*Healthcare and Medical Devices,* Vol. 79, 2023, 181–190
https://doi.org/10.54941/ahfe1003486

# Team-Based Text Analytics for Health Information Systems Learning

**Tim Arnold[1,2], Helen J. A. Fuller[1], Serge Yee[3], Ruth M. Reeves[1,4,5], and Seema Nazeer[3]**

[1]Human Factors Engineering and Health Solutions Management, Clinical Informatics and Data Management Office, Office of Health Informatics, Veterans Health Administration, Department of Veterans Affairs, Washington, DC, USA
[2]University of Michigan, College of Pharmacy, Ann Arbor, MI, USA
[3]Health Solutions Management, Clinical Informatics and Data Management Office, Office of Health Informatics, Veterans Health Administration, Department of Veterans Affairs, Washington, DC, USA
[4]Tennessee Valley Healthcare System, Department of Veterans Affairs, Nashville, TN, USA
[5]Vanderbilt University Medical Center Department of Biomedical Informatics, Nashville, TN, USA

## ABSTRACT

Learning about health information system usability, users, and their work can be challenging. Capturing and examining user perspectives is one approach for learning about users and usability of health information systems. When large numbers of perspectives are collected, readers can feel overwhelmed in trying to interpret the resultant large bodies of text. This phenomenon is sometimes referred to as information overload. In contexts of surplus text, people seek efficiencies for making sense of this information. Traditional strategies to address this challenge require manual thematic analysis and extracting key narratives to exhibit a concept or tell a story. Such approaches, however, may be time consuming, overwhelming, and require special expertise and resources. In this paper, we describe the initial design of a team-based text analytics approach in search of a usable middle ground in visualizing, interpreting, and sharing information from user survey comments about health information systems.

**Keywords:** Text analytics, Natural language processing, Human factors, Health information systems learning

## INTRODUCTION

In healthcare operations, we often rely on surveys to learn about users' experiences with their work systems. While some questions may be posed in structured ways that yield easily quantified results, questions that allow users to enter free text responses often result in more descriptive and nuanced data and thus are commonly used. Making sense of and coming to some shared meanings around the frequently abundant bodies of text from such surveys is often time consuming. A lack of resources and expertise may contribute to hesitation and indecisions on how or whether to analyze this type of text.

 

Because of challenges with analyzing text in operational settings, there may be organizational reluctance to capture narrative comments. Nonetheless, narrative comments can be a source of rich information that, with reliable and faster approaches for analyzing, may help with informing operational decisions and design efforts. In this paper, we describe using text analytics approaches for contributing to thematic analysis of users' comments to help with health information systems learning.

After a recent survey about a health information system, a team composed of members with expertise in human factors, informatics, patient safety, and clinical care explored several text analytic approaches as possible pathways to reduce the burden of reviewing comments about the system. Approaches included topic modelling, keyword extraction and creating word clouds, co-occurring word, and uniquely co-occurring word (u-coord) network visualizations, and text classifiers and nomograms that highlight top linguistic features for the trained classifier. The project team walked through example approaches and visualizations and decided on next steps.

Visualizations of word co-occurrence and uniquely co-occurring word networks and top linguistic features used to train a Naïve Bayes text classifier helped us envision possible class labels. Regular expressions were iteratively formulated consisting of some combination of words and stems as class definitions were formulated and extracts were repeatedly reviewed. Class formulation corresponded with refinement of regular expressions. Individual comments could be multi-labelled and not all comments were classified. Static visuals, text examples, regular expressions, and extract quantities were collected, presented, discussed, and refined with the team.

This work explored text analytic approaches to assist with interpreting a body of comments and to apply filtering techniques for addressing concerns of information overload to reduce hesitation with collecting and examining bodies of text. By reframing the task of interpreting user comment text as a filtering problem, we began to inquire into ways to review, label classes, and classify comments more quickly. With this goal of increasing general utility of this approach, including and fine-tuning text analytics approaches may help teams learn more quickly from survey comments about how users perceive working within health information systems. Finally, lowering thresholds for analyzing text may boost motivation for gathering free narrative text that may offer insight and rich information into vital viewpoints and language use across time.

## APPROACH

Multiple natural language processing (NLP) approaches and visualizations were explored with project team members. We walked through each approach and visualization while conceptualizing themes (Table 1). As we explored possible themes of interest, we reviewed context by reading and discussing comments. This progression helped the team select approaches and visualizations to share with a larger audience of stakeholders that included healthcare managers, informaticians, patient safety professionals, and clinicians.

