Text Analysis for Understanding Symptoms of Social Anxiety in Student Veterans

Morgan Byers, Vangelis Metsis
Texas State University
San Marcos, Texas 78666
meb298@txstate.edu, vmetsis@txstate.edu

Abstract

A significant portion of the veteran population suffers from PTSD, a mental illness that is often accompanied by social anxiety disorder. Student veterans are especially vulnerable as they struggle to adapt to a new, less structured college lifestyle. In order to assist psychologists and social workers in the treatment of social anxiety disorder we use machine learning to analyze transcribed interview text and apply topic modelling to highlight common stress factors for student veterans. The results detailed in this paper also have broader impacts in fields such as pedagogy and public health.

Introduction

Post-traumatic stress disorder (PTSD) affects 8 - 10% of returning veterans, and of those, another 7 - 13% also have social anxiety disorder (Trahan et al. 2019). Student veterans are particularly vulnerable because of the difficulty involved with transitioning to life on a college campus. This radical lifestyle shift, in combination with limited access to resources, significantly impacts student veterans' quality of life.

In order to improve the support network for student veterans, we developed a tool to assist social workers and psychologists in treating social anxiety disorder. Our goal was to provide both qualitative and quantitative information about social anxiety.

We gathered ten interviews with student veterans, transcribed them, and divided them into a set of responses. Each response was rated for anxiety on a scale of zero to three by three psychologists, and then those ratings were averaged and rounded to integers. Ultimately this resulted in four target classes, with class zero corresponding to no stress present in a response and class three corresponding to high stress present.

To begin classifying anxiety levels, we explored deep learning approaches before turning to traditional machine learning. We then experimented with several feature selection algorithms in order to both improve the performance of our classifiers and use the selected features as a form of qualitative data. Finally, we used topic modelling to shed light on common sources of stress for student veterans.

Methodology

To begin, ten student veterans from Texas State University were interviewed by a professional psychologist. The veterans' responses to each of the psychologist's questions were then transcribed. Next, three coders from the Texas State University School of Social Work rated each response for anxiety on a scale of zero (no anxiety present) to three (high anxiety present). Afterward, each of the coders' ratings were averaged and rounded to integers. The distribution of our target classes ended up skewing right, with only 18 out of 1,185 responses being labeled as class three.

The intraclass correlation coefficient (ICC) is a number between zero and one that quantifies how much our three coders agree with each other. In our case, the $ICC(3,3)$ was 0.678, which indicates moderate agreement between coders.

While rating each sentence, the coders also listened to a recording of the interviews. This resulted in some responses that were textually the same being labeled as different classes. For instance, the lone word 'yeah' appears as a response in our data set 68 times, and the target class for the word varies.

We first implemented a bidirectional long short-term memory (LSTM) network with attention. Since the input lengths of our responses vary greatly, an attention mechanism helps to mitigate the loss of information as it gets compressed along each time step (Bahdanau, Cho, and Bengio 2015). In addition, we experimented with two different word embeddings: word2vec (Maas et al. 2011) and Keras' built-in embedding layer.

Since our target is ordinal, we also approached our problem as a regression task. We used an LSTM network with a word2vec embedding layer. For this model, we did not round the average anxiety rating of each response to an integer class, instead leaving them as floating point values.

In order to overcome the limitations of our small data set, we chose to use a support vector machine (SVM) classifier with a radial basis function kernel. We prepared a pipeline to perform feature selection and then train the SVM on the extracted feature set. Our SVM ultimately performed quite well, which prompted us to try out other traditional machine learning approaches, like logistic regression. However, logistic regression did not perform as well. Although the classification accuracy was not low, analyzing the model's confusion matrix shows that logistic regression labeled nearly all instances as class one, our most frequent class.

