Opportunities for Natural Language Processing in Qualitative Engineering Education Research: Two Examples

Catherine G.P. Berdanier
Department of Mechanical Engineering, Penn State University
University Park, PA, USA
cgb9@engr.psu.edu

Eric Baker
Department of Computer Science, Penn State University
University Park, PA, USA
elb5465@psu.edu

Weiqin Wang
Department of Computer Science, Penn State University
University Park, PA, USA
wqw5186@psu.edu

Christopher McComb
School of Engineering Design, Technology, and Professional Programs, Penn State University
University Park, PA, USA
uum209@psu.edu

***Abstract*— This research full paper proposes opportunities to expand qualitative and textual data analysis using natural language processing (NLP), and demonstrates opportunities to use NLP in engineering education work in our presentation of two NLP-based projects as examples of how NLP can be used. The discipline of engineering education frequently employs qualitative data analysis techniques, but fundamentally, the analysis of large corpora of textual documents is limited by researcher's time. In this paper, we present a brief review literature of how NLP has been used in other qualitative research fields. This portion of the paper is aimed to provide a clear description of NLP for those unfamiliar with machine learning and natural language processing methods. The second part of the paper will provide two brief examples of how NLP is being employed in our research group. Example 1 is a study of engineering résumés, with the intention of being able to calculate the "disciplinary discourse density" based off the engineering language presented in engineering résumés, a technique validated in prior qualitative studies by hand. Example 2 is a genre analysis of engineering literature reviews, seeking to understand the ways in which sentences linguistically build into arguments, such that the task of writing literature reviews might be demystified. This paper will have a methodological impact for researchers attempting to use NLP methods to analyze qualitative data, articulating the opportunities and barriers in using these methods for engineering education research.**

Keywords—Natural language processing; machine learning; qualitative data

I. Introduction: An Orientation to Machine Learning and Natural Language Processing

Machine learning and natural language processing promise to continue to be an interesting interdisciplinary area. However, the discipline of engineering education has been slow to explore the role of machine learning and natural language processing to analyze large volumes of textual data. The goal of this paper is to introduce two small student-led projects exploring the use of natural language processing and machine learning techniques in analyzing qualitative data for engineering education research purposes.

Machine learning (ML) is a branch of computer science that uses statistical techniques to train the machine to learn from the data. Much machine learning research focuses more on prediction than inference; that is, the focus is on learning to predict something from training data sets with little emphasis on identifying correct underlying relationships. Prediction algorithms are often trained using an iterative approach in which fitting parameters are successively tuned to increase predictive accuracy. There are a variety of approaches to machine learning that can be applied based on the context, including artificial neural networks, deep learning, and support vector machines. **Artificial neural networks (ANN)** (or simply neural networks) analogically mimic the structure of biological neural systems in order to uncover complex relationships in the training data [1]. **Deep learning** refers to a multilayered (therefore, "deep") neural network approach that can handle increased levels of complexity in either a structured (classification) or unstructured (pattern recognition) format [2] **Support vector machines** are typically used for structured classification problems, and use training sets to uncover large categories by which additional inputs can be sorted [3].

Natural language processing (NLP) is a complementary field of study in which computers can be leveraged to analyze large amounts of useful data regarding human language (either spoken, or written). In recent years, NLP has begun to make frequent use ML algorithms to more efficiently analyze these large quantities of data. The applications of this technology are far-reaching across various disciplines, ranging from the ability to talk to electronic devices [4] to foundational English and rhetoric research [5]. The capabilities of NLP are continuously expanding as industries seek to automate tasks and access information, with some forecasts estimating that the NLP market will be work $22.3 billion by 2025 [6]. Although both machine learning and natural language processing algorithms are tightly coupled with foundational computer science research, researchers across disciplines (from engineering to linguistics) employ these methods in conducting other research.