One reason for walking through approaches and visualizations with the team was to share preferences for how to relay information while choosing

**Table 1.** NLP approaches explored.

| NLP Approach | Shared with Project Team | Planning to Share with Stakeholders | Use | Project Team and Stakeholder Feedback |
|---|---|---|---|---|
| Keyword Extraction | Yes | Yes | Review word usage and frequency | Expected analysis method and useful for overall impressions |
| Word Co-occurrence Network | Yes | Yes | Explore latent groupings and collect features for searching | Preferred visualization with reasonable threshold |
| Text Classification | Yes | Yes | Identify linguistic features most important to the classifier | Helpful for identifying distinctions between predefined classes |
| Corpus Viewer | Yes | Yes | Explore latent classes, review context, capture examples | Preferred visualization for reviewing context |
| Topic Modelling | Yes | No | Find linguistic patterns | Abstract latent topics |
| Word Clouds | Yes | No | Visualize frequency of words | "Busy" and lacked additional context |
| Concordance Viewer | Subset only | No | Visualize keywords in short context | Not initially insightful |
| Sentiment Analysis | Subset only | No | Detect strong emotions in comments | Not initially insightful, results were as expected |

examples that would help larger audiences follow along without distracting from the purpose of reviewing user comments for health system information learning.

Preliminary analysis looked at numerous approaches. Sentiment analysis lacked obvious and rapid insight that could be easily explained to those less familiar with these techniques. We will incrementally revisit these techniques to explore their usefulness for this work in the future. Focusing on techniques that highlighted clear themes, we initially set aside some while narrowing in on a collection of possibilities for further inquiry.

Approaches set aside for revisiting were not excluded but instead were deemphasized to facilitate meeting the immediate need. Topic models, word clouds, and varied word concordances did not resonate initially with the team or raised concerns about interpretability.

Topic modelling is an unsupervised or semi-supervised approach for identifying linguistic patterns in text (Ramage et al., 2009). Team reception of topic models was mixed, with team members reporting ambivalence when interactively exploring a common topic modelling methodology, Latent Dirichlet Allocation (LDA). Team members were concerned that trying to explain topic concepts could distract a larger audience such as stakeholders. Latent classes from unsupervised topic modelling approaches can be abstract and challenging to interpret. As we revisit topic modelling in the future as one possible approach for helping teams and audiences interpret survey comments, we will explore the possibility of a semi-supervised model with the added capability of constraining collections to line up with meaningful label classes (Ramage et al., 2011).

Word clouds are a common approach for visualizing terms in text that can help to summarize and promote visibility of word usage (Chuang et al., 2012). The team had mixed reactions to word clouds, explaining that they felt the visualizations were "busy" and lacked additional context around words for more meaningful interpretation. On the other hand, although in somewhat different contexts, Clough and Sen (2018) have highlighted the usefulness of word cloud visualizations and team members have observed that other analysts have found them helpful for conceptualizing the text while moving in between the parts and the whole when seeking to interpret larger bodies of text. This finding suggests that interactively revisiting the word cloud visualisation option and customizing to the analyst team's preferences for formatting and simplicity would be advisable.

The team found keyword extraction useful for forming overall impressions of word usage and frequency or significance. Moving back and forth between the keyword view and exploring interactive word co-occurrence network representations helps facilitate a form of interpretation. Adding another view, we employed text classification models using a bag-of-words associated Term Frequency-Inverse Document Frequency (TF-IDF) matrix and a Naïve Bayes classifier to identify words or linguistic features most important to the classifier.

We collected distinguishing features, reviewed the results with other NLP views as listed in Table 1, and plugged them into a corpus viewer (Demšar et al., 2013) to investigate within the context of complete comments. The corpus viewer helps with studying latent groupings, designing taxonomic structures for filtering, manually reviewing context, and capturing example stories for each class. The team iteratively designed the report before delivering to stakeholders.

## LESSONS LEARNED AND THEORETICAL CONSIDERATIONS

There are two layers of lessons learned to this work. One layer is interlaced throughout the approach section and describes technical details and specific preferences for the team-based text analytics process itself. The second layer of lessons learned, discussed in this section, considers the design of the team-based approach to NLP, the possibility of evaluating team learning, design frames and perspective (philosophical) base, and how findings can inform next steps.

Before we discuss lessons learned from the current effort, it is first valuable to consider past examples of the use of NLP in healthcare. NLP and text analytics have been commonly used for aiding in interpreting patient safety reports and processing natural language in clinical notes. Mull and

Nebeker (2008) suggested that advancements in NLP will contribute to finer clinical trigger tools. McKnight (2012) designed a semi-supervised classification method to group patient safety reports. Fong and colleagues (2017) described how they integrated NLP expertise into a committee reviewing patient safety reports. Reeves and team (2021) provided an account on how they modified an NLP system for a new healthcare environment to find social determinants of health.

