INFORMING THE CONSTRUCTION OF NARRATIVE-BASED RISK COMMUNICATION

by

Henry William King

A thesis submitted in partial fulfillment
of the requirements for the degree

of

Master of Science

in

Computer Science

MONTANA STATE UNIVERSITY
Bozeman, Montana

November 2019

ACKNOWLEDGEMENTS

This work would not be done without the help of many, many people. It was funded in large part by a grant from the National Science Foundation, Award #1635885. I would like to thank my advisor, Dr. Clemente Izurieta, for an excellent opportunity and even better guidance. I thank Dr. Mike Witte and Dr. Indika Kahanda for their help and instruction. I would like to acknowledge the help and inspiration of the other members of the Narrative Risk Communication group. Finally, I thank my family, who each supported me in a different but important way.

TABLE OF CONTENTS

TABLE OF CONTENTS – CONTINUED

LIST OF TABLES

LIST OF FIGURES

LIST OF ALGORITHMS

NOMENCLATURE

In this thesis, we will use the established Natural Language Processing labels for referring to examples of text. The smallest unit of measurement will be a term, which is one or more words treated as a single unit. A term count is the total unique occurrences of a term, either in a single document or over a corpus. A document is a collection of terms. A corpus is a collection of documents. The work in this thesis was centered around three corpora; Hero, Victim, and Villain.

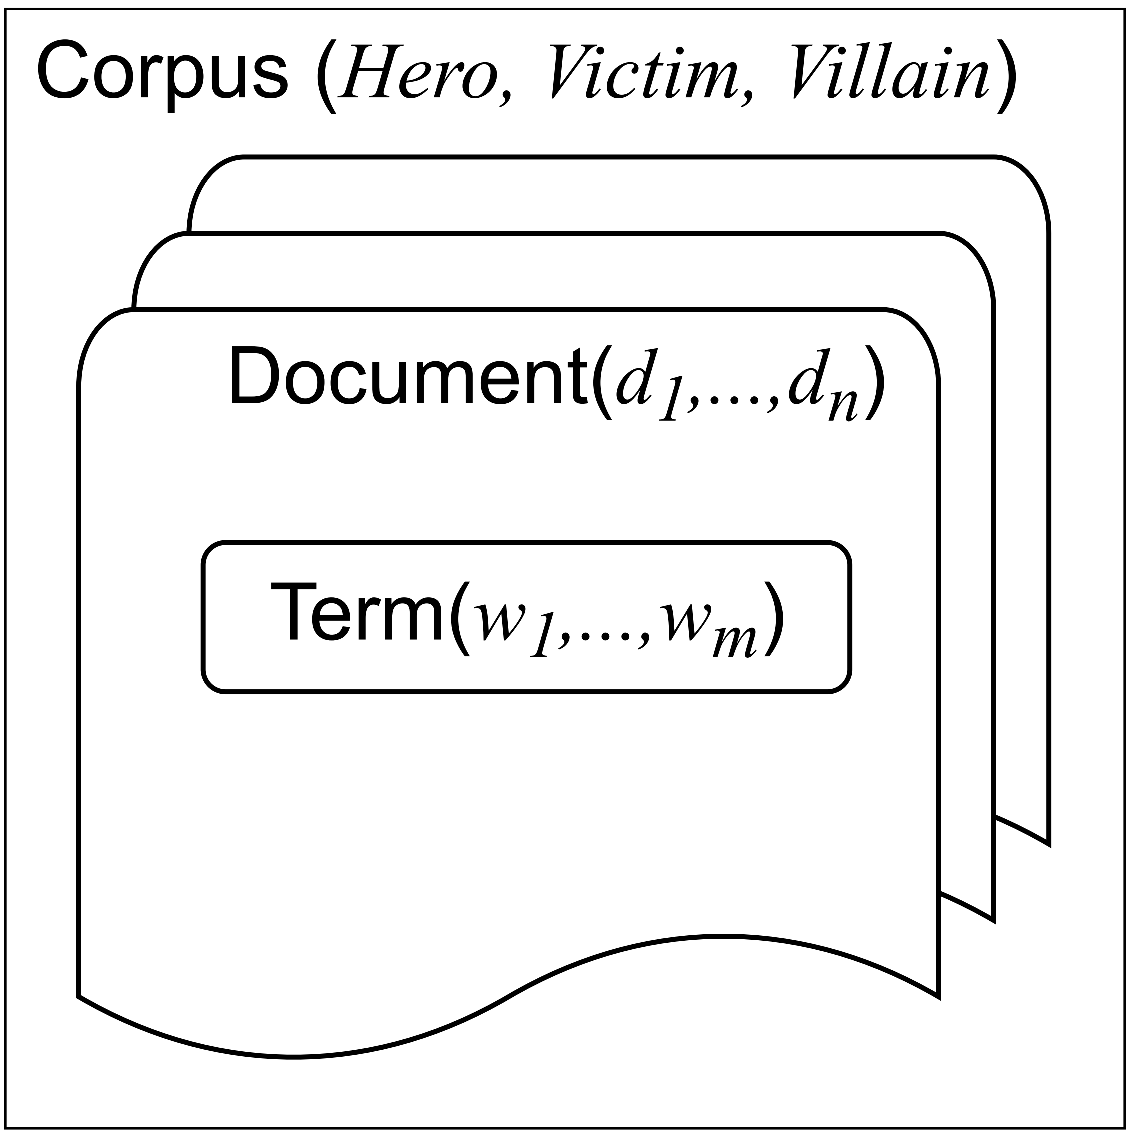

Figure 0.1: The hierarchy of a corpus

ABSTRACT

The current communication of flood risk by government agencies and the scientific community to the citizens living in the floodplain is ineffective. Using the Narrative Policy Framework (NPF), this communication can be enhanced through the use of Hero, Victim, and Victim to Hero character-based narratives. This thesis describes the methods used to inform users of the NPF to construct and test narratives using computational methods. Four natural language processing tasks are described; topic modeling, sentiment analysis, classification, and term frequencies. It was found that using the difference of transformed relative term frequencies produced an adequate vocabulary for each style of narrative. The narratives constructed from these vocabularies were used in work that sought to formalize the narrative construction process and in focus group studies which found that narrative-based scientific messages increased affective response versus traditional scientific messaging.

INTRODUCTION

When communicating scientific information and policy decisions to the public, scientists and government officials face the task of how best to clearly communicate precise information that may not necessarily be emotionally evocative. The information may even be arcane and dense to the target audience. This can result in the information being ignored, misunderstood, and possibly manipulated. To help ensure that information is being transmitted clearly and correctly, it is desirable to communicate in a way that resonates with your audience. In this thesis, we will examine text analysis methods for improving and informing the construction of such messages using language generated from the target audience.

Text analysis has been a pursuit of linguists, computer scientists, data managers, and more for a very long time. It is a broad term that encompasses most of the methods used to extract data from human generated text. Many of the areas covered in this work have early roots before the 20th century. As the means of collecting and storing textual data improved, so did the methods and interest analyzing it. There were many topics covered in this research that themselves contained many diverse and complex subgroups. This broad approach to gathering resources allowed us to see the many avenues taken by researchers performing similar tasks. Text analysis is wielded in different but specific ways dependent on the goals of the user.

This thesis was guided by a project that sought to improve the risk communication between government officials and citizens living in the flood plane of the Yellowstone River in South Central and Eastern Montana. Risk communication's primary goal is to reduce vulnerability by aligning the scientific communication of risk and the public's perception of that risk, therefore increase preparedness [14]. Preparedness includes making decisions such as purchasing flood insurance, fortifying your property against flooding, having an evacuation strategy, or keeping sandbags on hand. In particular, the project was concerned with the preparedness for an extreme flooding event where the probability of the event is low, but the potential consequences (economic, social, physical, emotional) are high. To increase preparedness, the project explored using narrative language in risk communication. A narrative is a rhetorical communication device, image, or story used to evoke an affective response [28]. Affective responses are emotional reactions (positive or negative) that will influence the behavior of the individual [66]. The general belief is that information about risk communicated through a narrative will generate an affective response. The affective response will in turn influence the behavior, in this case preparedness, of the individual citizen. Instead of taking the first part of that causal chain, that narratives will illicit an affective response, for granted, this project sought to validate that generated response through focus groups tested on sample narratives. The results of these tests informed the creation of a final mail survey containing narrative language that was mailed to the targeted communities.

Narrative Policy Framework

While there is no established theoretical method for constructing narratives, the Narrative Policy Framework (NPF) [28] does establish a method for identifying and testing different types of narratives. The NPF describes the basic elements of a narrative as characters, plot, setting, and moral of the story. However, there is an emphasis on the characters as the most fundamental units. These characters perform roles; villain, victim, and hero. The villain is defined as the agent doing the harm. The victim is defined as the agent who is harmed or fears being harmed. The hero is defined as the agent who provides promise of relief from the harm and a solution to the problem.

Narratives serve to generate affective responses in their audience. An affective response is some emotional reaction to the content. The idea being that the audience will be more engaged, invested, and concerned about the information being communicated when an emotion response is invoked. Narratives are said to also transport the audience into the story which in turn makes the audience more susceptible to persuasion. [28].

With the character roles in mind, we also wanted to test the affective response of communicating flood risk information in the form of Hero, Victim, Victim-to-Hero, and characterless narratives. In the narratives with character, an agent would play one of the three roles and their story would include the flood risk information, how they reacted to it, and how their actions increased or decreased the impact of the flood event. There was a secondary variable for each type, certainty or probability, that dictated how the inevitability of an extreme flood event was communicated. Certainty narratives portrayed flooding as definite and inevitable where probability narratives portrayed the risk in probabilistic terms. With a certainty and probability narrative of each of the four types, eight narratives were constructed in total. The affective response generated by these narratives was tested on focus groups held in the targeted communities and measured with a dial response Perception Analyzer [12].

Citizen Interview Transcripts

The target audience of the computationally enhanced narratives are residents of communities at risk of flooding along the Yellowstone River in Montana. In order to better communicate with this audience, an attempt was made to construct the narratives using language that they themselves would use when talking about flood risk and flood events. To that end, researchers conducted semi-structured in person interviews in the cities of Livingston, Miles City, and Glendive. These interviews contained questions concerning flood risk, experience, and preparedness. These spoken interviews were then transcribed by human hands to text transcripts. An example interview transcript is provided in Appendix B.

The narrative language from the transcripts was identified by researchers with domain knowledge using the Narrative Policy Framework. The narrative examples served as individual documents and were sorted into Hero, Victim or Villain corpora. These three corpora of natural language are the data sets for this project, referred to collectively as the flood corpora. These were coded with strong intercoder reliability, with "hero" coding have a Cohen's kappa of 0.8892 and "victim" having 0.8796 [62].

Documents of the flood corpora varied wildly in size and number. Some narrative examples were one sentence long, others multiple paragraphs. From smallest to largest, the Hero corpus contained 472 documents, the Victim 747, and the Villain 948. It should be noted that these discrepancies cannot lead to the corpora being called homogeneous and the overall flood corpora is too small to be considered highly representative of narrative language used by the target audience on the flood risk domain. Ideally the three subcorpora would be of roughly equal size and be a much larger sampling of the lexicon. You would also want a sample of language use outside of the flood domain from the same target audience in order to compare term usage against. These were concerns that guided us away from relying on more complex methods of text analysis, fearing that they would contain too much variance on such a small sample size to be reliable. We choose methods that returned reliable data at smaller sample sizes, with the potential to scale up when more information became available.

Motivations

This project was done by a multi-disciplinary team of researchers; hydrologists, economists, social and political scientists, geographers, and computer scientists. Our role was twofold; to enhance the creation of the sample narratives by illuminating as much as possible about the flood corpora, and to determine the vocabularies used to create the narratives for these tests. The task required that we use computationally enhanced techniques. We drew from established approaches in the Natural Language Processing (NLP) literature to find the best way to rank words used in the corpora to inform the construction of Hero, Victim, and Victim-to-Hero narratives. This was an iterative process that explored many avenues of text analysis to determine which methods would be most helpful to the group on this particular project.

Contributions: Enhancing Narratives

This thesis describes the data set used by the project in detail. We describe the methods used to collect, and represent the data for text analysis methods and in doing so describe current preprocessing steps that can be taken for natural language processing(NLP) projects. Then we describe in detail the four main areas of text analysis that were explored during the course of this project; Text Classification, Sentiment Analysis, Topic Modeling, and Term Frequencies. For each method, we describe the predominant methods and then go into detail on the methods used for this particular project. Finally, we discuss the narrative construction process and the impact the narratives had on affective response once that process was informed by computational methods.

6

TOOLS

Our primary role in this project was to find computational means of enhancing the narrative construction process for future scientists and government officials. Those researchers will most likely be political or social scientists, economists, hydrologists, and layman who may or may not be confident computer programmers. Although various tools were used to analyze the corpora, a particular emphasis was placed on tools and methods that would be relatively easy to access and use by those doing similar narrative construction projects in the future, with or without extensive knowledge of computer science.

The tools used at different stages of the process flow are depicted in Figure 2.1. The Hero, Victim, and Villain corpora came to us as multiple text documents. We then applied text preprocessing, discussed in Chapter 3, to prepare the data for the four text analysis methods shown in the middle column of Figure 2.1 and to create the Term Document matrices. These methods are outlined in detail in chapters 4, 5, 6, and 7. The output of each text analysis method is shown in the right hand column.

The Rstudio integrated development environment and the R programming language [71] were the primary tools used to research the computational enhancement of the narratives. Specifically, we used the `tm` [15] package, an R package for text mining and text analysis, to perform data preprocessing and to create the complementary data structures. The R package `qdap` [54] is designed to facilitate qualitative data analysis and natural language processing. It contains useful functions for manipulating data frames in R, as well as computing the frequencies and sentiment of words used in the narrative language segments. We used Microsoft Excel to visualize and communicate the data structures built in R to other group members. We constructed all visual plots using the R package `ggplot2` [77].