We chose to implement several feature selection algorithms, both to improve the accuracy of our SVM and also to provide qualitative data to psychologists and social workers. We implemented χ^2, mutual information, LASSO, recursive feature elimination, and tree-based feature selection. Each feature selection algorithm used provides some metric of attribute relevance e.g. a p-value for Chi-squared. Thus, for each feature selection algorithm we saved each keyword and its corresponding relevance metric in order to provide as much qualitative information as possible.

In addition to feature extraction, topic modelling allows us to find latent topics present in each interview. The results provided by topic modelling differ from those of feature selection since topic modelling works by aggregating related terms together while feature selection does not. The topic modelling algorithm we implemented is Latent Dirichlet Allocation (LDA).

Significant Results

Our bidirectional LSTM with attention and a word2vec embedding layer achieved a cross validated accuracy of 47.7% ($\pm$ 9.4%). However, this model struggled to overcome the imbalanced nature of our data set, even when the classes were weighted. Interestingly, the addition of an attention layer actually decreased our model's overall accuracy, but the f1-score for classes zero and two actually increased. The greatest increase in accuracy was caused by using a word2vec embedding layer.

The SVM achieved a cross validated accuracy of 59.1% ($\pm$ 7.2%). A notable aspect of the results is that although no model was able to correctly identify instances of the third class, the SVM labeled nearly all instances of class three as class two. Because of the ordinal nature of our data set, this misclassification is not as severe an error as our deep learning model classifying most instances of class three as class one.

Feature selection greatly improved the accuracy of our SVM. Chi-squared ended up providing the greatest boost to our model's accuracy. Additionally, feature selection highlighted words and phrases correlated with anxiety. For instance, mentions of "crowds" and "threats" appeared in the final feature set of many of our feature extraction algorithms.

Our topic model provided quite intuitive results. The final LDA model yielded five distinct topics, including subjects such as feeling seen, the fear of being in crowded places, and past military experiences. In combination with the powerful visualization tools available for LDA (Sievert and Shirley 2014), these results offer a keen look into areas of anxiety for student veterans.

Model	**Accuracy** (%)
Bi-LSTM	42.0
Bi-LSTM with attention	44.1
Bi-LSTM (word2vec)	48.4
Bi-LSTM with attention (word2vec)	47.7
LSTM regressor	45.0
Logistic Regression	48.3
SVM	59.1

Table 1: A comparison of model accuracy

Conclusion

Our results demonstrate that detecting and classifying stress levels from text alone is possible. In addition, we have shown that feature selection and topic modelling are viable methods of producing qualitative data about social anxiety. Our work produced tools that can help to support the decision making of psychology and social work professionals who are treating social anxiety.

The accuracy of our final models was limited by the small data set available to us. This data set was also heavily skewed, which caused most of our models to misclassify instances of our least common class.

There is much room for future work on this project. We will extend our project to take a multi-modal approach by incorporating the audio data from each interview with the transcription. This approach will allow us to offset the variance that was introduced by our judges giving different ratings to responses that are textually identical.

Beyond the immediate scope of our problem, our research is applicable to many other areas as well. For instance, automatic emotion recognition can assist teachers by identifying subject areas in which students don't feel confident.

Acknowledgments

This material is based upon work supported by the National Science Foundation under REU grant #1757893. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.

Maas, A. L.; Daly, R. E.; Pham, P. T.; Huang, D.; Ng, A. Y.; and Potts, C. 2011. Learning Word Vectors for Sentiment Analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 142–150. Portland, Oregon, USA: Association for Computational Linguistics.

Sievert, C.; and Shirley, K. 2014. LDAvis: A method for visualizing and interpreting topics. In *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*, 63–70. Baltimore, Maryland, USA: Association for Computational Linguistics.

Trahan, M. H.; Ausbrooks, A. R.; Smith, K. S.; Metsis, V.; Berek, A.; Trahan, L. H.; and Selber, K. 2019. Experiences of student veterans with social anxiety and avoidance: A qualitative study. *Social Work in Mental Health* 17(2): 197 – 221.