Literature Review

INTRODUCTION

The Digital Humanities is a field of study that came about in the late 1980s, concerned with handling vasts amounts of cultural heritage data - literary texts being the main form of this data. Just like with other forms of research, humanities research is having to adapt to a new “frontier” - a digital one that is. Google searches, online library databases, and email are all becoming increasingly relied upon as humanities research has adapted to work with new, modern outputs.

A subset of the Digital Humanities, computational literary studies (CLS) aims to digitally mine through texts in order to draw conclusions about genre, form, meaning, the evolution of words, etc (Da). Whereas before, these sorts of claims could only arise after people analyzed individual texts for themselves, the Digital Humanities aims to provide empirical evidence to human analysis, mining through entire corpuses at once. There are many different methods employed by CLS, any of which are starting to be used to enhance the digital humanities.

TEXT ANALYSIS

Text Analysis is one of the first, and most basic components of CLS. Text Analysis is about parsing through texts in order to obtain machine-readable data (https://www.ontotext.com/knowledgehub/fundamentals/text-analysis/). The ultimate goal is not necessarily for a computer to understand a text, but to gain specific information from the text and for the researchers to form conclusions from that extraction. However, this can be difficult because computers do not come loaded with the background knowledge that humans do when dealing with language. For example, if someone were to say, “The Panthers upset the Dolphins,” many people would recognize that the individual is talking about a football game. However, a computer, without having the same background information, could return several linguistically valid interpretations. Therefore, text analysis is a complicated method with a long history of development in the field.

The root of text analysis predominantly begins with psychological studies. One of the first examples of a type analysis applied to text was a study in the ‘80s by Walter Weitrub, a physician (Tausczik, 2010). Over the span of a decade, he hand-counted the pronouns that people would use in a speech or written text. From there, he concluded that the number of first-person singular pronouns a person would use (I, me, my, etc) were “reliably linked to people’s levels of depression.” Though this was one of the first reliable examples of a type of text analysis, it is also a rudimentary one. Now, with the advancement of technology and computer programs, methods of text analysis have progressed.

Text analysis tends to tokenize the text, or take the individual words, and then analyze the parts of speech of each token (ie is it an adverb/article/noun/etc?). Essentially, the document is transformed into data to then reach a greater conclusion about the individual text, and then multiple documents in comparison to each other. Alan Liu, in his article *Literary Studies in the Digital Age* concludes that “computationally assisted text analysis… is a way to experiment with literature to bring out, among other features, its latent social network and that of the characters in its imaginative worlds… After all, a concordance represents how even disjunct speakers share a sense of a word and so conjoin in a discursive structure that images a larger social structure” (2013).

METHODS OF NLP

The primary method that has been employed within CLS is Natural Language Processing (NLP). NLP differs from Text Analysis because, as opposed to just gathering the text as data, it performs a linguistic analysis on that text: helping a computer to “read” the data. NLP is a branch of artificial intelligence that helps computers learn how to understand, manipulate, and interpret human language. It is one of many strategies now being employed by CLS in order to branch the gap between the humanities and STEM fields.

One task within NLP is Semantic Textual Similarity (SLS), an evaluation tool for determining how similar two pieces of work are. Wang, Costellón, and Comelles explored STS in their study, which used STS to analyze the same dataset in two different ways: semantically and syntactically. They incorporated chunking, sentiment analysis, topic recognition (discussed further below), and the *hyponym-hypernym* structure in order to run their code. They used these different categories to assign a similarity score to various sentences, attempting to train computers to understand different phrases. They concluded that computers are very sensitive to lexical similarities. In another study, Underwood, Bamman, and Lee looked at the transformation of gender in English Fiction novels. Similarly to Wang, Costellón, and Comelles, they used syntax measurements. However, they were attempting to draw their own analysis based on those measurements, as opposed to having the computer decipher the meaning itself. Therefore, the authors used a different kind of NLP, graphing both the words correlated with female characters vs. male characters, and graphing the impact of the gender of the author on the novel’s characters, throughout time. The graphs allowed the authors to visually see how much space was provided to female characters in books, how many bestselling novels were written by men, the most common words used to characterize each gender, etc… After they mapped out associated terms over time, they drew a conclusion about how gender has evolved.

