

Applying NLP to Build a Cold Reading Chatbot

Peter Tracey†

School of Science, Engineering and Environment
University of Salford-Manchester,
M5 4WT
P.J.Tracey@salford.ac.uk

Mo Saraee

School of Science, Engineering and Environment
University of Salford-Manchester,
M5 4WT
M.Saraee@salford.ac.uk

Chris J Hughes

School of Science, Engineering and Environment
University of Salford-Manchester,
M5 4WT
C.J.Hughes@salford.ac.uk

ABSTRACT

Chatbots are computer programs designed to simulate conversation by interacting with a human user. In this paper we present a *chatbot* framework designed specifically to aid prolonged grief disorder (PGD) sufferers by replicating the techniques performed during *cold readings*. Our initial framework performed an association rule analysis on transcripts of real-world *cold reading* performances, in order to generate the required data as used in traditional rules based *chatbots*. However due to the structure of *cold readings* the traditional approach was unable to determine a satisfactory set of rules. Therefore, in this paper we discuss the limitations of this approach and subsequently provide a generative solution using sequence-to-sequence modeling with long short-term memory. We demonstrate how our generative *chatbot* is therefore able to provide appropriate responses to the majority of inputs. However, as inappropriate responses can present a risk to sensitive PGD sufferers we suggest a final iteration of our *chatbot* which successfully adjusts to account for multi-turn conversations.

CCS CONCEPTS

• Human-centered computing • Human computer interaction

KEYWORDS

Natural Language Processing (NLP), Association Rules, Apriori, Deep Learning, Chatbots, Grief, Cold Reading

1 Background

1.1 Motivation

This research is focused on providing comfort to patients suffering from grief, following another person's death. At these times many patients turn to mediums in order to receive a *cold reading*. This is often helpful because it allows the patient to find closure with the deceased. However often it is not possible for a patient to access a medium (either due to cost, or geographic limitations) and therefore a simulated *cold reading* provided through a *chatbot* can provide alternative comfort.

In order to enable a patient to achieve healthy grief they must achieve the completion of four grief tasks [1]. These tasks are:

- To accept the reality of the loss;

- To process the associated pain;
- To adjust to a world without the deceased;
- To find an enduring connection with the deceased in the midst of embarking on a new life.

However, there can be complications in completing these tasks which is commonly diagnosed as prolonged grief disorder (PGD), which occurs in approximately 10% of bereavements [1].

1.2 Current Approaches

Currently there are three main approaches to treating PGD: pharmacological; psychological; self-help.

Pharmacological treatment (such as the use of drugs) is effective at reducing depression symptoms but does nothing to target the underlying cause [1]. For many patients pharmacological treatment is not advised because it carries risks of dependence and can interfere with functions necessary for adaptation to loss.

Psychological interventions are a promising alternative, however according to a report by Mind [2] 10% of patients have been waiting for over a year for psychological therapy and over 50% have been waiting for over 3 months.

Therefore, many patients turn to self-help through a medium. Mediums are performers who purport to communicate on behalf of the deceased, using a process called *cold reading* [8]. *Cold reading* is the process wherein the medium makes probable assumptions called Barnum statements [8] about the client to infer knowledge about someone the client has lost. The reader claims that this knowledge has been imparted on them by the deceased, establishing a line of communication, which then allows the client to resolve their grief. Mediums use their *cold reading* skills to make a living, and therefore charge considerable fees to their clients, which has caused controversy due to allegations that these performers are taking advantage of other people's grief to make a profit. One way to avoid having to pay a living wage for a conversational service is to employ the use of a *chatbot*. A *chatbot* is a computer program designed to simulate conversation by interacting with a human user. In this case the patient would specifically be interested in a *griefbot* (a *chatbot* specifically designed for helping with grief).

