

BolVis: visualization for text-based research in philosophy

Citation for published version (APA):

van Wierst, P., Hofstede, S., Oortwijn, Y., Castermans, T. H. A., Koopman, R., Wang, S., Westenberg, M. A., & Betti, A. (2018). BolVis: visualization for text-based research in philosophy. In *Proc. 3rd Workshop on Visualization for the Digital Humanities (VIS4DH)*

Document license:
Unspecified

Document status and date:
Published: 01/01/2018

Document Version:
Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

BolVis: Visualization for Text-based Research in Philosophy

Pauline van Wierst¹
Rob Koopman⁴

Steven Hofstede²
Shenghui Wang⁴

Yvette Oortwijn¹
Michel A. Westenberg³

Thom Castermans³
Arianna Betti¹

ABSTRACT

Traditional research in philosophy consists for a large part in conceptual analysis and close reading of texts. This is a precise but time-consuming approach, in which the researcher focuses on one particular text passage or one philosophical concept from one or more works of an author. In this paper, we present BolVis, a visualization tool for text-based research in philosophy. BolVis allows researchers to determine quickly which parts of a text corpus are most relevant for their research by performing a semantic similarity search on words, sentences, and passages. It supports activities such as filtering, exploring the semantic context, comparing, performing close reading on selected passages, et cetera. Our approach enables in-depth analysis of texts at a significantly greater scale than is possible by traditional means, thereby allowing researchers to gain in speed without compromising on precision. We demonstrate the usefulness of BolVis by applying it to a corpus consisting of about 11,000 pages of the writings of the Bohemian polymath Bernard Bolzano (1781–1848). Our use case addresses an open question about Bolzano’s ideas concerning size equality for sets of natural numbers, and we show that the use of BolVis enabled us to find (at least a significant part of) the reason why he came to accept one-to-one correspondence as a sufficient criterion for size equality.

1 INTRODUCTION

Research in philosophy as it is traditionally done consists for a large part in conceptual analysis and close reading of philosophical texts. In this approach, the research focuses on one particular text passage or one philosophical concept from one or more works of an author. Since the texts are extremely complex and conceptually dense, the analysis in question routinely leads to hypotheses about conflicting interpretations of these passages or concepts, which are confirmed or rejected based on interpretation of the text [16]. This approach is generally very time-consuming, and poses severe limitations on the scale at which this type of research can be performed. Typically, only a small corpus is studied [1].

Despite the limitation of scalability, close reading of texts will remain an important research methodology in philosophy. However, we claim that this limitation can be mitigated by the use of computational tools. We present BolVis, a visualization tool which facilitates analysis of texts by philosophy researchers. BolVis allows researchers to determine quickly which parts of the text corpus are most relevant for their research, enabling in-depth analysis of texts at a significantly greater scale than is possible by traditional means.

In this paper, we will demonstrate how BolVis facilitates close reading and how it supports guided discovery. We put BolVis into use to address an open question concerning the ideas of the Bohemian polymath Bernard Bolzano (1781–1848), whose work

has been of fundamental importance for the development of Western logic and the foundation of the sciences such as mathematics and computer science [14]. We apply BolVis to a corpus consisting of about 11,000 pages of Bolzano’s writings, including books, manuscripts, and letters. Such a corpus is too large for a philosophy researcher working in isolation to adequately study in its entirety by close reading. The use case that we will describe investigates Bolzano’s ideas about *size equality* for sets of natural numbers – in particular, how Bolzano during his lifetime changed his ideas on size equality from the classical, Euclidean conception into a conception similar to the contemporary, Cantorian one.