The current work provides one such cross-cutting application in the utilization of machine learning and natural language processing to the analysis of qualitative engineering education data. The remainder of the paper is structured as follows: Two examples of the student projects employing NLP techniques are introduced, including a brief introduction to the respective motivation, data sets, methods, and findings for each example. Finally, the paper closes with recommendations for researchers interested in employing NLP techniques for qualitative human subject research, including a discussion on the inherent limitations of automating any qualitative data analysis methods.

II. LITERATURE REVIEW: NLP FOR QUALITATIVE DATA

Technological challenges and innovations continue to be researched from a fundamental computer science research perspective, often involving the development and testing of intensive statistical algorithms. The literature surrounding the mathematical algorithmic foundations are not addressed here, since reliable methods for ML and NLP are freely available through open source software. In this literature review, we focus on the use of NLP and ML for qualitative textual data analysis in educational contexts.

Natural language processing has proven to be useful across disciplines, opening doors to powerful data processing algorithms for those in and out of research. Previously, Kof [7], attempted to use natural language processing to extract text, analyze it, and automate the process for software development, and built classifications that established relationships between words, based on certain sentence constructions and other indicators [7]. Manual extraction of text from smaller documents was shown to be an effective means for scaling up to larger documents, although, it should be noted that manual work like this will continually be phased out in the future as developers create more NLP tools and algorithms.

In the domain of design research, NLP techniques have been employed frequently to assess the characteristics of text-based reports by human subjects. Latent Semantic Analysis (LSA) is a common technique that employs a bag-of-words model to assess the level of similarity between documents [8]. It has been used to analyze design team communication patterns [9, 10] and even in the assessment of the functional composition of design solutions [11]. LSA has also been utilized as a component in mapping analogical design spaces through the text of patents [12].

In engineering education in particular, NLP and ML techniques have rarely been employed. For example, Baba [13] presented fuzzy logic assessment systems for assessing English academic writing for engineering writing. In addition, researchers utilized latent Dirichlet allocation to assess differences between student communication and course content in online courses [14]. Most recently, researchers at Virginia Tech in engineering education are working use NLP to analyze student comments on course evaluations, and have demonstrated success using the "bag-of-words" approach to categorize student responses into "positive" and "negative" responses [15].

For NLP research dealing with interpreting language, one of the foundational struggles is for a machine to be able to understand and correctly interpret nuances, multiple meanings, and inferences in human language patterns that the human mind can easily interpret. Therefore, the development of appropriate training sets with qualitative data becomes difficult. As such, preprocessing and labeling data is a critical part of the research process before machine learning can be achieved.

III. EXAMPLE 1: NATURAL LANGUAGE PROCESSING IN THE ANALYSIS OF DISCIPLINARY DISCOURSE IN ENGINEERING RÉSUMÉS

A. *Summary of Motivation and Prior Work on Project*

The purpose of this ongoing project is to understand the ways in which "disciplinary discourse" is operationalized in engineering résumés. Traditionally, guidance on constructing a good résumé is relatively 'adisciplinary' with rules centered on formatting and structure, rather than on what content and words matter to a disciplinary audience. In past work, our research team has validated a framework for quantifying disciplinary discourse in engineering résumés, and have validated it on traditional one-page engineering résumés. This framework is based on descriptions of engineering competencies developed by the American Association of Engineering Societies (AAES) that translates well to rate more disciplinary specific competencies more highly than adisciplinary personal effectiveness competencies [16]. The framework for the competencies is structured into six tiers, which function as an *a priori* coding schema. Conveniently, the higher tiers correspond with higher disciplinary value. Therefore, it is possible to sum up the total "value" of each of the terms or phrases coded in order to calculate disciplinary discourse totals, density (per page of a résumé) and average disciplinary discourse quality (e.g. the average score of a coded phrase.) Résumés with a high quantity of disciplinary competencies will naturally score higher than a résumé filled with more general, lower-level codes. The AAES tiers are summarized in Figure 1. Method development and validation can be found in prior work [17-19].