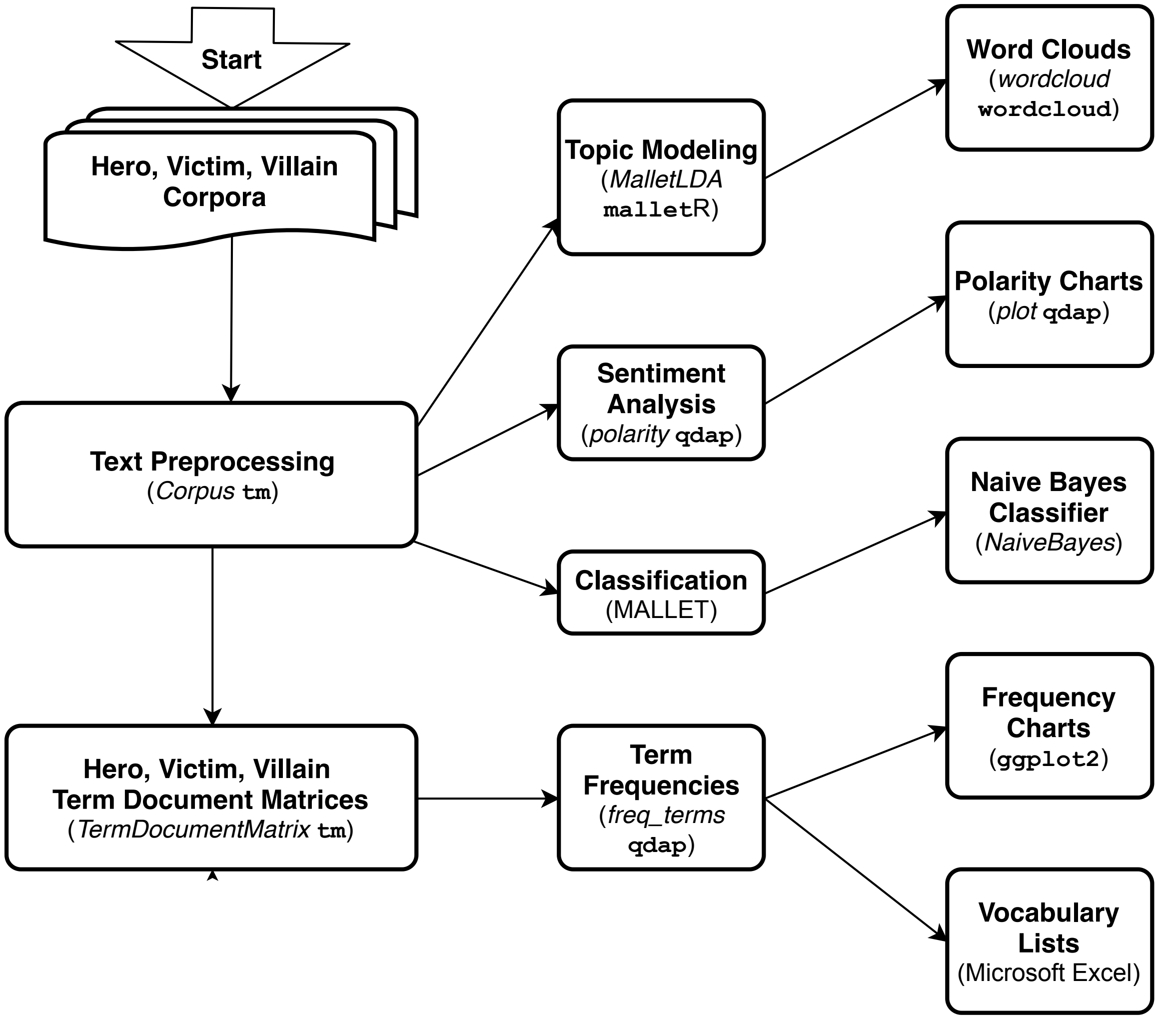

Figure 2.1: Information and task flow diagram. Each distinct tool or function is depicted in italics with the software package listed afterward where applicable.

In order to build the naive Bayes classifier, we used the Machine Learning for Language Toolkit (MALLET) [43]. This is a Java package that allows for the rapid building and testing of machine learning models of natural language. Its R complement, the `mallet` package [44], was used to generate the topic model. These topic models were visualized using the `wordcloud` R package [16].

PREPROCESSING

Natural language presents a combinatorial problem for computers, who view each character, word, sentence, and paragraph as an individual discrete feature for consideration. This high dimensionality can dramatically slow down text analysis. In order to speed up the downstream tasks, one usually performs some kind of preprocessing. Preprocessing is focused around reducing the number of features without losing too much information and reaching a vectorized representation for text analysis models.

Background

There are five predominant preprocessing methods; conversion to lowercase, character scrubbing, stop-word removal, stemming, and tokenization. While others exist, these methods are the most commonly used across NLP tasks. These steps are very important for feature extraction, feature selection, and the accuracy of the downstream tasks.

Lowercase conversion quickly reduces the number of features considered, as words like “he” and “He” would otherwise be seen as unique terms. The loss of information is minimal as most uppercase instances are dictated by the start of a sentence or an occurrence of a unique proper noun. That noun would remain a unique term regardless of case. Loss can happen when a word is a homonym with a proper noun, both terms then being treated as one. This is the case with names like “Art” and “Barb.” Lowercase conversion has been found to be effective in most situations regardless of the downstream task [74].

Character scrubbing removes any unhelpful terms from the corpus by scanning the documents at the character level, letter by letter. This includes punctuation and other irrelevant terms such as URL markers and numbers. The speaker is generally not able to vary these terms and characters through word choice and as such we usually do not want them to influence the resulting models. Similarly, stopword removal eliminates extremely frequent terms used in natural language such as “a”, “the”, and “that.” Again, because their presence in natural speech is dictated by grammar rules, not by word choice, they are unhelpful in the context of this analysis.

Stemming is a preprocessing step where words like “running” and “runner” would be stemmed to the root word “run.” It is usually considered an optional task because a large amount of information is lost. An alternative to stemming is lemmatization. Lemmatization reduces terms with similar meanings to the same term. For example, “am,” “is,” and “are” would reduce to “be.” These methods can dramatically decrease the dimensionality resulting in faster run times for particularly lengthy algorithms.

Tokenization, sometimes called segmentation, is how one determines the term length for consideration in downstream tasks. Term length refers to how many words are considered in a the smallest unit of the corpus. This can be anything from unigrams, with one word per term, to segmenting by paragraph. One approach to segmentation can simply be to gather all the text of the corpora into one document [30]. If you consider segmenting on the sentence or phrase level, you would need to alter your character scrubbing step to leave the punctuation that creates those boundaries intact. Tokenization transforms the documents into feature vectors, which hold the counts for each unique term in the corpus as a sequence of integers. These feature vectors are usually combined into a term document matrix, where there is a row for each unique term found in the corpus and a column for each document. Figure 3.1 shows an example of this representation. The cells in the term document matrix store the total number of occurrences of an unique term in a document. This is referred to as the term count or term frequency. This creates a very large and sparse matrix. From this representation, we can gather many interesting statistics about the corpus including the total number of unique terms in a corpus, the total number of words in a corpus, and the term count of a term over the whole corpus.

Document

Term	1	2	3	4	5	6	7	8
city	4	26	2	0	19	14	40	7
dike	9	11	21	14	20	35	32	30
flood	13	4	12	19	37	17	0	38
house	22	35	5	26	5	0	16	17
insurance	5	17	8	37	17	3	13	23
mean	5	14	16	15	38	27	24	29
one	12	9	24	6	3	10	25	18
people	38	26	40	29	12	36	15	29
right	5	39	23	4	18	22	31	29
river	31	1	1	25	36	9	22	17
think	29	4	39	16	8	34	34	11
water	31	26	25	4	10	29	26	32
way	28	31	27	31	5	25	3	3
year	38	2	39	24	14	2	37	15

Figure 3.1: An example term document matrix

Methods Implemented

From the 45 semi-structured interview transcripts, researchers pulled 472 Hero, 747 Victim, and 948 Villain narrative examples. These examples ranged from one sentence to a paragraph in length. This raw data was stored in three separate documents according to corpora, with each narrative example separated by headers of metadata generated by the program used to hand code the full collection of transcripts by narrative character, Nvivo [39]. The files were converted to plain text and read with the `readLines()` function from the base R package. A regular expression was used to find the headers and extract the language into new documents, one for each narrative chunk. Using the `tm` package in R, a Corpus object was created for each label by passing the directory of these new documents to the `Corpus()` function.

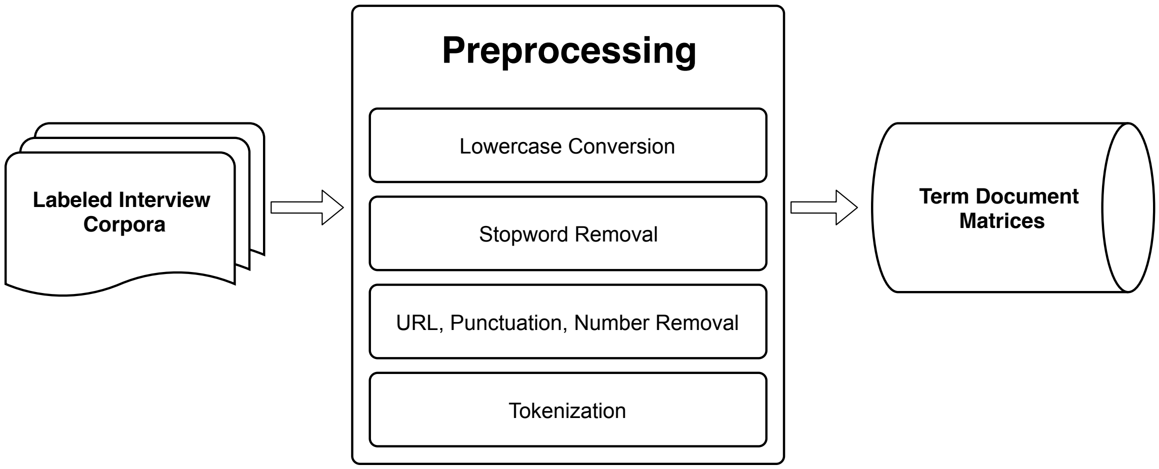

Figure 3.2: The preprocessing steps we applied to the flood corpora.

Each of the corpora went through four of the five preprocessing steps mentioned in the Background section above, shown in Figure 3.2. These were accomplished with the `tm_map` function, which takes in a Corpus object and a function to apply to that corpus. The text of all the documents was converted to lowercase and a variety of unhelpful characters were scrubbed from the data. Further, the `tm` package default list of 174 English stop-words was removed, as well as a custom list tailored to interview transcripts and the flood risk domain. The custom stop-words list included terms like uh, uhm, hmm which are important in vocal speech but are not relevant to our study. Stemming was considered in the analysis of the data, but upon examination of the results there was too much detail lost for our purposes. Words like "neighborhood" and "neighbors" held different connotations in the hero and victim corpora for example. The distinction held important weight in the flood lexicon. The researchers suggested this had to do with thinking of the collective, the "neighborhood" versus thinking of the individual, my "neighbors". Stemming removed these kind of distinctions. For tokenization, the documents were broken into unigram terms, with one word per term. Bigram and trigram models were explored, but due to the small size of the data set, they did not prove very useful. Unigrams also gave the most flexibility for word choice when constructing narratives, which would be done by human hands. By segmenting to unigrams, we created a so called "bag of words" representation. This representation assumes that the important relationship between the terms occurs at the document level, not the sentence level, as the intra-sentence associations are lost.

Table 3.1: The effect of preprocessing on the corpora.

Corpora Statistics					
Corpus	Documents	Unique Terms	Raw Word Count	Preprocessed Word Count	Reduction
Hero	472	3213	34481	15368	55%
Victim	747	3662	47224	20467	57%
Villain	948	4539	68017	29766	56%

Results

Table 3.1 shows how dramatic the impact of preprocessing can be. The steps we implemented did cause some minimal loss of information, such as losing the distinction between a proper noun and its homonym through lowercase conversion. But we achieved a reduction in corpora size of just over 50%. While the stopword list is only about $\frac{1}{10}$ the size of the number of unique terms list, those words make up over half the words in the corpora. However, this does not correlate directly to a 50% reduction in dimensionality as that is tied to the number of unique terms.

Tables 3.2 and 3.3 show the top terms before and after preprocessing. We find a much more informative list of frequent terms in 3.3, with river topping the list. River was ranked the 23rd most frequent term before preprocessing. These results show how much of our language is not dictated by word choice, and rather is controlled by grammar rules and social cues.

Table 3.2: The top 5 most frequent terms in the Hero corpus before preprocessing.

Term Rankings Before Preprocessing	
the	1979
and	1777
that	1105
you	945
they	618

Table 3.3: The top 5 most frequent terms in the Hero corpus after preprocessing.

Term Rankings After Preprocessing	
river	197
people	182
dike	158
think	143
one	136

After preprocessing, the natural language in the corpora was represented numerically as term-document matrices using the `TermDocumentMatrix` function. This served as the basic representation used to perform the downstream text analysis tasks of topic modeling, classification, sentiment analysis, and term frequency.

TOPIC MODELING

Topic modeling is an approach used to find the unifying themes that tie documents together within a corpus using quantitative measurements. These themes may or may not connect to abstract topics known before analysis. For instance, you have a corpus of high scoring customer reviews on your product. A topic model may find a theme you anticipate relating to a particular feature of the product, but may also find one relating to your customer service that you did not foresee. It may also find topics that don't particularly connect to abstract social constructs and instead represent the statistical relationships between terms in the corpus. In its primary use case, a topic model can be used to find the latent unifying concepts of a corpus too large for humans to read practically.

In our case, the researchers who conducted and coded the interviews had extensive knowledge of the text. We expected the topic model would serve two purposes. First, that it would reinforce the themes they were already aware of. They found that many interviewees were concerned with ice jams, economic damage, and governmental issues to name a few. It was hoped that the model would find topics that lined up with these ideas. Second, and perhaps the more interesting purpose, we hoped to find topics that might not have been apparent until sorted by narrative label. The researchers experienced the conversations as they happened but the model is seeing the text sorted by Hero, Victim, Villain. The goal was to find latent themes within the labels that would bring to light new information about the corpora.

Latent Dirichlet Allocation

In keeping with our primary concerns about the future users of our work, we chose to implement an Latent Dirichlet Allocation(LDA) topic model to explore what

1. Choose k, the number of topics
2. Choose $N \sim \text{Poisson}(\xi)$
3. Choose $\theta \sim \text{Dir}(\alpha)$
4. For each N words w_n:
   (a) Choose $z_n \sim \text{Multinomial}(\theta)$
   (b) Choose w_n from $p(w_n \mid z_n, \beta)$

Figure 4.1: The generative model for Latent Dirichlet Allocation

kind of impact a topic model could have on the narrative construction process. LDA was first outlined by Blei et al. [6] in 2003. They began by laying out the generative model, shown in Figure 4.1, that a LDA model assumes about how a document is created. The first assumption is that the document length N is determined by drawing from a Poisson distribution of document lengths. Then a topic mixture θ for the corpus is selected by drawing from a Dirichlet distribution $\sim Dir(\alpha)$ of topics. This is the probability of a topic occurring in a document from that corpus. The number of topics k is assumed to be fixed and known prior to the creation of the document. Then, for each N terms in a document, a topic z_n is drawn from θ. A term is selected from the multinomial probability conditioned on that topic $P(w_n \mid z_n, \beta)$. α and β are hyper-parameters specific to the corpus. α is the k dimensional vector containing the parameters for the Dirichlet distribution of topics. β is a $k \times V$ matrix containing the probabilities for every term in the corpus vocabulary V in every topic k.

In this way, their assumed generative model produces every document in the corpus by drawing a topic and then drawing a term from that topic's distribution over the fixed vocabulary. This is a mixed membership model allowing multiple topics to be present in the same document.

Collapsed Gibbs Sampling

In order to train an LDA topic model, the model needs to learn the Dirichlet parameter α and the matrix β. However, the only observed data are the words in the corpus organized by document. Working backwards from the terms, the model uses a collapsed Gibbs sampling method to infer those parameters for the corpus.

Gibbs sampling aims to construct a Markov chain that converges to the posterior distribution. For LDA, the posterior distribution is the topic distribution β. The algorithm steps through the chain multiple times, sampling and updating the topic assignments for each word in the corpus based on the topic assignments of all the other words at that time.

Algorithm 4.1 describes the collapsed Gibbs sampling algorithm [11] for LDA. Each term in the corpus is first assigned to a random topic. These assignments are stored in vector $\mathbf{z}$ of length N, the number of words in the corpus. Given the initial random assignments, track 3 counts; the number of documents assigned to each topic $n_{d,k}$, the number of times each word is assigned to each topic $n_{w,k}$, and the number of times a document is assigned to a topic n_k. For each word i, the topic assignments are updated based on the probability of a topic given a document and a word given a topic conditioned on the topic assignments of all the other words. These adjustments continue until the model converges or for a set number of iterations. Although convergence is theoretically guaranteed for Gibbs sampling, it cannot be calculated when it will occur. A minimum threshold can be used for the number of topic assignment changes per iteration.

