

Discerning Suicide Notes Causality Using Fuzzy Cognitive Maps

Ethan White

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

Lawrence J. Mazlack

Applied Computational Intelligence Laboratory
University of Cincinnati
Cincinnati, Ohio USA
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. The method involves an analysis of word frequencies that are used in a fuzzy cognitive map. The fuzzy cognitive map determines if there are suicidal tendencies. The method could have substantial potential in suicide prevention as well as in other forms of sociological behavior studies that also might exhibit their own identifying patterns.

Keywords—fuzzy cognitive maps; causality; suicide; language; natural language processing

I. INTRODUCTION

Computationally recognizing causality is a very difficult task. However, discovered causality can be one of the most useful predictive tools. One utilitarian outcome that causality can provide is the prediction of human behavioral patterns, either in a broad domain such as nations or religious groups, or in groups of individuals. One group of individuals that can be analyzed is the people who commit suicide. Describing the causalities involved is difficult to do quantitatively, so it must be represented qualitatively. The purpose of this work is to use fuzzy cognitive maps to discover and isolate root causal relationships based on words in suicide notes.

The long term goal of this work is to discover causal patterns in written material. The focus is on patterns in the frequency of words as opposed to grammatical structure. The objective of this work, which 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 suicidal 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:

A. 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 using discovered word frequency patterns can discover causal 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, then 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 the second aim.

B. 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.

C. 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 ensure that it can work on a variety of writing. This is also broken up into two steps as in the first two aims.

The first step is using the misspelled words as written, and the second step is using the corrected words. As with the second aim, 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.

II. BACKGROUND

A. Suicide research

Much of the research on suicide has been in two major areas: clinical studies and analysis of suicide notes. Clinical studies usually involve two types of studies. The first type of study is statistical analysis. This usually involves analyzing suicides based on age and gender as well as other demographics such as social and educational status [5]. This method has been traditionally a primary means of analysis in researching suicide [6]. For example, some studies have shown that nearly one-third of all suicides are found in the elderly over 60 years of age [1]. By using this technique, it is possible to discover what kinds of people are more likely to commit suicide and when. This allows researchers and suicide prevention organizations to identify particular trends in the demographics where they can focus their work in discovering ways to prevent such suicides.

The second type of clinical study involves psychological analysis of individual patients that are displaying suicidal thoughts and actions. This allows for discovering reasons for committing suicide on an emotional and psychological level. This method is done by studying people with suicidal thoughts in order to correlate emotional, mental, or physical state with suicide. An example would be to study the emotional state of people who are losing eyesight, who especially have a fear of losing sight, and how it correlates with suicidal thoughts [7]. If the thought processes of people who are undergoing blindness can be thoroughly understood, then it would be possible to anticipate and treat psychopathological symptoms before they lead to suicide [7].

The analysis of suicide notes has been primarily a study into the content of the notes in an attempt to discover the state of the persons thinking and thereby understand what circumstances and emotional states were driving the individual to commit suicide. These circumstances can be divided up into situations, relationships, emotional state, cognitive state, and demographic variables [3]. These different aspects are then analyzed to determine causes and motivations for suicide similar to the clinical studies approach.

B. Causality

Causality is a description of what causes a particular event or group of events in order to happen [10]. Causality is more than a statistical correlation as it contains both a direction, i.e. one event causes another event and not the other way around. Causality also has a particular order in a string of events that, if changed, would alter the result of the events [10]. If causality can be extracted from outside information, it would help intelligent systems in modeling its environment for decision making [11]. Also, causal relationships are more stable in describing the world than probabilistic models [12]. However, constructing causal relationships is difficult as often it is difficult to determine whether a particular event is necessary for the causal relationship or is merely coincidental [13]. Two methods used for modeling causality are text information extraction and graphs.

One of the challenging aspects of analyzing language is that language is highly complicated and often ambiguous; consequently making computer understanding difficult [14]. Extracting causal knowledge from text is of interest because the internet can provide a huge amount of text that is machine-readable. These documents have a great deal of information in them that can be used to establish causal models, e.g. analyzing newspaper articles in order to determine causal economic trends [15]. Another use in computationally extracting causal patterns from text documents is to answer questions by analyzing lexical and semantic constructions within the sentences [16].

Graph structures can be used for both modeling outside events, such as Mideast relations [9], as well as to describe causal relations in textual data [17]. The reason why these networks are powerful is that the graphs are made up of simple elements but create complex models that can be used to describe complex relationships [18]. One type of graph structure that can capture causality is a fuzzy cognitive map.