Table 1: AAES Engineering Competency Tiers [16]

Tier Number	Description
Tier 6: Job-specific Competencies	Occupation-specific competencies that demonstrate leadership and vision in accomplishing goals
Tier 5: Industry or Sector Functional Areas	Demonstration of specialized expertise
Tier 4: Industry-Wide Competencies	Demonstrate mastery of engineering foundations and fundamentals; includes demonstration of engineering ethics, economics, and global competencies
Tier 3: Workplace Competencies	Ability to work in teams; work with appropriate tools and technology (including engineering software packages)
Tier 2: Academic competencies	Foundational academic skills or coursework attributes; basic computer skills
Tier 1: Personal effectiveness competencies	Interpersonal skills and professionalism

As typical for qualitative data analysis, the coding process can become tedious and time consuming to gather any statistically large corpuses of engineering résumés. For this reason, we sought to explore the use of NLP techniques to analyze the engineering résumés by training a program to categorize the disciplinary discourse in engineering résumés in the same way that we did by hand.

B. *Description of the Corpus*

In order to collect a very large sample of engineering résumés without engaging in a large-scale data collection process, we decided to collect online engineering résumés that are publicly available through the online résumé repository *Indeed.com*. People who place their résumés on such websites intend for their résumés to be seen publicly and are therefore an easy resource by which to download hundreds of engineering résumés. We downloaded résumés for N=500 engineering résumés, and a subset of N=100 résumés were coded by hand to develop a training algorithm. We did not sample résumés based on experience level, gender, or discipline of engineering; however, we did not select people that work in engineering industries but are not working as engineers (for example, human resources personnel for engineering companies are not included in the corpus of résumés.)

C. *Natural Language Processing Methods*

The processing program for this project was created using Python 3 programming language. Initially, the use of Python's natural language processing (NLP) modules such as the Natural language toolkit (NLTK) were examined to determine whether or not they would aid in the processing of the text documents of résumés that were previously acquired. Considering the hand-coded training corpus, we determined that the majority of words coded for disciplinary discourse were verbs.

Knowing that there was mainly one type of speech that was dominant while processing these words meant that using certain toolkits for NLP might prove to be inefficient because the implementation of packages such as the NLTK would result in more time and resources used than what would result. As such, a python 'dictionary' was created in order to assign values to various words as determined by the coding schema from our group's previous study and the training algorithm.

The analysis program reads a text file (of a résumé), scans through all of the words and adds them to an "ordered list" of words and a "unique list" of words. The ordered list contains all original wording, converting all text to lowercase, and excludes any punctuation, numbers and special characters. The program then counts the frequency of words in this original list, which is stored for later use.

The unique list of words was created by scanning through both the original text document and the ordered list to make a separate ordered list of non-repeating terms. The unique list also references a dictionary of "skip-words" that the program is instructed to not code into a numerical category. This step is important because it eliminates a lot of unnecessary bulk found normally when returning a list of unique words and makes it more efficient for the researcher to check potential words that could be coded. Lastly, the program scans through the words in the unique word list and the dictionary of scored words. If the same word is present in both, it takes the frequency of that word from the original list and multiplies it by the score assigned in the dictionary. Afterward, the program tallies up an entire quantitative disciplinary discourse score for the whole résumé.

D. *Results and Discussion*

The use of this language processing algorithm has shown to be almost as accurate as the previous method of hand-grading the engineering résumés. The average accuracy for a résumé in against the hand-coded résumé is 92%, if we first ensure that the resumes are in the general format of a traditional engineering resumes (e.g., bullet points or limited text rather than multiple/long paragraph descriptions, a limitation that we discuss further later). We have also begun testing the algorithm on résumés outside of the training set, and have found consistently high levels of agreement between the researcher-coded résumés and the same résumés analyzed by the program. The algorithm tends to score resumes higher than the human coder, because it searches for words instead of interpreting the words through meaning as a human coder does. Provided the program is continually fed more data, it will be able to continue improving its accuracy. Future work on this research project includes extending the corpus to increasingly large sets of résumés that will offer samples sizes that will offer statistically significant results. We intend to use this method to also explore linguistic differences between online and traditional engineering résumés, and to continue exploring the engineering disciplinary discourse that engineers employ in their résumés.