The idea for *griefbots* [3] started in a 2013 episode of *Black Mirror*, titled "*Be Right Back*" [4] which told the story of Martha who loses her boyfriend Ash in a car accident. Martha then uses her instant messaging history

with Ash to recreate him virtually. In 2015, Eugenia Kuyda did much the same thing [5] by recreating her deceased friend Roman Mazurenko in the form of a *chatbot*. Following in her footsteps are Marius Ursache and James Vlahos, who founded *Eterni.me* [6] and *HereAfter* [7] respectively. Both services aim to virtually recreate deceased persons as a service by recording their experiences prior to their passing. This creates an accessibility issue for people who did not anticipate the passing of their loved ones and/or were unaware of the services and have therefore missed the window of opportunity to record their experiences. Kuyda averted this pitfall by using the method depicted in *Be Right Back* wherein instant messaging data forms the vocabulary for a *chatbot*. However not everyone uses instant messaging services, and if they do, they may use them sparingly or wish that their data is kept private after their death.

1.3 Proposal

Therefore, a new *griefbot* solution is required, one that does not necessitate the use of large volumes of instant messaging data or preparation prior to the deceased's passing.

This paper proposes to automate the *cold reading* process via a *chatbot*. Unlike contemporary *griefbots*, the *chatbot* would not need to use instant messaging data from the deceased nor would the deceased need to have preempted their passing and recorded their experience. Unlike mediums, the bot would not need to charge each user a living wage for its services.

2 Methods

2.1 Association Rules

Many *chatbots* are rules-based meaning that they consist of pattern-template pairs which need to be manually written, for example if one pattern was "how are you?" the corresponding template might be "I'm okay" [9].

We use association rule mining [10] to find potential *chatbot* rules. In particular, by using the apriori algorithm [11], we can find pairs of antecedents and consequents that could be used as patterns and templates respectively.

Other methods such as clustering and decision tree analysis were considered, but association rule analysis was determined to best suit the nature of the task.

2.1.1 Data

To find association rules for a *cold reading chatbot* we use the Archive of mediUm and cold Reader dAta (AURA) dataset [12] (provided for non-commercial use under a fair use license).

The dataset contains transcripts of readings conducted by mediums on the Larry King show [13] 1 transcript per episode, with a varying number of readings per episode.

The readings were conducted live and over the phone therefore no editing was used to embellish the effectiveness of the readers and no visual information was used in the readings.

A sample of the original text can be seen in Figure 1. The first step of our framework is to clean the data, to ensure it can be processed, as shown in figure 2.

We used 3908 lines of text from 273 readings.

```
R1
    KING: Let's take some calls
    hello?
    CALLER: Hello?
    KING: Hi.
    CALLER: Hi.
    KING: Go ahead.
    CALLER: Rosemary?
    ALTEA: Hello?
```

Figure 1: Sample of transcript from AURA dataset

```
R1    CALLER: Hello.
R1    READER: How are you?
R1    CALLER: I'm fine.
R1    READER: What's your question?
R1    CALLER: My question is -- I was just
R1    READER: How old was she?
R1    CALLER: Fifty.
R1    READER: Your daughter was 50?
R1    CALLER: Yes.
R1    READER: I'm actually seeing a lady w/
```

Figure 2: Sample of transcript after preparation.

2.1.2 Document-Term Matrix

After initial preparation our data is still not ready for association rule mining. To apply the apriori algorithm we need to transform the data into a document-term matrix.

First, we create a document-term matrix using lines spoken by the callers, where each row is a turn of speech, and the columns are n-grams up to 10-grams.

Secondly, we create a document-sentence matrix using lines spoken by the reader, wherein each row is a turn of speech and the columns are whole sentence responses.

We use binary weighting because the apriori algorithm operates on transactional datasets. For example, determining association rules for shopping habits requires the apriori algorithm to be applied on datasets where a given customer either did or did not buy a certain item.

We then add a prefix to the columns in each matrix, "C_" for columns in the caller matrix and "R_" for columns in the reader matrix. We then merge the matrices into a single matrix that we will apply the apriori algorithm to. By prepending our prefixes to the columns earlier, we can differentiate identical terms that appear in either the caller's input or the reader's response.

2.1.3 Parameters

There are certain parameters we need to set so that our results are not crowded with association rules that are statistically unsound.

2.1.3.1 Support

The "support" of a rule is a measure of how frequently a rule occurs within the dataset. We set the minimum support to only include rules which occur at least twice in the dataset.