C. Fuzzy Cognitive Maps

Fuzzy cognitive maps are a fairly recent development in the attempt to capture causality in a way that can be useful for studying the real world. These maps are directed weighted graphs that describe relationships between events. The nodes and edges together capture domain information allowing the fuzzy cognitive map to depict uncertain causal reasoning [8]. The individual nodes represent particular events and the directed weighted edges represent what kind of an influence each node has on the other nodes that it connects to. If node one has a positive influence on node two, then if the event of node one increases, the event of node two will also increase; if they are connected with a negative edge, then if node one increases, node two will decrease [9]. Each node in the graph is a fuzzy set that goes from 0% to 100% [9]. Also, each edge can be a set of concepts such as "a little" up to "a lot" or a range of numbers such as 0% to 100% as with the nodes [9] [2].

III. CREATING THE INITIAL FUZZY COGNITIVE MAP

A. Extracting patterns in general categories from written material

The first step in 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.

The data 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 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 investigator defined abstract general categories and sorted in order from most frequently used to least frequently used words. Each grouping was 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.

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




Figure 1. Word densities by percentage part 1

The results show that the greatest differentiation between the suicide notes and the non-suicide notes is found in the references to self and others as shown in Fig. 1, and, present tense as shown in Fig. 2. Also, according to the data, there is not a significant amount of misspellings and even the small amount that there is, does not show significant variation between suicide and non-suicide notes. However, these 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 as if it was analyzing a non-suicidal case. Therefore, the patterns have to be extracted on a word by word basis.




Figure 2. Word densities by percentage part 2

B. 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.




Figure 3. Word densities in self references by percentage

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.

On average each word has a specific affiliation to either the suicide notes or non-suicide notes. 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.




Figure 4. Word densities in others references by percentage

As before, there was a definite pattern with suicide notes having a large amount of references to the word “you” and the non-suicide notes having 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 are 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.

C. Developing the 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.




Figure 6. Initial fuzzy cognitive map

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 weights start at 0.5 or -0.5. This was making the simplifying assumption that all nodes have equal influence on the classification; this simplifying assumption is most likely not correct. However, by starting at this assumption, 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 in a computer program until the program either reaches a breaking point or until the graph reaches equilibrium. If the graph has reached equilibrium, then the final value of the suicide node is examined. If the value is over 0.80, i.e. over 80%, then the graph has determined the dataset to be suicidal. If the value is under 30%, then the dataset would be non-suicidal, and if the value is between 30% and 80%, then the classification is uncertain.

The starting values of the nodes were 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.




Figure 7. Final fuzzy cognitive map for testing

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 that made it similar to the suicide notes. 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.

IV. TESTING THE FUZZY COGNITIVE MAP

Once a fuzzy cognitive map had been developed that could accurately classify the four datasets, this design was tested against other collections of notes to see if the map could properly classify a random set of data.

A. 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 was a collection of suicide notes that contains 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 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. Before going straight into the word analysis, the results for the general categories were

compared with the original datasets. This was 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 8. Results of all seven datasets in general categories part 1




Figure 9. Results of all seven 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.

B. 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

V. FURTHER TESTING

One of the further tests done was an analysis based on gender distinction. A set of male and female notes were taken from the first set of suicide notes and separated to see

if gender was a factor in word frequency patterns. In addition, a separate set of male and female suicide notes were taken from [3]. None of these four sets of notes showed any significant difference in word frequency patterns and were correctly identified by the fuzzy cognitive map in Fig. 7. Two non-suicide note collections were then added. The first was from a personal blog colinhughes.blogspot.com that was correctly identified. The second collection was a set of personal letters from the 1820’s that were written by patients to Dr. James Carmichael. These were part of a collection from the University of Virginia found at carmichael.lib.virginia.edu. The fuzzy cognitive map was unable to distinguish these notes from the suicide notes. By performing a principle component analysis (PCA), Fig. 13, on the values of the values of the nodes in the map, it is easy to see why the Carmichael notes do not separate.




Figure 13. PCA of 13 notes

As can be seen in Fig. 13, the suicide and non-suicide notes are clearly separate from one another with the exception of the Carmichael notes which are found in with the cluster of suicide notes on the bottom left of Fig. 13. One possible explanation for why these particular notes are similar to suicide notes is that suicide notes generally take the form of a personal letter that is directed from one individual to another. This could cause particular word patterns to form that would be similar to other personalized letters. This could be especially true for the words found in the self references and references to others categories.

The problem was solved by adding more words from a new general category. Five more words were chosen from the past tense general category. The five words are “tried,” “got,” “told,” “wanted,” and “were.” By adding these words, there were enough distinct word patterns that were able to separate the Carmichael notes from the cluster of suicide notes. All that was left to be done was to find a new set of weights that would be able to separate the suicide from the non-suicide notes. After redistributing the weights, the new fuzzy cognitive map, Fig. 14, was able to successfully identify the Carmichael notes as non-suicidal.

