



UvA-DARE (Digital Academic Repository)

A Philosophical Perspective on Visualization for Digital Humanities

van den Berg, H.; Betti, A.; Castermans, T.; Koopman, R.; Speckmann, B.; Verbeek, K.; van der Werf, T.; Wang, S.; Westenberg, M.A.

Publication date

2018

Document Version

Final published version

Published in

3rd Workshop on Visualization for the Digital Humanities

[Link to publication](#)

Citation for published version (APA):

van den Berg, H., Betti, A., Castermans, T., Koopman, R., Speckmann, B., Verbeek, K., van der Werf, T., Wang, S., & Westenberg, M. A. (2018). A Philosophical Perspective on Visualization for Digital Humanities. In *3rd Workshop on Visualization for the Digital Humanities: Sunday, 21 October 2018 - Berlin, Germany : accepted papers VIS4DH*. <http://vis4dh.dbvis.de/papers/2018/A%20Philosophical%20Perspective%20on%20Visualization%20for%20Digital%20Humanities.pdf>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

A Philosophical Perspective on Visualization for Digital Humanities

Hein van den Berg* Arianna Betti* Thom Castermans† Rob Koopman‡ Bettina Speckmann†
Kevin Verbeek† Titia van der Werf‡ Shenghui Wang‡ Michel A. Westenberg†

ABSTRACT

In this position paper, we describe a number of methodological and philosophical challenges that arose within our interdisciplinary Digital Humanities project CATVIS, which is a collaboration between applied geometric algorithms and visualization researchers, data scientists working at OCLC, and philosophers who have a strong interest in the methodological foundations of visualization research. The challenges we describe concern aspects of one single epistemic need: that of methodologically securing (an increase in) trust in visualizations. We discuss the lack of ground truths in the (digital) humanities and argue that trust in visualizations requires that we evaluate visualizations on the basis of ground truths that humanities scholars themselves create. We further argue that trust in visualizations requires that a visualization provides provable guarantees on the faithfulness of the visual representation and that we must clearly communicate to the users which part of the visualization can be trusted and how much. Finally, we discuss transparency and accessibility in visualization research and provide measures for securing transparency and accessibility.

1 INTRODUCTION

In this position paper, we describe methodological and philosophical challenges that we have encountered within our interdisciplinary research project CATVIS¹. CATVIS is a collaborative effort between computer scientists, specifically researchers working on applied geometric algorithms and visualization; data scientists working at OCLC, a company that hosts *Worldcat*, the world’s largest database with information on library collections; and philosophers with a strong interest in the methodological foundations of computational approaches and visualization research. The goal of CATVIS is to provide visual tools and methods that allow librarians to manage and understand hundreds of millions of bibliographic records [3], and to develop visual tools and methods that aid philosophers in their research. More specifically, we aim to develop a cutting-edge visual analytics toolkit, to answer both the pressing needs of humanities researchers and concrete demands of the library industry. Our tools are intended to provide visual interfaces for: (1) data cleaning, clustering, and enrichment, (2) data analysis, and (3) intuitive and interactive (geographic) representation of search results [2]. Within CATVIS, we have had frequent discussions on philosophical and methodological aspects of visualization research. This paper provides an insight into these discussions.

In the three sections following we consider three sets of challenges that we were confronted with, and that we also take to be generalizable to Digital Humanities projects similar to ours. In Section 2 we discuss the need for proper evaluation, and offer a proposal for sound benchmarking of computational tools via ground truths

within (digital) humanities research. In Section 3 we argue for the necessity to provide provable guarantees on the faithfulness of the visual representation of data. Finally, in Section 4, we comment on the dual requirements of transparency and accessibility of the computational tools employed in Digital Humanities projects.

The challenges just mentioned are related to three issues in Digital Humanities projects as described by Rieder and Röhle [19], namely (i) the lure of objectivity, (ii) the power of visual evidence, and (iii) black-boxing. We take Rieder and Röhle’s discussion of each of these issues into account in the following sections. Here we wish to note that, in our view, our three sets of challenges are in fact three aspects of one single epistemic need: that of methodologically securing (an increase in) *trust*.