Algorithm 4.1: Collapsed Gibbs Sampling for LDA

Input: words $\mathbf{w} \in$ documents $\mathbf{d}$

Output: topic assignments $\mathbf{z}$ and counts $n_{d,k}, n_{k,w}$, and n_k

Randomly initialize z and increment counters

for each iteration **do**

  for $i = 0 \rightarrow N - 1$ **do**

    $word \leftarrow w[i]$

    $topic \leftarrow z[i]$

    $n_{d,topic}$ -= 1; $n_{word,topic}$ -= 1; n_{topic} -= 1

    for $k = 0 \rightarrow K - 1$ **do**

      $p(z = k \mid \cdot) = (n_{d,k} + \alpha_k)\frac{n_{k,w}+\beta_w}{n_k+\beta \times W}$

    end for

    $topic \leftarrow\sim p(z \mid \cdot)$

    $z[i] \leftarrow topic$

    $n_{d,topic}$ += 1; $n_{word,topic}$ += 1; n_{topic} += 1

  end for

end for

return $\mathbf{z}, n_{d,k}, n_{k,w}, n_k$

Using the output from this algorithm, the α vector for the Dirichlet distribution of topics and the topic by term probability matrix β can be easily calculated. From there, the topic model can now simulate its assumed generative model for the corpus.

Topic Model Applications and Background

One can use a trained topic model for a variety of interesting tasks. In the political science domain, a Bayesian Hierarchical topic model was used to model the communication between Senators and their constituents [18]. Similarly, a Structural Topic Model was used to identify latent themes within survey responses [55]. For an application in the humanities, LDA creator David Blei described a scholar organizing an archive of texts according to their domain knowledge and then tuning a topic model to that collection [5]. As theories are examined, the model is updated and used as additional evidence. He saw this iterative process as a way to highlight hidden structure of the text of those collections in the digital humanities. LDA has been used in an attempt to improve the ability of users to find new feeds to follow on Twitter by distilling the text within tweets to the language model [51]. Topic modelling was used as the basis for a comparison of the news topics distributed by Twitter compared to those distributed by the New York Times [83]. It was found that the topics covered were indeed different, with Twitter covering areas that received less traditional media attention as well as spreading important world news stories.

There have been many projects dedicated to improving the accuracy and validity of topic models since their inception. While the LDA model uses a unigram model, viewing the documents as a bag of words, there has been work showing that the results can be improved with bigram terms to incorporate word order [75]. It is clear that the generative model assumed above is not the actual process for generating human language. Lowe and Benoit [38] challenged this assumption and worked toward constructing a framework for determining how wrong a model is and validating its semantic meaningfullness.

Results

Using the `malletR` R package [44], a topic model was built using Latent Dirichlet Allocation for each corpus. After some tuning, we settled on 7 existing topics and 400 iterations of training. Although this is a subjective process, it produced topics that seemed both distinct from each other and relevant to the subject of flood risk.

As explained earlier, the topics themselves contain a list of every term that appears in the corpus and a value that indicates the probability of that term being in language of that topic. When sorted by that value, you can get an indication of the importance of a term to the topic. In order to visualize and communicate the results of the topic model to the group, the top one hundred words from each topic found in the model were displayed using word clouds. Figure 4.2 shows an example of a topic found in the Victim corpus visualized in a word cloud. The words in the topic are displayed by size and color. Larger size indicates importance, with color denoting tiers of words with similar importance to the topic.

Many of the topics found by the model did seem to mimic the language used in interview responses, although the models are not easy to interpret. It was difficult to evaluate the strength of a particular topic against the others, other than through subjective human measures. Most of the consideration comes down to questions like,"Does this look right given what we know about the domain?" Topic models output a term ranking for each topic that relates the prominence of every term to the topic. The difficulty comes in relating and visualizing this output for the qualitative human interpretation of the topics. It has even been found that the more statistically precise a topic model is on a corpora, the less semantically meaningful are the topics

Figure 4.2: Word cloud of topic within the Victim Corpus

that it generates [10]. So, even when they did appear similar to examples from the data set, they did not further our understanding of the data set. These issues made it difficult to apply the information in a way that is useful for constructing narratives. But, upon examination, the topics did appear to cover areas that were expected by the researchers who conducted the interviews. There were economic damage, natural occurrences, and social construct based topics, among others. This helped to validate the internal biases of the group.

SENTIMENT ANALYSIS

Digital text data is being generated at an increasing rate through internet avenues like customer reviews, social media, and blog posts. This data can quickly reach an overwhelming size if one wanted to extract information like customer preferences, opposition or support to a political idea, or reaction to a current event from all the available documents. It is no surprise that work has started towards an automatic and accurate way to determine these values using computers. This problem is defined as measuring the tone or emotion of a text segment and is referred to as sentiment analysis.

Background and Types of Opinion Mining

There are two major approaches to calculating emotional tone. The first involves classifying text segments in emotional bins. These bins are labeled by an emotion like happy, sad, anger, disgust, or fear. The second approach calculates a polarity score for the text segment, with the range centered at zero. Emotionally negative segments score below zero and emotionally positive segments score above zero. While both methods are considered under the umbrella of sentiment analysis, the second method is usually more specifically referred to as opinion mining while the first retains the sentiment analysis name [9]. Both approaches attempt to measure the sentiment of a text, and there is usually a further division in how that is accomplished, whether through dictionary look-ups or classification of sentiment through training a machine learning algorithm for the task. While these subfields remain fairly separate, there have been projects that blended the two. Singh et al. [64] explored the reactions to the demonetization of the 500 and 1000 rupee banknotes in India. They recorded their measurements into buckets like sentiment analysis but with positive and neutral labels like opinion mining.

Sentiment Analysis experienced a boom during the early 2000s thanks to the proliferation of text generated on the internet. In fact, one review found that 99% of the papers on this topic have been published since 2004 [46]. Early work in the field began to uncover the pitfalls of trying to calculate such an abstract concept from human language. Movie reviews were one of the most widely available and human scored sources of text available on the early web. Bo et al. [48] stated that sentiment required more understanding than topic based classification. They brought up the difficulty of measuring sentences without any negative words. Take this sentence for example, “How could anyone sit through this movie?” It is clearly negative but contains no directly negative terms. Rhetorical questions and term ambiguity present consistent difficulties for sentiment analysis.

Much of the early work centered around measuring sentiment as polarity, meaning some value on scale centered at zero. Hu and Liu [73] used WordNet, a word sense disambiguation tool, to determine the orientation of the adjectives in a sentence as positive or negative [23]. Then the sentence was scored by the counts of positive or negative opinion words. They went on to develop the Opinion Observer [37] system for calculating the polarity of customer reviews on different products segmented by different topics like size, weight, and price. These queries concerned the sentiment of phrases uttered about particular keywords. This has itself been explored as another pitfall of sentiment analysis. In Targeting Sentiment Expressions through Supervised Ranking of Linguistic Configurations [29], Kessler and Nicolov found that 14% of the references in a document about a particular idea were through pronouns. These phrases or sentences would not be returned when searching by keywords. This motivated work to resolve these coreferents [2]. By seeking to resolve coreferents like “lawsuit”, the “document” and “it”, they hoped to improve retrieval tasks and therefore the sentiment measurement over the entire corpus.

Many projects involved a supervised data set of opinion material that was fed into a machine learning classifier. The type of classifier has varied over time. Support vector machines, Naive Bayes, Maximum Entropy, were common in the beginning [48] [45]. These classifiers work well with linearly separable data and natural language is usually linearly separable [26]. Different data representations have been explored as well. Bo et al. [48] explored using binary vectors, 0 indicating absence of a term in a document and 1 indicating presence, for Naive Bayes and SVM models as opposed to vectors of terms counts. They found binary vectors to improve performance for the sentiment classification task. Symeonidis et al. [69] found that lemmatization, replacing repeated punctuation and contraction, and removing numbers also aided the classification task. Zhao [25] found similar improvement removing URLs, stopwords, and numbers.

Many packages have been developed to perform and run sentiment analysis from a scored lexicon of emotional terms. One such, SentiWordNet [3], uses a supervised lexicon of terms that have been assigned polarity values. This is a common occurrence in the field, but it does have its disadvantages. There is an assumption that the vocabulary is relevant and representative of your target domain. If you want a truly representative data set, researchers with specific domain knowledge would be required to score the lexicon. This can be difficult and costly to achieve, mainly due to the lack of available domain experts. With new tools available, studies like Warriner et al. [76] used the Amazon Mechanical Turk to score a lexicon of up to 14000 terms. While lexicons like this aim to be generally applicable to the common use of the language, it is recommended that you augment the lexicons to better fit your use in some way. These and other concerns about general sentiment lexicons have furthered research into automatic methods for building the lexicons without human subjective bias. In

Detecting Domain Dedicated Polar Words [63], the Chi squared goodness of fit test was used on a supervised data set of positive, negative, or neutral documents in a domain. The top 200 positive and negative words became the positive and negative sentiment lexicon for that domain.

Many sentiment analysis tasks involve user reviews of products that are accompanied by some additional information like a star rating or thumbs up or down rankings. These text documents are very useful for supervised text classification as they are labeled by the authors. These reviews can be used in interesting projects like studying the impact a bad review of an AirBnB room has on the price of the neighboring locations [34]. If one does not have a generous amount of labeled data, another method that has been explored is to use labeled polarity documents from one domain where you do have a large collection to inform measuring polarity in another [22]. Similarly, the majority of the labeled sentiment documents and lexicons are in English. Applying that work to other languages and cultures becomes a domain adaptation problem [9]. Of course, labeled data does come with biases. Guerini et al. [19] found gender based biases for negative words in the ANEW labeled dataset. This highlights the importance of knowing your data for sentiment analysis tasks.

It is clear that a particular sentence or phrase taken by itself is not necessarily indicative of the sentiment of the document as a whole. There has been many different attempts at incorporating additional information to account for it. Mullen and Collier [45] describe this as a continual challenge. They extracted value phrases from the text. They then gathered the average value of the phrases that came before and after the current value phrase and passed that information as an additional feature for the SVM classifier. They found increased accuracy when incorporating this data. Yang, Lin, and Chen [79] incorporated mood labels from web blogs into their classification model. Documents were scored into classes labeled HAPPY, JOY, SAD, or ANGRY.

Recent endeavours in sentiment analysis involve pulling in data from multiple media sources like film, audio, images, and text. This multimodal sentiment analysis has some interesting challenges to overcome, such as how to combine information from different mediums into vectors for Recurrent Neural Networks or Generative Adversarial Networks [40] to modeling correlation between different modalities via density matrices with Quantum-inspired Multimodal Representation [82].

Motivations

We wanted to use sentiment analysis methods to learn more about the emotional tone of language used in the three sub-corpora. Given that score, the goal was to characterize the corpora against each other and use that characterization to impact word choice when writing from a particular corpus point of view. We wanted to answer questions like is Victim language more negative than Hero language in general? If so, should we use the negative victim language to make our constructed narratives feel more authentic?

Knowing that we were going to evaluate the narratives using a Perception Analyzer [12] test on a focus group, we chose to pursue opinion mining for this task. The Perception Analyzer test involves turning a dial based on how you are reacting to what you are hearing. Higher values indicate a positive reaction, and turning the dial downward indicates a negative one. This score begins centered at 50. We believed that the quantitative polarity score would be more analogous to the perception analyzer score, and therefore more relevant to our study.

qdap's Polarity Function

Sentiment Analysis has been described as a "big suitcase" research field [8], meaning it is an NLP task that can encompass many smaller NLP tasks within it. In our case, these smaller tasks included preprocessing and text segmentation. We choose to do the analysis at the sentence level, where each sentence in the corpus would receive a polarity score. In the Rstudio programming evironment, we used the `qdap` package's polarity function. In order to preserve as much of the original grammar as possible, the corpora went through a slightly different preprocessing stage for this analysis. Stop-words and punctuation were not removed and there was no need to represent the corpus as a term-document matrix. Instead, using another `qdap` function, `sentSplit`, every sentence in the corpus is pulled out for analysis.

Using `qdap`'s `polarity` function, the sentences in the corpora were scored to a degree of positivity or negativity according to the following algorithm. First, the sentence is scanned for words that appear in the polarity frame argument, which is a combination of the default dictionary of flagged words [23] and the additional words provided by the researchers. For each word found within the polarity frame, a context cluster is created around it. A context cluster contains the 4 words prior and the 2 words after a tagged word, unless there is a comma before the tagged word. Then, only the words found after the comma go into the context cluster. The words in the context cluster are tagged according to how they modify the tagged word and are referred to as valence shifters. They can be neutral, negators, amplifiers, or de-amplifiers. The tagged word is first weighted by its value from the polarity frame, being 1 or -1. Then its weight is modified by the number and value of the valence shifters within its context cluster. We used the default, 0.8, for the amplifier/deamplifier value c. The valence shifter modification values are calculated as follows.

$$\text{Amplification, } x_{i^A} = \sum (w_{neg} \cdot w_{i^a}) \tag{5.1}$$

$$\text{De-amplification, } x_{i^D} = \max(x_{i^{D'}}, -1) \tag{5.2}$$

$$x_{i^{D'}} = \sum (-w_{neg} \cdot x_{i^a} + x_{i^d}) \tag{5.3}$$

$$\text{Negation, } w_{neg} = \left(\sum x_{i^N}\right) \mod 2 \tag{5.4}$$

These values are then combined to generate the value for a particular context cluster.

$$x_{i^t} = \sum((1 + c \cdot (x_{i^A} - x_{i^D})) \cdot w(-1)^{\sum x_{i^N}} \tag{5.5}$$

The values for all the context clusters found within the sentence are then summed and divided by the square root of the number of words in the sentence, n.

$$\text{Sentence Polarity} = \frac{\sum_{t \in T} x_{i^t}}{\sqrt{n}} \tag{5.6}$$

Results

Figures 5.1, 5.2, and 5.3 display the polarity scores for the three corpora. The upper bar chart shows the polarity for each sentence in the corpus in order. Each vertical bar indicates a sentence from the corpus. Blue indicates negative polarity and red indicates positive polarity. The lower scatter plot chart shows the scores for each sentence, centered at zero. Each red dot is a sentence in the corpus placed at its polarity score, left being negative, right being positive. The red dots are jittered vertically for better visibility. The black dot indicates the mean polarity score of the corpus.

Table 5.1 shows the cumulative polarity scores for each corpus. The average

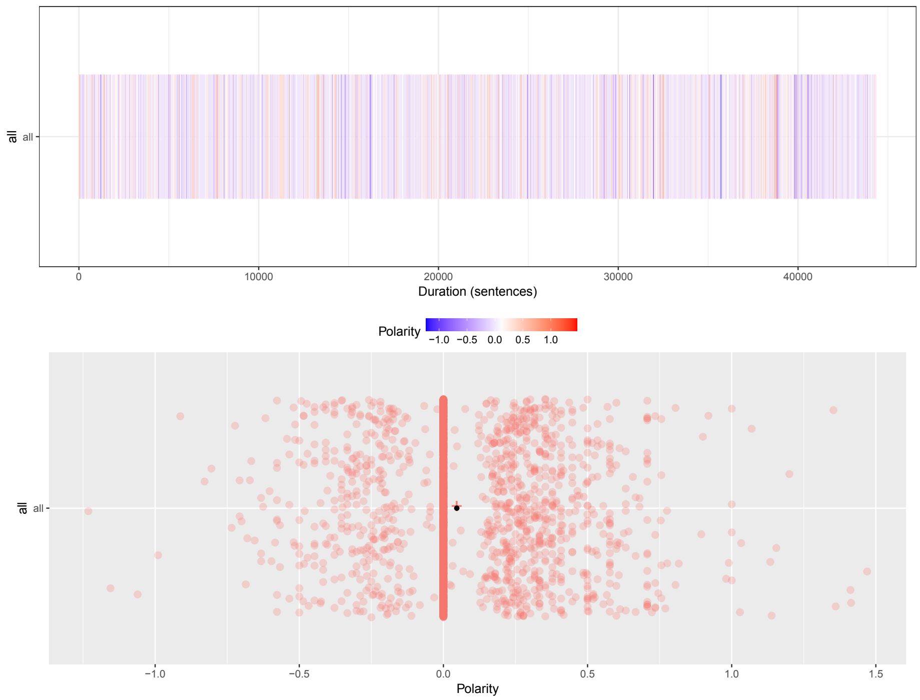

Figure 5.1: Polarity scores for every sentence in the Hero Corpus.

polarity is calculated by adding the polarity scores for each sentence in a corpus and dividing by the total number of sentences. The standard mean polarity score is calculated by dividing the average polarity by the standard deviation.