Thus, NLP has been the subject of wide-ranging analytical work in healthcare. Adding to this chronicle, we share lessons learned, design perspectives, and theoretical considerations from and about socializing and operationalizing team-based NLP practices with respect to both their use and their continual health services development to examine user comments for health information systems learning.

## Socialization: Team-Based Approach

Text-based analytics and NLP work seems primarily focused on solving problems and providing some product or service methods, tools, pipelines, or answers to questions. Implementing a text-based analytics design into daily team-based work for operational purposes is less understood. There are multiple purposes for encouraging experiences and gaining skills in this space, and doing so may generate value for the analyst team in different ways. First, formulating work with human teams around NLP techniques secures time for consistent practice, building skills, and contributing to organizational maturity. Second, the process of co-creating analysis methods helps to include multiple perspectives into the process design. This participatory design can help with visibility of system status, explainability, and generating conversations around capabilities and limitations of these NLP methods. Third, coming to each project with cooperative and multi-layered learning in mind may encourage early, consistent, and appropriate adoption of these practices in the development process. Finally, viewing these work efforts as collaborative learning opportunities may also shed light on two important and yet seemingly mysterious disciplines and sets of practices—human factors and text analytics/NLP—both of which are important for studying, designing, and evaluating information work systems. Next, we look at a set of properties for helping to evaluate learning that may occur from team-based exploration of NLP-based visualizations to help classify and interpret text.

## Evaluation: Restructuration Based

To further understand the possibility of learning from this approach, we consider evaluating based on a concept by Wilensky and Papert (2010) called restructuration. Evaluating our approach based on properties of structuration or representational framings used to convey knowledge could be an initial step toward understanding the impacts and risks of implementing text-based analytics into daily team-based work and the downstream consequences. The properties of structuration are *power*, *cognitive*, *affective*, *social*, and *diversity*. We highlight a few of these properties here. *Affective* properties point to favorable aesthetics or to the appeal of the knowledge representations. Initial observations from collective engagement suggest that NLP-based visualizations and network representations have overall positive affective properties, which we can design further based on human factors principles and practices. *Cognitive* properties consider how simple a new structuration is to learn, drawing from connection to existing knowledge. Our approach fits with pre-existing knowledge used to perform thematic analysis based on code book development and categorizing into preselected classes/taxonomies and text classification approaches based on a trained model. We take it one step further by using visualizations and network diagrams (Figure 1) of select linguistic features to aid in the learning process.

## Theory: Classification & Context

According to Rosch (1978), humans have a tendency to classify in order to access maximal information while limiting cognitive effort. Perhaps the number of taxonomies and classification systems in healthcare alone suggests that people find value in the process of creating classification labels and taxonomies or it is a necessary step for making sense of large bodies of text and the text-based representations of knowledge. Also, it may be that the iterative act of classification is required for the process of interpretation by going in between pieces of text and greater context surrounding each piece. When classes are predefined or coders select from existing taxonomies, text is assigned to pre-existing groups. This reduces the work necessary to formulate labels. However, certain boundaries are pre-established and work entails assigning text to these existing classes. But the very nature of pre-defining eliminates the step of formulating class labels, and people are confined to making comments that fit within predefined classes, which requires them to decide what the class means while juggling the ambiguity of language. A constraint on the use of predefined classes is that they may work well for some applications but can be problematic for others.

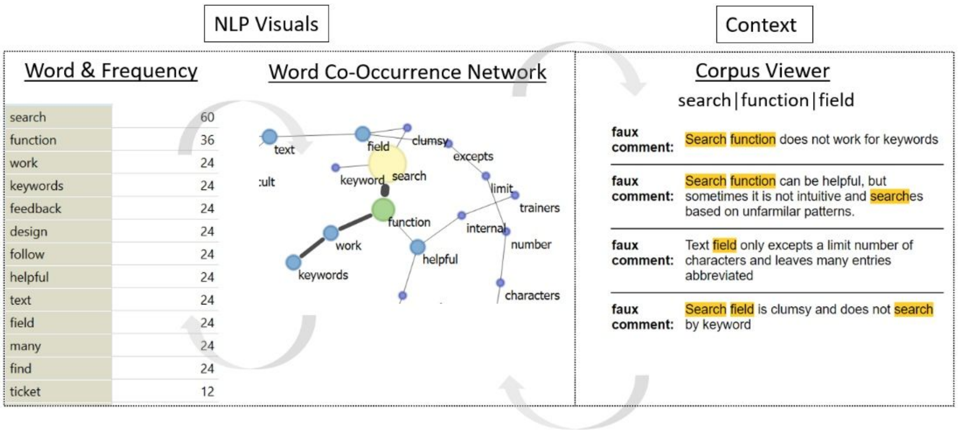


**Figure 1:** Simple depiction of NLP approach for reviewing user comments. Visualizations created using orange: data mining toolbox in python (Demšar et al., 2013).

