GazeGIS: A Gaze-based Reading and Dynamic Geographic Information System

Laura G. Tateosian, Michelle Glatz, Makiko Shukunobe, and Pankaj Chopra

Abstract— Location is an important component of a narrative. Mapped place names provide vital geographical, economic, historical, political, and cultural context for the text. Online sources such as news articles, travel logs, and blogs frequently refer to geographic locations, but often these are not mapped. When a map is provided, the reader is still responsible for matching references in the text with map positions. As they read a place name within the text, readers must locate its map position, then find their place again in the text to resume reading, and repeat this for each toponym. We propose a gaze-based reading and dynamic geographic information system (GazeGIS) which uses eye tracking and geoparsing to enable a more cohesive reading experience by dynamically mapping locations just as they are encountered within the text. We developed a prototype GazeGIS application and demonstrated its application to several narrative passages. We conducted a study in which participants read text passages using the system and evaluated their experience. Evaluations indicate a positive reception of this new reading paradigm. Lastly, we present a case study to test its application for intelligence analysis and discuss how experts in this domain envision it use.

Index Terms—Eye-tracking, GIS, geoparsing, dynamic mapping, gaze-contingency

1 Introduction

We are proposing an idea to change the way we read narratives. Most narratives we read on a daily basis have something in common. They reference toponyms (place names), the names of cities, countries, rivers, and islands, and the names of these places themselves carry their own connotations. The name of a Swiss village invokes a mountainous scene, a reference to Siberia implies the frigid temperatures there, the story taking place on a Balinese island suggests a backdrop of Hinduism, and an event in a Canadian state occurred in a first-world economic environment. Familiar place names offer underlying context for the story we are reading. Still, many place names are unfamiliar. If we do not take the time to interrupt the flow of reading to look at these places on a map, we will miss this locational context. If we do stop to seek a map, by the time we get back to the reading, we have lost track of where we were reading. In fact, we may have even forgotten what the story was about.

Our idea to change the way we read narratives combines reading with a dynamic geographic information system. This application would provide a map and corresponding geographically derived information on the fly as the reader encounters toponyms in the text. A device would sense when a person is reading a place name in the text and display a map of the place. Pictures of the site would also appear along with related information, like population, current weather conditions and currency valuation, political climate and dominant religion, or even proximity and route to the nearest airport. The map and geo-information would be introduced in a way that allows the reader to access the supplementary material without losing their place in the text. The map and images would help us to retain and comprehend the story arc. Studies have shown improvement in recall and comprehension of text content when accompanying maps are provided [2, 20, 27]. This could change the way we read online news and e-books, the way school children study political science and geography, the way officers study military history, the way information analysts consume reports, and the way we plan our next vacation.

We demonstrate this idea with a prototype called GazeGIS, a gaze-based reading and dynamic geographic information system. GazeGIS couples geoparsing and Web mapping with an eye-tracking device. Eye-tracking devices have begun to be used for interactive reading systems in other contexts. Our innovation is combining this approach with geoparsing and GIS. To contextualize our contribution and provide a backdrop for our prototype implementation, we discuss related work in geoparsing and eye-tracking (Section 2). Then, we explain the system design and implementation in Section 3. To see how our users interact with our system and how they perceive it, we conducted a study, which we describe in Section 4. Short passages were used in the study so that reading would take only a few minutes. The content in these passages could be considered leisure reading. We were also interested to learn how domain experts would use our system for reading or writing reports to rapidly analyze complex scenarios. As a case study, we obtained a mock report from an analyst in the intelligence community and applied GazeGIS to the report. Section 5 discusses insights from this experience. This is followed in Section 6 by our conclusions and plans for future work.

2 Related Work

Geoparsing algorithms enable place names to be extracted from text and geolocated. One approach for implementing our idea is to combine this geoparsing cabability with eye-tracking technology for eye-based interaction. When discussed in the context of related work, this idea appears to be a natural extension of previous work conducted in these areas.