The concept of trust has recently been subject to intense philosophical reflection and has also been applied to contexts that deal with information and communication technologies (dealing with the concept of *e-trust* [22]). Trust is usually defined as a relation between people. Thus, interpersonal trust is taken to be the dominant paradigm of trust [18]. Trust is defined as “an attitude that we have towards people whom we hope will be trustworthy, where trustworthiness is a property, not an attitude” [18]. A condition for trustworthiness is that the trustworthy person “is competent and committed to do what s/he is trusted to do” [18]. We believe that this condition can also be applied to visualizations. We claim that the increase in methodological soundness afforded by reliable evaluations based on ground truths that humanities scholars themselves create (Section 2), representational accuracy provided by clear and robust algorithms (Section 3) and transparency and accessibility of the tools (Section 4) are together necessary conditions for methodological trustworthiness within Digital Humanities projects.

Related Work. There have been prior discussions of Digital Humanities projects that partly touch upon the problems we discuss. These discussions typically provide a fresh perspective on the value of visualization for the Digital Humanities. Hinrichs and Forlini [13] claim that visualization should not be seen merely as a means to an end, but constitutes a research process in its own right. They claim that visualizations provide new perspectives on data and new modes of knowledge production. Abdul-Rahman et al. [1] similarly stress that visualizations provide new modes of knowledge production. In an insightful article on poetry visualization, they show visualizations that “help literary scholars to make observations more effectively, to stimulate different interpretations, and to visually evaluate interpretations” [1]. Jänicke [14] claims that visualizations typically provide a new perspective on cultural heritage data, whereas Jänicke et al. [15] claim that a visualization should be able to “(1) confirm existing hypotheses, (2) refine humanities scholars’ research questions, (3) offer new ways of answering research questions, (4) negotiate quantitative and qualitative interpretation of the underlying text corpus, and (5) trigger new research questions”. Finally, Coles [10] argues that the Digital Humanities must incorporate Humanities-based values and methods, whereas Hinrichs et al. [12] discuss several case studies of Digital Humanities projects where visualizations were used as tools to enrich humanities scholars research practices, or as tools for answering research questions. They emphasize that research styles and thinking can change due to the collaboration

*University of Amsterdam
[h.vandenberg | a.betti]@uva.nl

†TU Eindhoven
[t.h.a.castermans | b.speckmann | k.a.b.verbeek | m.a.westenberg]@tue.nl

‡OCLC Research
[rob.koopman | titia.vanderwerf | shenghui.wang]@oclc.org

¹<https://www.win.tue.nl/CatVis/>. Supported by the Netherlands Organization for Scientific Research (NWO) under project no. 314.99.117.

between computer scientists and humanities scholars.

We agree with the authors just mentioned that visualizations provide new modes of knowledge production. Thus, we do not deny that visualizations, for example, enhance existing research practices, provide new perspectives on data, refine research questions, and provide new research questions. These are all beneficial aspects of visualizations. In this paper we argue, however, that in order to secure trust in visualizations, visualizations must be able to represent research evidence in a reliable way – first of all by (aiding in) replicating known hypotheses or interpretations (including negative ones). This does not mean that visualizations, or computational tools in general, are only beneficial if they show the expected. Visualizations are beneficial in the many ways described by the authors cited above. However, our concern in this paper is trust. If visualizations cannot be trusted, that is, if they cannot be reliably used as (faithfully representing) evidence, they can only be beneficial by accident.

2 BENCHMARKING AND GROUND TRUTHS

Rieder and Röhle note that humanities scholars might become interested in computers under the belief that computers can help attain results with a higher epistemological status, that is, objective and unbiased results, of the kind produced in the mathematical and natural sciences (“the lure of objectivity”). Rieder and Röhle argue that this belief is fallacious: “questions of bias and subjectivity, which the computer was thought to do away with, enter anew on a less tangible plane – via specific modes of formalisation, the choice of algorithmic procedures, and means of presenting results” [19, p. 73]. We generally agree on this last caveat, but we unpack and address the worry along a somewhat different route. There are at least two ways to talk about “more objective” results. We think that computational procedures can be said to yield “more objective” results if certain conditions are met, namely when all parties involved in the research are clear on what is exactly computed, in the presence of a robust assessment of error, of transparent and accessible algorithmic procedures, and of sound evaluation methods: in short, when the three sets of challenges we address are met in an adequate way.

