EuroVis Workshop on Visual Analytics (2016)
N. Andrienko and M. Sedlmair (Editors)

Visualization of Latin Textual Variants using a Pixel-Based Text Analysis Tool

Bharathi Asokarajan[1], Ronak Etemadpour[2], June Abbas[1], Sam Huskey[1], Chris Weaver[1]

[1]The Digital Latin Library Project, The University of Oklahoma
[2]Department of Computer Science, Oklahoma State University

Abstract

One of the most important activities of Latin scholars is to analyze fragmentary copies of a Classical text and assemble an annotated reconstruction as a conjecture about its original form. We have developed a pixel-based visual text analysis tool to help Latin scholars visualize the evolution of historic copies and analyze the details of alterations and errors introduced in transcription. Coordination of pixel-based visualizations with focus+context navigation across multiple views allows compact representation of text variation across scales of text structure. This approach helps scholars validate the accuracy of variations and assess subtle differences across fragmentary copies as well as past reconstructions. In this paper, we describe the central design features of the tool that help scholars analyze the density and distribution of variants by interacting with text at the granularities of words, lines, and pages simultaneously. We present the results of a user study on our initial multiple view focus+context design and discuss how the results motivate a more visually integrated focus+context design using tiered views.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Interfaces and Presentations]: User Interfaces-Graphical user interfaces (GUI)—Interaction Styles

1. Introduction

Our knowledge of the Latin writings of antiquity is based on reconstructions of texts, not originals. Classics scholars conjecture how an original looked by interpreting and collating the variations in a relatively small collection of fragmentary copies as well as the past conjectures of other scholars. Substantial variation in spelling, phrasing, and even replacement or omission of entire passages is common. Exploring and analyzing the details and patterns of variation is central to Classics scholarship, pedagogy, and even routine reading and language learning by Latin students and devotees. Classical Latin texts are at the core of western cultural heritage.

On the Digital Latin Library project [DLL], we work with Classics scholars and their professional societies to develop new techniques and tools that apply visual analytics to their scholarship. Within the broad agenda of visual analytics [TC06], our focus is on developing new representation and interaction techniques to help scholars study patterns in textual provenance. The state-of-the-art in Classics consists of human entry of word-level variations across large, sparse, ad hoc spreadsheets. We work on integration of pixel-based and focus+context approaches as a way to let scholars simultaneously analyze fragmentary influences at different scales of text structure in these *collation tables*. Scholars publish their conjectured reconstructions in printed *critical editions* with an *apparatus* of often cryptic footnotes derived from collation tables. We aim to provide visual techniques to support dissemination of the details of scholarly reasoning in a complementary, collaborative, and broadly accessible form.

We evaluated a prototype tool designed to reveal patterns of variation across *witnesses*—manuscript copies, earlier editions, and other sources—for particular *lemmata*—words or segments—in a text. A focus+context design [SZ00] supports contextual navigation across scales of aggregated text structure. The pixel-based representation allows focusing on text variation at the level of an individual lemma or an entire line or page. We used the evaluation results to assess the usability of: (1) aggregating variations in lemmata, lines, and pages in coordinated pixel-based views, and (2) encoding variation by mapping edit distance into stepped hue gradients. We concluded that the design is effective for reading details of variation for individual lemmata/lines/pages, and that using separate views to display context is effective but somewhat inefficient, suggesting a more integrated layout.

© The Eurographics Association 2016.

DOI: 10.2312/eurova.20161119

delivered by EUROGRAPHICS DIGITAL LIBRARY
www.eg.org diglib.eg.org

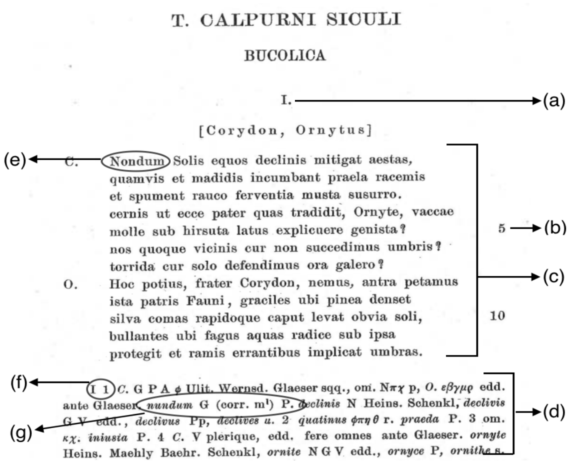

Figure 1: *The first page of Giarratano's 1910 edition of* Calpurnius Siculus, *including (a) poem numbers, (b) line numbers, (c) text, and (d) apparatus. An apparatus entry (f) lists past variants (g) with one chosen for use in the text (e).*