2 RELATED WORK

There are many visualization tools for text analysis that support close and distant reading, see Jänicke et al. [8] for a comprehensive overview. They distinguish between single text analysis, parallel text analysis, and corpus analysis. The tools for *single text analysis* quite often provide summary visualizations of the text via tag clouds, depict the structure, and have support for close reading via highlighting and visualization of annotations or other additional information. Tools for *parallel text analysis* display similarities and differences between similar works or passages. Finally, *corpus analysis* tools produce statistics about the texts in the corpus. In this area, topic model visualization has received much attention (see [4–6, 15], for example). Such systems support distant reading by presenting an overview of the corpora based on text mining techniques and clustering and are therefore less suitable for our needs. Brehmer et al. [3] present an application aimed at investigative journalists, which focuses on exploratory content-based analysis of large document collections. Their approach constructs a hierarchical clustering of all documents based on TF-IDF scores. The clusters are labeled with keywords extracted from these scores. Users can navigate the tree, inspect clusters and individual documents, retrieve documents via keyword search, and they can assign tags to documents. The limitation of their approach is that only keywords can be searched, whereas we require direct search of the full text.

Our approach aims at forming hypotheses about conflicting interpretations of a passage or concept, and hence we require a combination of single text analysis and corpus analysis. There are not many approaches that combine distant and close reading to suit our needs. The Metatation tool [13] bridges the gap between close and distant reading by providing literary analysis experts a system that supports their workflow of reading, annotating, hypothesis generation, and re-reading. While the focus is on supporting the close reading activities, the annotations can be used as distant reading queries to identify relevant connections to other texts in large corpora. Since Metatation was designed for analyzing literary texts and focuses on linguistic features, it cannot be easily adapted or used for our purpose; we focus on what they refer to as *active reading*.

3 BACKGROUND/REQUIREMENTS

We use a corpus that consists of a substantial part of Bolzano’s complete oeuvre written in German. The corpus includes texts in various genres and from different phases of Bolzano’s thoughts on various topics, including mathematics. All material was originally published as volumes of the modern collection of Bolzano’s writings, the *Bernard Bolzano Gesamtausgabe* (BGA). The corpus was

¹ University of Amsterdam, e-mail:

[a.betti | y.oortwijn | p.m.a.vanwierst]@uva.nl.

² TU Eindhoven, e-mail: s.e.hofstede@student.tue.nl.

³ TU Eindhoven, e-mail: [t.h.a.castermans | m.a.westenberg]@tue.nl.

⁴ OCLC Research, e-mail: [rob.koopman | shenghui.wang]@oclc.org.

scanned, converted to plain text format, and corrected to 99% accuracy by a professional service for use in our research. The texts are segmented at the sentence level, and we have worked out and annotated manually the correspondence between pages, works, and the hierarchical structure of each work: chapter, subchapter, section, et cetera. Each work is assigned to a genre (book, manuscript, letter).

Text analysis on the corpus starts from a query (a word or a string of words), which (is taken to) correspond(s) to a concept. Relevant parts of the corpus are then identified by performing a semantic similarity search. These parts can be individual sentences, passages, sections, or chapters, to name a few, depending on the segmentation granularity of the corpus (sentences in our case). These search results form the basis of further analysis, involving activities such as filtering, exploring the semantic context, comparing, performing close reading on selected passages, et cetera. We identified the main analysis tasks that are performed in text-based philosophical research, and derived the following requirements. BolVis must:

- R1** make clear in which book, chapter, section, paragraph, and sentence each search result can be found in the entire corpus;
- R2** show for multiple search queries whether there are similarities, overlap, or co-occurrences in the results on each level;
- R3** allow performing a query within the context of a previous result;
- R4** enable access to the local context of each result;
- R5** support selection of passages for comparison and close reading;
- R6** allow ordering and filtering the search results by various criteria.

4 BolVis

Semantic similarity search in BolVis is based on ARIADNE, a tool developed at OCLC Research [9, 10]. ARIADNE operates on a dimension-reduced semantic matrix that is built from sentence-level co-occurrence statistics of the corpus and embeds words in a semantic space [10]. Using the weighted average of the word embeddings [9], sentences from the corpus or any strings of words can also be embedded in the same semantic space. This way, it is possible to calculate the similarity between words, sentences or any strings of words. As ARIADNE operates on unannotated text, it can work directly with the nineteenth-century German of the Bolzano corpus.

Starting from a query, either a single term, a phrase, or a sentence, ARIADNE returns an ordered list of sentences which are most related to the query, i.e., its closest neighbors in the semantic space (see Figure 1). The semantic space in this case accommodates words, sentences and passages. This way of searching is different from exact search, which relies heavily on the occurrences of the query terms.