As to the other sense of “more objective”, namely the sense in which one might attempt to make the very object of humanities research free from bias and subjectivity, however, we see no way to attain objectivity. The very result of this kind of research is interpretation, a reading of the output data according to a certain perspective. The only way to deal appropriately with bias and subjectivity in this sense, we say, is to make the perspective in question – namely the one behind the interpretation – as explicit as possible. In our opinion, the need to make interpretive bias and subjectivity explicit is connected to the issue of providing ground truths in the humanities, in the way we explain next.

A frequently recurring issue within our CATVIS project is the lack of ground truths in the (digital) humanities that can be used to benchmark or evaluate data processing output and its visualization. This lack is unfortunate, because only if we can properly evaluate the reliability of the tools employed can we properly trust those very tools. In one of our projects, we aimed to provide a proper visual representation of the output of the ARIADNE tool developed by OCLC Research [16, 17]. ARIADNE embeds natural language terms, bibliographic entities and records in a multi-dimensional semantic space in which their distances (representing similarities) could be computed. Within this space, ARIADNE identifies, for a given *focus entity*, a set of meaningful closest neighbors. We developed the tool SOLARVIEW [8] which enables users to visually explore this space.

To evaluate SOLARVIEW, we decided to let expert users evaluate the visualization of a philosophy dataset of (terms standing for) entities extracted from over 1.7 million bibliographic records of articles from about 300 philosophy journals. The evaluation was conducted by the senior philosophers in our team, who have intimate knowledge of philosophy as a discipline. The evaluators did not

have access to any ground truth or golden standards for the dataset in question, however, nor was a specific research question formulated to guide the exploration. The philosophers did not know the data well enough, nor, given the size, could they in fact know it well enough from traditional research (they would be in the – highly unrealistic – latter situation if they had studied at least a representative part of 1.7 million articles). This resulted in problems. The philosophers did not and could not know whether abnormal configurations in the visualization were due to a bug or due to an issue with the data – that is, benchmarking SOLARVIEW was impossible.

In particular, it was impossible to distinguish whether (i) there was a bug in the ARIADNE tool code, or (ii) ARIADNE operated on incomplete, dishomogeneous, or corrupted data, or (iii) the results of ARIADNE were trustworthy but the visualization was not for any number of reasons. For example, when searching for philosophical topics, very often the same author, Dale Jacquette (1953-2015), would appear as a neighboring entity. Dale Jacquette was a prolific contemporary author on many different philosophical topics. However, one of our philosophers thought it was strange that the visualization of ARIADNE would display Jacquette’s name so often. She thought the result was probably due to a bug or any of the reasons quoted above having to do with the working of ARIADNE. One of our other philosophers tried to make sense of the result and thought that maybe the results did faithfully represent the data. In the end, we could not decide whether the result made sense or not.

The lack of ground truths in the (digital) humanities is a known problem. In text-based humanities, scholars often deal with conflicting interpretations of a text, a circumstance that makes it difficult to settle on ground truths. For instance, philosophers, especially historians of philosophy, often write long papers on one particular passage of a philosophical work or on one philosophical concept employed in the work of an author. These papers engage with multiple conflicting interpretations of a passage or concept and deal with multiple shades of meaning, making the creation of ground truths for particular interpretations of passages or concepts a non-trivial task. The problem transfers immediately to the Digital Humanities, for the very notion of ground truth requires a manually curated reliable standard against which to measure the results of a computational procedure: if there are no such suitable ground truths in the humanities, how can we ensure reliability of Digital Humanities tools?