This particular study illuminates a number of inherent limitations of the use of NLP in analyzing qualitative data. Firstly, format is not analyzed in the engineering résumés, which is one way in which traditional résumés often differ, and can be used to show professionality. As a computer program simply reads text files and analyzes the words involved within, there is no way to interpret subjective aspects such as design that humans find valuable when evaluating the résumé of a potential employee. In addition, online engineering résumé repositories such as Indeed.com often include information that is irrelevant to a traditional "on-paper" résumé, such as willingness to work outside the U.S. In addition, people describing employment history in online contexts often write much more about their background, history, and trajectory than in traditional résumés, because they are not "limited" by the typical one-page length guideline. Often, users will also describe their companies in depth, which cause inaccuracies within the program, as the program cannot distinguish between words describing the company and accomplishments or job responsibilities of the person her/himself. However, for online résumés that share formatting with traditional résumés (e.g., bullet-point oriented lists, brief descriptions of job responsibilities), the automated program performs as well as a second researcher, indicating future potential for this method.

IV. EXAMPLE 2: NATURAL LANGUAGE PROCESSING AND MACHINE LEARNING IN GENRE ANALYSIS OF ENGINEERING LITERATURE REVIEWS

A. Summary of Motivation and Prior Work on Project

The purpose of the project in this example is to explore NLP for conducting genre analysis as it has been applied previously by our group for engineering literature reviews. Genre analysis is a method by which the rhetorical patterns of a text can be categorized, a technique that was formalized by Swales [20] as he proposed a structure by which rhetorical moves are categorized and comprise subcategorical rhetorical steps. Previously, other researchers have used this technique to analyze a variety of texts, including research articles, academic research proposals, and introduction sections. Prior work by our research group and our colleagues [21, 22] has analyzed engineering literature reviews for the purposes of more effectively teaching engineering students how to effectively build arguments within the literature review. However, this process is quite time intensive, and it is of interest to the research team to be able to gather rhetorical patterns from very large numbers of text, both from published articles and for students as they are learning to write a literature review.

The qualitative data analysis process and the coding schema has been described in previous work, but is outlined here for context. In a given literature review as text to be analyzed, we seek to categorize each sentence of the literature review into its respective purpose in the argument. In past work, we allowed each sentence to perform more than one function at a time, to support rhetorical theory; however, in our task to simplify the process, we sought to categorize each sentence into one rhetorical move. The unit of analysis in the coding procedure is at the sentence level. In addition, we did not seek to code into subcategorical rhetorical "steps," within each of the three moves, working to be able to simply teach a machine to categorize real sentences within a literature review into three main rhetorical moves, shown in Table 2.

Table 2: Rhetorical moves for sentence classification algorithms

Rhetorical Move [21, 22]	Description
Move 1: Announcing the Importance of the Study	Claim the importance of the topic through a statement of context (not usually literature-based); describe how the research context affects humans; or identification of a specific technical problem to be solved
Move 2: Preparing for the Present Study	Discuss the state of the field/current findings (citing literature); establish a gap or challenges in literature; acknowledge proposed solutions to the problem and or failures or previous solutions/approaches; identify benefits of a new/different approach to the problem; or familiarize readers with scientific background (not establishing context, but preparing the reader to interpret the current study).
Move 3: Introducing the Present Study	State the objective, purpose, hypothesis or need specifically for the present research study; intended outcomes; impact; attributes, capabilities, or benefits; and/or novelty of the study

In our study, we seek to conduct supervised learning techniques on text corpus to conduct sentence classification. By using the machine learning approach, we could give structure to unstructured data, which is called a "sequence to sequence" process. This technique has typically been applied to tasks such as topic classification or semantics analysis. In our study, we simply endeavor to categorize each sentence of a literature review into one of the Moves.