2. Related Work

Pixel-based visualization is an approach popularized by Keim, et al. [Kei00]. It displays large data sets in a compact format by representing each data point as a single pixel or small region of pixels. The color of pixels is typically used to represent a quantity or category. In text visualization, color might represent word count or a part of speech (POS) category. In our case, color represents similarity between a variant and its lemma in the base text. Representing the entire data set as pixels allows users to scan for patterns of variation across lemmata and witnesses in the reconstructed text. The VarifocalReader [KJW*14] similarly focuses on visual exploration of innate hierarchical structure in text.

In the digital humanities there is increasing interest in software tools that apply visualization to text scholarship. Juxta [Jux] integrates several simple visualizations including heat maps to depict the density of variants across pages in text. The Interactive Timeline Viewer [MKF*02] allows comparison of several texts to identify differences among variants. TRAViz [JGF*15] provides a graph-based visual representation of variants for purposes of comparison. While such tools help in recognizing patterns and density for a given text metric, they lack features for viewing these metrics at different levels of text structure such as pages, lines and words. The visualizations used are also limited in how many variants can be compared, typically not more than seven. Comparing a larger number of variants—which in our case is often 20 or more—is difficult in existing tools. Our pixel-based focus+context approach is designed to address these issues of granularity and scalability in text variation. We used Improvise [Wea04] to design and implement the coordinated multiple views and dynamic queries in the tool.

3. Tool Design

Our analysis of the tool uses Giarratano's 1910 critical edition of *Calpurnius Siculus* [TC10]. Figure 1 shows the first page of the edition. The base text and critical apparatus are in the conventional format of printed editions. We converted the first poem of the edition into an XML format, following Text Encoding Initiative (TEI) conventions [TEI], and used XPath queries to convert the XML into tabular data for visualization. The reconstruction of the first poem in the work involves approximately 1200 variants across 7 pages, with an average of 6 lemmata for each of the 94 lines.

Through in-depth interviews with Latin scholars, we identified some of the major tasks that they perform while studying textual variation. They identify and characterize patterns of variation over a text for a given witness and also patterns of variation over different witnesses for a given lemma. They focus on the apparatus, using the base text as context, and referring to the location of lemmata in the base text to do that. To support these tasks, the pixel-based text analysis tool (Figure 2) displays the base text along with four pixel-based views: an overview plus three detail views that aggregate variation at the levels of lemmata, lines, and pages.

The overview displays lemmata along the horizontal axis, in the order encountered in the apparatus, which generally corresponds to the reading order of words in the base text. The vertical axis lists all witnesses having variants incorporated somewhere in the edition. The central grid represents the presence and degree of variation for each lemma and witness as a colored pixel. Along the horizontal axis, lemma color alternates between black and red to indicate successive pages. This helps in tracking a given lemma during panning or after glancing at other views. The base text view displays the reconstructed Latin text in its natural form. Hovering over a lemma in the overview highlights the corresponding word/segment in the base text view.

Pixel coloring plays a vital role in depicting similarity between variants, including for special types of variation such as omission. We use a case-sensitive Levenshtein edit distance [Lev66] to calculate similarity between a lemma and its corresponding variant for each witness. For instance, the lemma *Nondum* ("not yet") in the base text has a variation *nundum* (a scribal error) with distance two. Distances are normalized over the entire text and mapped into a five-level univariate color scheme from ColorBrewer [HB11]. Omissions are mapped into a distinct color (purple). The overview colors the "pixel" grid cell at $(lemma, witness)$ whenever the apparatus notes some variation of that lemma by that witness. Brighter pixels indicate higher lexicographic similarity and hence smaller edit distances. When a column of pixels is represented by the same color, it indicates that many witnesses agree on a particular variant. Conversely, a column of pixels with varying colors indicates a lot of variation across sources, suggesting substantial past disagreement that may warrant ongoing scholarly attention.

© The Eurographics Association 2016.