Their methods allowed them to quantitatively quickly recognize historical variation in apparently stable categories like gender.

CRITICAL VISUALIZATION

Another method within CLS is data network visualization, which provides a visual component to the relational network of texts. Critical visualization is most useful when working with multiple texts or online databases, emphasizing the relations between large corpuses. In *On Close and Distant Reading in Digital Humanities: A Survey and Future Challenges*, Janicke, Franzini, Cheema, and Schuermann explore different visualization approaches and analyze them as tools of distant reading (viewing global features of texts), a concept introduced by Franco Moretti. They explain that, “A visualization displays summarized information of the given text corpus, facilitat[ing] distant reading. The process of transforming such information into complex representations can be based upon a large variety of data dimensions” (2015). Some of these transformations can be heat maps, timelines, tag clouds, and graphs. Heat maps are frequently the method used to display textual patterns, timelines are used to chart temporal information, tag clouds encode the number of occurrences of a word in a given selection, and maps represent some sort of geographical space (whether that be the distance between the fictional physical spaces mentioned in the book, or locations of each author, etc).

DATA TRANSFORMATION METHODS

Critical Visualization is, at its core, a method to make information easier to comprehend. However, in order to successfully create a data network visualization, the raw data itself (ie the corpus of texts) must be pre-processed in order to be inputted into the visualization network.

Janicke, Franzini, Cheema, and Schuermann conducted a second study (2017) on different types of visualizations and concluded that the most common transformations are:

1. Tokenization and Normalization - types of NLP
2. Sequence Alignments
3. Part-of-Speech Tagging (POS)
4. Named Entity Recognition (NER)
5. Topic Modeling

They go on further in their article to explain how these transformations can be combined with close reading techniques, and what are the critical aspects in developing each type of visualization.

NETWORK ANALYSIS

Other types of computation - which often combine NLP and visualization - can also be relied upon when studying the humanities. A study conducted by David Elson, Nicholas Dames, and Kathleen McKeown used network analysis, a method of CLS that is rare, even for standards of the digital humanities. To reach their goal, a better understanding of the representation of social interaction in nineteenth century European fiction, they relied upon statistical methods of analysis in order to test 60 novels and build a social network. network, which focuses on conversations, “Vertices represent characters (assumed to be named entities) and edges indicate at least one instance of dialogue interaction between two characters over the course of the novel. The weight of each edge is proportional to the amount of inter- action” (2010). They then went on to conclude, with 95% precision, that the narrative voice (ie first or third person) trumps the setting of the novel when determining various features of a social network. Although their work was exactly what Da criticized when she explained,

> Capturing this kind of complication—or capturing network complexity by studying a network whose degree distribution of nodes to links is neither arbitrary nor regular but obeys some other mathematical law—is not the same as saying that network diagrams

of who is talking to whom in Shakespeare can capture the complexity of connections in Shakespeare or character discourse.

However, this study still shows a potential path in CLS that could be refined and developed in the future.

Another example of constructing a social network using CLS methods is Peter S. Bearman and Katherine Stovel's *Becoming Nazi: A Model for Narrative Networks"* (2000). They focus on first-person narratives from 1934 of people explaining why they became a Nazi. They mapped each narrative by coding each distinct element of the story as a node, and then representing the author's explicit connections between those distinct elements as arcs to connect the nodes. By mapping out the narrative, they were able to recognize that the narrative about *becoming* a Nazi was twice as dense as the narrative of *being* a Nazi. Their visual representation showed that becoing a Nazi is locally dense, globaly sparse, knotty, and redundant. By comparing this to the much less dense, and much less connected social network of *being* a NAzi, the authors conclude that "the self disappears. Nothing is left to operate the story, except for a mechanical agency." Their computational analysis of the social network allowed the authors to tap into a unique part of Nazi narratives.

TOPIC MODELING

Topic Modeling, a technique as previously mentioned is frequently used in data visualization and in NLP, gauges rhetorical or thematic patterns that influence a document. At Rutgers University, Andrew Goldstone and Ted Underwood used topic modeling to review thousands of articles of literary scholarship (2014). They used *probabilistic topic modeling*, creating an algorithm seeking to infer meaningful groups of words statistically. They admit that building the algorithm requires subjective analysis on their part - as Da would criticize - but then explain that is actually *why* an algorithm is valuable; their assumptions are explicit in their parameters. By analyzing topics across research articles, and then showing the change in frequency of certain topics over time, the authors were able to reveal “quiet changes… that can slip under the radar for many reasons… but can turn out to be both numerous and intellectually significant.” They conclude that topic modeling can be a useful tool to articulate changes in scholarship that might be too gradual to have previously noticed.