2.1 Geoparsing

Parsing place names within structured data files such as tables is simply a matter of cleaning the data and discerning the pattern (e.g., it's easy to parse county names, if you know they are in the third column of a table). A different approach is needed for unstructured text like news articles or novels. Advances in natural language processing enable place names within unstructured text to be automatically detected and geolocated. The process of identifying toponym instances within a text and subsequently assigning a coordinate to each name is referred to as *geoparsing*. The first part of this process is implemented with a named-entity identifier and a *gazetteer* (a database of place names and related information). Once something is identified as the name of a place, the toponymic homonym problem may need to

- *Laura G. Tateosian is with the Center for Spatial Analytics at North Carolina State University. E-mail: lgtateos@ncsu.edu.*
- *Michelle Glatz is with the Center for Spatial Analytics at North Carolina State University. E-mail: mlglatz@ncsu.edu.*
- *Makiko Shukunobe is with the Center for Spatial Analytics at North Carolina State University. E-mail: mshukun@ncsu.edu.*
- *Pankaj Chopra is a Research Scientist. E-mail: chopra.pankaj@gmail.com.*

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014; date of publication xx xxx 2014; date of current version xx xxx 2014. For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

be resolved. Toponymic homonyms are different places that share the same names, as in SpringField, OH and Springfield, RI. The geoparsing algorithms usually make some assumption such as proximate locations being clustered within text, or the occurance of other small town names in the text, or if another place name within the same border is mentioned. Failing other cues, they may be ranked by population on the assumption that larger places are mentioned more frequently. In [21], Leidner and Lieberman review the workflow and challenges involved in geoparsing.

In the past, geoparsers were primarily propietery software. Recent open source projects, such as GeoDict, Unlock Text, and CLAVIN, have made this technology more widely available. As a result, researchers have begun adapting geoparsing algorithms to handle specific unstructured media, such as news articles and microtext [12, 15] and specific content domains such as historical texts and classic literature [26, 17]. Though our work does not target a particular media or content domain, the system could be modified to be tailored for specific content.

Geoparsing along with mapping tools have also begun to be conceived of as a means for visualizing collections of documents and multilingual texts [1, 24]. Geoparsing underlies the digital maps project, DIGMAP, which uses a map as part of an interactive user interface for searching digital resources [22]. Our prototype is currently designed to display a single document at a time. In the future, we would like to add the capability to explore databases of documents.

2.2 Eye Tracking

Eye-tracking has a long history of application in both diagnostics and interaction [11]. From the 1970's, eye-trackers have been used in controlled experiments to record eye movements and focus as users perform a cognitive task or inspect a stimulus, such as a work of art, an advertisement, or a Web page. Data visualization and cartographic design choices are also being evaluated in this way [9, 8]. In these experiments, the data is recorded to be analyzed afterward. Whereas, applications for real-time eye movement consumption appeared, from the early 1990's onward, with eye trackers acting as devices for interacting with applications [7, 18]. For example, gaze focus has been used as a pointer to make selections in virtual environments [29] and to act as a typing device for people with a loss of movement [19]. These tools require the user to actively adjust eye movements to control their environment. Applications such as ours that attend to the user's natural eye movements to trigger timely responses are termed as gaze-contingent.

An early gaze-contingent reading application used the reader's gaze path to drive zoom and voice narration for items of interest in the novella, "The Little Prince" [7]. In the mean time, diagnostic eye-tracking systems have been used extensively to study how people read in relation to cognition [25]. This work provided the foundation for several interactive eye-tracking applications that react, in real-time, to what the user is reading and how they are reading it. The GWGazer Reading Assistant is a remedial reading application that uses gaze to determine when the user hesitates over a word and provides assistance by speaking the word [28]. SUITOR harnesses gaze information to infer the reader's interest and automatically finds and displays relevant information in a scrolling display at the bottom of the screen [23]. The iDict provides automatic translation help when users appear to be having difficulty reading in a foreign language [16]. In Text 2.0, words or phrases are associated with sound effects that are played or images that are displayed when this selection is read. Additional information like translations or explanations are presented when the user's attention indicates difficulty. Also, if skimming is detected, the document display is altered such that words that are likely to be skipped are faded [5, 6, 4]. The Text 2.0 system features general aids to reading comprehension. Like our system, these systems also use figures to supplement the text. They do not, however, focus on geographic content.