We have built a client-server application on top of ARIADNE, making BolVis easily accessible to users via their browser. The user interface consists of four main components, as shown in Figure 3. We will briefly discuss the functionality of each component now, using the situation shown in Figure 3 as a running example.

Users kick off by entering a query. BolVis is aimed at expert users who have a starting point in mind; it does not target supporting exploration without prior knowledge by laymen.

Hierarchy overview. This overview is populated upon submitting the first query. It shows the entire corpus that is loaded, as an expandable list. Top level items are works, below that are chapters, then subchapters or sections, et cetera. Not all works have the same number of levels, BolVis does not impose restrictions on it.

To the left of every list item, we show discs corresponding to the queries entered by the user (in the same order). This set visualization [11] shows if a query result occurs in the work, chapter, et cetera, represented by that list item (filled disc) or not (empty disc). This allows users to quickly spot works, or parts of works, which are of interest (**R1** and **R2**). The overview can be sorted to show items

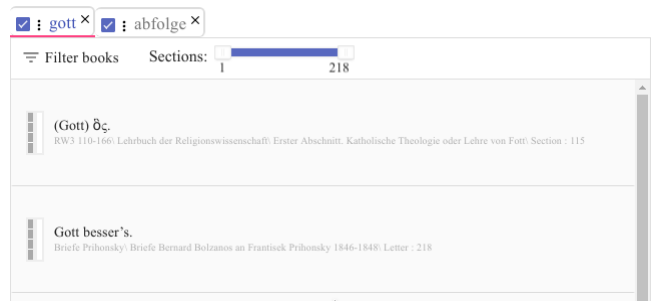


Figure 1: Closeup of the panel with query results (for *Gott* and *Abfolge*). Each query is represented by a tab, allowing users to quickly switch between the search results of each query. The 4-part bars (to the left of the results) indicate how closely related the queries in the tabs are to the results. The figure shows that while *Gott* is closely related to both, *Abfolge* is not related to either.

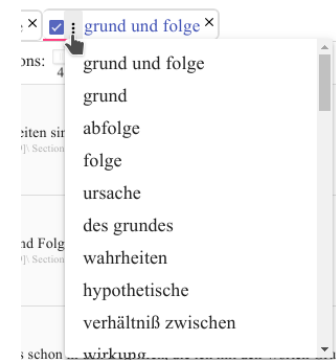


Figure 2: BolVis encourages guided discovery by suggesting queries that are semantically close to queries performed earlier by users.

with most hits first (**R6**). It can additionally be sorted and filtered by year of publication (**R6**); this metadata is present for most items.

Results list. The passages most closely related to the query are displayed, in order, in the results list. The bucketed *score* of results is also displayed in 4-part bars. This shows that for example in Figure 1, even though *Gott* scores higher than *Gott besser's*, both results are highly relevant—as one would intuitively expect.

Results can be filtered to include only those from selected works and/or sections (**R6**). This allows users to narrow their search scope and focus on a subset of the corpus that has their interest.

Users can mouse over results to reveal three buttons (normally hidden to reduce clutter), as illustrated in Figure 3. From top to bottom, these enable (i) querying for the result, enabling guided discovery; (ii) adding the result to the favorites list (see below), and (iii) viewing the result in the original work, to see it in context (**R4**).

Favorites list. In order to support **R5**, users can mark query results as favorites. The interaction with items in this list is the same as for the results list. The favorites can be exported to CSV using the download button, for further processing in another application.

The last requirement BolVis needs to support is **R3**. This is implemented via a popup menu on the tabs that represent queries, as illustrated in Figure 2. The menu presents a list of semantically close queries, and enables the user to engage in guided discovery.