SENTIMENT ANALYSIS

Sentiment Analysis, another method discussed previously, is an additional way in which researchers can computationally analyze texts. When analyzing a work published by Anderson and McMaster in 1982, Evgeny Kim and Roman Klinger concluded that sentiment analysis is necessary because the “emotional tone of a story can be responsible for the reader’s interest” (2018). They further discussed the implications by saying that Anderson and McMaster’s “suggested that by identifying emotional tones of text passages one can model affective patterns of a given text or a collection of texts, which in turn can be used to challenge or test existing literary theories.”

There are two major types of sentiment analysis: document sentiment analysis and fine-grain sentiment analysis (Kim and Klinger, 2018). Document sentiment analysis aims to classify an opinion document as holding either a positive or negative opinion. It makes the implicit assumption that any document is from a single opinion holder and expresses opinions on one entity (Liu, 2015). By gauging the opinion, document analysis aims to understand the underlying positive or negative sentiment implied by the opinion. Fine-grain sentiment analysis, on the other hand, can extract multiple phrases and assign a polarity scale to the sentiment (Kim and Klinger, 2018). Therefore, it can be more informative than the one-note document analysis.

Although Sentiment Analysis has been used rather frequently within CLS, it is a multifaceted problem with many different components (as discussed in the chapter *The Problem of Sentiment Analysis,* in the book *Sentiment Analysis* by Bing Liu). It once again relies on a subjective measure, however, instead of using subjective measures to gauge objective fact (like the topic of a certain passage), it is attempting to gauge another subjective measurement: the *opinions* expressed by the text. This creates another layer to the challenge of a quantitative approach to the humanities, when one implicitly encodes bias. In *Examining Gender and Race Bias in Two Hundred Sentiment Analysis Systems* (2018), the authors display how NLP can perpetuate biases and vary in quality of analysis based on the race or gender of the speaker. The authors compiled 8640 English sentences and then assign an intensity score to the sentiment (ager, joy, fear, sadness, and sentiment/valence intensity) of the sentence, making sure that each sentence contained one race or gender-associated word. They discovered that in both gender and race, the systems would consistently give higher scores (ie more intense) scores to females, and African Americans, with African American females being perceived as the most intense in anger, fear, and sadness of the group.

The risk of sentiment analysis as a method is perhaps the most obvious of this review. By grading associated emotions within a sentence, machines can be coded to quantify their bias, scoring different races and genders unfairly.

CRITICISM AND RESPONSES TO NLP

However, CLS is not without its detractors. Nan Z Da, an English professor at the University of Notre Dame, explains in her paper "The Computational Case Against Computational Literary Studies" that NLP fails to provide meaning to texts that human readers can, it simply looks for patterns throughout texts. Within the study, she walks through various studies in which the authors utilized NLP to study different texts and concludes that, "Computational literary criticism is prone to fallacious overclaims or misinterpretations of statistical results because it often places itself in a po- sition of making claims based purely on word frequencies without regard to position, syntax, context, and semantics." She even points to one study, which analyzed changes in Augustine's *Confessions,* that showed that the language of 3 of its chapters is statistically different than the other chapters. She responds though that quantitative approach fails to take into account that Augustine completely shifts his subject, meaning that inherently different words will show up. However, this begs the question: would it be possible to run some sort of NLP program on something other than the primary text itself; for example, could one mine through text of secondary sources, collecting data on the *analysis* of the primary text, as opposed to the primary text itself?

David Mimno touched on this question in his paper "Computational Historiography" (2012). By using JSTOR as a public database from which to mine data, he concluded that it was possible to parse through secondary sources using the bag-of-word technique, an algorithm that analyzes the frequency of unique words as opposed to the sequence or order of words. Although he conceded that this sort of "distant reading" does not match the depth of analyzing primary texts, he suggested that the two methods might not be mutually exclusive; both "close" reading and "distant" reading could help a reader gain a better understanding of a body of literature.