B. Description of the Corpus

To develop a training corpus on which to explore automated genre analysis, we collected a set of 200 literature reviews from recent volumes of Journal of Propulsion and Power, a journal published through The American Society of Mechanical Engineers. Though the corpus encompassed only recent articles (from years 2017-2018) published in one journal, there are several subdisciplinary areas and subjects of research that can be published in this journal, such that the technical content of the articles widely varies. Before any NLP/ML approaches were explored, each sentence in the Introduction section of these articles (encompassing the literature review) was coded by hand by the research team in order to develop the training corpus. Before beginning any data analysis, all files were converted from .pdf files into text files, and separated by sentence, storing the sentences in a .csv file with the corresponding rhetorical category.

C. Application of NLP Methods

In order to apply Deep Learning techniques for text corpus, the first goal is to preprocess the data. The normal preprocessing for text classification includes stemming, lemmatization, tokenizing and part of speech tagging. The purpose of stemming and lemmatization is to convert the words into their cleanest form. The most common stemming algorithm is Porter's algorithm [23], which removes the ends of words to normalize the text data. For instance, by using Python NLTK library, the stemming result for 'studied,' or 'investigating' returns 'studi' and 'investigat' [24]. Lemmatization is process to return the words into the original form. For instance, the word 'studies', 'studied' will both transfer to 'study' after the Lemmatization step.

After the Stemming and Lemmatization process, the next step is to tokenize the sentence into words and conduct part of speech tagging. Tokenizing refers to the process of converting a sentence into individual words for vectorization purposes. We used the Python NLTK tool "word_tokenize" for this process. Part of speech (POS) tagging refers to the process of give the words tag based on the context and dependencies of the sentence. The POS tagging process is important for the automatic content analysis work especially for sentiment analysis [25]. The preprocessing steps enable the text data to be fed into the deep learning neural network training.

We chose a "long short-term memory" (LSTM) neural network in order to try to capture long-term dependencies within the textual data. Within a neural network approach to machine learning, there are many layers of learning to conduct the approximation work. We selected parameters to comply with our interpretations of the qualitative data. For example, we used a sigmoidal activation function for our non-linear model and to conduct the multiclass classification we use the soft-max as the output layer. These decisions were made in order to accommodate the assumption that our data was non-linear and to aim for a multiclass classification system. The layers in the middle of the process seek to learn the parameters such as weights and bias from each input and output. The neural network was trained using the root mean square propagation (RMSprop) algorithm [26].

D. Results and Future Work

Deep learning requires a large amount of data in order to be effectively trained, much larger than the training corpus in this exploratory project. The model is able to recognize some patterns through the center word such like 'moreover', 'based on', 'therefore' and so on to conclude the topic. However, the dependencies of the content have not yet been revealed in our model. Likely, the size of the corpus is not large enough to train sufficient numbers of layers. In addition, the model still is dependent on content to determine the rhetorical moves, instead of using the arrangement of various words/parts of speech in order to determine purpose, which is part of our goals for future work.

V. Discussion and Recommendations for Exploring NLP Techniques

In this section, we reflect on these ongoing projects to discuss some of the relevant considerations that we have learned, in order to help other researchers beginning to use NLP or ML techniques. First, there are a variety python modules, as well as software in other languages that would be useful for anyone looking to utilize the power of machine learning or natural language processing. One of the most popular software packages available for testing machine learning algorithms is TensorFlow, which was created by engineers and researchers from the Google Brain Team and is open-source. As for natural language processing, Python's NLP toolkits are also open source, and have a great deal of online support through forums and toolkits.

Next, before selecting methods by which data will be categorized, it is important to understand the problem being faced, and to make appropriate choices for processing based on assumptions about the data set. For example, in the genre analysis example, we assume that the sentences are related to each other but not dependent on each other, and we also assume that all the words in a sentence cannot simply be analyzed as words in order to categorize for rhetorical purpose (a highly difficult task.) However, in the résumé example of NLP, because the same words were emerging across disciplines through hundreds of résumés, we assumed that we could indeed analyze by word and not by a more complex unit of analysis.