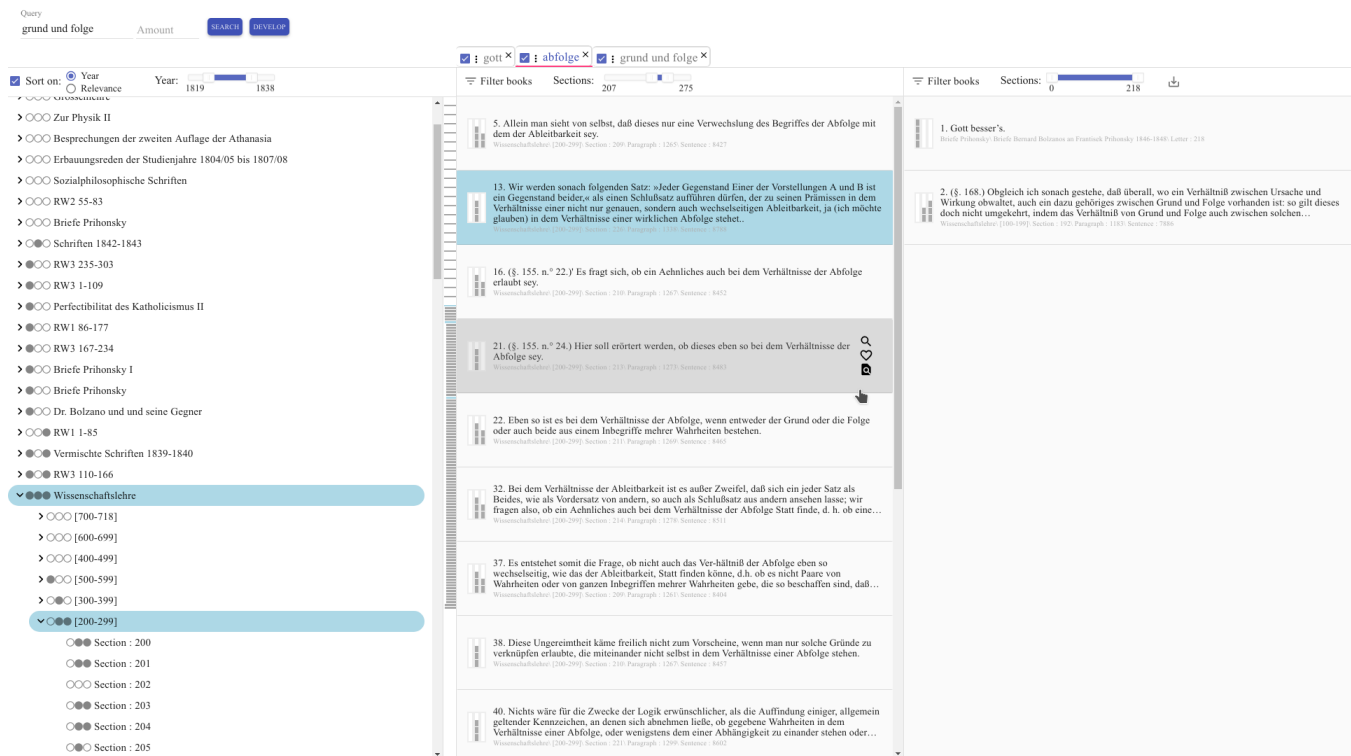


Figure 3: Overview of BolVis. The query input and search button are located in the top left. Below that, three panels evenly share the available space. From left to right, they are the *hierarchy overview*, the *results list* and the *favorites list*. The result highlighted in blue has been clicked. The hierarchy overview shows where that result occurs in the corpus, which is also displayed in the scrollbar of the overview. Mousing over a result triggers display of buttons that allow users to (i) query for that result; (ii) add the result to their favorites; and (iii) view the passage in the source text, to see the result in the context of the original work. All panels have sorting and/or filtering controls near the top.

5 USE CASE

Here we describe how BolVis aids in philosophical research, that is, how BolVis can be used in addition to traditional methods (close reading) to answer text-based questions in (the history of) philosophy. Our research question concerns Bolzano’s position within the development of *set theory*, the current foundational framework for mainstream mathematics. A crucial prerequisite for the development of set theory has been the adoption of *one-to-one correspondence* as a sufficient criterion for equality of size (according to this criterion, two sets A and B are of equal size just in case for every element a in A there is one element b in B , and vice versa). It is known that Bolzano rejected this criterion as sufficient for equality of size for a long time, but at some point in his life changed his position and embraced it – at least where it comes to sets of natural numbers [12]. However, it is not known exactly *why* Bolzano came to change his mind on this. We aim to identify Bolzano’s reasons for adopting one-to-one correspondence as a sufficient criterion for equality of size. Understanding why Bolzano – who rejected this criterion for so long – came to embrace one-to-one correspondence as a sufficient criterion for equality of size, will help us to evaluate the widespread view that one-to-one correspondence is the only “right” (in the sense of mathematically coherent) criterion for equality of size (see e.g. [7]).