2.1.3.2 Confidence

Confidence is measured by the support of a rule over the support of its antecedent. Therefore, confidence is the conditional probability of the consequent, given the antecedent. We disregard rules for which the confidence is <50%, because a *chatbot* following these rules will be wrong more often than it is right.

2.1.3.3 Lift

“Lift” is measured by the support of a rule over the product of the support of the antecedent and the support of the consequent. Therefore, it is the ratio of the support of the rule to the expected support if the antecedent and consequent were independent of each other. The higher the lift, the greater the dependence between the antecedent and the consequent. If the lift value is <1, the antecedent and consequent are inversely dependent upon each other, and therefore we disregard rules which have lift <1.

2.1.4 Results

```
{C_saw_him} => {R_i saw him }
{C_you_saw_him} => {R_i saw him }
{C_you_saw} => {R_i saw him }
{C_greatgrandmother} => {R_yes}
{C_seeing} => {R_yes }
{C_good_evening_sylvia} => {R_yes }
{C_evening_sylvia} => {R_yes }
```

Figure 3 Association rules generated using methods described in this paper.

From the dataset we processed we only found 7 association rules which fit our hyperparameters as shown in Figure 3. While we could find more rules if we used less stringent hyperparameters, these rules would not be statistically sound to use in our *chatbot*.

To build a fully conversational system we need to consider tools beyond traditional rules-based *chatbots*, and thus in the next section we describe a generative model using deep learning techniques.

2.2 Deep Learning

Deep learning is named after the structure of models which are built with many layers of artificial neurons, connected together into a deep network.

2.2.1 Model

Artificial neural networks (ANN) mimic the complex interactions between neurons in biological systems and with enough data and training time, they can learn to generate new outputs given previously unseen inputs.

Recurrent neural networks (RNN) develop the concept of ANN’s by taking the hidden layers of one ANN and using them as the input for another, then repeating the process as many times as required. Encoding data in this way captures sequential properties, i.e. parsing a sentence word by word through an RNN encodes the order of the words and thus each word carries with it the context with which it was used.

Sequence-to-Sequence neural models take the RNN concept and enhance it by using 2 RNN’s, one as an encoder, to which input sequences are parsed, and another as a decoder, from which output sequences are generated. Typically, this system would be used for translation e.g. English to French, but the same process works for query and response pairs.

Long Short-Term Memory networks are another addition to the RNN, whereby incorporating memory cells and gates can negate the vanishing gradient problem in RNN’s, where older parts of a sequence are forgotten the longer the sequence becomes.

2.3.5 Data

In addition to the AURA dataset, we use the *Cornell Movie-Dialogs Corpus* [14] (provided for non-commercial use under a fair use license). This is to give our model the capacity for general conversational ability, upon which the ability to give *cold readings* will rely on. For both the Cornell and AURA datasets, we remove input-output pairs where either the input or output is longer than 25 characters. This will make our model more efficient and improve overall performance as deep learning models can struggle with longer sequences.

For the Cornell dataset we also remove unique input-output pairs. This is to avoid the model learning responses which were only appropriate in a single specific context. We also convert each dataset to lowercase and remove punctuation and bind the two datasets together into a single dataset by their rows.

To use the datasets in our model, we parse a tokenizer over both datasets. This builds a vocabulary wherein each word is represented by a unique token. For our target data, we use one-hot encoding where each word is represented by a list of 0’s and a single 1 at the digit which corresponds to that word.

2.3.6 Training

We train a sequence-to-sequence model with long short-term memory on the combined corpora for 200 epochs with a batch size of 4.

These small batch sizes are chosen because of the relatively small volume of data available, and we find that 200 epochs are sufficient for comprehensible responses to begin emerging.