We think that the problem can be obviated by requiring that the humanities scholars performing evaluations are always asked to do an initial evaluation of the tools on data (i) they know well, and (ii) for which they themselves are able to create ground truths. For philosophy in particular, ground truths need to be created that concern particular philosophical interpretations of a concept or position. How can we do this? We claim that (conceptual) *models* (more precisely their concrete realizations in texts), as defined by Betti and van den Berg [4, 5], can function as ground truths, as long as they receive previous preliminary testing in traditional research. Betti and van den Berg’s models are abstract constructs set up within the humanities, that model concepts as relational complexes with stable elements, representing conceptual continuities, and variable elements, representing conceptual discontinuities. Differences between two concepts are then e.g. modeled as (dis)similarities between the (stable and variable) elements of two different models [6]. An example of a model in our sense is the Classical Model of Science, which codifies an influential traditional ideal of science accepted by many past philosophers. According to this model, a science is a system S satisfying the following conditions:

- (1) All propositions and all concepts (or terms) of S concern a specific set of objects or are about a certain domain of being(s).
- (2a) S contains so-called *fundamental concepts* (or terms).
- (2b) All other concepts (or terms) occurring in S are *composed of* (or are *definable from*) these fundamental concepts (or terms).
- (3a) There are in S a number of so-called *fundamental propositions*.

- (3b) All other propositions of S follow from or are grounded in or provable / demonstrable from these fundamental propositions.
- (4) All propositions of S are true.
- (5) All propositions of S are *universal* and *necessary* in a sense.
- (6) All propositions of S are known to be true. A non-fundamental proposition is known to be true through its *proof* in S .
- (7) All concepts or terms of S are *adequately known*. A non-fundamental concept is adequately known through its composition (or definition) [11].

Models such as the Classical Model of Science can be used to guide traditional qualitative close reading and conceptual analysis, but can also be fruitfully used for investigations that focus mainly on quantitative large-scale pattern finding [6, 20]. Models can also be implemented computationally as *ontologies* that can guide the computational exploration of texts [5, 6, 23]. By setting up the Classical Model of Science, we have reduced thousands of complex and conceptually dense pages of texts to a manageable and informative model. We have gathered detailed information from traditional research on the fact that certain texts or philosophers accept the concept of axiomatic science fixed by the Classical Model of Science: we know the extent of that acceptance, and the subtle variations in which it presents itself. This information – the model itself and the way it is embodied (or not embodied) in texts across time – can thus be taken as a ground truth for the benchmarking of Natural Language Processing (NLP) tools and visualizations. The idea is that the computational tools should point towards an interpretation of a philosophical text that conforms (or does not conform) to our interpretative model of such a text [5, 6]. If the computational methods match the outcome (or model) we expect, we know that they are reliable in aiding Digital Humanities enterprises.

It may be objected that models in our sense are problematic because they incorporate scholars’ interpretative biases. However, what is distinctive of models in this sense is that they make biases fully explicit and revisable [4]. In our opinion, this is the best and only way to counter interpretative biases in text-based humanities. The fact that models make interpretative assumptions fully explicit and revisable distinguishes our approach from much traditional work, which tends to use implicit interpretative assumptions. The advantage of using models is thus that philosophers and other text-based humanities researchers have a clear, concise and transparent articulation of a self-made interpretative framework that they can use to properly evaluate computational methods in the Digital Humanities. Such proper evaluations further the trustworthiness of the scientific results obtained by computational methods, and the reliability of these computational methods themselves.

The method for evaluation sketched above involving ground truths in the sense explained implies that we first apply computational tools and methods to rather small, human-sized corpora that the domain experts are intimately familiar with, and in order to reply to research questions the humanities scholars know the answer to. Once we have tested our computational tools on small corpora that are well-known by the experts, on research questions with known answers, and have thus established that the tools are reliable, we can apply our computational tools to larger corpora that domain experts are less familiar with, and to unanswered research questions.

A downside to this approach is that it works only for cases in which top-down interpretative knowledge from the field is available. If such knowledge is unavailable, or if we want to test computational tools without a presupposed layer of interpretation in the form of a model, we should always start by either (i) acquiring top-down interpretative knowledge, or else (ii) limiting the research to tools that give provable guarantees (Section 3), and are transparent and accessible (Section 4). Cases like (ii) might for instance be cases in which an NLP step on a large-size corpus is necessary, or any case of big data manipulation.