The Hero corpus did display a slightly higher standard mean polarity and average polarity than the other two corpora. However, this difference was very small. Each of the three plots show a large number of sentences at neutral polarity, with a collection

Table 5.1: The Polarity scores for each corpus.

Corpus	Total Sentences	Avg. Polarity	Standard Deviation	Std. Mean Polarity
Hero	2683	0.042	0.221	0.0189
Victim	4224	0.003	0.196	0.013
Villain	5779	0.004	0.213	0.019

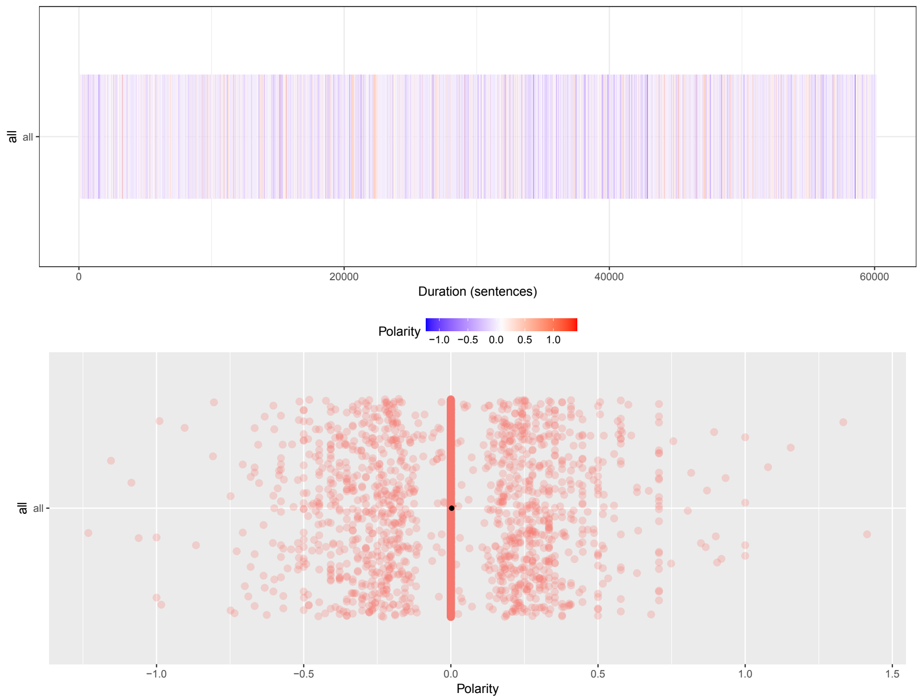

Figure 5.2: Polarity scores for every sentence in the Victim Corpus.

of sentences slightly negative and slightly positive. It is difficult to conclude any meaningful difference in polarity from these findings. This may be the result of our supplemental dictionary being too small, where too many words from the flood domain are not being considered specifically polarizing by the model. We may have received different results by segmenting in different ways as well, such as grouping all the language into one segment per corpora and comparing at that macro level. The standard mean polarity compares the variability of the three corpora. Again, they are quite close suggesting that each corpora exhibits a similar variability of emotional language across its particular range of sentiment.

While sentiment analysis was useful in gauging the polarity of one text segment over another, the narratives had to convey certain information regardless of the polarity of the words needed to communicate it. This made using sentiment analysis

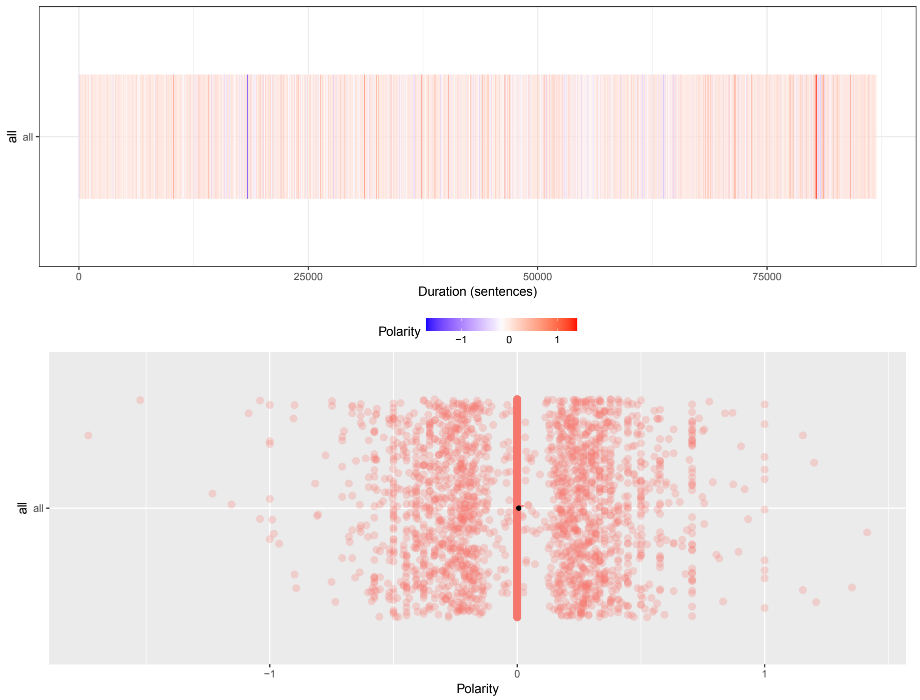

Figure 5.3: Polarity scores for every sentence in the Villain Corpus.

awkward for the narrative construction process. We expected Hero narratives to start generally neutral or slightly negative, and then grow to positive over the course of the story. Victim narratives would trend negative throughout. The Hero to Victim narrative was expected to start negative and become more positive as the narrative progressed. With these expectations in mind, we were able to score the draft narratives as they were developed with our model and used the scores to guide their construction. We generated polarity scores for each narrative, and displayed the scores of the individual sentences within the narratives for the group. A selection of these scores are included in Appendix A.

CLASSIFICATION

Human coding the flood corpora was a time consuming task. Using the NPF, the researchers needed to read all of the the interview transcripts, closely looking for segments that fit Hero, Victim, and Villain character narratives. This became a bottleneck for many of the other members of the group, who wanted to use the narrative examples for their own analysis. While this stage of the overall project was specifically focused on improving the narrative construction process, the group was considering future work that would make use of news media articles about flooding gathered from the web. This represented an extremely large potential data set and it would be one where human coding would not be cost or time effective. An automated classifier of narrative language would speed up the coding process and allow us to gather narrative examples from much larger collections of documents. To that end, we explored machine classification.

Classification is a machine learning task where unlabeled examples are given class labels by a computer model [20]. In a supervised setting, the class labels are known before hand. A data set of human labeled examples is provided and the model is trained on the features of the examples found within. In an unsupervised setting, the model examines the features of the unlabeled data and tries to discover groups of similar examples. This is usually referred to as clustering. As we had a well curated and labeled data set, we would be pursuing the supervised classification task.

Background

Text Classification methods followed the popularity trends of machine learning techniques. Early methods like naïve Bayes, support vector machines (SVM), and k-means clustering have been used for a long time [17]. These linear and probabilistic techniques were some of the first attempts at classification, then called automatic indexing [41]. More recently, SVMs have been used to speed up data collection in the political science realm by filtering results returned from a database of text documents [13]. As research into neural networks waxes and wanes, so do the methods for text classification that employ them. Neural Networks have been used to infer membership in labeled corpora [33] and to infer sentiment in text [61] [70].

Naïve Bayes has been fit to Yahoo! shopping data [49], detect spam email [80], and used for information retrival [36]. Much of the research into text classification has been attempts to improve the classification accuracy by manipulating the input. Projects have looked at representing the term counts with TF-IDF scores or other normalized methods instead of actual counts [81]. Other projects have examined the impact of different preprocessing methods on classification accuracy [67]. Kim et al. [32] proposed including different weighting heuristics to help aid the classification process. The data set plays an important role in classification, particularly in the supervised setting. Caliskan-Islam et. al [7] found that biases in the human labeled corpora manifested in the algorithm trained upon it.

Recently, projects have looked at combining classification algorithms together [24]. This technique is particularly interesting in the semi-supervised classification setting, where the data set is only partially labeled by hand. An algorithm is progressively trained on the labeled portion and use to label some amount of the unlabeled portion. Nigam et. al [47] combined Estimation Maximization and naïve Bayes in this fashion.

Naïve Bayes

For this project, the three sub-corpora made up the three classes of the training set; Hero, Victim, Villain. Each document of coded narrative language within was considered a single labeled example. Each term used within the narrative language was a feature, with their term counts as values. The ultimate goal was to classify unlabeled examples constructed by the group or articles scraped from the web. During the training phase, we used 10-fold cross validation [17] to hold out labeled examples and test the accuracy of the model. We used the classification accuracy, f1 score, precision, and recall of the classifier as performance metrics on this task.

Classifying the data required an approach that could correctly label examples of each of the three different types of documents and deal with a very large number of potential features. Considering the vast number of features, we chose to construct a naïve Bayes (NB) classifier using the MALLET toolkit. Naïve Bayes has been used successfully on NLP tasks in the past. In particular, the multinomial variation of naïve Bayes that MALLET [43] includes have been found to perform better on large vocabularies and examples of different lengths [42]. Naïve Bayes also falls in line with our concerns for the future audience of our work, being relatively easy to implement when compared to other text classifiers such as deep neural nets.

Naïve Bayes infers the class based on the conditional probabilities of the features in the input being present in each of the different classes and the relative probability of each class appearing. NB is an extension of Bayesian theory that assumes conditional independence of the features [65]. This "naïve" assumption allows the classifier to perform well when the number of features in very large, such as in NLP tasks where it is commonly in the thousands.

In order to train a NB model, a probability table is constructed for each class. The table holds the probability of a feature value occurring within that class, for all the features in the training set. In an NLP setting the features are each unique term, so the table holds the likelihood of a term being present in class of documents. In our case, the classes were the three corpora; Hero, Victim, and Villain. NB also requires learning the probabilities of the classes occurring within the training set, shown below where N_k is the number of documents of class k and N_C is the total number of documents in the training set.

$$P(\mathrm{c}_k) = (\frac{\mathrm{N}_k}{\mathrm{N}_C}) \tag{6.1}$$

Once the the likelihoods of the classes and the features within the classes are memorized and stored in the tables, the model calculates the probability of an unlabeled example with the following equation:

$$\mathrm{c}_{MAP} = \underset{\mathrm{c}_j \in \mathrm{C}}{\arg\max}\, P(\mathrm{c}_k) \prod_i P(\mathrm{a}_i|\mathrm{c}_j) \tag{6.2}$$

Where $P(\mathrm{a}_i|\mathrm{c}_j)$ is the probability of the term a_i being present in class c_j and $P(\mathrm{c}_k)$ is the probability of class c_k occurring in the training set. The unlabeled example is given the class with the *maximum a posteriori* hypothesis, denoted here as c_{MAP}.

The text was passed to the classifier as preprocessed documents, with stop words, numbers, punctuation, and capitalization removed. We did not stem as it has been found to be harmful for classification accuracy in some cases [4] [72]. These were given to the classifier as labeled examples. Naïve Bayes then assigns a log likelihood score for each label to a document given the frequency of that label within the training set, and the frequency of the features and terms within the examples of that label. The unlabeled example is assigned the label with the highest score.

Results

Our classifier was able to achieve a mean accuracy of about 70% using 10 fold cross-validation during training but achieved a fairly poor score on the test data of

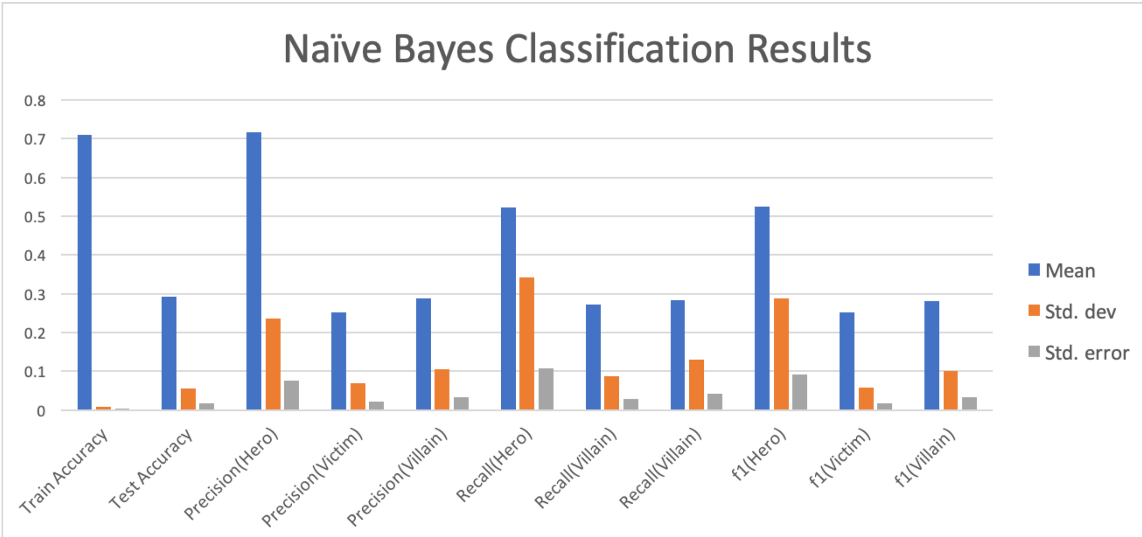

Figure 6.1: The classification results after 10-fold cross-validation. Each metric shown on a 0 to 1 scale.

about 30%. This implies that the model is likely to generalize fairly poorly. If we were to pursue naïve Bayes further, we would look into further methods of preprocessing, like stemming, that have been shown to improve classificaiton performance.

The accuracy was not reasonably close to human performance on this task so classification did not provide an adequate benefit to justify pursuing it further at this stage of the project. This work provides a good basis for the project in the future if the number of researchers with domain knowledge diminished or the data set grows in size. In these cases, even a classifier that works at 70% would save time and money. It could be used as a first pass over a large data set, the results of which could then be examined by human eyes.

Issues with Text Classification

Limitations of this approach were uncovered. The first issue was that the measurable features, term frequencies, did not necessarily divide the classes naturally.

It appeared that the classes were not well separated by language use. Using the same terms in similar relative frequencies to one another, one could produce a narrative of any of the three classes. This is a function of human language and the fact that the division of the classes was derived from an abstract social construct, narrative characters. The second issue deals with the length of the examples. They had a wide range of different lengths that may have skewed the classifier to favor the classes that contained longer examples as the longer a narrative, the more likely it was to contain a particular feature.

Similarly, this brought up issues about segmentation of future unlabeled examples. The current supervised data set contains examples that range from one sentence to many paragraphs in length. Even given a reliable classifier, its most valuable use would be on unlabeled text to identify narratives without the need for human coding. But how would one determine the text segmentation length on said unlabeled text? Narratives don't necessarily need to follow conventional textual boundaries like sentences, paragraphs, or pages. Do you submit an entire news article? Do you break by punctuation, like submitting text between quotation marks? Do you assume the unlabeled examples are of one document and you break them into examples every 500 words? These were questions that needed to be answered before any rigorous classification work was pursued further.