Our research has a traditional starting point: the clearest evidence that Bolzano at some point had adopted one-to-one correspondence as a sufficient criterion for equality of size for sets of natural numbers is a letter which he wrote to a pupil a couple of months before his own death [2, 12]. In this letter, Bolzano rejects a specific argument which he gave in one of his main works – namely, in the *Wissenschaftslehre* (henceforth, WL), in which he develops a logic which is to serve as a foundation for mathematics. The argument

is given in WL §102, so our research starts by analyzing this section, with the aim of understanding what exactly Bolzano in the end, witnessing the letter, came to reject.

Understanding WL §102 is hard, for example because Bolzano uses many technical terms, some of which are currently not commonly used in philosophy or mathematics anymore, and others seem to have (slightly) changed their meanings. Therefore, it is crucial to our aims to clarify the meaning of the terms that Bolzano uses in this section. What Bolzano aims to show in WL §102 we understand from the section’s title: “*Keine endliche Menge von Maßen genügt, die Weiten aller Vorstellungen zu messen* (No finite set of measures suffices to measure the width of all ideas)”. We know what Bolzano means with “ideas”: ideas are abstract objects which are the building blocks of propositions, where the latter are the bearers of truth and falsity and serve as the meanings of sentences (see e.g. WL §48, cfr. [14]). But we are unsure of what exactly he means with “to measure” and the “width” of ideas. We use BolVis to find passages which help us determine the meanings of these (technical) terms, i.e. the way in which they are used by Bolzano in the corpus.

We use BolVis to query “Weite” and see that it occurs in four works, but a closer look at the results reveals that in only two works he uses it in the relevant sense – in the other cases, he uses it in a non-technical sense when arguing that a certain definition is too broad. We learned that Bolzano introduced the concept in the WL, mentions it once more in the first section of the introduction to the *Größenlehre* (henceforth GL), and then discards it. Bolzano defines *Weite* in GL §1 in the following way:

Von dem Begriffe der Höhe läßt sich mit Nutzen noch jener der Weite als ein höherer unterscheiden. Unter der Weite einer Vorstellung verstehen wir nämlich eine Größe,

welche die Menge aller ihr unterstehenden Gegenstände mißt, und wir sagen sonach, daß ein Paar Vorstellungen gleichweit sind, wenn sie gleichviele Gegenstände haben, und daß die eine weiter sey als die andere, wenn jene mehr Gegenstände hat als diese, unangesehen, ob die Gegenstände der Einen auch zugleich Gegenstände der anderen sind oder nicht.

From this passage and the other query results in which Bolzano uses the term in the relevant sense, we learn that the width of an idea is the *quantity* (*Größe*) of the set of objects to which that idea refers. For example, the idea *root of the equation* $x^2 - 5x + 6 = 0$ has a width of 2, for it refers to two objects: 2 and 3 (WL §117).

We turn to the second term in need of clarification: “to measure” (*messen*). The fourth result in the query for this term is exactly the title from the section under consideration, WL §102. We click on this result to query this passage and obtain a list of other passages from the corpus that are semantically similar to it. The first three results consist of other sentences from WL §102. Result number 8 is remarkable: it is a sentence from WL §93 saying that infinite sets cannot be measured. The reason this result is remarkable, is that in the body of the section where our research started, WL §102, Bolzano presents a series of infinite sets and talks about their measurement – implying that they *can* be measured.

