Discovering Causality in Suicide Notes Using Fuzzy Cognitive Maps

Ethan White

Applied Computational Intelligence Laboratory
University of Cincinnati
Cincinnati, Ohio 45221
whitee4@mail.uc.edu

Lawrence J. Mazlack

Applied Computational Intelligence Laboratory
University of Cincinnati
Cincinnati, Ohio 45221
mazlack@uc.edu

Abstract

An important question is how to determine if a person is exhibiting suicidal tendencies in behavior, speech, or writing. This paper demonstrates a method of analyzing written material to determine whether or not a person is suicidal or not. The method involves an analysis of word frequencies that are then translated into a fuzzy cognitive map that will be able to determine if the word frequency patterns are showing signs of suicidal tendencies. The method could have significant potential in suicide prevention as well as in other forms of sociological behavior studies that might exhibit their own identifying patterns.

Introduction

Computationally recognizing causality is a difficult task. However, discovered causality can be one of the most useful predictive tools. This is because understanding causality helps in understanding the underlying system that is driving the causal relationships [Steyvers, 2003]. One utilitarian outcome that causality provides is the prediction of human behavioral patterns either in a broad domain such as nations or religious groups or in groups of individuals. One such group of individuals that can be analyzed is those people who commit suicide. Suicide is one of the top three causes of death for 15-34 year olds [Pestian, 2010]. Therefore, suicide is a very pertinent topic for study. One of the ways to study suicide is to study the notes that were left behind by the ones who committed suicide [Leenaars, 1988]. Using these notes, a linguistic analysis can be performed that causal relationships can be extracted from. However, describing the causalities involved is difficult to do quantitatively, so previous causal analysis has mostly been qualitative. In contrast, this work considers causal suicide analysis using a quantitative method. This work uses fuzzy cognitive maps to discover and isolate root causal relationships based on words in suicide notes from people who take their own lives.

The long term goal of our work is to discover patterns within written material that may indicate causal relationships in human behavior. The focus of this work is based on patterns in the frequency of words as opposed to grammatical structure. The objective of this research, that will be the first step toward the long term goal, is to analyze a specific human behavioral pattern, i.e. suicide, in the form of suicide notes in a way as to contrast it with non-suicide notes. The central hypothesis is that human behavioral patterns can be extracted from word frequencies in written material, and that these patterns can be represented using fuzzy cognitive maps.

To test the central hypothesis and accomplish the objective of this research, three specific aims are pursued:

Discover and extract patterns in written material in order to produce an initial fuzzy cognitive map to describe causality

The first step toward this aim is the analysis of suicide notes according to the working hypothesis that the causal patterns can be discovered by finding word frequency patterns. This will be done for both the original written material and a set of the same data with spelling corrections.

The reason for making the distinction between spelling errors and corrected errors is that misspellings in suicide notes could have patterns that are exclusive to such writings as opposed to other written material. If, on the other hand, it turns out that misspellings are not significantly tied in with either suicide or non-suicide notes then notes that have had their spelling corrected will not be considered in the analysis. Only the original notes will be used in developing the fuzzy cognitive map.

The second step is an analysis of non-suicide notes based on the same working hypothesis. Again this has to be done for both the original and corrected versions of the data. Once this analysis has been done, the frequency patterns of the data will be used to produce an initial fuzzy cognitive map for analysis in aim two.

Perform rigorous testing on the fuzzy cognitive map on the original data and make adjustments where necessary

Once the first aim has been accomplished and the patterns discovered are converted to a fuzzy cognitive map, then testing must be performed in order to ensure that the map will be able to tell the difference between the suicide notes and the non-suicide notes that were originally tested.

Again, as in the first aim, this must be broken up into testing the original data and the data with spelling corrections. These must be further divided up into testing groups of notes and testing individual notes. This will show how sensitive the fuzzy cognitive map is to the amount of data available. Once the map has been altered to a point where the results are acceptably reliable then aim three will be performed.

Perform rigorous testing on the fuzzy cognitive map based on different material

Once the cognitive map is able to distinguish between the two original data sets used to build it, the map must be able to find the patterns in different written sources to make sure that it can work on a variety of writing. This is also broken up into two steps as in aim one and aim two.