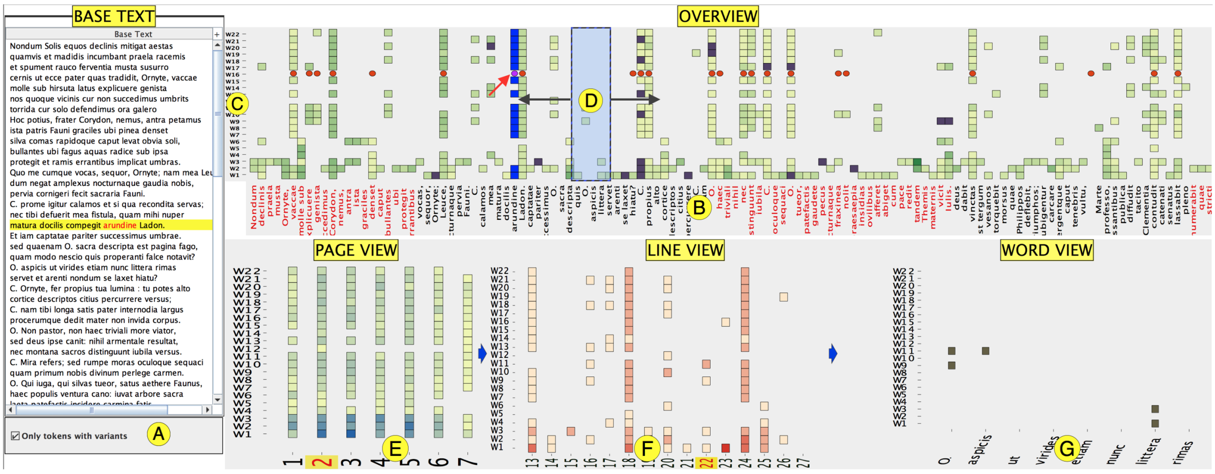

Figure 2: *The tool for analyzing variation in classical Latin texts. The layout consists of: a base text view with a global filter option (A); an overview of variation between lemmata (B) and witnesses (C); a draggable lens to select a subsequence of lemmata (D); and a hierarchy of drill-down views that summarize variation at the level of pages (E), lines (F), and words (G).*

The three detail views similarly display variation in a drill-down hierarchy from pages to lines to individual lemmata. The page view shows all pages. Selecting a page number determines the range of lines shown in the line view. Selecting a line number filters the lemma view to show only the lemma columns on that line. Unlike the overview and lemma ("word") views, pixel color in the page and line views represents the total number of variants for each witness. To aid drill-down navigation, a draggable horizontal lens in the overview highlights corresponding page numbers in the page view. Through a combination of panning, hovering, and drill-down interactions, scholars can examine pixel location and coloring to analyze interesting patterns and outliers in the relationships between lemmata and witnesses. There are often runs of lemmata that have few or no lexicographic variants, which results in a sparse distribution of pixels. The tool provides a check box to display only lemmata that have at least one lexicographic variant. Key tasks that a scholar can perform using the tool include:

- **Examine the pattern of variation over lemmata.** Hover over a pixel (Figure 2, red diagonal arrow) to highlight (in red) all variations of lemmata for a particular witness.
- **Examine the pattern of variation over witnesses.** Hover over a pixel to highlight (in blue) all variations between witnesses for a particular lemma.
- **View a lemma in context.** Brush pixels to highlight the corresponding lemmata in the base text view.
- **Drill down to a passage of interest.** Drag and stretch the lens in the overview to highlight page and line numbers in the detail views. This lets scholars view lemmata in complementary ways: in direct succession in the overview to look at fine-grain patterns, and in aggregated succession in the detail views to look at coarse-grain patterns.

4. User Study

Following the guidelines described in [Car08], we designed a user study to assess the effectiveness and usability of our text visualization tool. Our goals were to assess the overall usability of the tool, the effectiveness of pixel coloring in identifying patterns, and the ease of use in performing queries using focus+context across multiple views. To define representative user tasks, we interviewed scholars to identify typical questions that they raise when examining printed critical editions by hand. To formulate an experimental hypothesis, we identified three major visual analysis tasks, each associated with one of our three goals. We grouped these tasks into *synoptic tasks* as explained in [AA05] and [EMdSP*15]. The **interaction-related (I)** tasks were for investigating ease in using selection, panning, and zooming features to perform a query on the data. The **perception-related (P)** tasks were for investigating the effectiveness of pixel color mapping to perform grouping and pattern identification related queries on a given data set. The **navigation-related (N)** tasks were aimed at verifying the usability of the multiple coordinated view visualization features, particularly the distributed focus+context organization. Our hypothesis is that *the performance of the pixel-based text analysis tool is the same across all three task groups*. So if the *alternate* hypothesis holds, we can identify the best and worst performing groups, which helps inform features to include and add in future designs.