TERM FREQUENCIES

To a human, there are many characteristics of natural language that could serve as features to discern between documents or corpora. We can consider very abstract concepts like tone, cadence, passive or active voice, sarcasm, humor, or even vocabulary level. These concepts are social and cultural constructions that add meaning to the base definitions of the words we choose. But if we examine how we come to make these judgments about the language we read or write, we are making these decisions based off of the terms in the text. How often do they each appear, how do they appear next to each other within sentences or paragraphs. This is the motivation behind term frequency analysis. The foremost reliably countable features of corpora are words [31] and we want to use the quantifiable term frequencies as the basis for more complex and abstract text analysis methods.

Background

Given that term frequencies are some of the most overt features of natural language, term frequency analysis is a fairly old field of text analysis. Early work was focused around the best methods for indexing a collection document to best facilitate information retrieval on that corpora. Using precision and recall(how many of the results returned truly fit the query and how many of the total documents in the corpora that fit the query respectively) as heuristics, the Cornell Group began work into term ranking systems in the late 1960s and early 1970s. They imagined every document as a vector in n-dimensional space where n was the number of unique terms in the corpora. If you could visualize the ends of these vectors creating a point cloud of the document space, similar documents should be densely packed. Their goal was to determine how to find the terms that reduced the space density of the document when assigned as indexing values to documents. They called these terms discriminators, and gave each term a discrimination value [60].

Using discrimination value as a comparative metric, the Cornell Group and others began looking term ranking systems. One of their early findings related to the document frequency of the terms. When plotted against average discrimination rank, document frequency created a U-shaped graph. This indicated that terms at the head and tail of the document frequency rankings were on average poorly positioned as discriminators relative to the terms that appeared more selectively in the corpus.

At the same time Jones [68] noticed that penalizing terms by their document frequency created a ranking system that seemed appropriate for information retrieval. This is frustratingly common for many of the most accepted methods in term frequency analysis. The attempts to tie term rankings systems to strong theoretical backing mainly focusing on tieing it to Shannon's theory of information [56]. In spite of these attempts, there is no theoretically proven best term ranking system. Salton and Yang of the Cornell Group stated about term rank weighting systems that they "may be useful under some circumstances, but that it cannot be guaranteed to perform well in all environments." [58] Jones' method was at the time called term specificity but we refer to it today as inverse document frequency or IDF. The Cornell group later went on to combine the IDF method with term frequency measurements to create the widely used TF-IDF term ranking system used today [59]. TF-IDF measures term frequency but also suppresses words that appear broadly in the corpus, meaning they appear in a high number of documents. Equation 7.1 shows how TF-IDF is calculated for term t in d documents where N is the total number of documents in the corpus and N_t is the number of documents t appears in. Term Frequency is the normalized word count of the corpora.

$$\text{TF-IDF}_{t,d} = \text{Normalized Term Frequency}_{t,d} \cdot \log(\frac{\text{N}}{\text{N}_t}) \qquad (7.1)$$

While TF-IDF does not display theoretically backed properties like some probability based term ranking systems. Those systems rely on relevance information of the terms in the target documents of information retrieval as well as the non-target documents. It has been shown that this relevance information will reduce down to TF-IDF in well defined conditions [57].

While a document grows in length, each terms frequency has the potential to grow higher as well with each new word added. Moving into the 1990s, term normalization factors began to reach prominence in term ranking systems in order to account for the growing amount of digital text available for consideration [56]. Many new applications of term frequency rankings came about as well. One of these, WebWatcher [27], used TF-IDF to compare words in web pages to previous search queries and guide users from web page to web page.

With larger and larger corpora, the need for automatic corpora comparison grew. Many collections of documents had reached points beyond which it would be practical for a human to read and determine their similarities and differences to other corpora. In a big data situation, corpora comparison usually comes down to generating some similarity measurement between the two corpora. For instance, how similar is a robust and representative corpus of British English to a similarly robust and representative American English corpus. This is commonly determined by using a test known as the χ^2 test or Mann-Whitney-Wilcoxon Test [30]. Other projects view corpora comparison as a psychological modeling problem and rank their methods against human judges of document similarity [35]. Both of these tests compare results on measurements of the features of the corpora. The features are the terms and their values are frequencies, either as counts of occurrences or as binary values of present or not present (0 or 1). While these stats are useful for comparing two corpora that have been classified as different, there is an underlying attempt to determine which of the features are truly distinct. Kilgarriff postulated that "any difference in the linguistic character of two corpora will leave its trace in a difference between their word frequency lists" [30]. As with most data analysis situations, "No Free Lunch" still applies. Rayson [52] states that where expected frequencies are low, more complex methods like χ^2 should be avoided as they become unreliable.

Methods Implemented and Results

As discussed in Section 3, frequency doesn't necessarily indicate importance or distinctiveness of a term. Many term ranking schemes attempt to account for this fact, measuring other characteristics like document frequency and co-location with other frequent terms. Frequent terms that appear in almost every document of the corpus can be thought of as a function of the domain, and not necessarily the speaker's word choice. These are terms that a speaker is compelled to use when speaking about flood risk, but not necessarily distinct and important terms. We believed that these broadly used terms would appear in the narratives naturally and would not inherently endear our audience to the message. But as importance does leave its trace in frequency, the goal was to find frequent terms that define narrower regions of the flood risk lexicon used by the target audience.

Ultimately, word frequencies were the most informative text analysis for the creation of narratives. In building the narratives, we needed to best relay the message in language that would resonate with the intended audience. Measuring the term counts across the corpora would hopefully bring the conscious word choice decisions that the interviewees made to light. Then, our constructed narratives would mirror those choices when communicating risk back to the residents of the flood plane. Word frequency measurements on actual transcripts of that target audience provided the exact vocabulary the audience would use to communicate these messages themselves.

We had a very specific question to address. Given a human researcher constructing a narrative, that researcher has domain knowledge, experience with the target audience and their language, and the ability to think abstractly about the subject. But they also have a subjective bias. We needed to provide a vocabulary list that would serve as a check on the subjective bias of the researcher when constructing narratives from a particular point of view.

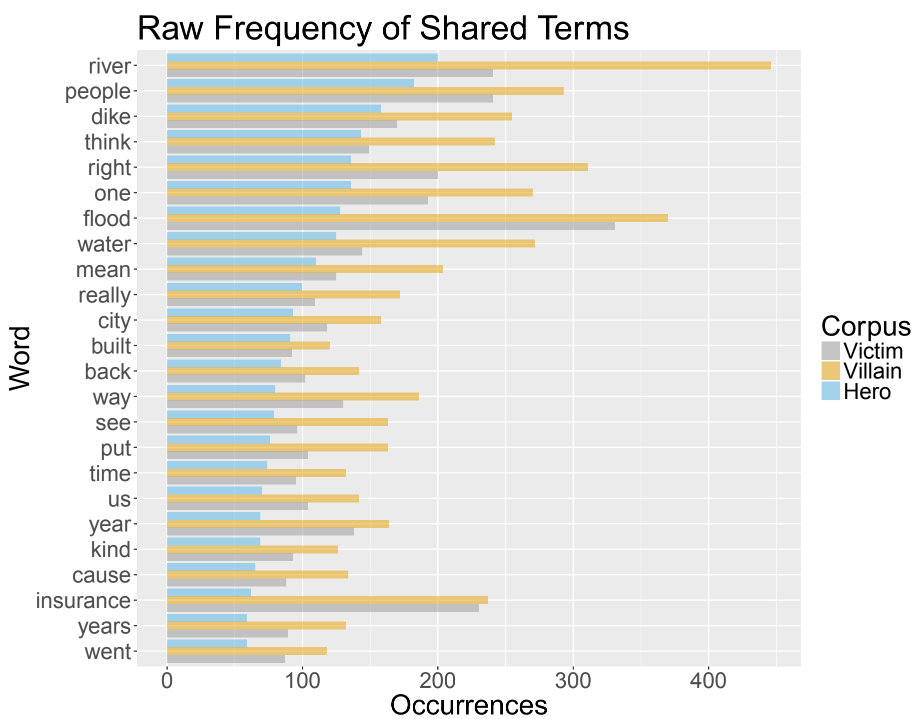

Figure 7.1: The raw frequency of shared terms colored by corpus.

There is no best frequency weighting system, it is corpora and application dependent. In our case, finding that ranking system involved trying many of the existing systems and comparing the results.

The first frequency ranking system was generated from the term-document matrix. A simple term count ranking of the terms in the corpus for each label. This is calculated by taking the row sums of the matrix. These lists were approximately 5000 terms long. This allowed for the comparison of the most frequently used words in each corpus against each other. Figure 7.1 shows the terms with the highest frequency in the Hero corpus and their raw frequencies in the other corpora. We found two flaws

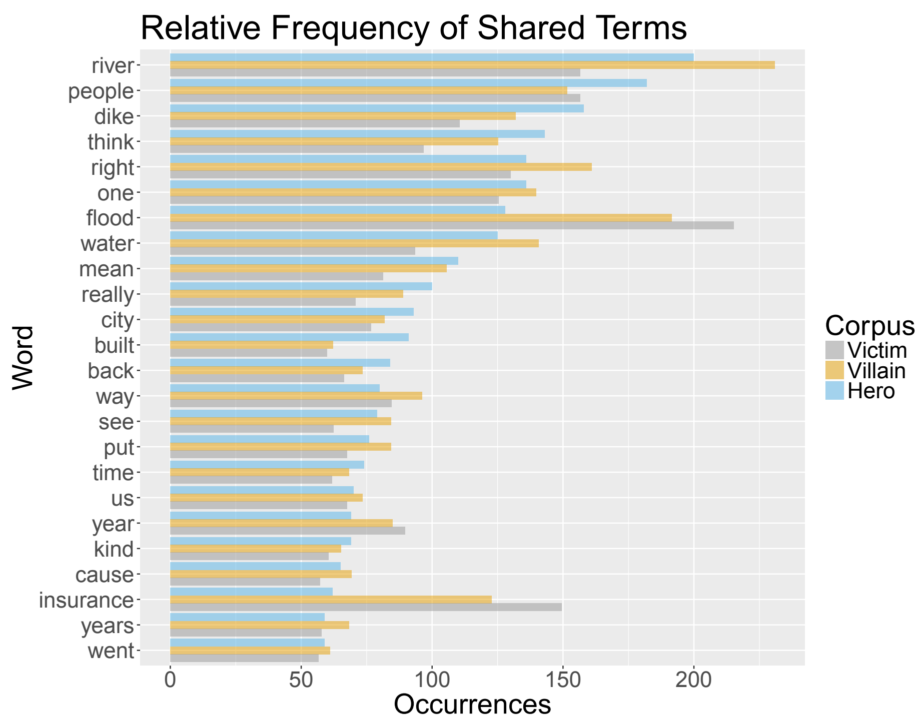

Figure 7.2: The relative frequency of shared terms colored by corpus.

with this approach. First, the corpora are different sizes. The Hero corpus is 472 documents with 15365 words, the Victim corpus is 748 documents with 20448 words, and the Villain corpus is 947 documents with 29768 words. Thus, simple term count comparisons are skewed because the term counts for a particular word are higher in corpora with higher word counts. This is fixed with normalization, where the term counts are divided by the word count of the corpus. Figure 7.2 shows the normalized or relative word frequencies, again sorted by the Hero corpus frequencies.

The second flaw is the assumption that frequency necessarily equals importance. If this were true, we would not remove stop-words and “the” would be one of the most important words in the English language. Although most stop words were removed in the preprocessing step, they were common English stop-words. This left a less apparent, context specific list of stop-words in the documents. These were words that constitute the common vernacular of the interviews. They are important terms in that they are the unifying context of the transcripts, but they dont help distinguish Hero from Victim because they appear broadly in each corpus. To attempt to address this issue, we used the Term Frequency Inverse Document Frequency method. In practice, it was a transformation applied to the values in the term-document matrices created during preprocessing. This new ranking is shown in Figure 7.3.

This was an improvement over the previous frequency lists. TF-IDF is designed to identify the important words in the corpus for downstream tasks such as topic modeling or information retrieval. However, a specific list of words was necessary. The list needed to emulate the frequency of the majority of the interviews and the suppression of broadly used terms was concerning in this context. Another concern was that a single interview could potentially skew the results on the list. For instance, if one person used a term a large number of times, that word would be given too much importance relative to the rest of the corpus. TF-IDF would exacerbate this, rewarding it for appearing in fewer documents. Given the small number of interviews(i.e, 45), this seemed a valid concern. TF-IDF became a supplemental measure, but not the primary word ranking system.

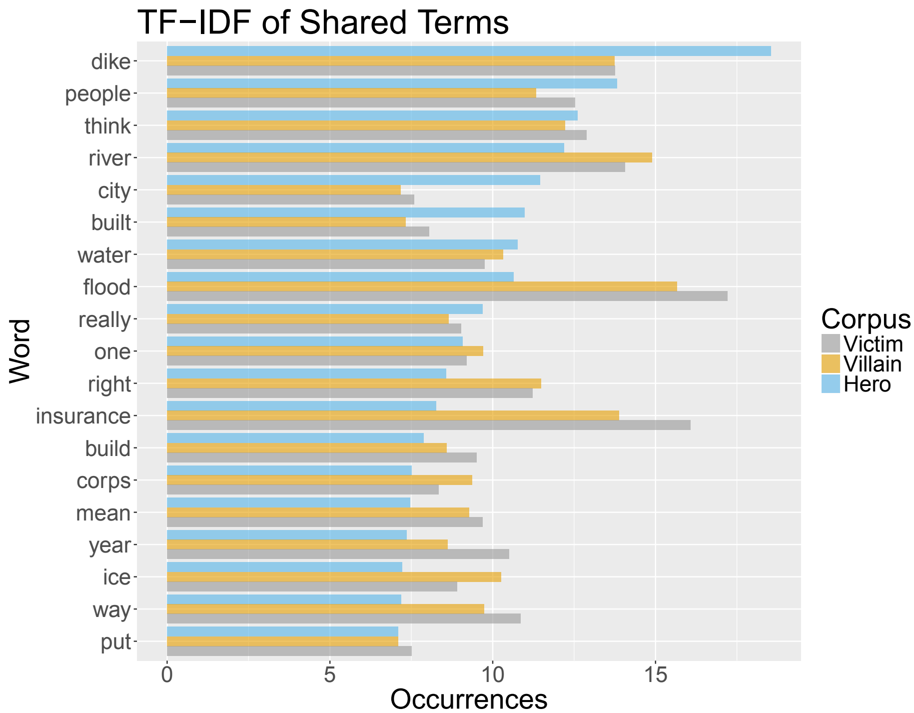

Figure 7.3: The TF-IDF of Shared Terms colored by corpus.

Transformed Relative Frequency

In order to address one interview skewing the results, the text was reorganized into three new Hero, Victim, and Villain corpora. Each document in the new corpora represented all the hand coded language of that type from a particular interview. For example, document three in the Hero corpus holds all the Hero coded language from interview number three. This representation led to two new frequency measures. The document frequency measurement (Dfc), which is the number of interviews a term appeared in and the transformed relative frequency (TRF). TRF is the relative frequency of terms within the new interview based corpora normalized by the size (i.e.,word count) of that corpus and transformed by taking the square root. Dfc was used by the researchers to check against one interviewee skewing the results, and TRF became the definitive word frequency ranking system within a corpus for narrative construction.