3 PROVABLE GUARANTEES

Why should provable guarantees matter to users of visualizations? We think that guarantees are a necessary condition for users to trust visualizations. Rieder and Röhle note that in general visualization has a schizophrenic nature: images are taken to provide powerful argumentative evidence despite presenting condensed or reduced information. For example, many visualizations have to reduce informational dimensions. These reductions introduce a focus that “renders certain interpretations and explanations more plausible than others” [19, p. 74], which should make us cautious to unreflectively assign too much evidential weight to visualizations. Visualizations must hence prove themselves worthy of trust, and visualizations are trustworthy (competent and committed to what they are trusted to do) if they provide a faithful representation of the data. Now, if a visualization provides provable guarantees on the faithfulness of the visual representation, then the user has a *reason* for trusting the visualization. Hence, provable guarantees are important for users, even if users are typically not aware of this fact.

In the context of computer science research there are many types of provable guarantees. Typically computer science researchers are interested in guarantees on the running time of an algorithm or on the quality of the output. More specifically, computer scientists are likely to care about the optimality of the solution, that is, how optimal is the encoding of the data in a specific computed visualization with respect to the chosen design space. However, these are not the type of provable guarantees users would typically worry about and neither do they need to. The biggest, and in some sense only, concern of the users is that what they see in the visualization accurately reflects the data. For a visualization to support evidence-based humanities research, the degree to which a visualization reflects the underlying data needs to be quantifiable. Therefore, the visualization needs to be equipped with provable guarantees on the faithfulness of the visual representation.

Digital Humanities research often involves high dimensional data sets. Visual representations are only useful if they reduce these high-dimensional data spaces to a more comprehensible abstraction. Any meaningful abstraction will, by design, make it impossible to faithfully represent all data and the resulting visualization will give only a distorted view. The same limitations apply in general to all types of complex (and not necessarily high-dimensional) data. However, these inherent limitations are not a reason per se to abandon all attempts at faithful visualization representations with guarantees. Specifically, in the context of evidence-based humanities research, it is necessary to identify as large a subset as possible of the salient information to represent faithfully. This calls for more research in abstraction techniques which provide partial provable guarantees. At the same time, it is of paramount importance to clearly communicate to the users which part of the visualization can be trusted and how much. Visual encodings should offer users visual cues that allow them to become aware of distortions in the visual encoding of data. Such visual cues (see [7–9, 21]) thus provide users with a reason to properly assign evidential weight to a visualization and hence, ultimately, trust the visualization.

The philosophers in our team have in the past used visualizations that did not provide provable guarantees. In a paper on applying computational tools to the works of the philosopher and mathematician Bernard Bolzano (1781-1848), the philosophers used a Gephi visualization (see Figure 1) of the similarity between paragraphs of Bolzano’s main work, the *Wissenschaftslehre* [23]. The visualization shows a cluster of paragraphs that are about the concept of *analyticity*, which was the focus of the investigation. The philosophers knew the data well and expected the paragraphs about analyticity to be related to each other (since they are similar). However, the visualization indicates that these paragraphs are not only similar to each other but at the same time very different from the remaining paragraphs. This is doubtful, since the analyticity paragraphs are in




Figure 1: A Gephi visualization of paragraph similarity in Bolzano's *Wissenschaftslehre*. Top left: a close-up of the top right, focusing on the analyticity paragraphs.

fact very similar to other paragraphs. There can be multiple causes for this discrepancy between knowledge and representation. Either the distance metric did not capture similarity between paragraphs as expected, or the visualization did not represent the salient distances faithfully. In the absence of provable guarantees on the visualization, it is impossible to determine the underlying cause of the discrepancy. As a result, the visualization as a whole cannot be trusted, effectively negating any positive effect of using these computational tools as reliable methodological devices in the philosopher's research. Note that the visualization can still be said to benefit the paper, as it stresses visually a known fact (that certain paragraphs are saliently similar): the point is however that it does so by accident. As a faithful representation of evidence for phenomena in the data, the visualization is a form of cherry picking, e.g. (fallacious) confirmation bias: the reader is supposed to take as significant one particular aspect that confirms previous knowledge (the similarity between a group of paragraphs), and ignore other aspects contradicting that knowledge (the dissimilarity between that group of paragraphs and all other paragraphs).