4.1. Experimental Setting

We conducted a pilot study with the participation of four graduate students and one faculty member on the DLL project. Together with assessment of preliminary results, we aimed to test the duration of tasks, amount of training re-

© The Eurographics Association 2016.

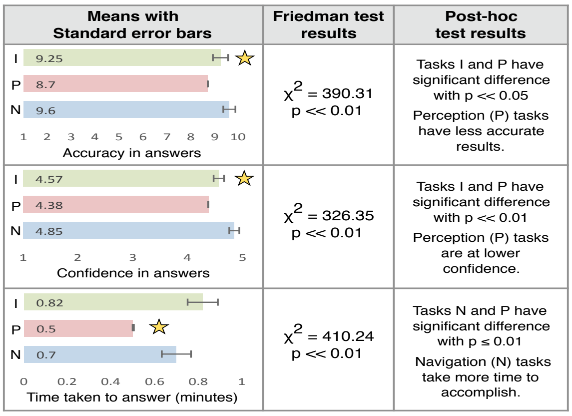

Figure 3: *Accuracy, confidence, and time taken results for tasks in the three task groups (I, P, N), with mean, error bars and p-value. A star indicates the best performing task group.*

quired to perform various tasks for a particular user group, and robustness of the study interface. The pilot study suggested interesting trends in support of the alternate hypothesis. The actual study was conducted over several days with 14 participants. All were undergraduate or graduate students with normal vision. Sessions were individual, included a brief training phase, and engaged five questions in each of the three task groups. Participants used typical interaction techniques with mouse, keyboard, and a 21.5" display. Think-aloud before and after each task performance captured participants' approach to queries and their observations about the task. We collected qualitative feedback about the design and usability of the tool at the end of each session.

4.2. Study Results and Analysis

In Figure 3 we summarize the performance measurements and analysis results for the I, P, and N task groups. We used three performance metrics: **accuracy** (in %), specifically the correctness of answers based on ground truth, normalized; **confidence level** in answers, on an increasing scale from 1 to 5; and **time taken** (in minutes) to perform a task. A check for data normality using the Shapiro-Wilk test revealed that all three measurements had a non-normal distribution. To further check the statistical significance of data, we applied a non-parametric approach using the Friedman rank sum test to compare group means. All three metrics had $p \ll 0.01$, which indicates a significant difference in group means.

A post-hoc analysis using a pairwise Wilcoxon signed rank test helped us to identify the best/least performing task group for each metric. With respect to accuracy and confidence level, I and P tasks had a significant difference in means (accuracy: $Z = -1.35$, $p \ll 0.05$, confidence: $Z = -3.82$, $p \ll 0.01$). Task I results were the most accurate and involve the highest confidence. Task P results displayed less accuracy and confidence in answers. The primary reason behind this is the effort involved in identifying patterns in colored pixels and the rotated orientation of text along the horizontal axis of the overview. To better understand the lower performance in P tasks, we checked qualitative feedback elicited from participants at the end of each session. Eight out of 14 users mentioned that, although they were able to identify patterns in colors easily, they were not certain of the correctness of their answers. Providing sorting features along the vertical axis, to rearrange pixels in order of color/similarity, might be one way to overcome this issue. Text orientation is a common design tradeoff in visualization tools, and remains a problem. When we consider time taken as a measure of efficiency, N and P tasks both showed significant difference in means ($Z = -3.09$, $p \leq 0.01$). P tasks were less time consuming; color patterns can be readily identified due to the compact representation of data points. N tasks needed more time to accomplish because they require traversing multiple views with a variety of drill-down steps.

There were several notable highlights in the qualitative feedback. Participants found interaction, especially panning and zooming in the overview, particularly useful when there is a close-call in distinguishing pixel colors. Brushing a pixel to highlight the corresponding pixels in the same column and row proved helpful for interpreting the corresponding distributions of variations. We plan to use this study approach to compare our design with other variant visualizations and existing practice. The initial study aimed primarily to understand the tasks for which the pixel-based focus+context combination is promising. Based on the results, we are evaluating a new tool with an integrated multi-tier design inspired by the Perspective Wall [MRC91]. We anticipate that the new design improves overall performance in all three task groups.

