UX Methods in the Data Lab: Arguing for Validity

Trinity C. Overmyer
Purdue University
West Lafayette, IN, USA
trinity@purdue.edu

ABSTRACT

Based on a case study at a national science lab, this paper outlines strategies for demonstrating the value of technical communication research when working with subject matter experts in both data and domain sciences. This paper argues that technical communicators can find common goals with experts in highly technical fields using UX methods as the medium of identification, particularly in the realm of scientific computing. This paper outlines a qualitative usability method created for data scientists at the lab to validate highly specialized scientific visualization applications.

KEYWORDS

technical communication, user experience design, user testing, scientific visualization, big data

ACM Reference Format:
Trinity C. Overmyer. 2019. UX Methods in the Data Lab: Arguing for Validity. In *SIGDOC '19: The 37th International Conference on the Design of Communication Proceedings, October 4-6, 2019, Portland, OR*. ACM, 6 pages. https://doi.org/10.1145/3328020.3353954

1 INTRODUCTION

Technical communicators often inhabit a liminal space between computational systems-centered STEM fields and human-centered goals. The boundaries between these two stances are unproductively constructed in ways that can put technical communicators epistemologically at odds with the high-tech environments in which they work, and vice versa. This tension plays out in very real ways when technical communicators work in data-centric fields where they must argue for the value of human-centered, qualitative methods among their quantitatively-driven colleagues [21].

In the midst of this apparent divide lies common ground. Data visualization experts and technical communicators are both ultimately concerned with meaning: how to make it, how to curate it, and how to communicate it effectively. Of course in systems-centered work—defined here as activity primarily concerned with the functionality of computational systems—meaning can take many forms. Semantics and coding languages, computational workflows, and concerns with how inputs are transformed into outputs are only a few ways that machines participate in meaning-making. Particularly with the rise of neural networks and deep learning, it can appear as if high-tech environments are moving further and further away from considering human meaning making in these computational processes.

One of the fundamental "big data" challenges deals with the growing gap between the amount of data able to be processed and the ability for humans to make meaning out of it [11]. Essentially, there is too much data and too little cognitive bandwidth of the experts who read it. This is particularly true in scientific work that requires multiple runs of petabytes of data and large-scale simulations. On the systems side, significant amounts of funding and expertise is poured into parallel computing [12] and *in situ* data processing strategies [See section 2.2, *In Situ Visualization*]. Both of these processes save computational time and power, but research in these areas has also begun to build strategies that consider human productivity and meaning [16] as part of the complex knowledge economy.

This gap between data and meaning making forces usability and visualization experts to forge new computational methods that process larger amounts of data in shorter time frames, and to design interfaces that attend to the ways humans make knowledge with data [1]. As user advocates [17] [18] and professionals poised to understand arrangement, invention, and discourse as epistemic [19] [13], technical communicators can find ways to enter data-centric areas of research and industry where such professionals are sorely needed. Because the end goal of scientific computation is always to create or uncover meaningful information, and such meaning is co-constructed by computers, data ecologies, and scientists, human users still remain central to the shared concerns of technical communication and scientific data work.

Over the course of ten weeks embedded in a team of data scientists and experts in computational visualization at a national science lab, a method was developed for conducting qualitative user testing with domain scientists to better understand fit and usability of custom-made scientific applications. In doing so, it became clear that technical communication research and qualitative UX methods could be insightful components of the team's development process, even in such quantitatively-driven fields.

This paper first outlines some cursory knowledge of large-scale scientific data visualization, and then goes on to discuss one particular research group and the method developed for eliciting qualitative user feedback. Additionally, this paper considers strategies for demonstrating the value of technical communication research when working with subject matter experts in both data and domain sciences by finding common goals, using UX methods as the medium of identification.

SIGDOC '19, October 04–06, 2019, Portland, OR

ACM ISBN 978-1-4503-6790-5/19/10. . . $15.00
https://doi.org/10.1145/3328020.3353954

2 DATA SCIENCE FOR TECHNICAL COMMUNICATORS

Data science is a burgeoning field that has not yet reached maturity. Academic programs are emerging in universities around the country while data buzzwords infiltrate all manner of business and industry job descriptions. In fact, in 2015, the National Research Council called for nation-wide educational initiatives that focus on digital curation and "meaningful use" of digital objects, specifically, quantitative big data [5]. While STEM programs may easily house the technological skills required for such data initiatives, technical communication as a field has an opportunity to explore what meaningful data practices might entail.