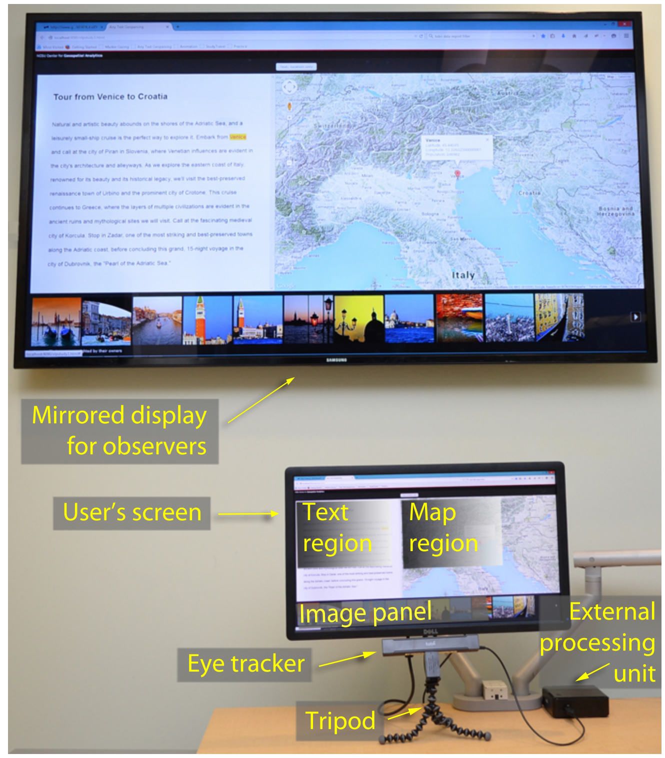

Fig. 1. The physical setup of GazeGIS in the lab: A seated reader views the application on a desktop monitor with an eye tracking device centered below the screen. The wall-mounted display mirrors the primary display, allowing others to observe the interactions.

3 GazeGIS 1.0 Design and Implementation

We implemented our idea in an application we call GazeGIS. GazeGIS consists of a gaze-contingent Web page application for reading text documents while an eye tracking device computes the reader's point of gaze. We use a layout widely employed within military history books where the story often hinges on the location of commanders and troops [14, 10]. In books such as this, passages of text are accompanied by a corresponding map on each facing page, so that the text can refer to important locations pictured opposite. Analogously, we split the main part of a web page with reading on the left and a map on the right. Of course, in the digital medium, we can use a single Web map and move or mark the map to reflect place names being referenced. To supplement the maps, we also display photographs to provide some information about the appearance of the named places. We display the images in a panel that runs across the bottom of the screen, beneath both the text and map. This approach follows a popular design for Web pages such as Esri Story Maps, which display text, maps, and images [13].

The resulting GazeGIS web page is divided into three regions, with a box for text on the left, a box for a map on the right, and an image panel along the bottom. The system augments the reading content with geographic information in the form of maps and images to orient readers to the local scenery and geographic location of a place just as their eyes encounter the location's name within the text. When the system detects that the user's gaze has reached a place name within the text, the map pans to display and tag that location, and photos taken near that location appear in the image panel. At the same time, the place name within the text is highlighted, so that readers may inspect the map and images without losing their place. Gaze-contingent behavior was only implemented for the text region of the web page, not for the map and image panels, as we found that this could disrupt the flow of reading. The reader can pan and zoom the map or browse the images with the mouse. More details on hardware, software, and gaze interaction are described next.

3.1 Hardware

The primary hardware component of GazeGIS is an X2-60 Compact Eyetracker from Tobii Technologies. This 60HZ eye tracking system is mounted below a 27-inch desktop monitor with 1920 x 1080 resolution. The tracker comes with adhesive for mounting on the monitor frame. We choose not to use this method of attachment in order to preserve portability. Our solution is to hold the tracker in place with a flexible cell phone tripod, a Gorillapod Grip. Mirroring the screen on a secondary overhead 65-inch Samsung wall-mounted monitor facilitates demonstrations for groups and allows researchers to observe usage unobtrusively (Fig. 1).

3.2 Software

We combined several open source software libraries and proprietary APIs to create GazeGIS. For geoparsing, we selected the open-source package, CLAVIN-REST, for its modifiability and ease of use with Web mapping packages and other services (CLAVIN stands for Cartographic Location And Vicinity INdexer). CLAVIN-REST extracts place names from unstructured text using an entity extractor and gazeteer (by default the GPL licensed Stanford Named Entity Recognizer and the GeoName world gazeteer). Additional information is returned along with the place names (latitude/longitude, population, alternative names by different languages, etc.). Fuzzy matching is used to find misspelled place names and context-based heuristics are used to disambiguate topographic homonyms. The heuristics rely on the population of a location, the geographic proximity to other place names in the text, the text-based distance to other place names within the text, and political boundaries [3].