The first step is using the misspelled words as written, and the second step is the corrected words. Also, as in aim two, this must be tested for both individual notes and for groups of notes to determine if the amount of data affects the outcome. If satisfactory results have not been attained, then the new data must be factored into the fuzzy cognitive map until the results are reliably accurate. Then aim three must be performed again using a different source of data.

Creating the Initial Fuzzy Cognitive Map

Extracting patterns in general categories from written material

The first step to accomplishing the first aim and developing the initial fuzzy cognitive map was to analyze the written data of both suicide notes and non-suicide notes. The group of suicide notes that were studied consisted of notes written by those that successfully committed suicide. The non-suicide notes consist of three sets that are approximately the same size as the number of words used in the group of suicide notes.

All three sets are taken from informal sources, i.e., each source represents a natural human form of communication as opposed to magazine articles, professional journals, and other such written works. The first set is a collection of various product reviews extracted from www.Amazon.com. This sample set was taken from a number of different products over a range of different ratings that ranged from the highest rating of five stars to the lowest rating of one star. The second set is a collection of notes from a private blog at archbishop-cranmer.blogspot.com. This is different from the amazon.com data because it represents an individual instead of a group of people. The final set comes from a political website called www.biggovernment.com. This set contains more specific topics than are covered by random product reviews on amazon.com and random notes from an individual.

All of the words were grouped into abstract general categories and sorted in order from most frequently used to least frequently used words. Each grouping is defined by how dense they are by percentage compared to the entire dataset, i.e., how frequently each group is used in a given set of data. The categories used are references to self, others, financial terms, medical terms, religious terms, negative and positive words, and misspelled words. The densities of these categories are shown in Fig. 1.

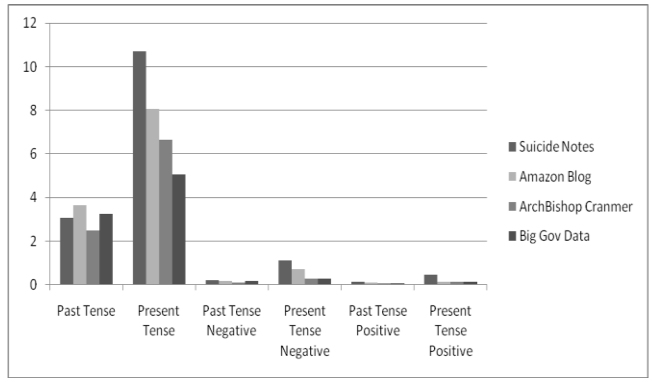

Figure 1. Word densities by percentage part 1

In addition to these categories, past tense and present tense words are also included along with their corresponding negative and positive references as shown in Fig. 2.

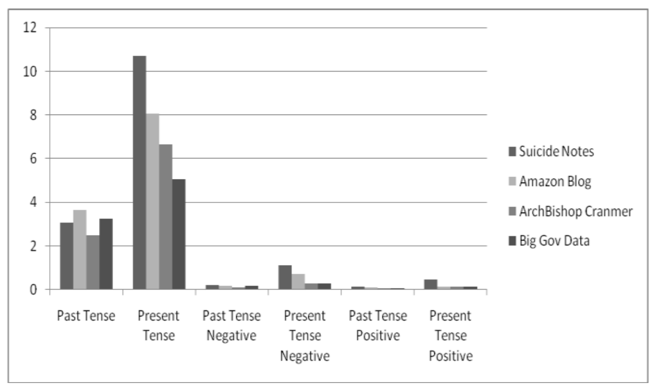

Figure 2. Word densities by percentage part 2

The results show that the greatest differentiation between the suicide notes and the non-suicide notes is found in three main categories that are references to self and others in fig. 1 and present tense in Fig. 2. Also, according to the data, there is not a significant amount of misspellings and even the small amount that is, does not show significant variation between suicide and non-suicide notes. Since the misspellings are not significant, they will not be considered in the analysis of the data. The three main categories are chiefly dominated by the set of suicide notes. This means that there would be no nodes in the fuzzy cognitive map that would push the final result toward a non-suicidal classification if it was analyzing a non-suicidal case. Therefore, the patterns have to be extracted on a word by word basis.