When formulating class labels around any given context based on NLP visualization, the team examined patterns and came to some momentary agreements on how the team would describe boundaries in what we are referring to as systems. In an NLP-based inductive labelling exercise, label choices were informed by interdisciplinary teams mapping the representations of computational language models to contexts. This process of text interpretation has the benefit of identifying latent patterns that may be missed through other methods. Beyond that, the team simultaneously worked through how to describe something as a system, through fashioning boundaries and describing external environments that will eventually inform decisions and next steps.

## Theory: Systems Thinking

Human tendencies to classify are addressed in systems theory in the central concept of *boundaries*. Boundary placement can be contentious, and there can be many ways observers might conceptualize something described as a system (Mingers, 2014). Rasmussen (2002) walks us through Niklas Luhmann's unity of difference and the three elements "That (1) which is distinguished from something else (2) and the distinction itself (3)." The distinction itself or the boundary conceptualization is vital in determining how a problem is studied and the methodology used, which in turn can inform modelling, design, and evaluation efforts. Pre-specified taxonomies or codes often require some domain adaptation and pre-emptively draws boundaries based on perceptions of a party removed from the context of work. This is in contrast to internal team-based NLP guided analysis of text comments drawn from actual workers in the systems.

These considerations of boundaries and how they are drawn bring to light the importance of interdisciplinary teams in exchanging perceptions of systems boundaries and features of environments that also take place during team-based analysis work. How boundaries are drawn will likely impact information interpretation and communication while influencing the type of and justification for future work.

## Empirical Considerations

Analysis of and sharing information intended to help with deciding on and planning for future work can be challenged and side tracked, sometimes for good reason but sometimes unduly due to differences in interpretations about methods of analysis and perception of data strength. Mingers (2014) draws attention to this topic suggesting the importance of an evidence base while pointing out the limitations of data and analyses. An additional challenge to text analytics and NLP techniques are the audience's overall comfort levels with the methods and results. The methods may seem complex, and certain results appear fuzzy to interpreters, which can trigger initial suspicion and scrutiny. Once data and analytics discussions go down the rabbit hole of methodological argument, rebounding to a state that approaches some shared interpretation or even explainability can be difficult or next to impossible. Also, people sometimes prefer the convenience of a summary and simple approach to reporting on data and analytics that seemingly works initially, but then after reflection they find that they need greater context. The present work faced these challenges; for future work we are adapting our technique and designing messaging to address these concerns with plans to evaluate receptivity. We also consider some additional questions that may help us study this space: What can the text analysis and data tell us? What can it not tell us? What are the limitations of the text and analysis? How might we think about the text and how might it be supplemented with existing information and which methods could help generate missing information?

### Action and Next Steps

The primary goal for reviewing recent comments about health information systems is informing decision-making and next steps for improvement and safety activities. Information gathered from survey comments with other relevant data can provide background and incentive for future studies, modelling, design, and evaluation activities. Comment analysis can inform trade-off decisions and help us understand workflow nuance. Unconstrained by pre-existing classes, this analysis identified both what can go wrong and what the requirements are for things to go right. Fuller and colleagues (2023) lay out an action plan that could be supported by the analytic capabilities outlined in this paper. Future work should focus on designing analytic capabilities that inform decisions and consider trade-offs when selecting from multiple methods for designing and problem solving.

## DISCUSSION

The purpose of this work was to review and interpret users' survey responses and comments about health information systems. Thematic analysis is a widely used approach for analyzing text from surveys and patient safety reports. Often a predefined code book helps teams to categorize comments. This can be a time consuming and laborious process. The task of categorization is focused on a specific discipline, domain, or set of viewpoints that distinguishes between one thing and another. The interpretation of the text is based on possibilities open to the reader just as the written text is based on the possibilities open to the survey respondent. Differences defined through code construction and drawn through the act of text classification may tell us more about the actualizations of possibilities open to and boundary construction of readers and/or writers (Rasmussen, 2004). In this work, it is important to describe a third grouping of people, drawing a distinction between the interpreters generating the analysis and the interpreters in the larger audience.