Given these two – at least apparently – contradicting passages, we want to know whether or not infinite sets can be measured according to Bolzano. We query “Mengen messen” (*sets and to measure*) and we obtain some sentences of WL §102, some (seemingly) irrelevant sentences, and several sentences in which Bolzano is talking about *comparing* (*vergleichen*) infinite sets with each other. We repeat the search with slightly different query terms: “Mengen Größe” (*sets and quantity*); again, in almost all sentences that we obtain Bolzano speaks about comparing (*Verhältniss, Vergleichung*) infinite sets, and in no one of them he speaks about measuring them or about their quantity. We keep in mind that BolVis does not do an exact search, and so there might be passages in which Bolzano does speak about measuring infinite sets which did not turn up in our query results. However, we feel at this point confident enough to suppose that Bolzano held that infinite sets *cannot* be measured, but that they can merely be compared among each other.

At this point it is a mystery why Bolzano in WL §102 is measuring infinite sets, whereas he generally held that infinite sets cannot be measured. We turn back to WL §102 in order to look for a clue. In this section, Bolzano is considering a series – in modern terminology, a *nested sequence* – of which the first term is the set of all natural numbers, the second term the set of all squares of natural numbers, the third all fourth powers, and so forth. Bolzano claims that for each term j of this series, the corresponding set is *infinitely larger* than the set corresponding to its successor $j + 1$ – for example, the set of natural numbers is infinitely larger than the set of squares –, and since this series has infinitely many terms, it proves the claim made in the title of the section: that no finite set of measures suffices to measure the width of all ideas. However, we now recognize that in this argument, Bolzano actually *does not* measure infinite sets: he bounds the sets by introducing some largest number N and then shows that for every N the j^{th} set is larger than the $j + 1^{\text{th}}$ set; given that, as he argues, we can take this number N larger and larger, this result must hold in general. In other words, in the argument in WL §102 Bolzano measures – or better: attempts to measure – infinite sets by making them finite.

Bolzano later wrote in the letter to his pupil that he came to see that there are *just as many* natural numbers as squares, because for every natural number there is a square number (namely, the square of that number), and for every square there is a natural number corresponding to it (namely, its root) – in other words, there is a one-to-one correspondence between the set of natural numbers and the set of squares [2]. In that letter, Bolzano explicitly rejects the idea

underlying the argument in WL §102, that there are more natural numbers than squares. For, as he writes: “The false result was due to an unjust inference from a *finite* set of numbers, namely which do not succeed the number N , to *all of them*” [2]. We now understand what he must have meant: it was a mistake to (attempt to) measure infinite sets by bounding them above. Thus, from our analysis of WL §102 and our discovery that Bolzano held that infinite sets cannot be measured, we may conclude that at least part of the reason that Bolzano came to accept one-to-one correspondence as a sufficient criterion for equality of size is that he came to see that the argument that he gave in WL §102 is false.

6 CONCLUSION

The use case showed that BolVis aids our text-based philosophical research in a substantial way. We used BolVis to quickly find relevant passages for our research and we went back and forth from BolVis to the original text to do our conceptual analysis. In this manner, we discovered the information – that is, that Bolzano held that infinite sets cannot be measured – necessary to answer our research question: to understand why Bolzano later in his life rejected his former position and came to accept one-to-one correspondence as a sufficient criterion for equality of size of sets of natural numbers.

However, to get a grip on exactly how useful BolVis is, and how it can be improved, is not easy. A general problem for philosophers using digital tools such as BolVis is how to evaluate these tools. For example, when presented with the results of a query, it is easy enough to come up with a justification for these results and give a reason why these passages are relevant given the query. But we cannot be certain that BolVis shows *everything* relevant, nor are we certain that we queried for the right terms, checked enough results, et cetera. It must be acknowledged, however, that such issues also occur in the traditional close reading setting.

Part of the usefulness of BolVis consists in the possibility to visualize results of only a part of the corpus. For example, we can search only in Bolzano’s later, more mature works, or we can investigate what Bolzano writes about *infinity* in his mathematical works, but not in his religious works. In our use case, ordering the results chronologically made it easy to see that the last time Bolzano used the term “Weite” was in the *Größenlehre*. Such chronological ordering of the search results fosters new hypotheses (to be addressed in later research): does the disappearance of *Weite* have something to do with Bolzano’s embracing of one-to-one correspondence?