Extracting patterns from specific words in written material

The three best places to gather words that can provide varying reliable patterns are the groups for self references, references to others, and present tense. These groups contain the most references than any other kind and, therefore, the words in these categories are most likely to be found in a random set of notes to be analyzed and classified as suicidal or non-suicidal. The densities for these words, however, are not based on how many of each word is used in the entire dataset but rather on how many of each word is used in the group it occupies. Upon further analysis of the three groups, there were a number of words that proved to have either distinct suicidal influences or distinct non-suicidal influences. All words that had small percentages over all four datasets or did not vary significantly between suicide and non-suicide were removed from consideration. Fig. 3 shows the final results for references to self.

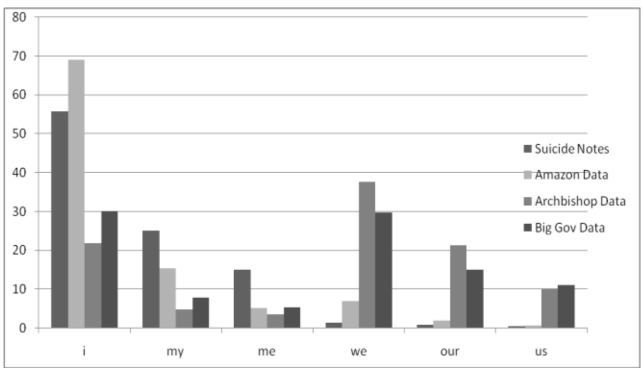

Figure 3. Word densities in self references by percentage

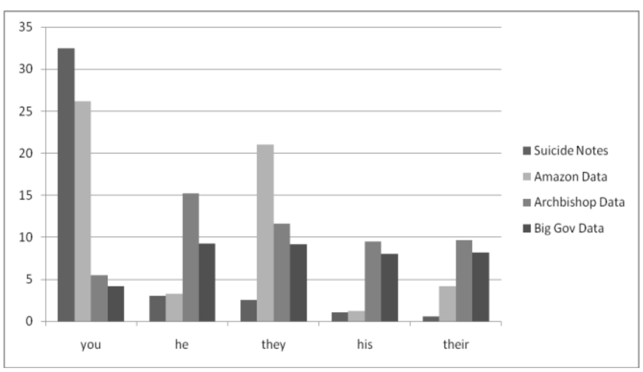

Figure 4. Word densities in others references by percentage

On average each word has a specific affiliation to either the suicide notes or non-suicide. However, the Amazon.com data shows definite anomalies in the words I, we, our, and us as compared with the other two non-suicide collection of notes. However, the apparent pattern is that suicide notes have more singular self references, i.e. I, my and me, while non-suicide notes seem to have more group self references, i.e. we, our, and us. Fig. 4 shows the results for references to others.

Again, there is a definite pattern with suicide notes have a large amount of references to the word "you" and the non-suicide notes have larger references to "he", "they", "his", and "their". Again, the Amazon.com data shows anomalies being similar to the suicide data in the word "you" but showing a great deal more influence in the word "they". The final results for present tense words is shown in Fig. 5.

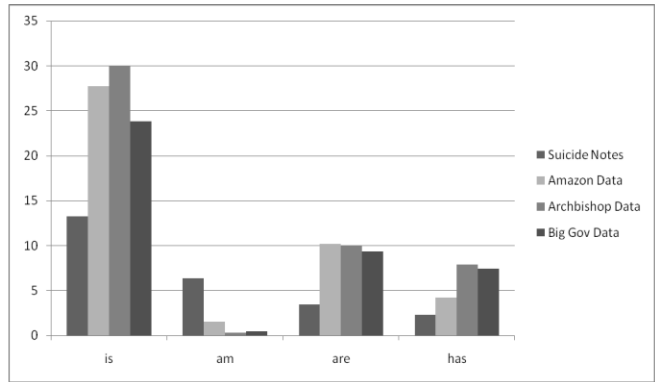

Figure 5. Word densities in present tense by percentage

In this group, the Amazon.com data acts similarly to the other non-suicide data except that the percentage for the word "has" is a little low, although not entirely problematic.