All the term rankings systems; RF, TF-IDF, Dfc, Ifc, and TRF were collected into a spreadsheet that held the values for each term that appeared in the corpus. The researchers who would construct the narratives were primarily concerned with the distinctions between the Hero and Victim corpora as the narratives to be constructed were Hero, Victim and Victim to Hero. Having now ranked the terms by use in each corpus, we found many of top terms were still shared between the Hero and Victim corpora. To handle this, we subtracted the TRF of a term in the Victim corpus from the TRF of the term in the Hero corpus. This created a scale in the range of about -20 to 20. The positive scores were Hero terms, the negative were Victim terms, and the near zero terms were either low frequency or evenly shared between both corpora. Using this scale, we could now rank words by their importance to each corpora. The group took the head and tail (top and bottom 4%) of the list to create the Hero and Victim vocabularies. These vocabularies were used to construct the narratives for the perception analyzer focus group tests.

IMPACTS TO INTERDISCIPLINARY RESEARCH

As discussed in previous chapters, this work served a larger project that has produced, as of writing, two major works. In this section we will discuss how our work impacted those projects and some of their findings.

Draft Narrative Analysis

Having identified the Hero and Victim vocabularies, we began to use them to construct narratives. The narrative construction process was iterative, with each member of the group raising concerns about the current drafts from the perspective of the field they represented.

Our role shifted to providing feedback about the draft narratives as they progressed. We used the previously explained NLP text analysis methods to provide information about the drafts. Each draft was scored for polarity, using sentence by sentence measurements and then returning an average for the narrative as a whole. The draft narratives were classified using the naive Bayes classifier to check that they were still representative of the coded corpus according to automated means. Unsurprisingly, they scored very highly as a member of their particular corpus. Naive Bayes is a frequency based model, so using the terms that appeared most frequently within a label should generate that label unless the labels share too many features. We also provided some terms that were highly associated with the select characters from the narratives. These terms had a higher probability to co-occur with the characters in the coded text segments within the same corpora. We presented frequency information for each term as it appeared in the corpus. This included TF-IDF score and frequency rank of the terms relative to the draft narrative and relative to the corpus. The feedback provided informed the narrative construction.

An example of this feedback is provided in Appendix A.

Additionally, narratives were also given a visual word use signature in the form of a histogram if they contained character driven narrative language. Again using the Rstudio environment, words that appeared within a narrative were plotted by HTRF-VicTRF term ranking, centered at zero.

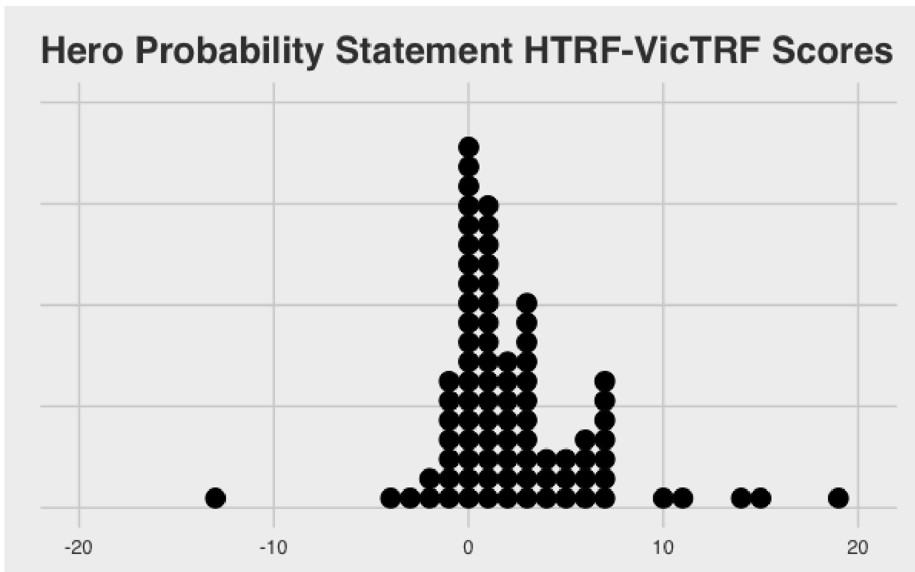

Figure 8.1: Distance scores of a Hero narrative

Figures 8.1 and 8.2 show the distance scores of words used in two of the final constructed narratives. The first plot shows a bias to the positive side of the HTRF-VicTRF scale, where the Hero language lies. The second plot shows a more even spread of Victim and Hero language. This narrative turns from Victim to Hero as it progresses. If a plot returned undesirable behavior, such as a Hero narrative containing too much Victim language, the group considered adjusting their word choices accordingly.

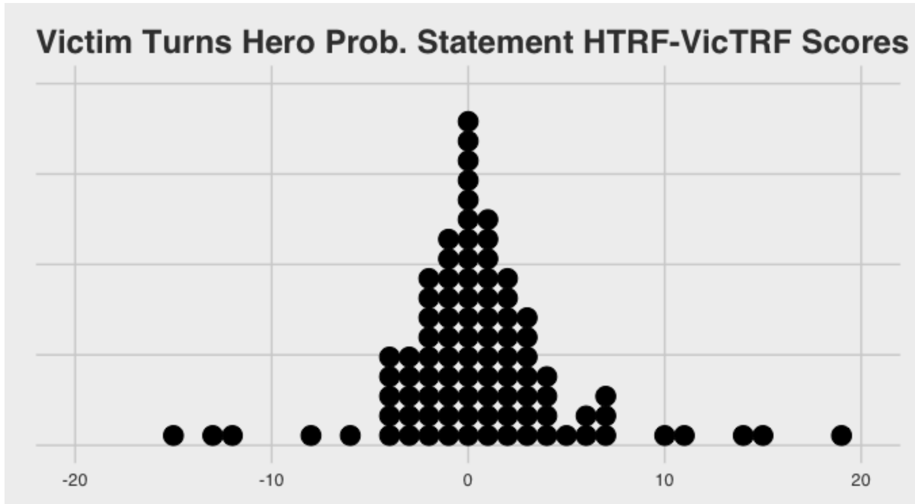

Figure 8.2: Distance scores of a Victim to Hero narrative

The Impact of Narratives on Affective Response

In order to test the impact of certainty, probability, and character-based narratives on affective response, eight risk communication messages were constructed according to the variables laid out in Figure 8.3. Each message contained segments for a Flood Definition and Science Information. The Science Information segment was either certainty based (An extreme event will happen even if we don't know exactly when) or probability based (an extreme event is likely to occur in the future, 100-year flood, 30-year flood). For each type of Science Information, four messages were constructed, a traditional science message and three character-based narratives. The character-based narratives would embed the traditional message segments in a narrative for each style Hero, Victim, and Victim-to-Hero. This added two additional segments to the messages. Problem framing introduced the character(Hero, Victim, Victim-to-Hero), and Characters in Action showed the actions taken by the characters according to their emphasis. Heroes emphasize that the audience is capable of preparing for flood risk. Victims emphasize the negatives outcomes from flooding.

<table>
<tr><th></th><th>Certainty</th><th>Probability</th></tr>
<tr><td>Conventional</td><td colspan="2">Flood definition, Science Information</td></tr>
<tr><td>Hero</td><td colspan="2" rowspan="3">Flood definition, Problem Framing, Science Information, Characters in Action</td></tr>
<tr><td>Victim</td></tr>
<tr><td>Victim-to-Hero</td></tr>
</table>

Figure 8.3: Each type of risk communications constructed for the experiment. The cells show the segments contained within each message in the order that they appeared.

Victim-to-Heroes emphasize that negative outcomes can be reversed by the audience members. Hero narratives were constructed by drawing from the Hero vocabulary identified in Chapter 7 and avoiding the Victim terms. Victim narratives were done in the opposite manner and Victim-to-Hero pulled from the Victim vocabulary in the beginning and the Hero vocabulary in the end. This was first laid out in Narrative-based Risk Communication: A Lingua Franca for Natural Hazard Messages [50], which was presented at the 2018 Midwest Political Science Association's conference in Chicago. This work was primarily focused on the narrative construction process but did seek to address certain intuitions we had about the different narrative types.

In order to measure the affective responses of the different narratives, 12 focus group studies measured the second by second readings of the audience using Perceptions Analyzer dials [12]. The average reading for each second was taken across all participants and examined. The results from the focus group dial readings are shown in Figure 8.4. It was found that the traditional science messages produced flatter contours when these averages were graphed compared to the Hero and Victim-to-Hero narratives. Hero language evoked significantly more positive responses than science language, and it was significantly more positive than victim language. Conversely, Victim language was not found to be significantly more negative than science language. This work was elaborated on and expanded in Characters matter:

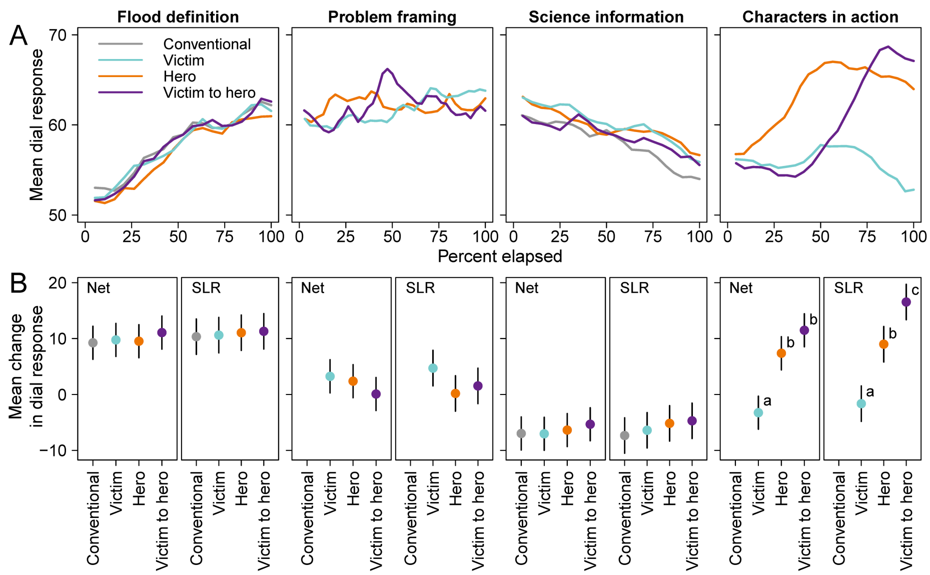

Figure 8.4: The average dial readings across the message segments, as well as the T_{Net} and T_{SLR} for each communication type. Courtesy of Shanahan et al. 2019 [62]

How narratives shape affective responses to risk communication [62]. This work was focused on two questions. The first, does affective response to probabilistic language differ from affective responses to certainty language. Second, does narrative language influence affective responses to hazard preparedness messages? These questions were approached through further analysis of the focus group dials readings using three linear mixed-effect models, standard deviation(S.D.), the net change per segment (T_{Net}), and the slope of the simple linear regression of the dial reading over elapsed message time(T_{SLR}). It was found that there was no significant difference in the measures of affective response between certainty and probability messages. It was also found that all narrative message types had a larger S.D. than that of the traditional non-narrative science messages. This was particularly pronounced among the Characters in Action segments of the narratives, with Hero and Victim-to-Hero leading compared to Victim. This work was accepted by the journal PLOS ONE [1] in 2019.

THREATS TO VALIDITY AND FUTURE RESEARCH

This section discusses some of those methods sorted by how they could help cover potential threats to the validity of the work described in this thesis according to the definitions of Wohlin [78].

Internal

Internal threats to validity are unanticipated relationships between variables. We did not control for differing lengths of examples between the classes for the classification task. This could bias the classifier towards the class containing longer examples as the longer an example, the more opportunity a term has to appear in it. It is possible that future narrative examples may simply be grouped into one document by character and then segmented by word count to create the individual documents. This would normalize all the documents to the same length and cover this internal validity.

External

External validity is concerned with applying our findings outside of the specific case study. As stated in the Chapter 2, we chose methods that would work well on smaller data collections and scale up as more information was collected. As it stands now, the data set represents the language use of a specific community of people. We would continue to gather more documents over time and add to the corpora to reach a more representative and robust data set. That will make the more complex methods viable and reliable. This includes collecting the larger national media and federal agency flood risk language corpora from the internet for comparative purposes. A more robust data set would be easier to generalize to other communities that are also at risk of flooding.

The topic models were graded subjectively, but also not rigorously validated among the group. If a stronger case for topic models arises in the future of this project, some method of adequately interpreting and validating the model would need to be developed.

Construct

Construct validity refers to the meaningfulness of measurements and that both independent and dependent variables are represented correctly in the study. It is possible that the ideas of Hero, Victim, and Villain narrative language are only being understood in terms of their relations to the other narrative styles and to the non-narrative language still concerning the flood domain. If we could compare the narrative-based corpora to a so called "normative corpus" [53], we could use that to examine the difference in sentiment and term frequencies that set off character based flood risk language from the everyday language used by the target audience. Similarly, one of our concerns during this project was surrounding the difference between spoken language and written language. In this project, we used transcripts of spoken language as the data set but were constructing written narratives. Many of the curated corpora in NLP research are of written material, especially if taken from online sources. It would be interesting to study the difference between the spoken and written language use of the target audience.

Concerning classification, only one type of classification was explored, naïve Bayes. If we examined other classification models, neural nets in particular, we may find an even stronger classifier for this task.

Content

Content validity refers to "the degree to which elements of an assessment instrument are relevant to and representative of the targeted construct for a particular assessment purpose" [21].

One of the sections that showed the most promise was sentiment analysis. If we could continue to expand and update the flood risk domain specific polarity terms, the accuracy and usefulness of that data would also improve. This is an area that simply requires human eyes and input. One of the drawbacks of the sentiment analysis approach comes in verifying the accuracy of the results. It is difficult to trust the sentiment analysis scores of sentences and corpora as being truly representative of the emotional tone of the language in the flood risk context without examining them, which at a certain scale defeats the purpose.

While relative frequency does accurately capture the most prominent terms of the corpora, these may not actually be the distinguishing terms that define the idea of Hero, Victim, and Villain language. These ideas remain abstract and difficult to measure quantitatively. While we did construct vocabularies according to the difference transformed relative frequencies for each concept, more research could be done towards covering other aspects of these ideas as they are represented in the text of the transcripts.

CONCLUSION

We have laid out the four text analysis methods we explored toward the creation of a narrative construction process. That process would enhance the risk communication of extreme flood events from scientists and government officials to citizens living in the flood plane of the Yellowstone River in Montana. We have laid out our methods for preprocessing, document organization, and text analysis to help future researchers parse out meaning from large and unwieldy corpora. After examining the results of topic modeling, machine classification, and sentiment analysis we found that a difference in transformed relative term frequencies generated the best information to construct computationally enhanced vocabularies for this narrative construction use case. These vocabularies define the differences in word choice between the three different corpora. The Hero and Victim vocabularies were used to construct Hero, Victim, and Victim-to-Hero narratives that were embedded with conventional risk communication. After testing the affective response in focus group studies, it was found that the narrative-based risk communication generated a larger affective response than that of conventional risk communication. Our results were used in two publications, "Narrative-based Risk Communication: A *Lingua Franca* for Natural Hazard Messages?" [50] shown at the Midwest Political Science Association's annual conference in 2018 and "Characters matter: How narratives shape affective responses to risk communciation" [62] published in the journal PLOS ONE. These findings have furthered the study of narratives as more effective means of risk communication and the impact that natural language processing methods can have on constructing character-based narratives.

REFERENCES CITED

[1] PLOS ONE. `https://journals.plos.org/plosone/`, 2019.

[2] J. Atkinson, G. Salas, and A. Figueroa. Improving opinion retrieval in social media by combining features-based coreferencing and memory-based learning. *Information Sciences*, 299:20–31, 2015.

[3] S. Baccianella, A. Esuli, and F. Sebastiani. Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In *Lrec*, volume 10, pages 2200–2204, 2010.