2.3.7 Results

Following the training process, our *chatbot* is able to successfully respond to simple messages as seen in Figure 4.

```
Human : hello
Bot : hello
Human : are you okay
Bot : i think so
```

Figure 4 Demonstration of general conversational ability

Although our *chatbot* has not fully developed the ability to replicate a *cold reading* or the use of techniques such as Barnum statements, it shows significant promise by responding appropriately. This is both succinct, and

appropriate to situations a medium might find themselves in, such as a client enquiring after their deceased father (Figure 5) or grandmother (Figure 6).

```
Human : i lost my father
Bot : i know hes right
Human : can you hear my father
Bot : i yeah
Human : is he okay
Bot : yes
Human : is he proud of me
Bot : yes
```

Figure 5 Chatbot responding to situation where the user has lost their father

```
Human : my grandmother died
Bot : yes
Human : can you hear her
Bot : yeah
Human : what does my grandmother say
Bot : hello
Human : can you tell her i say hello
Bot : okay
```

Figure 6 Chatbot responding to situation where the user has lost their grandmother

Currently our *chatbot* doesn't always behave as we expect. For example, Figure 7 demonstrates the *chatbot* responding in ways that might be upsetting to a PGD sufferer. Furthermore, Figure 8 demonstrates how often the *chatbot* responds entirely nonsensically. While the approach is remarkable, we are limited by the data available in the domain of *cold readings*. More data could allow for more complex relationships to be learned and even more complex model structures.

```
Human : whats heaven like
Bot : nothing
Human : whats hell like
Bot : you
```

Figure 7 Unexpected but nonetheless interesting results from *chatbot* responding to questions about the afterlife

```
Human : how are you
Bot : i
Human : is my dad watching me
Bot : is it
Human : are there any messages for me
Bot : is i this okay
```

Figure 8 Examples of *chatbot* failing to produce meaningful results

3 Discussion

We proposed a *cold reading chatbot* with rules chosen based on the results of association rule analysis on transcripts of *cold reading* performances. Our association rule analysis only managed to produce 7 rules. These rules were insufficient for a conversational system to function

and therefore we tried for a generative approach using a sequence-to-sequence model with long short-term memory. The generative *chatbot* responds appropriately to a greater variety of inputs than our rules-based model would have, however it still fails to respond appropriately to many inputs which would most likely be dangerously frustrating to a PGD sufferer.

We suggest that this is because each response in an actual *cold reading* is based on more context than the most recent utterance by the client. Therefore, we suggest another attempt at making a *cold reading chatbot*, with more data and a more complex architecture which encodes the context of previous utterances made over the course of each conversation, in order to account for all the context in multi turn conversations.

ACKNOWLEDGMENTS

We thank Dr. Mahsa Saracee for contributing their medical expertise to this paper.

References

- [1] Paul A Boelen and Geert E Smid. 2017. Disturbed grief: prolonged grief disorder and persistent complex bereavement disorder. *British Medical Journal*.
- [2] Mind. 2013. We still need to talk. A report on access to talking therapies.
- [3] Eduard Fosch Villaronga. 2019. "I Love You," Said the Robot: Boundaries of the Use of Emotions in Human-Robot Interactions. *Emotional Design in Human-Robot Interaction, Human-Computer Interaction Series*.
- [4] Owen Harris. 2013. Be Right Back.
- [5] Casey Newton. 2016. The Verge. <https://www.theverge.com/a/luka-artificial-intelligence-memorial-roman-mazurenko-bot>.
- [6] Marius Ursache. 2015. Medium. <https://medium.com/@mariusursache/the-journey-to-digital-immortality-33fcbd79949>.
- [7] HereAfter Inc. 2020. HereAfter AI: Your Life Story & Voice, Forever. <https://www.hereafter.ai/>.
- [8] Denis Dutton. 1988. The Cold Reading Tehnique. *Experientia*. 326-32.
- [9] Shafquat Hussain, Omid Ameri Sianaki and Ababneh Nedal. 2019. A Survey on Conversational Agents/Chatbots Classification and Design Techniques. *Web, Artificial Intelligence and Network Applications. Advances in Intelligent Systems and Computing*.
- [10] Rakesh Agrawal, Tomasz Imielinski and Arun Swami. 1993. Mining association rules between sets of items in large databases. *ACM SIGMOD Record*. <https://doi.org/10.1145/170036.170072>
- [11] Rakesh Agrawal and Ramakrishnan Srikant. 1994. Fast Algorithms for Mining Association Rules. *Proceedings of the 20th VLDB Conference*. Santiago, Chile.

- [12] Justin Smith. 2005. Statisticool.
<http://www.statisticool.com/aura.htm>.
- [13] Larry King. 1978. Larry King Show.
- [14] Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics, ACL 2011.