According to the National Science Foundation (NSF), data science is defined by its focus on "the processes and systems that enable the extraction of knowledge or insights from data" [2]. It is no surprise that a systems-centered approach is highlighted in their definition, considering data science's close coupling with computer science and statistics. Just as importantly, if not more so, is the NSF's focus on insight and knowledge making as a key component of the field. Data science is a combination of data mining and curation, computation, and visualization, in ways that help make data meaningful to a human viewer.

While big data has been critiqued for its tendency to flatten and decontextualize information, it is no more and no less than curated collections of granular information, which then must be read, arranged and interpreted with the help of computation. A data set, no matter the size, is an object of digital curation. An archive. It resists finitude and invites exploration and interpretation [23]. Data are not objects, but rhetorical processesâĂŤmutually constitutive relationships between bits of information, technology, communication, and human knowledge practices.

For example, early 20th century climate data was recorded manually but when scientists, data collection methods, and instrumentation changed over time, so did the ways researchers formatted and organized their data [7]. Historic data cannot be re-collected using contemporary methods, and such data does not get automatically updated in kind or format when new techniques emerge. The process of curation and formatting takes large amounts of time and human effort. The data used in much of big data work, but especially in climate science, is generated over long periods of time, by many different people with a wide range of training and expertise.

Our current knowledge of climate science is essentially built on historical traces, a process Edwards terms, "data friction" [7]. The work of constructing present or future knowledge from large or historic data consists of a lot of piece work that requires interpretation and various forms of inscription, whether it be in collecting, curating, or interpreting a dataset. Data are never raw—they are productions [8]. At each point in its life cycle, data is constructed through interpretation.

2.1 The Role of Visualization

Visualization is a key component for working with data, but especially when working with data that are too large to make sense of without computation. It is more than just a step in the sensemaking process when it comes to so-called big data. Visualization is the fulcrum upon which humans and computers collaborate to identify relationships and make insight.

There are clear implications for technical communication scholarship when it comes to visual rhetoric, infographics and data visualizations that are created to help audiences make meaning. Often, technical communication work focuses on public-facing visualizations and visual information for nonexpert users. However, the users for most scientific data visualizations are the scientists themselves, not an outside audience. Nonetheless, visualization formats and platforms have to be designed rhetorically, with an audience in mind, and scientific visualization applications still need to be designed with communication and meaning in mind, even if it is only for a single expert user.

2.2 In Situ Visualization

Large scale physics simulations and other visualizations require high performance computing (HPC), which is becoming more and more accessible for scientists with enough clout and grant money. However, even with access to the most powerful supercomputers in the world, massive scientific datasets can take months and hundreds of iterative runs between data input and visual output. While loss of time is a factor in scientific research, the cost to rent time on the machines, as well as the ecological impact of running and cooling such massive computers are additional bottlenecks that HPC communities are currently attempting to address [6].

For many experimental scientists, the most problematic issue is not waiting months for their data, as much as waiting for data only to find that there were malfunctions with experimental instrumentation or methods that rendered the data unusable [22]. In this instance, current post-processing methods cost more than compute time. For researchers who collect data in less accessible laboratories or environments, or for those whose experiments require considerable preparation time, unusable data output can set research back months or even years.

Currently, researchers are working to forgo post-processing and move toward in situ visualization, where vast amounts of data may be collected and rendered as a graphic almost instantaneously, rather than being filtered through months of post-processing. In situ methods are particularly exciting for experimental scientists whose data collection relies on short, intensive experiments that may cost thousands to run and months to plan. With growing computational capacities and the race to exascale, in situ data visualization could be a key way to address issues of environmental impact, the growing lack of data storage, and the overflow of data that does not have the human time and sensemaking capacity to match.

3 DESIGNING SCIENTIFIC APPLICATIONS FOR IN SITU VISUALIZATION

This paper is based on research conducted while embedded as a rhetorician at a national science laboratory over the course of two and a half months of ethnographic methods and one-on-one interviews collected with a group of 15 domain and computational scientists, statisticians, physicists and data visualization experts [14].