Developing the Initial Fuzzy Cognitive Map

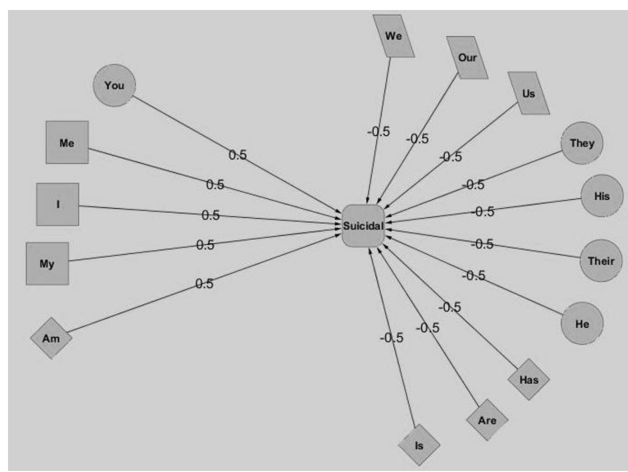

Figure 6. Initial fuzzy cognitive map

The fuzzy cognitive maps consist of a series of connected nodes that will represent the words being used from Fig. 3-5. These words will in some way connect to a suicidal node that will determine the classification of the dataset. The simplest graph that can be constructed, is for the suicidal node to be central with all word nodes connected to only that one node as shown in Fig. 6.

The node roles in the graph are indicated by the shapes of the nodes. The square nodes are the words that are the singular self references. The parallelograms are words that are plural self references. The circles are words that are references to others. Finally, the diamond shaped nodes are present tense words.

Each of the edges has a weight between -1.00 and 1.00 that is attached to it to determine how much influence and what kind of influence a particular node has on the suicidal node. The nodes on the left of Fig. 6 are all the nodes that are associated with suicide notes and thus have a positive influence, while all the nodes on the right represent non-suicide notes and are therefore negative in their influence.

All of the initial edge weights are arbitrarily set to start at 0.5 or -0.5. This would be true if all nodes would have equal influences on the classification; these starting values are expected to change. However, by starting with these values, it can be determined whether or not the general structure of the map is good or bad.

Each of the nodes starts at a particular value between 0.00 and 1.00 and then the graph is allowed to iterate by a computer program until the graph reaches equilibrium or until enough time has shown that it will never reach equilibrium. If the graph has reached equilibrium, then the final value of the suicide node is examined. If the value is over 0.50, i.e. over 50%, then the graph has determined the dataset to be suicidal. If the value is under 50%, then the dataset would be non-suicidal, and if the value is at 50%, then the classification is uncertain.

The starting values of the nodes are determined by normalizing the data in the particular group, e.g. in Fig. 5, all four datasets would be normalized according to the archbishop result for the word "is". This means that about 30% is the new 100% which all other values are compared to within that group. Fig. 3 and 4 would have their own number for normalization. The starting number for the suicidal node is 0.00 because it is assumed that there is no initial influence from this node.

The final results for the fuzzy cognitive maps for each dataset were not entirely successful. The Amazon.com data was particularly unsuccessful because of its anomalies which made it similar to the suicide notes. This means that the nodes in the graph do not have the same influence. Therefore, in order to determine if this map structure can distinguish between the datasets correctly, a set of weights must be found that can find the dividing line. By using machine learning techniques (supervised learning), it was discovered that there is a set of weights which allows the fuzzy cognitive map to correctly classify each dataset. The graph with its final weights is shown in Fig. 7.

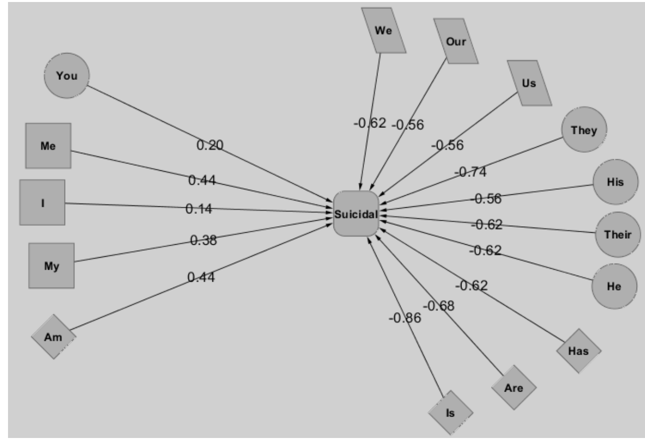

Figure 7. Final fuzzy cognitive map for testing

Testing the Fuzzy Cognitive Map

Now that a fuzzy cognitive map has been designed that can accurately classify the four datasets, this design must be tested against other collections of notes to see if the map can properly classify a random set of data.