To visualize the geographic information extracted from the text, GazeGIS uses a Web map and an online geotagged photo database. There are a number of web mapping APIs, both proprietary and open source. GazeGIS is built using a Google API which has abundant documentation. For imagery, we compared Flickr and Panoramio. Though Flickr hosts a greater number of images than Panoramio, Flickr also contains many personal images that are not fitting for our purposes. Hence, GazeGIS subscribes to Google's Panoramio, a geolocation-oriented photo-sharing site hosting millions of geotagged images. To receive related image feeds for selected toponyms, we filter images with a bounding box centered on a toponym's coordinates.

To communicate between eye tracker and browser, GazeGIS uses an open source Text 2.0 Framework package, named gaze.io [6]. This interface enables event notification based on the user's gaze position. The package provides several functions to obtain gaze data from the eye tracker. We invoke a function to report an event as soon as the reader encounters a place name (the reader is not required to dwell on the location name). A small trigger radius around the element is specified to avoid simultaneous gaze-over event triggering for proximate place names within the text layout.

3.3 Gaze Interaction

The GazeGIS display updates are driven by user gaze feedback. Our javascript/HTML Web page application first allows the user to select a text document to read. The system then geoparses the selected document to identify toponyms and creates an HTML version of the document contents, inserting HTML tags within the HTML to mark the toponyms. Then, the HTML document is loaded into the Web page application and the preprocessing is complete. Next, the feedback loop commences. The eye tracking system analyzes the infrared video image of the eye and computes the coordinates of the gaze-point (the screen position of the viewer's gaze) and sends this to our GazeGIS application. The application tracks the reader's gaze as it passes over the text. When the user‘s gaze reaches a toponym identified by the geoparser, the display is updated (Fig. 2). The place is marked on the map, local imagery is displayed in an image panel in the browser, and the place name is highlighted within the text. This toponym remains highlighted until the reader returns to the text and then encounters another toponym within the text. This behavior is designed to enable users to inspect the map and images without losing their place in the reading.

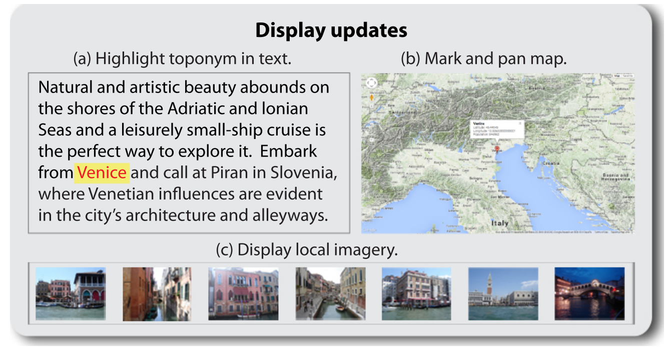

Fig. 2. When a viewer's gaze falls on a toponym, the GazeGIS display updates. A video demonstration is available at `http://tinyurl.com/p5qy5q9`.

4 User Study

We wondered how this idea would be received by users and how they would use the system. Here we explore some preliminary questions. Would readers enjoy the experience of reading with GazeGIS or would they find the map and images distracting? Would they spend time looking at the maps and images or just focus on the text? Would highlighting and mapping toponyms assist in recall of these place names? To answer these questions, we conducted a study in which we invited participants to try the system.

For our study, we selected two text passages for participants to read within our application. The first passage, S1, describes a historical event and the second one, S2, discusses travel. We selected topics to be accessible for a general audience and portray real-world usage scenarios. S1 and S2 are also suitable for the study due to their brevity and because they contain a few place names that might not be familiar to southeastern U.S. undergraduate students. S1 is a 6 sentence, 144 word, excerpt from a Wikipedia page, describing Paul Revere's midnight ride to warn of the British Army approach. This refers to six place names in the greater Boston area (Boston, Charles River, Lexington, Somerville, Medford, and Middlesex County). S2 is a passage from a travel brochure, touting a Mediterranean cruise docking in Italian and Croatian ports. S2 refers to 12 distinct place names, two of which appear twice in the passage, yielding a total of 14 tagged toponyms (Venice, Croatia, Adriatic Sea, Piran, Slovenia, Italy, Urbino, Crotone, Greece, Korcula, Zadar, and Dubrovnik, with Venice and Adriatic Sea appearing twice).