5. Conclusion

In this paper, we described the key features of an initial pixel-based focus+context tool design for analysis of detailed variation in reconstructions of classical Latin texts. Visually correlating witnesses and grouping them into categories appears likely to help scholars analyze and propose improvements to the detailed historical provenance of these texts. Our user study verified our hypothesis about the accuracy and confidence of detailed questioning of text that can be achieved using the tool. The focus+context organization compliments the pixel-based representation and proves helpful for efficient data exploration. As a part of development in the ongoing Digital Latin Library project, this combination is slated to be included alongside storyline and other visualization techniques to provide a full-featured desktop application for rich visual analysis of classical texts.

6. Acknowledgments

This work was supported as part of the Digital Latin Library project by a grant from the Andrew W. Mellon Foundation.

© The Eurographics Association 2016.

References

[AA05] ANDRIENKO N., ANDRIENKO G.: *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005. 3

[Car08] CARPENDALE S.: *Information Visualization: Human-Centered Issues and Perspectives*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, ch. Evaluating Information Visualizations, pp. 19–45. 3

[DLL] Digital Latin Library. `http://digitallatin.org/`. Accessed: 2016-02-19. 1

[EMdSP*15] ETEMADPOUR R., MOTTA R., DE SOUZA PAIVA J. G., MINGHIM R., DE OLIVEIRA F., CRISTINA M., LINSEN L.: Perception-based evaluation of projection methods for multidimensional data visualization. *Visualization and Computer Graphics, IEEE Transactions on 21*, 1 (2015), 81–94. 3

[HB11] HARROWER M. A., BREWER C. A.: ColorBrewer.org: An online tool for selecting colour schemes for maps. In *The Map Reader: Theories of Mapping Practice and Cartographic Representation*, Dodge M., Kitchin R., Perkins C., (Eds.). John Wiley & Sons, Chichester, 2011, ch. 3.8, pp. 261–268. 2

[JGF*15] JÄNICKE S., GESSNER A., FRANZINI G., TERRAS M., MAHONY S., SCHEUERMANN G.: TRAViz: A visualization for variant graphs. *Digital Scholarship in the Humanities 30*, suppl 1 (2015), i83–i99. 2

[Jux] Juxta. `http://www.juxtasoftware.org/`. Accessed: 2016-02-19. 2

[Kei00] KEIM D. A.: Designing pixel-oriented visualization techniques: Theory and applications. *IEEE Transactions on Visualization and Computer Graphics 6*, 1 (January-March 2000), 59–78. 2

[KJW*14] KOCH S., JOHN M., WORNER M., MULLER A., ERTL T.: VarifocalReader – in-depth visual analysis of large text documents. In *IEEE Transactions on Visualization and Computer Graphics* (2014), pp. 1723–1732. 2

[Lev66] LEVENSHTEIN V.: Binary codes capable of correcting deletions, insertions and reversals. In *Soviet Physics Doklady* (1966), vol. 10, p. 707. 2

[MKF*02] MONROY C., KOCHUMMAN R., FURUTA R., URBINA E., MELGOZA E., GOENKA A.: Visualization of variants in textual collations to analyze the evolution of literary works in the cervantes project. In *6th European Conference on Research and Advanced Technology for Digital Libraries (ECDL* (2002), pp. 638–653. 2

[MRC91] MACKINLAY J. D., ROBERTSON G. G., CARD S. K.: The perspective wall: Detail and context smoothly integrated. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (1991), CHI '91, pp. 173–176. 4

[SZ00] STASKO J., ZHANG E.: Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *Information Visualization, 2000. InfoVis 2000. IEEE Symposium on* (2000), pp. 57–65. 1

[TC10] TITUS CALPURNIUS SICULUS, CAESAR GIARRATANO: *Calpurnii Et Nemesiani Bucolica*. 1910. 2

[TC06] THOMAS J. J., COOK K. A.: A visual analytics agenda. *IEEE Comput. Graph. Appl. 26*, 1 (Jan. 2006), 10–13. 1

[TEI] Text Encoding Initiative. `http://www.tei-c.org`. Accessed: 2016-02-19. 2

[Wea04] WEAVER C.: Building highly-coordinated visualizations in Improvise. In *Information Visualization, 2004. INFOVIS 2004. IEEE Symposium on* (2004), pp. 159–166. 2

© The Eurographics Association 2016.