In this paper, we applied BolVis to the works of only one philosopher: Bernard Bolzano. In the future, we plan to extend this use of BolVis to the works of more than one philosopher. We expect BolVis to be useful also to identify conceptual relations between the works of different philosophers, and to aid in analyzing the development of philosophical and mathematical concepts through history.

ACKNOWLEDGMENTS

The Netherlands Organization for Scientific Research (NWO) is supporting A. Betti, P. van Wierst, Y. Oortwijn and T. Castermans under project no. 314.99.117 and A. Betti and P. van Wierst also under project no. 277.20.007.

REFERENCES

- [1] A. Betti, H. van den Berg, Y. Oortwijn, and C. Treijtel. History of Philosophy in Ones and Zeros [accepted; edited copy approved by authors on Jun 30, 2018]. In M. Curtis and E. Fischer, eds., *Methodological Advances in Experimental Philosophy*, Advances in Experimental Philosophy. Bloomsbury, London, 2018.
- [2] B. Bolzano. Vermischte philosophische und physikalische Schriften 1832-1848. In J. Berg, ed., *Bernard Bolzano Gesamtausgabe 2A 12/2*. 1978.
- [3] M. Brehmer, S. Ingram, J. Stray, and T. Munzner. Overview: The design, adoption, and analysis of a visual document mining tool for

- investigative journalists. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2271–2280, 2014. doi: 10.1109/TVCG.2014.2346431
- [4] Y. Chen, L. Wang, M. Dong, and J. Hua. Exemplar-based visualization of large document corpus. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1161–1168, 2009. doi: 10.1109/TVCG.2009.140
- [5] W. Cui, S. Liu, L. Tan, C. Shi, Y. Song, Z. Gao, H. Qu, and X. Tong. Textflow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2412–2421, 2011.
- [6] J.-D. Fekete and N. Dufournaud. Compus: visualization and analysis of structured documents for understanding social life in the 16th century. In *Proc. ACM conference on Digital Libraries (DL'00)*, pp. 47–55, 2000. doi: 10.1145/336597.336632
- [7] J. Ferreirós. *Labyrinth of Thought: A History of Set Theory and Its Role in Modern Mathematics*. Birkhäuser, Basel, 1999.
- [8] S. Jänicke, G. Franzini, M. F. Cheema, and G. Scheuermann. On Close and Distant Reading in Digital Humanities : A Survey and Future Challenges. In *Eurographics Conference on Visualization (EuroVis) (2015) – STARs*, pp. 1–21, 2015. doi: 10.2312/eurovisstar.20151113
- [9] R. Koopman, S. Wang, and A. Scharnhorst. Contextualization of topics: Browsing through the universe of bibliographic information. *Scientometrics*, 111(2):1119–1139, 2017.
- [10] 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. ACM, 2015.
- [11] A. Lex, N. Gehlenborg, H. Strobel, R. Vuillemot, and H. Pfister. Upset: Visualization of intersecting sets. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1983–1992, 2014. doi: 10.1109/TVCG.2014.2346248
- [12] P. Mancosu. Measuring the size of infinite collections of natural numbers: Was Cantor’s theory of infinite number inevitable? *The Review of Symbolic Logic*, 2(4), 2009.
- [13] H. Mehta, A. Bradley, M. Hancock, and C. Collins. Metatation: Annotation as Implicit Interaction to Bridge Close and Distant Reading. *ACM Transactions on Computer-Human Interaction*, 24(5):1–41, 2017. doi: 10.1145/3131609
- [14] E. Morscher. Bernard Bolzano. In E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2014 ed., 2007.
- [15] D. Oelke, H. Strobel, C. Rohrdantz, I. Gurevych, and O. Deussen. Comparative exploration of document collections: a visual analytics approach. *Computer Graphics Forum*, 33(3):201–210, 2014.
- [16] H. van den Berg, G. Parra, A. Jentzsch, A. Drakos, and E. Duval. Studying the history of philosophical ideas: Supporting research discovery, navigation, and awareness. In *Proceedings of the 14th International Conference on Knowledge Technologies and Data-driven Business*. ACM Digital Library, 2014. doi: 10.1145/2637748.2638412