Fifty-four Parks, Recreation, and Tourism Management students ranging in age from 18 to 26 and one professor (over 45) participated in the study. Twenty-eight of our participants were male and 27 were female. Only three of the participants had previously used an eyetracker, five had GIS experience, and eight readers had traveled to the Mediterranean.

When the participants arrived for the study they were welcomed to our laboratory and given an explanation of the purpose of the study and a brief description of our application. Once they signed the informed consent agreement, they were asked to complete a pre-study biographical data sheet. Next, the participants were seated in a non-adjustable, non-swivel, stationary chair and a nine-point eye tracker calibration was explained and executed. After this, the participants were told that once the eye-tracking session started, instructions about how to proceed would be displayed on the monitor. They were also told that once they finished with each step, they should use the 'Esc' key to proceed to the next screen. Participants were asked to relax, read at their own pace, and look at whatever interested them. They were not given a time limit. The first element that appeared was a set of instructions. Next, passage S1 was displayed in our GazeGIS application in the text panel. When a user decided to proceed, another set of instructions was displayed and finally, passage S2 was displayed. When they finished experimenting with our system, they were given a survey to rate their experience and their perception of the application. They were asked to assign a value from 1 (strongly disagree) to 5 (strongly agree) for each question. The participants were also asked questions to assess their recall about the information presented in passage S2.

4.1 User preference evaluation

In our post-session survey we also asked questions to determine whether displaying spatial information related to the text was perceived as helpful. In particular, did the participants like the content and display format and was the dynamic update based on what was being read positive or distracting. The results of these survey questions for our 55 participants are shown in Table 1. Participants liked the map display and image panel and thought the content was appropriate. In addition, the dynamic screen update with geographic information relevant to the text being read was thought to be helpful and not distracting. Participants expressed some interest in being able to control the reading mode by toggling the dynamic behavior. Overall, the application was very positively received by the participants.

Table 1. Mean survey results, μ, with standard deviation, σ. Responses range from 1 (Strongly Disagree) to 5 (Strongly Agree).

Question	μ	σ
Integrating the eye tracker enhanced the reading experience.	3.89	0.71
The dynamic display of geographic information was helpful.	4.09	0.69
The geographic detail was displayed at the appropriate time.	3.73	0.92
The content of geographic detail was appropriate.	4.20	0.44
The display of geographic information was distracting.	2.35	0.86
I liked the image panel.	4.15	0.70
I liked the map display.	4.07	0.76
I would have liked a way to disable/enable the dynamic display of geographic details.	3.13	0.87
If this was an app, I would use it.	3.93	0.83

4.2 Toponym recall

As part of the post session survey, we asked the participants to list the place names they remembered from the readings. The participants had no prior knowledge that they were going to be asked for this information. This question was asked to assess if displaying geographical information would help participants to recall the names of the places they were reading about. The average recall rate of the place names that were highlighted and displayed on the map with GazeGIS was 18% with some participants recalling one-half to two-thirds of these locations. Additional research involving a control group is needed to understand how dynamic geographic displays influence retention. Asking followup questions that require mentally accessing information made available by maps, such as route planning or judging relative distances, may reveal if this design has a positive effect on geographic recall.

4.3 Gaze patterns

To investigate how participants interacted with the system, we used the Tobii Studio screen capture function to record study participants' eye movements during the study reading. Through this, we observed some interesting behaviors. Some readers peered at the map and images each time they updated. A small minority of readers seem to focus only on the reading and only looked at the map and image features after reading. Some users seemed to read through the text a second time to 'play' with the interface and review the places of interest. Though fixation analytics are not the focus of this paper, analyzing the sequence and duration of gaze visits to the text, map, and image panel regions may reveal additional information about usage patterns.

5 Case Study

Our user study applied GazeGIS to educational and leisure-related documents and user feedback was positive, lending support to the potential for GazeGIS to be adopted for day-to-day text consumption. However, these casual uses are only a part of our focus for developing GazeGIS. We are also interested in how our system can support critical sense-making tasks. For example, intelligence community analysts study events in which only part of the story is known. As a crucial component for interpreting events, analysts compose a narrative report to share with fellow analysts. To explore this capability, we used a report on the Malaysian Airlines Flight (MH-17) that exploded over Ukrainian airspace in July 2014.