This interdisciplinary research team was created to develop an application for two materials scientists, in order for them to visualize their experimental data. Rather than acting as clients, the two scientific users co-led the team, alongside a data visualization researcher. Like most scientific software development, there was a clear exigence and use-case for the proposed application. Unlike simulation science, which employs physics principles to develop models of physical behavior, this visualization application was designed for use during time-constrained and very expensive experiments in a laboratory.

3.1 PowerPoint as Data Manager

Until this team was assembled, researchers' workflows operated along these lines: Schedule specialized lab for experiment (lab might only be free two days out of the year)

- Schedule specialized lab for experiment (lab might only be free 2 days out of the year)
- Plan and prepare experiment (6+ months)
- Run experiment (36 hours straight)
- Post-processing of data (2+ months)
- Data Analysis

The interface for the final data analysis step consisted of hundreds of individual images, arranged in a single PowerPoint file. A slide deck was the primary way scientists worked with and analyzed their data. Months of planning, hundreds of thousands of dollars, and countless hours of education and research all culminated in a PowerPoint file.

The experiments in question can only take place at one of three locations in the country. Because time in the facility was competitive and due to the massive cost of each experiment (in the hundreds of thousands of dollars), research plans were designed and revised meticulously for several months or a year prior to each run. Experiments lasted 36 hours nonstop, and the attending research team spent that all 36 of those hours on the lab floor. The researchers had to take a single shot with the lab equipment, collect measurements, recalibrate, and shoot again, in a constant loop for the entire 36 hours.

The problem with their workflow and the exigence for developing the new application was that they could not see or access any of the experimental data during the experiment to be sure that the equipment was calibrated properly and the shots were aligned correctly. In fact, because of the massive size of the data generated, researchers did not have access to this data for months after the experiment during computational data post-processing. They could not confirm that they designed the right shots or collected usable data during lab time, and therefore sometimes spent months and hundreds of thousands of dollars to rerun the experiment.

To replace the slide deck, the team worked to develop an in situ visualization workflow and attending interface, whereby scientists could see their data visualized in real time during experimentation, which meant that these researchers would also, for the first time, have the ability to intervene in the experiment. If the data was unclear, or instruments needed reset, researchers would be able to monitor the environment and make changes while they were still conducting the experiment, rather than 6-12 months down the road.

3.2 Qualitative Vs. Quantitative Usability

Data visualization scholarship and technical papers usually require some sort of usability testing, which often takes the form of quantitative online quizzes and tasks completed by a random labor pool. These usability tests are facilitated by services such as Amazon Mechanical Turk, where random participants take quizzes and receive micropayments for each user test, along the lines of $.01 per completed task. Such tasks may be geared toward testing graphical perception [9], cognitive activities [3], and other micro-functions in and around data visualization and representation.

Mechanical Turk is a widely accepted tool for studying cognitive tasks in visualization [9] and even psychology and sociology [15]. Hundreds and sometimes thousands of participants perform micro tasks associated with a particular data or visualization hypothesis, making Mechanical Turk a tool for oxymoronic large-scale micro user testing. Quantitative random task-based testing (along with computational remains the most accepted way to demonstrate validity in visualization scholarship.

Based on an informal analysis of scientific visualization scholarship, of the 47 papers published from the 2018 IEEE SciVis conference, 16 of them used some kind of qualitative process to validate their argument. Nearly a quarter of them (10) used direct feedback from the one or more scientists for whom they created the application. Although qualitative user-testing and validity arguments do exist in scientific visualization scholarship, quantitative random task-based testing (along with computational performance metrics) remains the most accepted way to demonstrate efficacy when designing a novel visualization application. Even though crowd-sourced inquiry such as this is becoming more accepted, it still leaves wide gaps in how developers might understand the more situated and complex human processes as they pertain to knowledge making with and through quantitative data. But professionals in high-tech environments often do not understand or appreciate the value of qualitative methods [21] for providing insight into quantitative work like data analysis and computer science.

Those who work directly with domain scientists are often unable to perform user testing at the same scale as Mechanical Turk studies, since the techniques and platforms being tested are often custom made for small research teams and their specific goals and data. Such was the case for the project team described above, which consisted of two target users. When we consider users' experiences, not as single tasks, but as “ecosystems of activity" [17], it becomes even clearer that random quantitative user testing, while convincing in some fields, cannot replicate the kind of information generated from working directly with end users to design an application from the concept stage to the end product.