As important as knowing what data to incorporate is to know what *not* to incorporate, as processes can become unnecessarily complicated. Corpora can be managed in a multitude of ways, all with varying characteristics. In the résumé example, there was no need to characterize the words by part of speech, so the NLTK toolkit was not used to analyze the résumés. If we had not been comparing the words in the résumé to pre-existing values though, the NLTK toolkit would have provided the necessary tools to categorize the text in other ways. In developing a training corpus, researchers are urged to consider that if an appropriate training set has not been developed by others, that hand-coding and developing those mechanisms is highly time intensive, as is formatting data properly to be able to be fed into an algorithm.

Lastly, it is important to remember that though machine learning and NLP is a powerful tool, as with all methods, it is limited by the scope and decisions that the researchers make. Therefore, machine learning and NLP data is in no way bias- or value free; the machine can only learn from data fed by the researcher, and similarly, the machine will do only what the researcher asks it to do. It is highly important for any researchers using NLP or machine learning to stay engaged through the methodological decision-making in a process, and to interpret results within the confines of the research decisions, research question, and underlying theory.

Deep reliance on research topic-specific theory will become especially important if qualitative human subject data is analyzed, as one of the philosophical tenets of qualitative research is the deep relationship that a researcher has with the qualitative data, and the inferences that a human mind can reach in making data analysis and qualitative coding decisions. However, perhaps with large enough training sets, and highly validated coding schemas, deductive analysis using a priori codes might be accomplished on very large qualitative data sets using machine learning and natural language processing.

We see a great deal of potential for the use of NLP and/or ML techniques in analyzing qualitative data for engineering education and other social science and human subjects research. Primarily, one of the main limitations of qualitative research is that it is highly time intensive to analyze data from even small numbers of participants. While the use of NLP methods does not excuse the researcher from deeply engaging in the data, if an algorithm were trained to analyze the data from the paradigm of the research group, it may be useful to expand the sample sizes available to be analyzed through qualitative research methods. If trained accurately, the use of NLP could increase the consistency and accuracy of coding qualitative data, eliminating the room for human error and other factors, so long as the algorithm being used is designed correctly.

VI. Conclusion

In conclusion, this paper presents two examples of ongoing work exploring the use of natural language processing and machine learning in engineering education research. Our findings indicate that it is accurate as well as potentially convenient to incorporate the use of NLP and ML within different aspect of the field, after sufficient training sets are developed by which to train algorithms. Lastly, we provide recommendations for other researchers considering exploring NLP techniques, with a focus toward exploring the future of an intersection between ML/NLP and qualitative research methods.

REFERENCES

[1] McCulloch, Warren S., and Walter Pitts. "A logical calculus of the ideas immanent in nervous activity." *The bulletin of mathematical biophysics* 5.4 (1943): 115-133.

[2] Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems*. 2012.

[3] Cortes, Corinna, and Vladimir Vapnik. "Support-vector networks." *Machine learning* 20.3 (1995): 273-297.

[4] Lamere, Paul, et al. "The CMU SPHINX-4 speech recognition system." *IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP 2003), Hong Kong*. Vol. 1. 2003.

[5] Hogenboom, A., Frasincar, F., De Jong, F., & Kaymak, U. (2015). Using rhetorical structure in sentiment analysis. *Communications of the ACM*, *58*(7), 69-77.

[6] Tractica, Inc (2017) "Enterprise Applications for Natural Language Processing Technologies (Processing, Understanding, Generation) Software and Systems: Market Analysis and Forecasts," https://www.tractica.com/research/natural-language-processing/

[7] Kof, Leonid. (2004). Natural Language Processing for Requirements Engineering: Applicability to Large Requirements Documents.