The MH-17 report models work that would be used in agencies concerned with intelligence research. The plane was shot down, but it was not definitively known by whom or with what intention. (Additional information has since been uncovered, but this was an initial report.) The report provides background on the region's political climate in the months leading up to the event. The narrative encodes key geographic content in both toponyms and latitude/longitude coordinates. Location is also given in relative terms. E.g., contact was lost with the MH-17 plane 50 km from the Russian-Ukraine border. Places-of-interest (e.g., airports) and aircraft direction and speed are also reported.

The implications of this report are complex. A number of the spatial elements and other components have visual analogues that could aid analysts in rapidly reading or writing a report expressing such complexities. To discuss these ideas with expert users, we loaded the simulated report into GazeGIS and showed it to potential users from the intelligence community. These experts found GazeGIS appealing for its potential to streamline workflow. For example, an analyst under time pressure would not need to interrupt report writing or reading to look up the location of an unfamiliar city or to map a geographic location given in latitude/longitude coordinates. As demonstrated in the MH-17 case, this information can be key to drawing conclusions about events. Where the debris was found, where ground control lost contact and boundaries of the no-fly zone were specified in geographic coordinates. In these cases, mapping these points can't simply be bypassed as a casual reader might do. The expert users also reinforced the utility of automatically displaying pertinent imagery, scenes of the places being discussed. Sourcing in-house image databases would also be advantageous for their domain application. They also said that mapping non-geographic location names, places-of-interest such as airport or museum names would be informative. There was also interest in adding more GIS features. They suggested displaying an overview map alongside the main map. They also suggested overlaying additional GIS layers to the map display. As an example, cloud cover and other weather conditions on the day of the MH-17 incident could influence opinions of whether the shooters intended to target a commercial airline or thought it was a Ukrainian military plane.

This idea of a novel means for interacting with narrative and GIS was well received by these experts. The detailed feedback we received implies that they are readily able to envision using this system to their advantage. Encouraged, we have begun to investigate some extensions to our system that could incorporate these suggestions.

6 Conclusions

We have presented a novel idea to change the way we read narratives by using attention-contingent displays to augment the text with geographic information. We implemented the GazeGIS tool using geoparsing and eye tracking technology to dynamically update map and image displays to provide a reader with easy access to pertinent geographic information. User experience feedback from laymen was strongly positive.

We exposed our system to expert analysts and studied a compelling use case scenario. These experts indicated interest in using the system for their work. Additional information and visual analytics can be added to the display to support specific domains. In this vein, some directions for future work are adapting the system for geographic literacy development and other educational purposes, for travel planning, or to facilitate efficiently summarizing the spatial elements of a document and generating maps for reports. Geographic context is fundamental for comprehending the nuances of the narratives we routinely encounter in our consumption of ebooks, news, and other electronic reading. With systems like GazeGIS, gaze-contingent behavior can provide pertinent geovisualizations in a timely and convenient fashion.

Acknowledgments

The authors wish to thank Tonya Adelsperger for her significant contributions to our research.This material is based upon work supported in whole or in part with funding from the Laboratory for Analytic Sciences (LAS). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government.

References

[1] R. Abascal-Mena, E. Lopez-Ornelas, and J. S. Zepeda. Geographic information retrieval and visualization of online unstructured documents. *International Journal of Computer Information Systems and Industrial Management Applications*, 5:089–097, 2012.

[2] R. R. Abel and R. W. Kulhavy. Maps, mode of text presentation, and childrens prose learning. *American Educational Research Journal*, 23(2):263–274, 1986.

[3] Berico Technologies. Open source software for geotagging unstructured big data - CLAVIN. All Things Open, 2013.

[4] R. Biedert, G. Buscher, and A. Dengel. The eyebook–using eye tracking to enhance the reading experience. *Informatik-Spektrum*, 33(3):272–281, 2010.

[5] R. Biedert, G. Buscher, S. Schwarz, J. Hees, and A. Dengel. Text 2.0. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, pages 4003–4008. ACM, 2010.