[4] L. D. Baker and A. K. McCallum. Distributional clustering of words for text classification. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 96–103. ACM, 1998.

[5] D. M. Blei. Topic modeling and digital humanities. *Journal of Digital Humanities*, 2(1):8–11, 2012.

[6] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning research*, 3(Jan):993–1022, 2003.

[7] A. Caliskan-Islam, J. J. Bryson, and A. Narayanan. Semantics derived automatically from language corpora necessarily contain human biases. pages 1–14, 2016.

[8] E. Cambria, S. Poria, A. Gelbukh, and M. Thelwall. Sentiment analysis is a big suitcase. *IEEE Intelligent Systems*, 32(6):74–80, 2017.

[9] E. Cambria, B. Schuller, Y. Xia, and C. Havasi. New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems*, 28(2):15–21, 2013.

[10] J. Chang, J. Boyd-Graber, C. Wang, S. Gerrish, and D. M. Blei. Reading tea leaves: How humans interpret topic models. In *Neural Information Processing Systems*, 2009.

[11] W. M. Darling. A theoretical and practical implementation tutorial on topic modeling and gibbs sampling. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 642–647, 2011.

[12] Dialsmith. Perception analyzer. `http://www.dialsmith.com`.

[13] V. D'Orazio, S. T. Landis, G. Palmer, and P. Schrodt. Separating the wheat from the chaff: Applications of automated document classification using support vector machines. *Political Analysis*, 22(2):224–242, 2014.

[14] D. Dransch, H. Rotzoll, and K. Poser. The contribution of maps to the challenges of risk communication to the public. *International Journal of Digital Earth*, 3(3):292–311, 2010.

[15] I. Feinerer and K. Hornik. tm: Text mining package, 2017.

[16] I. Fellows. wordcloud: Word clouds, 2014.

[17] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, NY, USA:, 2nd edition, 2008.

[18] J. Grimmer. A bayesian hierarchical topic model for political texts: Measuring expressed agendas in senate press releases. *Political Analysis*, 18(1):1–35, 2010.

[19] M. Guerini, L. Gatti, and M. Turchi. Sentiment analysis: How to derive prior polarities from sentiwordnet. 2013.

[20] H.-q. Han, D.-H. Zhu, and X.-f. Wang. Semi-supervised text classification from unlabeled documents using class associated words. In *Computers & Industrial Engineering, 2009. CIE 2009. International Conference on*, pages 1255–1260. IEEE, 2009.

[21] S. N. Haynes, D. Richard, and E. S. Kubany. Content validity in psychological assessment: A functional approach to concepts and methods. *Psychological assessment*, 7(3):238, 1995.

[22] Y. He, C. Lin, and H. Alani. Automatically extracting polarity-bearing topics for cross-domain sentiment classification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 123–131. Association for Computational Linguistics, 2011.

[23] M. Hu and B. Liu. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177, New York, NY, USA, 2004. ACM.

[24] A. Jain and J. Mandowara. Text classification by combining text classifiers to improve the efficiency of classification. *International Journal of Computer Application*, 6(2), 2016.

[25] Z. Jianqiang. Pre-processing boosting twitter sentiment analysis? In *Smart City/SocialCom/SustainCom (SmartCity), 2015 IEEE International Conference on*, pages 748–753. IEEE, 2015.

[26] T. Joachims. Text categorization with support vector machines: Learning with many relevant features. In *European conference on machine learning*, pages 137–142. Springer, 1998.

[27] T. Joachims, D. Freitag, and T. Mitchell. Webwatcher: A tour guide for the world wide web. In *IJCAI (1)*, pages 770–777. Citeseer, 1997.

[28] M. D. Jones, M. K. McBeth, and E. A. Shanahan. *Introducing the Narrative Policy Framework*, pages 1–25. Palgrave Macmillan US, New York, 2014.

[29] J. S. Kessler and N. Nicolov. Targeting sentiment expressions through supervised ranking of linguistic configurations. In *ICWSM*, 2009.

[30] A. Kilgarriff. Using word frequency lists to measure corpus homogeneity and similarity between corpora. In *Fifth Workshop on Very Large Corpora*, 1997.

[31] A. Kilgarriff. *Comparing Corpora*, volume 6. International Journal of Corpus Linguistics, 2001.

[32] S.-B. Kim, K.-S. Han, H.-C. Rim, and S. H. Myaeng. Some effective techniques for naive bayes text classification. *IEEE transactions on knowledge data engineering*, 18(11):1457–1466, 2006.

[33] S. Lai, L. Xu, K. Liu, and J. Zhao. Recurrent convolutional neural networks for text classification. In *Twenty-ninth AAAI conference on artificial intelligence*, 2015.

[34] A. Lawani, M. M. R. Reed, T. Mark, and Y. Zheng. Reviews and price on online platforms: Evidence from sentiment analysis of airbnb reviews in boston. *Regional Science Urban Economics*, 2018.

[35] M. D. Lee, B. Pincombe, and M. Welsh. An empirical evaluation of models of text document similarity. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 27, 2005.

[36] D. D. Lewis. Naive (bayes) at forty: The independence assumption in information retrieval. In *European conference on machine learning*, pages 4–15. Springer, 1998.

[37] B. Liu, M. Hu, and J. Cheng. Opinion observer: analyzing and comparing opinions on the web. In *Proceedings of the 14th international conference on World Wide Web*, pages 342–351. ACM, 2005.

[38] W. Lowe and K. Benoit. Validating estimates of latent traits from textual data using human judgment as a benchmark. *Political Analysis*, 21(3):298–313, 2013.

[39] Q. I. P. Ltd. Nvivo qualitative data analysis software, 2018.

[40] N. Majumder, D. Hazarika, A. Gelbukh, E. Cambria, and S. Poria. Multimodal sentiment analysis using hierarchical fusion with context modeling. *Knowledge-Based Systems*, 161:124–133, 2018.

[41] M. E. Maron. Automatic indexing: An experimental inquiry. *J. ACM*, 8(3):404–417, 1961.

[42] A. McCallum and K. Nigam. A comparison of event models for naive bayes text classification. In *AAAI-98 workshop on learning for text categorization*, volume 752, pages 41–48. Citeseer, 1998.

[43] A. K. McCallum. Mallet: A machine learning for language toolkit, 2002.

[44] D. Mimno. mallet: A wrapper around the java machine learning tool mallet.

[45] T. Mullen and N. Collier. Sentiment analysis using support vector machines with diverse information sources. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, 2004.

[46] M. V. Mntyl, D. Graziotin, and M. Kuutila. The evolution of sentiment analysisa review of research topics, venues, and top cited papers. *Computer Science Review*, 27:16–32, 2018.

[47] K. Nigam, A. McCallum, and T. Mitchell. Semi-supervised text classification using em. *Semi-Supervised Learning*, pages 33–56, 2006.

[48] B. Pang, L. Lee, and S. Vaithyanathan. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*, pages 79–86. Association for Computational Linguistics, 2002.

[49] D. Pavlov, R. Balasubramanyan, B. Dom, S. Kapur, and J. Parikh. Document preprocessing for naive bayes classification and clustering with mixture of multinomials. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 829–834. ACM, 2004.

[50] E. Raile, H. King, E. Shannahan, J. McEvoy, C. Izurieta, N. Bergmann, R. Ready, A. M. Reinhold, and G. Poole. Narrative-based risk communication: A lingua franca for natural hazard messages. *76th Annual Midwest Political Science Association, Chicago IL*, 2018.

[51] D. Ramage, S. Dumais, and D. Liebling. Characterizing microblogs with topic models. In *Fourth international AAAI conference on weblogs and social media*, 2010.

[52] P. Rayson, D. Berridge, and B. Francis. Extending the cochran rule for the comparison of word frequencies between corpora. In *7th International Conference on Statistical analysis of textual data (JADT 2004)*, pages 926–936, 2004.

[53] P. Rayson and R. Garside. Comparing corpora using frequency profiling. In *Proceedings of the workshop on Comparing corpora-Volume 9*, pages 1–6. Association for Computational Linguistics, 2000.

[54] T. W. Rinker. qdap: Quantitative discourse analysis package, 2017.

[55] M. E. Roberts, B. M. Stewart, D. Tingley, C. Lucas, J. LederLuis, S. K. Gadarian, B. Albertson, and D. G. Rand. Structural topic models for openended survey responses. *American Journal of Political Science*, 58(4):1064–1082, 2014.

[56] S. Robertson. Understanding inverse document frequency: on theoretical arguments for idf. *Journal of documentation*, 60(5):503–520, 2004.

[57] G. Salton and C. Buckley. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523, 1988.

[58] G. Salton and C.-S. Yang. On the specification of term values in automatic indexing. *Journal of documentation*, 29(4):351–372, 1973.

[59] G. Salton, C.-S. Yang, and C. T. Yu. Contribution to the theory of indexing. Report, Cornell University, 1973.

[60] G. Salton, C.-S. Yang, and C. T. Yu. A theory of term importance in automatic text analysis. *Journal of the American society for Information Science*, 26(1):33–44, 1975.

[61] A. Severyn and A. Moschitti. Twitter sentiment analysis with deep convolutional neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 959–962. ACM, 2015.

[62] E. A. Shanahan, A. M. Reinhold, E. D. Raile, G. C. Poole, R. C. Ready, J. McEvoy, C. Izurieta, N. T. Bergmann, and H. W. King. Characters matter: How narratives shape affective responses to risk communication. *PLOS ONE*, 2019.

[63] R. Sharma and P. Bhattacharyya. Detecting domain dedicated polar words. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 661–666, 2013.

[64] P. Singh, R. S. Sawhney, and K. S. Kahlon. Sentiment analysis of demonetization of 500 & 1000 rupee banknotes by indian government. *ICT Express*, 4(3):124–129, 2018.

[65] S. S. Skiena. *The Data Science Design Manual.* Springer, 2017.

[66] P. Slovic and E. Peters. Risk perception and affect. *Current directions in psychological science*, 15(6):322–325, 2006.

[67] F. Song, S. Liu, and J. Yang. A comparative study on text representation schemes in text categorization. *Pattern analysis applications*, 8(1-2):199–209, 2005.

[68] K. Sparck Jones. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21, 1972.

[69] S. Symeonidis, D. Effrosynidis, and A. Arampatzis. A comparative evaluation of pre-processing techniques and their interactions for twitter sentiment analysis. *Expert Systems with Applications*, 2018.

[70] D. Tang, B. Qin, and T. Liu. Document modeling with gated recurrent neural network for sentiment classification. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 1422–1432, 2015.

[71] R. C. Team. R: A language and environment for statistical computing.

[72] M. Toman, R. Tesar, and K. Jezek. Influence of word normalization on text classification. *Proceedings of InSciT*, 4:354–358, 2006.

[73] P. University. About wordnet, 2010.

[74] A. K. Uysal and S. Gunal. The impact of preprocessing on text classification. *Information Processing & Management*, 50(1):104–112, 2014.

[75] H. M. Wallach. Topic modeling: beyond bag-of-words. In *Proceedings of the 23rd international conference on Machine learning*, pages 977–984. ACM, 2006.

[76] A. B. Warriner, V. Kuperman, and M. Brysbaert. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, 45(4):1191–1207, 2013.

[77] H. Wickham. ggplot2: Elegant graphics for data analysis, 2009.

[78] C. Wohlin, P. Runeson, M. Hst, M. C. Ohlsson, B. Regnell, and A. Wessln. *Experimentation in software engineering*. Springer Science & Business Media, 2012.

[79] C. Yang, K. H.-Y. Lin, and H.-H. Chen. Emotion classification using web blog corpora. In *Web Intelligence, IEEE/WIC/ACM International Conference on*, pages 275–278. IEEE, 2007.

[80] H. Zhang and D. Li. Nave bayes text classifier. In *Granular Computing, 2007. GRC 2007. IEEE International Conference on*, pages 708–708. IEEE, 2007.

[81] M.-L. Zhang, J. M. Pea, and V. Robles. Feature selection for multi-label naive bayes classification. *Information Sciences*, 179(19):3218–3229, 2009.

[82] Y. Zhang, D. Song, P. Zhang, P. Wang, J. Li, X. Li, and B. Wang. A quantum-inspired multimodal sentiment analysis framework. *Theoretical Computer Science*, 2018.

[83] W. X. Zhao, J. Jiang, J. Weng, J. He, E.-P. Lim, H. Yan, and X. Li. Comparing twitter and traditional media using topic models. In *European conference on information retrieval*, pages 338–349. Springer, 2011.

APPENDICES

APPENDIX A

DRAFT NARRATIVE FEEDBACK

Text Analysis Feedback

Original Narrative:

On the Yellowstone River there are many reasonable way you and your neighbors can work together to help protect from flooding. While many people in the river communities of the Yellowstone feel saved and protected by their levee, people also have a healthy respect for the power of this awesome river. Engineers and local government provide the benefit of technical assistance for levees, dams, and bank protection measures, but in the face of an extreme flood event, even the best engineered solutions may not help. What was thought to be a 100 year flood is now considered to be a 33 year flood, meaning that 3 times in the last 100 years, the streamflow gauge at has exceeded the 100 year flood. Furthermore, there are likely to be much larger floods than previously anticipated, time the magnitude of the 2011 flood. Taking care to work and prepare for an extreme flood event will likely improve economic stability in the face of an extreme flood event.

Classification

The log-likelihood of Hero/Victim/Villain according to our naïve Bayes classifier trained on the coded narratives

Hero	**0.999021678**	Victim	1.88E-10	Villain	9.78E-04

Word Associations

Some highly associated words from the Hero Corpus of some select nouns from the narrative.

(>.6 correlation lower limit)

Dam-

ahold(probably hold), blank, dredge, electricity, film, furnish, heaviest, impressive, marvelous, wonderful

Yellowstone-

bay, integrating, key, melts, muddy, sluffing, mud

Neighbor-

backtrack, flipside, responds, ordered

Engineers-

corps

(>.5 correlation lower limit, no data until that point)

Levee-

elevator, grading, height, recognizing, wisdom

Economic-

viable, cancel, egg, gal, groans, nest, angers, astronomical, headed, essence, arguing, coastal

(>.4 correlation lower limit, no data until that point)

Government-

reasonable, brad, accomplish, analysis, appreciation, biologist, bunchgrasses,

Word Frequency:

Comparison to the Hero Corpus Top 500 TF-IDF, sorted by Frequency in the Narrative. Can't run TF-IDF on Narrative, as its only 1 document, but provides an interesting comparision to the important words in the Hero Corpus. For every word in the narrative, if it appeared on the Top 500 TF-IDF ranking of the Hero Corpus, it made this list. Sorted by the Frequency Rank in the Narrative.

Figure A.1

	Narrative		Hero Corpus	
Freq. Rank in Narrative	**Term**	**Total Occurrences**	**TF-IDF**	**TF-IDF Rank in Hero Corpus**
1	flood	7	10.640402	8
2	event	3	2.229471	256
3	river	3	12.200831	4
4	year	3	7.365685	25
5	face	2	1.529351	431
6	help	2	9.592944	10
7	many	2	3.954626	93
8	people	2	13.821726	2
9	work	2	4.063322	89
10	yellowstone	2	3.345539	126
11	also	1	2.192516	266
12	assistance	1	2.042421	298
13	bank	1	3.316346	131
14	benefit	1	1.613004	405
15	best	1	2.187254	267
16	care	1	3.745954	104
17	considered	1	1.789543	361
18	engineers	1	5.740979	40
19	even	1	4.637472	68
20	feel	1	2.759282	191
21	flooding	1	6.912737	30
22	floods	1	2.868569	177
23	government	1	3.194218	138
24	last	1	2.525277	213
25	levee	1	8.171024	17
26	local	1	2.465161	219
27	may	1	2.304131	244
28	much	1	4.407539	76
29	neighbors	1	8.10287	18
30	prepare	1	2.868538	178
31	protect	1	1.377643	490
32	protected	1	1.44551	469
33	protection	1	1.819328	354
34	saved	1	3.292344	134
35	thought	1	2.457427	223
36	time	1	6.398795	34
37	times	1	3.179441	139
38	together	1	3.909021	95
39	way	1	7.196228	27
40	will	1	4.976345	56
41	years	1	6.654316	32

Same comparison to the Hero Corpus Top 500 TF-IDF, sorted by TF-IDF in the Hero Corpus. How many "important" words from the Hero Corpus did we hit?