[8] T. K. Landauer and S. T. Dumais, "A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge," Psychological Review, vol. 1, no. 2, pp. 211-240, 1997.[9] A. M. Agogino, S. Song, and J. Hey, "Triangulation of Indicators of Successful Student Design Teams," *International Journal of Engineering Education*, vol. 22, no. 3, pp. 617-625, 2006.

[10] A. Dong and A. M. Agogino, "Text Analysis for Constructing Design Representations," Artificial Intelligence in Engineering, 11, 65-75, 1997.

[11] J. Moss, K. Kotovsky, and J. Cagan, "The role of functionality in the mental representations of engineering students: some differences in the early stages of expertise," Cognitive Science, 30, 1, pp. 65-93, Jan. 2006.

[12] Fu K, Chan J, Cagan J, Kotovsky K, Schunn C, Wood K. The Meaning of "Near" and "Far": The Impact of Structuring Design Databases and the Effect of Distance of Analogy on Design Output. ASME. J. Mech. Des. 2013;135(2):021007-021007-12. doi:10.1115/1.4023158.

[13] Baba, A. F., Cin, F. M., & Ordukaya, E. (2015). Intelligent fuzzy assessment system for English Academic writing in engineering education. *International Journal of Engineering Education*, *31*(1 (A)), 83-93.

[14] Lim S, Tucker CS, Jablokow K, Pursel B. Quantifying the Mismatch Between Course Content and Students' Dialogue in Online Learning Environments. ASME. International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, *Volume 3: 19th International Conference on Advanced Vehicle Technologies; 14th International Conference on Design Education; 10th Frontiers in Biomedical Devices* ():V003T04A016. doi:10.1115/DETC2017-67339.

[15] Soledad, M., Grohs, J., Bhaduri, S., Doggett, J., Williams, J., & Culver, S. Leveraging institutional data to understand student perceptions of teaching in large engineering classes. In *Frontiers in Education Conference (FIE)* (pp. 1-8). IEEE. (2015).

[16] Amer. Assoc. Eng. Soc. Eng. (2016). Competency model. [Online]. Available: http://www.aaes.org/model

[17] Berdanier, C.G.P., McCall, M. and Mike, G. (2016). "A degree is not enough: Promoting engineering identity development and professional planning through the teaching of engineering resume writing," presented at the IEEE Frontiers Educ. Conf., Erie, PA, USA, Oct 12–15, 2016.

[18] Berdanier, C.G.P., McCall, M. and Mike, G. (2016). "Resumes in the development of undergraduate engineering identity: A genre analysis with teaching implications," presented at the IEEE Prof. Commun. Conf., Austin, TX, USA, Oct. 3–5.

[19] Fillenwarth, G. M., McCall, M., & Berdanier, C. (2018). Quantification of Engineering Disciplinary Discourse in Résumés: A Novel Genre Analysis With Teaching Implications. IEEE Transactions on Professional Communication, 61(1), 48-64.

[20] Swales, J. (1990). Genre analysis: English in academic and research settings. Cambridge University Press.

[21] Berdanier, C. G. (2016). Learning the language of academic engineering: Sociocognitive writing in graduate students (Doctoral dissertation, Purdue University).

[22] Lenart, J. B., & Berdanier, C. G. (2017). Optimizing a genre analysis framework to investigate engineering literature reviews. In Professional Communication Conference (ProComm), 2017 IEEE International (pp. 1-6). IEEE.

[23] Porter, Martin F.1980. An algorithm for suffix stripping. *Program* 14 (3): 130-137.

[24] Bird, Steven, Edward Loper and Ewan Klein (2009),Natural Language Processing with Python. O'Reilly Media Inc.

[25] Das, O., & Balabantaray, R. C. (2014, June). Sentiment Analysis of Movie Reviews using POS tags and Term Frequencies. Retrieved April 19, 2018, from https://pdfs.semanticscholar.org/d7b3/c50daf478667e22bf2d2f8a76df386586757.pdf

[26] http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf

.