Mimno's technique ran into issues with data mining across translations, which he claimed was one of the "key challenges of working with humanities scholarship." Yuri Bizzoni, Marianne Reboul, and Angelo del Grosso, responded to this by comparing over 150 French translations of the *Odyssey*, using a novel text analysis method called text alignment, gauging how far the translations vary from the original Greek text. By examining the translations digitally, they were able to provide a quantitative analysis to back qualitative theories; their findings emphasized how omissions and consistencies can emphasize culture shifts over time and sociological trends.

CONCLUSION

These are just some of the pioneering studies in Computational Literary Studies. Each can provide a lens to see how the digital humanities is a field that needs to be further explored. With the possibility of digitally mining corpuses of texts, the humanities can reach new corners in research, combining data and quantitative analysis, what has previously only been thought of as a STEM field. Methods that combine both close reading with distant reading offer seemingly the greatest potential, drawing upon the more traditional approaches in the humanities but combining it with methods of study most typically associated with other disciplines.

Bibliography

1. Bearman PS, Stovel K. Becoming a nazi: A model for narrative networks. Poetics. 2000;27(2):69-90. https://www.sciencedirect.com/science/article/pii/S0304422X99000224. doi: 10.1016/S0304-422X(99)00022-4.

2. Bizzoni Y, Reboul M, Grosso AD. Diachronic trends in homeric translations. *Digital Humanities Quarterly*. 2015;11(2). https://gup.ub.gu.se/publication/227104.

3. Da NZ. The computational case against computational literary studies. *Critical Inquiry*. 2019;45(3):601-639. http://dx.doi.org/10.1086/702594. doi: 10.1086/702594.

4. Elson DK, McKeown K, Dames NJ. Extracting social networks from literary fiction. . 2010.

5. Goldstone A, Underwood T. The quiet transformations of literary studies: What thirteen thousand scholars could tell us. *New Literary History*. 2014;45(3):359-384. https://muse.jhu.edu/article/558875. doi: 10.1353/nlh.2014.0025.

6. Jänicke S, Franzini G, Faisal Cheema M, Scheuermann G. On close and distant reading in digital humanities: A survey and future challenges. . 2015.

7. Jänicke S, Franzini G, Cheema MF, Scheuermann G. Visual text analysis in digital humanities. *Computer Graphics Forum*. 2017;36(6):226-250. https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12873. doi: 10.1111/cgf.12873.

8. Kim E, Klinger R. A survey on sentiment and emotion analysis for computational literary studies. . 2018. https://arxiv.org/abs/1808.03137.

9. Kiritchenko S, Mohammad SM. Examining gender and race bias in two hundred sentiment analysis systems. *Computer Science*. 2018. https://nrc-publications.canada.ca/eng/view/object/?id=c0ba249c-450d-4a6a-b313-0c9168998b8e

10. Liu A. From reading to social computing. . 2013.

11. 1. Liu B. *Sentiment analysis.* New York: Cambridge University Press; 2015. http://bvbr.bib-bvb.de:8991/F?func=service&doc_library=BVB01&local_base=BVB01&doc_number=027715231&sequence=000002&line_number=0001&func_code=DB_RECORDS&service_type=MEDIA.

12. Mimno D. Computational historiography. *Journal on Computing and Cultural Heritage (JOCCH)*. 2012;5(1):1-19. http://dl.acm.org/citation.cfm?id=2160168. doi: 10.1145/2160165.2160168.

13. Tausczik YR, Pennebaker JW. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*. 2010;29(1):24-54. https://journals.sagepub.com/doi/full/10.1177/0261927X09351676. doi: 1.1177/0261927X09351676.

14. Underwood T, Bamman D, Lee S. The transformation of gender in english-language fiction. *CA: Journal of Cultural Analytics*. 2018. http://gateway.proquest.com/openurl?ctx_ver=Z39.88-2003&xri:pqil:res_ver=0.2&res_id=xri:ilcs-us&rft_id=xri:ilcs:rec:abell:R05730837.

15. Wang C, Castellón I, Comelles E. Linguistic analysis of datasets for semantic textual similarity. *Digital Scholarship in the Humanities*. 2019. doi: 10.1093/llc/fqy076.