General Category Testing

Three more datasets were used for testing. These consist of two sets of suicide notes and one non-suicide with each one only a fraction the size of the original four datasets. The first data set is a collection of suicide notes that contain some notes from the original suicide note collection as well as new ones. This was obtained from the website www.well.com/~art/suicidenotes.html and is labeled as suicide notes 2 in the analysis. The second set is a collection of suicide notes or the last words from famous actors, poets, and musicians labeled as suicide notes 3 in the analysis that was obtained from the website www.corsinet.com/braincandy/dying3.html. The final dataset is a collection of non-suicide notes from the private blog gregmankiw.blogspot.com.

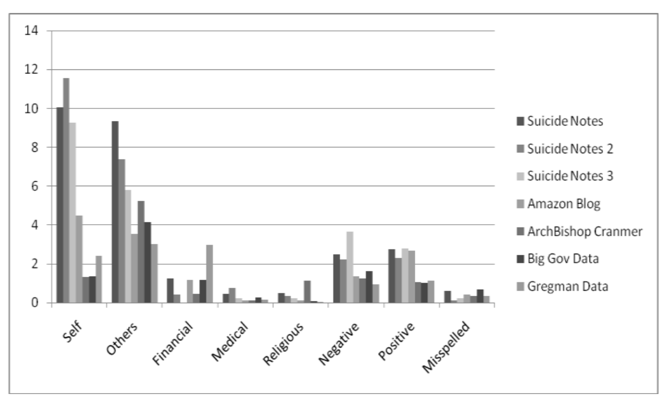

Figure 8. Results of all 7 datasets in general categories part 1

Before going straight into the word analysis, the results for the general categories should be compared with the original datasets. This is for the purpose of making sure that all of the datasets are following a predictable pattern. Fig. 8 and 9 show the results for each of the general categories.

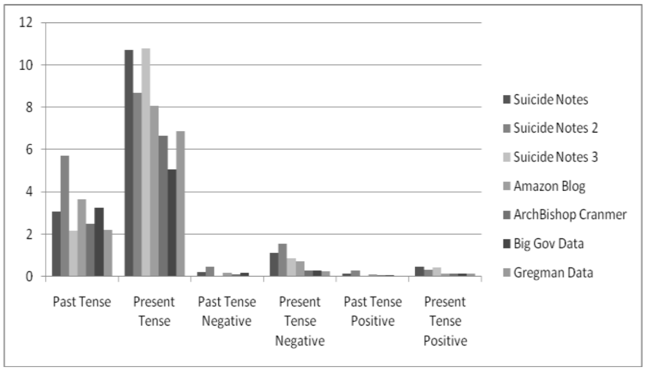

Figure 9. Results of all 7 datasets in general categories part 2

As can be seen from Fig. 8 and 9, the three new datasets follow similar patterns in both the suicide and non-suicide cases. Since there were no significant differences, then the specific word analysis could begin.

Specific Word Analysis

The final results for the word analysis are shown in Fig. 10, 11, and 12. As can be seen from the graphs, the three new datasets follow the same pattern for their respective classification with the exception of suicide notes 3 which produces some anomalies in the form of very large values in Fig. 12 for the words "is" and "are" which are very close to non-suicide patterns. Each of the new cases was normalized into the starting values for the nodes of Fig. 7. Each time, the fuzzy cognitive map accurately identified each dataset as either suicidal or non-suicidal. These findings suggest that this fuzzy cognitive map design is somewhat robust in that it was able to handle a random relatively small collection of suicide notes, i.e. suicide notes 3, and correctly identify them as such even with the non-suicidal like behavior that were found in Fig. 12.