4 A USABILITY METHOD FOR SPECIALIZED APPLICATIONS

Because the accepted quantitative, anonymous user testing frameworks would not work for an application specially designed for two users, qualitative user testing was considered. Even though some types of qualitative validity tests are accepted in science visualization scholarship (as noted above), including use cases and interviews with scientists, computer scientists and other technical experts often do not have methodological background in such methods and techniques. It can be difficult for these teams to spend the time conducting interviews, let alone learning qualitative techniques that can take years to cultivate.

Based on this professional context, a structured usability method was designed to be 1) accessible for technical staff, 2) low in time commitments beyond the application design process, 3) portable for numerous development projects, and 4) useful in understanding the real needs of end users. UX design frameworks hinge on active consideration and participation of actual users throughout the development process. While the technical staff generally struggled to get on board with qualitative measures, they were already fluent in designing systems with user needs in mind. They may not have had an understanding of possible procedures, but they understood the value of working closely with scientists and attending to specific workflows. Scientific software and applications are highly specialized and often their development is funded as a component of a larger scientific project. Therefore, having a more structured process that visualization experts could follow to engage their users more systematically was useful, particularly when publishing on these novel applications and positioning user feedback to validate their designs.

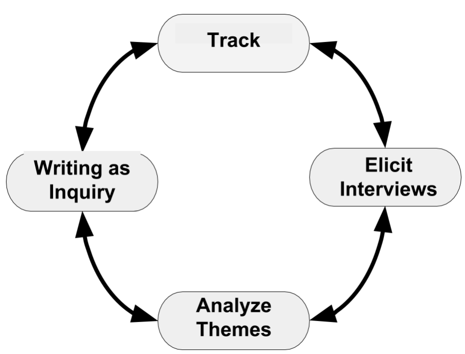

Figure 1: Basic structure for qualitative user elicitation

4.1 Track, Elicit, Analyze, Write

The method developed for this project consists of four processes, which can and should occur in tandem throughout the design and assessment process:

- Track challenges, changes and constraints throughout the process
- Elicit interviews
- Analyze themes
- Write

Though the method outlined below may be seem straightforward or overly simplified for qualitative researchers, the process was developed specifically to attend to the constraints and needs of scientific data visualization researchers unfamiliar with these techniques.

4.1.1 Track. Making note of users' domain challenges and visualization constraints, as well as logging notes from all communications throughout the design process allows visualization experts to implement and track primary users' feedback without adding excessively to preexisting workloads.

Characterize the problem. Clearly note domain characterization of the problem and how it may translate from existing domains or applications. Additionally, track growth and any changes between early conceptions of problems and solutions, and how they shift throughout the process.

Workflow Considerations. Make note of decisions and feedback that occur throughout the process in communications and meetings, rather than relying on post-production interviews and write ups. This list will guide final analysis and allow the gathering of user data iteratively, laying the framework for a better application.

4.1.2 Elicit Interviews. When possible, interviews with users should be conducted, in addition to tracking the collaboration throughout the project. The goal of interviews is to let the user speak at length about the way applications are used and useful in their work, and to talk about the problems it addressed and the challenges that should be considered in the future.

Workflow Considerations. You may conduct short interviews iteratively throughout, rather than long interviews at the end of the process. Interviewees should spend time with the application open during the interview to spark feedback and demonstrate their assessment.

4.1.3 Analyze Themes. Analyzing qualitative data can be, in its simplest form, identifying themes in the data that provide insight into user needs and application benefits. The goal of this step is twofold: 1) to identify areas of interest early in the collaboration, and 2) to begin crafting a narrative and evidence of the application's efficacy.

Workflow Considerations. Bring all tracking notes and interviews together and begin to identify themes, but do not limit themes to what is expected, and do not discount outliers. Use themes as headings in a document to organize your user data. These sections will form the basis for the narrative in your eventual write-up.

4.1.4 Writing as Inquiry. Writing helps the writer make cognitive connections and form understanding in ways that cannot be accomplished otherwise, and should be done iteratively, just like research and design.

Workflow Considerations. Collaborative writing and iterative drafting should not be left as the final step, but should begin early in the process and honed throughout. The text is not a timeline of events. The goal is to create a cohesive narrative using the thematized data, so the order in which themes are discussed and the connections between each theme is crucial to the audience's understanding of the application.