Finally, two more sets of personal letters were added to datasets for a total of 15. The first set was a collection of

letters from Alexander Bell to his wife. These were most similar to the Carmichael data. The last dataset was a collection of personal letters from a soldier during the Spanish American War. Both sets came from the digital collection at the Library of Congress. The final results are shown in the PCA of the data, Fig. 15.




Figure 14. New fuzzy cognitive map




Figure 15. Final results for all 15 datasets

With the addition of the five past tense words, the Carmichael notes have shifted away from the suicide cluster and are almost at the y-axis at $y \approx -40$. Immediately to the right is the collection of Alexander Bell's letters. Finally, the Spanish American War letters are found in the cluster on the right hand side of the graph near $x \approx 75$.

VI. CONCLUSION

The results 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. Testing needs to be done on age based collections. There is the possibility that notes from adolescents, for example, are different from adults. Another such testing

would be an analysis of suicidal ideation or intent to commit suicide [1] which may or may not result in an attempted suicide. Also, further tests can be done from other collections of suicide notes and non-suicide notes that may result in further adjustments to the fuzzy cognitive map.

This work is creative and original because it employed 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 can further the understanding of causality discovery in text 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 may allow the ability to predict the consequences of actions from military activities, governmental restructuring or societal rebuilding [2] [4]. In the context of this work, fuzzy cognitive mapping was used to analyze writing and potentially it may be used to predict suicidal acts and allow possible intervention that could save lives.

REFERENCES

- [1] S. Barnow and M. Linden, "Suicidality and tiredness of life among very old persons: Results from the Berlin Aging Study (BASE)," *Archives of Suicide Research*, 1997, pp. 171-182.
- [2] B. Kosko, "Fuzzy Cognitive Maps," *Academic Press, Inc.*, vol. 24, 1986, pp. 65-75.
- [3] A. A. Leenaars, *Suicide Notes Predictive Clues and Patterns*, Windsor Ontario, Canada: Human Sciences Press, Inc., 1988.
- [4] L. Mazlack, "Approximate Representations In The Medical Domain," *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence*, Toronto, August 31 – September 3, 2010.
- [5] A. S. Richardson, H. A. Bergen, G. Martin, L. Roeger and S. Allison, "Perceived Academic Performance as an Indicator of Risk of Attempted Suicide in Young Adolescents," *Archives of Suicide Research*, 9, 2005, 163-176.
- [6] B. L. Paulson and R. D. Everall, "The Teen Suicide Research Project," *The Alberta Journal of Educational Research*, 47, 2001, 91-94.
- [7] D. Leo, P. A. Hickey, G. Meneghel and C. H. Cantor, "Blindness, Fear of Sight Loss, and Suicide," *Psychosomatics*, 40, 1999, 339-344.
- [8] M. Hagiwara, "Extended Fuzzy Cognitive Maps," *IEEE International Conference on Fuzzy Systems*, San Diego, CA, March 8-12, 1992, 795-801.
- [9] B. Kosko, 1993, *Fuzzy Thinking: The New Science of Fuzzy Logic*, Hyperion, New York.
- [10] G. D. Weber, "Discovering Causal Relations by Experimentation: Causal Trees," *Proceedings of MAICS*, 1997, 91-98.
- [11] J. Pearl and T. S. Verma, "A Theory of Inferred Causation," *Second International Conference on the Principles of Knowledge Representation and Reasoning*, 1991, 441-452.
- [12] S. Stromsten, "First-Order Causal Knowledge," 2003.
- [13] D. Buchsbaum, A. Gopnik, T. L. Griffiths and P. Shafto, "Children's Imitation of Causal Action Sequences is Influenced by Statistical and Pedagogical Evidence," *Cognition*, In Press, Corrected Proof, 2011.
- [14] K. Chan and W. Lam, "Extracting Causation Knowledge from Natural Language Texts," *International Journal of Intelligent Systems*, 20, 2005, 327-358.
- [15] H. Sakaj, S. Sekine and S. Masuyama, "Extracting Causal Knowledge Using Clue Phrases and Syntactic Patterns," *Proceedings of the 7th International Conference on Practical Aspects of Knowledge Management*, Springer-verlag, Berlin, 2008, 111-122.
- [16] R. Girju, "Automatic Detection of Causal Relations for Question Answering," *Proceedings of the ACL 2003 workshop on Multilingual Summarization and Question Answering*, 2003, 76-83.

- [17] C. S. G. Khoo, S. Chan and Y. Niu, "Extracting Causal Knowledge from a Medical Database Using Graphical Patterns," *Proceedings of the 38th Annual Meeting on Association for Computational Linguistics*, Hong Kong, Association for Computational Linguistics, Stroudsburg, PA, USA, 2000, 336-343.
- [18] W. Buntine, "A Guide to the Literature on Learning Probabilistic Networks from Data," *IEEE Transactions on Knowledge and Data Engineering*, 8, 1996, 195-210