[6] R. Biedert, J. Hees, A. Dengel, and G. Buscher. A robust realtime reading-skimming classifier. In *Proceedings of the Symposium on Eye Tracking Research and Applications*, pages 123–130. ACM, 2012.

[7] R. A. Bolt. A gaze-responsive self-disclosing display. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 3–10. ACM, 1990.

[8] A. Brychtova and A. Coltekin. An empirical user study for measuring the influence of colour distance and font size in map reading using eye tracking. *The Cartographic Journal*, 2014.

[9] M. Burch, N. Konevtsova, J. Heinrich, M. Hoeferlin, and D. Weiskopf. Evaluation of traditional, orthogonal, and radial tree diagrams by an eye tracking study. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2440–2448, 2011.

[10] D. M. Detweiler and D. Reisch. *Gettysburg: the story of the battle with maps*. 2013.

[11] A. Duchowski. *Eye tracking methodology: Theory and practice*, volume 373. Springer Science & Business Media, 2007.

[12] C. DIgnazio, R. Bhargava, E. Zuckerman, and L. Beck. Cliff-clavin: Determining geographic focus for news. *NewsKDD: Data Science for News Publishing, at KDD*, 2014, 2014.

[13] I. Environmental Systems Research Institute. Story maps: Everyone has a story to tell. `http://storymaps.arcgis.com/en/`. Accessed: 2015-07-24.

[14] V. J. Esposito and J. R. Elting. *A military history and atlas of the napoleonic wars*. Greenhill Books, 1999.

[15] J. Gelernter and S. Balaji. An algorithm for local geoparsing of microtext. *GeoInformatica*, 17(4):635–667, 2013.

[16] A. Hyrskykari, P. Majaranta, and K.-J. Räihä. From gaze control to attentive interfaces. In *Proceedings of HCII*, volume 2, 2005.

[17] L. Isaksen, E. Barker, E. C. Kansa, and K. Byrne. Gap: A neogeo approach to classical resources. *Leonardo*, 45(1):82–83, 2012.

[18] R. J. Jacob. What you look at is what you get: eye movement-based interaction techniques. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 11–18. ACM, 1990.

[19] R. J. Jacob and K. S. Karn. Eye tracking in human-computer interaction and usability research: Ready to deliver the promises. *Mind*, 2(3):4, 2003.

[20] R. W. Kulhavy, W. A. Stock, M. P. Verdi, K. A. Rittschof, and W. Savenye. Why maps improve memory for text: The influence of structural information on working memory operations. *European Journal of Cognitive Psychology*, 5(4):375–392, 1993.

[21] J. L. Leidner and M. D. Lieberman. Detecting geographical references in the form of place names and associated spatial natural language. *SIGSPATIAL Special*, 3(2):5–11, July 2011.

[22] J. Machado, G. Pedrosa, N. Freire, B. Martins, and H. Manguinhas. User interface for a geo-temporal search service using digmap components. In *Research and Advanced Technology for Digital Libraries*, pages 483–486. Springer, 2009.

[23] P. P. Maglio and C. S. Campbell. Attentive agents. *Communications of the ACM*, 46(3):47–51, 2003.

[24] B. Pouliquen, M. Kimler, R. Steinberger, C. Ignat, T. Oellinger, K. Blackler, F. Fuart, W. Zaghouani, A. Widiger, A.-C. Forslund, et al. Geocoding multilingual texts: Recognition, disambiguation and visualisation. *arXiv preprint cs/0609065*, 2006.

[25] K. Rayner. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin*, 124(3):372, 1998.

[26] C. Rupp, P. Rayson, A. Baron, C. Donaldson, I. Gregory, A. Hardie, and P. Murrieta-Flores. Customising geoparsing and georeferencing for historical texts. In *Big Data, 2013 IEEE International Conference on*, pages 59–62. IEEE, 2013.

[27] N. H. Schwartz and R. W. Kulhavy. Map features and the recall of discourse. *Contemporary Educational Psychology*, 6(2):151–158, 1981.

[28] J. L. Sibert, M. Gokturk, and R. A. Lavine. The reading assistant: eye gaze triggered auditory prompting for reading remediation. In *Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 101–107. ACM, 2000.

[29] V. Tanriverdi and R. J. Jacob. Interacting with eye movements in virtual environments. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 265–272. ACM, 2000.