This method acts as a basic guide for data visualization researchers to construct qualitative assessment of their designs, based on their real, primary users. Additionally, this outline hinges on values shared by both technical communication and data visualization, such as the value of user input and expertise, the goals of sensemaking, and creating and identifying opportunities for action.

5 DISCUSSION

Because the research team was led by the primary users, they ended up employing codesign and iterative UXD techniques, even though that was not their goal. The users chose their team of experts and engaged the collaborators in months of a "structured learning phase," [14] where team members took turns educating the others on each of their own expertise and disciplinary perspectives as it related to the task at hand. Structured learning was followed by a long span of weekly all-hands meetings and small working groups. It was in these regular meetings that the majority of the team members saw the "real work" of the project being done. The team's knowledge was forged from incremental articulation work of codesigning, which is the recalibration work that gets things back on track [4]. Sensemaking and the codesign of platforms for data insights are both kinds of articulation work that require a process of continually wading in and then adjusting course.

Likewise, the in situ interface was designed purposefully for data sensemaking to have the same possibilities for exploration and adjustment. By creating a digital space for researchers to move data components around and connect them freely, the in situ platform lets the researcher make meaning from more complex and relational data. It provides a paratactic, non-linear environment where the researcher's process of arrangement produces narratives that aid sensemaking for the researchers themselves and their eventual audiences. Data visualization has taken an acute interest in narrative, both in terms of public audiences and understanding, and because of the power of narrative for data workers and researchers to understand their own data [20] [10].

As SciVis researchers develop platforms and techniques, their goal is to help scientists explore and understand their data and its potential. Novel visualization platforms do not rely on rote procedural steps. Instead, interfaces are created that emphasize the possibilities in data. They are built around interfaces that allow multiple views at once, and the ability to interactively rearrange data to find patterns and ask questions. Whether they use our terminology or not, scientific visualization is heavily invested in invention and arrangement in their research and platform development.

And yet, one of the drawbacks of the proposed method is precisely that it is so procedural, and could easily be deployed too rigidly or in contexts that do not lend themselves to these methodologies. When practitioners learn techniques but are unfamiliar with the methodological concerns involved in them, fissures between intent and efficacy can easily appear. However, until technical environments begin to more readily see the value that qualitative researchers and technical communicators can add, simplified methods such as this one can help computational researchers begin to practice more effective user-engaged design. As user experience design concepts and techniques filter in and become more accepted in the scholarship of highly technical fields, there will continue to be professional openings for technical communicators who have expertise in these areas.

Identifying shared values and concerns between data science and technical communication is a key tool for obtaining and maintaining a critical place in interdisciplinary research, which can be a vehicle for more diverse technical communication research. This case and the development of the method has uncovered a few key shared values between the two fields that might hold promise for future collaboration.

Shared values of data science and technical communication include:

- Valuing user needs and constraints, while also seeing past what the user asks for to what the user needs.

- Bridging how information is communicated and able to be arranged across diverse audiences and expertise levels.

- Using narrative as a sense making tool, which provides different types of insight and uncovers relationships among possible arguments and understandings [20].

- Facilitating users' abilities to investigate, explore and understand data, in order to use it to take action.

When technical communicators enter high tech fields or new workplaces, we expect, at best, some amount of explaining what our expertise consists of beyond writing document, or at worst, arguing that our expertise has value. This does mean, however, that technical communicators have the opportunity to make new arguments and to redefine the boundaries of our work in each new space. Even in high tech, quantitatively-driven workplaces, our human centered focus—the insights we can gain from concentrating on the human in technological processes—can be our biggest asset.

Data science is ripe for new pathways to reconfigure user experience design concepts. Often, data visualization experts do not have outside domain science expertise, therefore, they have to work closely with their users to develop successful processes that lead to the specific kinds of insight needed in each field. The field of scientific visualization is particularly open to insights about usability and user testing, because their scholarship requires it, and yet very few practitioners have education or training related to user studies. While usability can sometimes be relegated to the background work in high tech industries, in data science and visualization, it is considered a necessary part of valid research. By demonstrating in both professional practice and academic work that user experience design and knowledge making do fall within the purview of professional and technical communication, we can create openings for trained technical communicators to explore fruitful research alongside data-centered science.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Ben Simms. This work was partially completed at Los Alamos National Laboratory, supported by the U.S. Department of Energy. The publication has been assigned the LANL identifier LA-UR-19-26829.