4 ACCESSIBILITY AND TRANSPARENCY

Accessibility and transparency of computational tools are necessary conditions for methodological trustworthiness. On the topic of accessibility Rieder and Röhle say "possibilities for critique and scrutiny are related to technological skill: even if specifications and source code are accessible, who can actually make sense of them?" [19, p. 76]. This lack of true accessibility is a common problem in our work, especially if the tools and methods developed for humanities researchers are developed by computational experts. The computational experts have the algorithmic and mathematical skills to develop – and hence also to understand – the algorithms that are applied to humanities data. However, sometimes the humanities experts lack the background and do not have the necessary technological skills to fully understand these same algorithms.

The only solution we see to this problem is to let humanities experts work closely together with the computational experts. Humanities scholars should become as literate as possible when it comes to understanding computational tools, while conversely computational experts should try to understand the domain problems with which humanities scholars are concerned; ideally, all tools should be co-developed in full. This means that the computational experts must be able to clearly explain how certain technologies

work. They should also have a deep awareness of the humanities data to which these technologies are applied. However, it also means that humanities scholars must become technologically literate. They should study the (mathematical) methods that are applied and ideally should also develop an understanding of programming. Only in this way can humanities scholars develop a deep understanding of the methods that are employed in their research.

To further the technological literacy of the humanities scholars, the philosophers in our team sometimes try to write the technical and methodological sections of Digital Humanities publications. This ensures that they fully understand the methods employed. Furthermore, this practice also helps to explain technologies to a humanities audience that is not familiar with the computational tools used.

Rieder and Röhle note that tools and methods used in the Digital Humanities are often not transparent. Here transparency is defined as "our ability to understand the method, to see how it works, which assumptions it is built on, to reproduce it, and to criticize it" [19, p. 75]. They state further that "some of the approaches computer science provides us with are positively experimental, in the sense that the results they produce cannot be easily mapped back to the algorithms and the data they process" [19, p. 76]. For example, machine learning often produces outputs that are unanticipated and that are very difficult for the user to reconnect to the inputs (ibid.).

This lack of transparency is intrinsically difficult to overcome. In our project we have found two different approaches to be helpful. First of all, we test our tools on small data sets which our expert users know very well. This allows users to accurately evaluate how the output of a tool relates to the input provided to the tool. If this relation is clear and the visual representation is deemed to be faithful to the data, then we have more reason to trust the tool when it is applied to bigger datasets which the users know less well. A second, less common, approach is the use of synthetic data. Synthetic data can be constructed in a variety of ways and provides reliable ground truths against which the faithfulness of a given visual representation can be evaluated. It is important that such synthetic data is not biased towards certain representations and that it allows the user to examine all aspects of a visual encoding equally well. As such this calls for further research in the construction of suitable synthetic data sets for visualization in Digital Humanities. At the current time qualitative evaluations based on synthetic data are not fully accepted in the visualization community. We argue that experiments with well-designed synthetic data are an important tool to make the inner workings of algorithms more transparent and to ultimately increase trust in visualizations.

5 CONCLUSION

In this position paper, we have described a number of methodological and philosophical challenges that arose within our interdisciplinary Digital Humanities project CATVIS. The three challenges we have described concern the need for proper evaluation of computational tools on the basis of ground truths, the need to provide provable guarantees that the data are faithfully represented by visualizations, and the need for transparency and accessibility of the computational tools employed in Digital Humanities projects. These challenges must be met if users are to trust visualizations. In order to provide provable guarantees that data is faithfully visualized, we require the identification of a large subset of salient information which must be represented faithfully. This requires more research into abstraction techniques which provide provable guarantees. To further transparency for visualization in the Digital Humanities we have argued that we must evaluate visualizations on the basis of synthetic data, even if the use of synthetic data is not yet fully accepted in the visualization community. This proposal calls for further research in the construction of suitable synthetic data sets for visualizations.