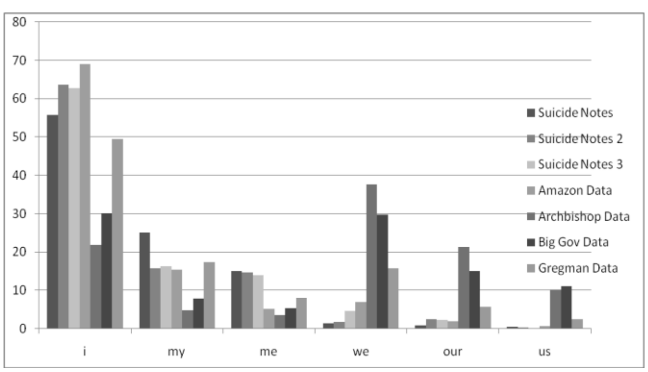

Figure 10. Word densities in self references by percentage

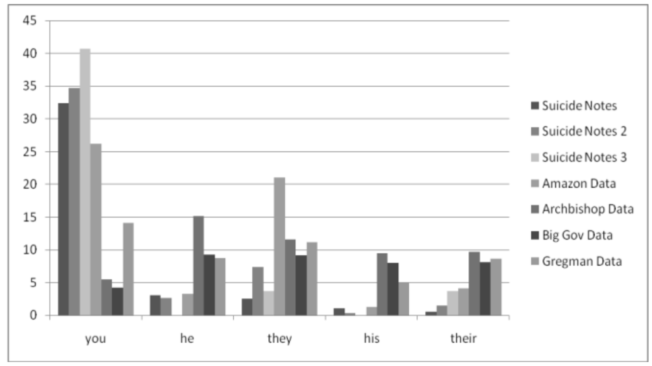

Figure 11. Word densities in others references by percentage

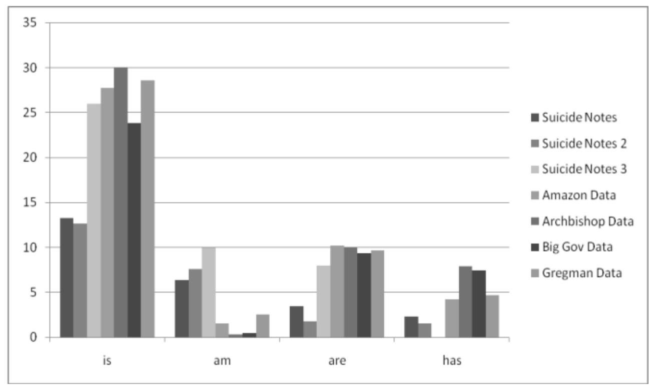

Figure 12. Word densities in present tense by percentage

Conclusion

The results of this research appear to provide strong evidence that it is possible to differentiate between suicidal behavioral patterns and non-suicidal patterns. Further testing must be done in order to ensure that this method can be used in all given situations. One such testing would be an analysis of suicidal ideation or intent to commit suicide [Barnow, 1997] which may or may not result in an attempted suicide. Also, further tests can be done from other collections of suicide notes as well as other sets of non-suicide notes.

This research is creative and original because it employs the use of fuzzy cognitive maps based on word frequencies in order to define human behavioral patterns. It is expected that the results of this research will further the understanding of causality and the prediction of human behavior. The broad application and positive impact of this work is a further development in the techniques for capturing causal relationships. Identification of causal relationships allows the ability to predict the consequences of actions from military strategies, governmental restructuring or societal rebuilding [Kosko, 1986] [Mazlack, 2010]. In the context of this research, fuzzy cognitive mapping is used to analyze writing and potentially to predict suicide cases allowing possible intervention that could save lives.

References

Barnow, S. and Linden, M. 1997. Suicidality and tiredness of life among very old persons: Results from the Berlin Aging Study (BASE). Archives of Suicide Research: 171-182

Kosko, B. 1986. Fuzzy Cognitive Maps. Academic Press, Inc. vol. 24: 65-75

Leenaars, A. A. 1988. Suicide Notes Predictive Clues and Patterns. Human Sciences Press, Inc. Windsor Ontario, Canada.

Mazlack, L. August 31 – September 3, 2010. Approximate Representations In The Medical Domain. Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence.

Pestian J., Nasrallah H., Matykiewicz P., Bennett A., and Leenaars A. 2010. Suicide Note Classification Using Natural Language Processing: A Content Analysis. Biomedical Informatics Insights: 19-28.

Steyvers M., Tenenbaum J. B., Wagenmakers E., and Blum B. 2003. Inferring causal networks from observations and interventions. Cognitive Science Society, Inc.: 453-489.