REFERENCES

[1] Chris Argenta, Jordan Benson, Nathan Bos, Susannah BF Paletz, William Pike, and Aaron Wilson. 2014. Sensemaking in Big Data Environments. In *Proceedings of the 2014 Workshop on Human Centered Big Data Research*. ACM, 53.

[2] Francine Berman, Rob Rutenbar, Brent Hailpern, Henrik Christensen, Susan Davidson, Deborah Estrin, Michael Franklin, Margaret Martonosi, Padma Raghavan, Victoria Stodden, and Alexander S. Szalay. 2018. Realizing the Potential of Data Science. *Commun. ACM* 61, 4 (March 2018), 67–72. https://doi.org/10.1145/3188721

[3] Michelle A Borkin, Azalea A Vo, Zoya Bylinskii, Phillip Isola, Shashank Sunkavalli, Aude Oliva, and Hanspeter Pfister. 2013. What makes a visualization memorable? *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2306–2315.

[4] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.

[5] National Research Council et al. 2015. *Preparing the workforce for digital curation*. National Academies Press.

[6] Jack Dongarra, Steven Gottlieb, and William TC Kramer. 2019. Race to Exascale. *Computing in Science & Engineering* 21, 1 (2019), 4–5.

[7] Paul N Edwards. 2010. *A vast machine: Computer models, climate data, and the politics of global warming*. Mit Press.

[8] Lisa Gitelman. 2013. *Raw data is an oxymoron*. MIT press.

[9] Jeffrey Heer and Michael Bostock. 2010. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 203–212.

[10] Jessica Hullman and Nick Diakopoulos. 2011. Visualization rhetoric: Framing effects in narrative visualization. *IEEE transactions on visualization and computer graphics* 17, 12 (2011), 2231–2240.

[11] Rob Kitchin. 2014. Big Data, new epistemologies and paradigm shifts. *Big data & society* 1, 1 (2014), 2053951714528481.

[12] Vipin Kumar. 2002. *Introduction to parallel computing*. Addison-Wesley Longman Publishing Co., Inc.

[13] Janice M Lauer and Kelly Pender. 2004. *Invention in rhetoric and composition*. Parlor Press LLC.

[14] Trinity Overmyer. 2018. Analysis and Strategies of a Successful LDRD Collaboration. *Los Alamos National Laboratory* LA-UR-18-29988 (2018).

[15] Gabriele Paolacci and Jesse Chandler. 2014. Inside the Turk: Understanding Mechanical Turk as a participant pool. *Current Directions in Psychological Science* 23, 3 (2014), 184–188.

[16] John M Patchett, Boonthanome Nouanesengsy, James Paul Ahrens, Michael Kenneth Lang, David Honegger Rogers, Jennifer Kathleen Green, Francesca Samsel, Giovanni Antonio Cone, and Hans-Jurgen Hagen. [n. d.]. Delivery of In Situ Capability to End Users. ([n. d.]).

[17] Liza Potts and Michael J Salvo. 2017. *Rhetoric and Experience Architecture*. Parlor Press.

[18] Michael J Salvo. 2001. Ethics of engagement: User-centered design and rhetorical methodology. *Technical communication quarterly* 10, 3 (2001), 273–290.

[19] Robert L Scott. 1967. On viewing rhetoric as epistemic. *Communication Studies* 18, 1 (1967), 9–17.

[20] Edward Segel and Jeffrey Heer. 2010. Narrative visualization: Telling stories with data. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 1139–1148.

[21] Patricia Sullivan and Rachel Spilka. 1992. Qualitative research in technical communication: Issues of value, identity, and use. *Technical Communication* (1992), 592–606.

[22] Sven C. Vogel, Chris M. Biwer, David H. Rogers, James P. Ahrens, Robert E. Hackenberg, Drew Onken, and Jianzhong Zhang. 2018. Interactive visualization of multi-data-set Rietveld analyses using Cinema:Debye-Scherrer. *Journal of Applied Crystallography* 51, 3 (May 2018), 943–951. https://doi.org/10.1107/s1600576718003989

[23] Susan Wells. 2002. Claiming the archive for rhetoric and composition. *Rhetoric and composition as intellectual work* (2002), 55–64.