REFERENCES

- [1] A. Abdul-Rahman, J. Lein, K. Coles, E. Maguire, M. Meyer, M. Wynne, C. R. Johnson, A. Trefethen, and M. Chen. Rule-based visual mappings—with a case study on poetry visualization. In *Computer Graphics Forum*, vol. 32, pp. 381–390, 2013.
- [2] A. Betti, T. Castermans, B. Speckmann, H. van den Berg, and K. Verbeek. GlamMapping Trove. In *Proceedings of the 18th VALA*, 2016.
- [3] A. Betti, D. Gerrits, B. Speckmann, and H. van den Berg. GlamMap: visualising library metadata. *Proceedings of the 17th VALA*, 2014.
- [4] A. Betti and H. van den Berg. Modelling the history of ideas. *British Journal for the History of Philosophy*, 22(4):812–835, 2014.
- [5] A. Betti and H. van den Berg. Towards a computational history of ideas. In *CEUR workshop proceedings*, vol. 1681. 2016.
- [6] A. Betti, H. van den Berg, Y. Oortwijn, and C. Treijtel. History of Philosophy in Ones and Zeros. In *Methodological Advances in Experimental Philosophy*, Advances in Experimental Philosophy. Bloomsbury, 2018, forthcoming.
- [7] M. Brehmer, M. Sedlmair, S. Ingram, and T. Munzner. Visualizing dimensionally-reduced data: interviews with analysts and a characterization of task sequences. In *Proceedings of the 5th BELIV workshop*, pp. 1–8. ACM, 2014.
- [8] T. Castermans, K. Verbeek, B. Speckmann, M. A. Westenberg, R. Koopman, S. Wang, H. van den Berg, and A. Betti. Solarview: Low distortion radial embeddings with a focus. *IEEE Transactions on Visualization and Computer Graphics*, 2018, forthcoming.
- [9] J. Chuang, D. Ramage, C. D. Manning, and J. Heer. Interpretation and trust: Designing model-driven visualizations for text analysis. In *ACM Human Factors in Computing Systems (CHI)*, pp. 443–452, 2012.
- [10] K. Coles. Show ambiguity: Collaboration, anxiety, and the pleasures of unknowing. In *Workshop on Visualization for the Digital Humanities*, 2016.
- [11] W. R. De Jong and A. Betti. The classical model of science: A millennia-old model of scientific rationality. *Synthese*, 174(2):185–203, 2010.
- [12] U. Hinrichs, M. El-Assady, A. J. Bradely, S. Forlini, and C. Collins. Risk the drift! Stretching disciplinary boundaries through critical collaborations between the humanities and visualization. In *Workshop on Visualization for the Digital Humanities*, 2017.
- [13] U. Hinrichs and S. Forlini. In defense of sandcastles: Research thinking through visualization in DH. In *Proceedings of the conference on Digital Humanities*, 2017.
- [14] S. Jänicke. Valuable research for visualization and digital humanities: A balancing act. In *Workshop on Visualization for the Digital Humanities*, 2016.
- [15] S. Jänicke, G. Franzini, M. F. Cheema, and G. Scheuermann. Visual text analysis in digital humanities. In *Computer Graphics Forum*, vol. 36, pp. 226–250, 2017.
- [16] R. Koopman, S. Wang, and A. Scharnhorst. Contextualization of topics: browsing through the universe of bibliographic information. *Scientometrics*, 111(2):1119–1139, May 2017. doi: 10.1007/s11192-017-2303-4
- [17] R. Koopman, S. Wang, A. Scharnhorst, and G. Englebienne. Ariadne’s thread: Interactive navigation in a world of networked information. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 1833–1838. 2015.
- [18] C. McLeod. Trust. In E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2015 ed., 2015.
- [19] T. Rieder and B. Röhle. Digital methods: Five challenges. In *Understanding Digital Humanities*, pp. 67–84. Springer, 2012.
- [20] A. Sangiacomo. Modelling the history of early modern natural philosophy: The fate of the art-nature distinction in the dutch universities. *British Journal for the History of Philosophy*, 2018, forthcoming.
- [21] J. Stahnke, M. Dörk, B. Müller, and A. Thom. Probing projections: interaction techniques for interpreting arrangements and errors of dimensionality reductions. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):629–638, 2016.
- [22] M. Taddeo and L. Floridi. The case for e-trust. *Ethics and Information Technology*, 13(1):1–3, 2011.
- [23] P. van Wierst, S. Vrijenhoek, S. Schlobach, and A. Betti. *Phil@Scale: Computational methods within philosophy*. In *CEUR workshop proceedings*, vol. 1681. 2016.