The larger audience is less likely to read deeply into the comment text because of the sheer volume. Therefore, this team's task, in general, is to act as intermediaries by filtering, reading, and interpreting text while designing an approach to identify, format, and share usable and actionable information with a larger audience. To this end, a common task of text analysts is designing translational representations of the intended boundaries described in the text to aid in decision making and help strategize about work efforts. Future work might include a system to map medical terms to the UMLS such as CLAMP or cTAKES to identify the hierarchical positions of extracted clinical terms within the terminological taxonomies and ontologies determined by these tools, and identify relationships, potentially generating additional ontologies or partially merging to existent vocabularies and other ontological objects.

## CONCLUSION

As we reflect on the various methods and techniques open to us in text analytics, we should consider how to arrive at translational representations of text that speak to the analysis team as well as the target audience and provide insightful and actionable information. Future work will focus on the perspectives and needs of analysis team members and larger audiences using the information while considering the usability and risks of text-based analytics approaches. This work aligns with human-centered and participatory design principles for co-creating systems and processes.

This text analytics work feeds into other healthcare systems learning, evaluation, and design work, so it is important to also account for needs associated with the users and designers of such systems. In addition, we should consider the feedback loops present in such work. Analysis of text will help us better understand our systems and inform safety efforts, and theories of safety will influence how we perform our analysis by determining how we categorize and filter information and how we think about the boundaries of the systems. Text analytics methods that are accessible and usable by study teams and stakeholders can support health information systems learning, which should lead to better and safer designs.

## ACKNOWLEDGMENT AND NOTES

We thank those working at the Veterans Health Administration for their commitment to patient safety. There were no relevant financial relationships nor any source of support in the forms of grants, equipment, or drugs. The authors declare no conflicts of interest. The opinions expressed in this article are those of the authors and do not necessarily represent those of the Department of Veterans Affairs or the United States federal government. The auto reader completion time was approximately 28 minutes for reading this manuscript. The completion time was approximately 33 minutes when including acknowledgements and references. Auto reader completion times may vary.

## REFERENCES

Chuang, J., Ramage, D., Manning, C., & Heer, J. (2012, May). Interpretation and trust: Designing model-driven visualizations for text analysis. *In Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 443–452).

Clough, P. D., and Sen, B. A. (2018) Evaluating tagclouds for health-related information research. *In Proceedings of the 13th International Symposium on Health Information Management Research – ISHIMR 2008.*

Demšar, J., Curk, T., Erjavec, A., Gorup, Č., Hočevar, T., Milutinovič, M., & Zupan, B. (2013). Orange: data mining toolbox in Python. *The Journal of Machine Learning Research*, 14(1), 2349–2353.

Fong, A., Harriott, N., Walters, D. M., Foley, H., Morrissey, R., & Ratwani, R. R. (2017). Integrating natural language processing expertise with patient safety event review committees to improve the analysis of medication events. *International journal of medical informatics*, 104, 120–125.

Fuller, H. J. A., Arnold, T., O'Neil, M., Wilson, D. (Accepted for publication). Activities to Promote Resiliency During Health Information System Transitions.

McKnight, S. D. (2012). Semi-supervised classification of patient safety event reports. *Journal of patient safety*, 8(2), 60–64.

Mingers, J. (2014). *Systems thinking, critical realism and philosophy: A confluence of ideas*. Routledge.

Mull, H. J., & Nebeker, J. R. (2008). Informatics tools for the development of action-oriented triggers for outpatient adverse drug events. In AMIA Annual Symposium Proceedings (Vol. 2008, p. 505). *American Medical Informatics Association*.

Ramage, D., Manning, C. D., & Dumais, S. (2011, August). Partially labeled topic models for interpretable text mining. *In Proceedings of the 17th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 457–465).

Ramage, D., Rosen, E., Chuang, J., Manning, C. D., & McFarland, D. A. (2009, December). Topic modeling for the social sciences. *In NIPS 2009 workshop on applications for topic models: Text and beyond* (Vol. 5, pp. 1–4).

Rasmussen, J. (2004). Textual interpretation and complexity. Radical hermeneutics. *Nordic Studies in Education*, 24(3), 177–193.

Reeves, R. M., Christensen, L., Brown, J. R., Conway, M., Levis, M., Gobbel, G. T., & Chapman, W. (2021). Adaptation of an NLP system to a new healthcare environment to identify social determinants of health. *Journal of biomedical informatics*, 120, 103851.

Rosch, E. (1978): Principles of categorization. In: Rosch, E. & Lloyd, B. B. (eds.): *Cognition and Categorization*. Hillsdale, N. J.: Lawrence Erlbaum Associates, 27–48.

Wilensky, U., & Papert, S. (2010). Restructurations: Reformulations of knowledge disciplines through new representational forms. *Constructionism*, 17, 1–15.