Figure A.2

Narrative			Hero Corpus	
Freq. Rank in Narrative	**Term**	**Total Occurrences**	**TF-IDF**	**TF-IDF Rank in Hero Corpus**
8	people	2	13.821726	2
3	river	3	12.200831	4
1	flood	7	10.640402	8
6	help	2	9.592944	10
25	levee	1	8.171024	17
29	neighbors	1	8.10287	18
4	year	3	7.365685	25
39	way	1	7.196228	27
21	flooding	1	6.912737	30
41	years	1	6.654316	32
36	time	1	6.398795	34
18	engineers	1	5.740979	40
40	will	1	4.976345	56
19	even	1	4.637472	68
28	much	1	4.407539	76
9	work	2	4.063322	89
7	many	2	3.954626	93
38	together	1	3.909021	95
16	care	1	3.745954	104
10	yellowstone	2	3.345539	126
13	bank	1	3.316346	131
34	saved	1	3.292344	134
23	government	1	3.194218	138
37	times	1	3.179441	139
22	floods	1	2.868569	177
30	prepare	1	2.868538	178
20	feel	1	2.759282	191
24	last	1	2.525277	213
26	local	1	2.465161	219
35	thought	1	2.457427	223
27	may	1	2.304131	244
2	event	3	2.229471	256
11	also	1	2.192516	266
15	best	1	2.187254	267
12	assistance	1	2.042421	298
33	protection	1	1.819328	354
17	considered	1	1.789543	361
14	benefit	1	1.613004	405
5	face	2	1.529351	431
32	protected	1	1.44551	469
31	protect	1	1.377643	490

Polarity

Figure A.3

Position in Narrative	Word Count	Polarity	Positive Words	Negative Words	Sentence
1	21	0.8728716	c("reasonable", "work", "help", "protect")	–	On the Yellowstone River there are many reasonable way you and your neighbors can work together to help protect from flooding.
2	30	0.9128709	c("saved", "protected", "healthy", "respect", "awesome")	–	While many people in the river communities of the Yellowstone feel saved and protected by their levee, people also have a healthy respect for the power of this awesome river.
3	34	0.5144958	c("benefit", "assistance", "protection", "best", "help")	–	Engineers and local government provide the benefit of technical assistance for levees, dams, and bank protection measures, but in the face of an extreme flood event, even the best engineered solutions may not help.
4	32	0.1767767	exceeded	–	What was thought to be a year flood is now considered to be a year flood, meaning that times in the last years, the streamflow gauge at has exceeded the year flood.
5	18	0	–	–	Furthermore, there are likely to be much larger floods than previously anticipated, time the magnitude of the flood.
6	24	0.8164966	c("work", "prepare", "improve", "stability")	–	Taking care to work and prepare for an extreme flood event will likely improve economic stability in the face of an extreme flood event.

Polarity of Narrative **Figure A.4**

Total Sentences	Total Words	Avg. Polarity	Std. Deviation Polarity	Std. Mean Polarity
6	159	0.5489186	0.387296	1.41731

Polarity of each Sentence in order, followed by each Sentence by Polarity
Figure A.5

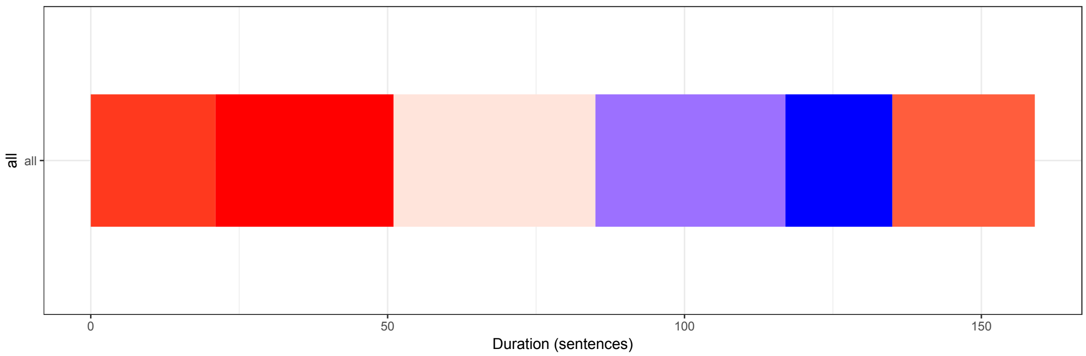

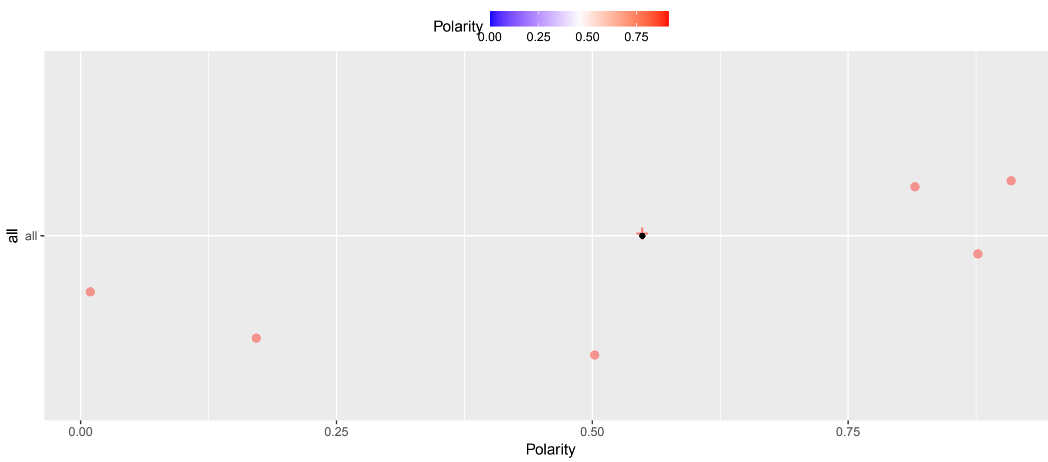

Hero Corpus Polarity for Comparison

Figure A.6

Total Sentences	Total Words	Avg. Polarity	Std. Deviation Polarity	Std. Mean Polarity
3301	45264	0.04633499	0.2269424	0.2041707

Polarity of each sentence in order, followed by each Sentence by Polarity

Figure A.7

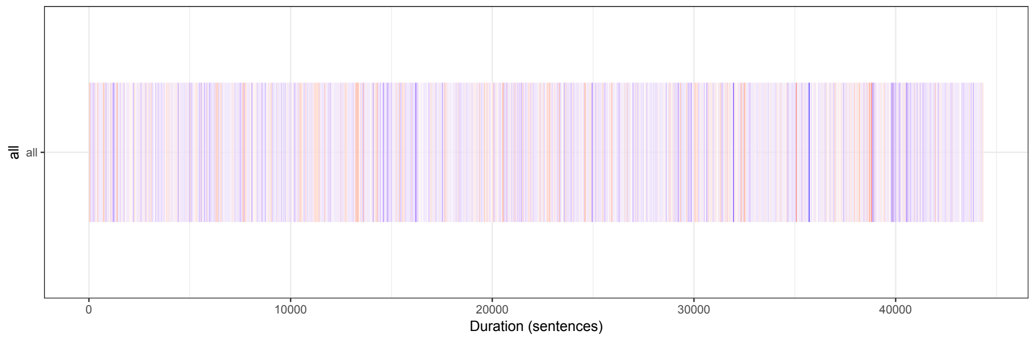

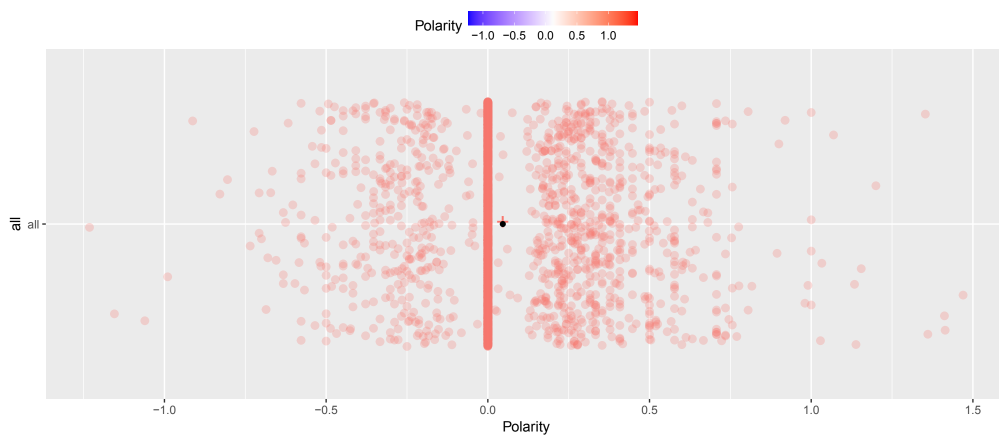

APPENDIX B

CITIZEN INTERVIEW TRANSCRIPT

R1: Well, yeah, there's uh … I would say that one of the tricks that's played on us is our loss of memory, it's just … and it's, I remember being frustrated with my mother about that and now I'm hearing if from my son: "Mom, I already told you that," and I'm, like, "Okay, you can be kind about it." You know.
R2: It's just the problem with memory.
I: [Chuckles]
R2: They're so old and so wise and so intelligent. Our file banks are so large that it takes longer for our (so CPU?) …
I: [Chuckles]
R2: To dial through all of the information and find a piece we want. Now you would think that I knew where Whitehall is, and I do and that I [could have found] Three Bridges in my head but I couldn't or Three Forks in my head and I'm stilling going Three Bridges, 'cause Twin Bridges ––
I: Twin Bridges, right, right, right ––
R2: –– is down at Whitehall.
I: –– right, I'm following you.
[Collective laughter]
I: That's funny.
R2: So it's gotta go through all this CPU (stuff?) –– my computer's just slower. Now these kids have faster computers and they don't have as much stuff in there to look through, so it happens quicker. (Don't think I'm losing?) my stuff, it's all there.
I: Right, well, I might––I might need to call you back and quote you to my son on this because I'm supposed to, like, I don't know. Just kidding.
[Collective laughter]
I: Alright, so this––I'm Liz Shanahan, and what is today? The 15th of February and I'm here talking with (redacted) and––Do you go by (redacted) or (redacted)? Or doctor?
R2: No, I don't go by doctor. There's a long story behind that one, too.
I: Okay.
R2: (Redacted) is fine.
I: (Redacted). Okay, excellent. Well, thank you very much. And, my name and number are on, it's on that sheet that you have too.
R1: Okay.
I: Uhm, 'cause we will definitely be potentially following up with further questions or also, uhm, you know when we've done our study if you're interested in what we have found, we can get back to you and share with you what we've found. So, the first couple questions are really just about, they're very general. They're not flood related questions. Uhm, but it is about what it is that you, maybe you personally or –– and, or you as a community value about the Yellowstone––and maybe the Tongue and the Yellowstone River –– rivers in your community.
R2: Value? [Laughs]
R1: [Laughs]
I: Yeah.

R2: I —— I don't know. (I'm not?)——Miles City has forgot that the river's there really for so many purposes.
R1: There's a swimming hole for the kids.
I: Mmhmm.
R1: The Tongue was anyway.
R2: When —— when we moved to town, that was the flood side. Nobody could go over there. I mean, it was really the bad side of town. If you crossed Milwaukee tracks you were in the floodplain. Well … You know, nobody went over there. I mean [cough] it was kind of a division point to a lot of Miles City social economical ...
I: Mmhmm.
R2: … umbrella, which really aren't real.
I: What —— tell me more, they're not real. Like …
R2: I can go to the north side and I can find shanties and I can go to the south side and I can find shanties. I can find shanties all over Miles City.
I: Mmhmm.
R2: Find beautiful homes all over Miles City.
I: Mmhmm.
R2: Some of them right down next to a dike that've been there forever. In fact, we own a little house right next to a dike and it's got a basement. It's never been flooded, it's never been wet.
I: Mmhmm.
R2: Uhm, we own property elsewhere. Same kind of story, so …
I: And so when you say, so then, the value is… what you, well how did you put it? That Miles City kind of ignores the river or that it's, it's in the, in that floodpl ——
R2: It's in that floodplain. It's in that poor section of town. It's in the, the ——
I: —— floodplain. Okay. Hmm. And that's a myth?
R2: And it's a myth.
I: Mmhmm.
R2: Uh, it's not a myth in terms of value of the land down there's real cheap.
I: Mmhmm.
R2: But, even that has changed in the last twenty years. It's just —— to some degree.
I: Mmhmm, mmhmm. So there may be some value for some recreation or swimming of some kind, uhm but ——
R1: Swimming, fishing, boating.
I: Mmhmm. There's no irrigation, then, that comes up of the … No economic value then?
R2: Well, TY —— the TY out of the Tongue River (becomes?) TY ditch and it goes across.
I: Uh-huh.
R2: Uh, and it's a major, uh, irrigation system, so for the (inaudible) farmers and ranchers, yes, it's a major impact on agricultural and, uhm, I'd say really that's … the biggest part of …
I: Is the TY different than the Slough?
R2: Yes.

I: Okay. I have to look at a map and look at that then.
R1: Of course you have to remember years ago they built communities near rivers because that was their means of transportation, and that hasn't been a means of transportation for I don't know how many years.
I: Mmhmm.
R2: The TY ditch comes off thirteen miles up the Tongue River.
I: Oh, okay.
R2: And comes across, they've got a diversion dam on the Tongue River.
I: Mmhmm.
R2: That goes across a little highland and then all of it comes down through the whole valley. And then of course the Yellowstone on this end is where the overflow goes to when it's done. I mean, it all flows through and gets used and back into the Yellowstone, and in the current Tongue River, it was altered somewhere and I don't know where in history, but it was altered privately because the Tongue River actually meandered all over Miles City. The last (hanging?) in Miles was here (in late?) and the railroad tracks which was the bridge over the Tongue River.
I: Huh.
R2: [Laughter]
I: So you're a historian, too, I see.
R2: [Laughter] I read a lot, but uh, you know, the bowl, which used to be down by the high school was a big hole that they played football in, which was the Tongue River bed. A number of years ago they filled it in and made it level.
I: Mmhmm.
R2: So, that area is no even longer even seen. You can see places of what used to be the river and some of back neighborhoods, and it's really surprising, front of the neighborhood that here's that ground level and the back drops sixteen feet, all in the period of a hundred feet of property, 'cause it's right on what…
I: Was.
R2: … was a river. So the Tongue River, uh, current dike was privately built, is my understanding, that the dike——the Army Corps of Engineers certified it and never had anything to do with it. Uh, well we have tried over the years to do various things to improve it. The Army Corps' also stopped us from improving it.
I: Because of?
R2: Who has jurisdiction.
I: Hmm. Who has jurisdiction? The Army Corps, Feds, or the town, who needs to protect itself?
R2: The town does itself the protection and then the Feds say, well that's not adequate, but they wouldn't help us when we were doin' it. Now, I wasn't around in those days when that dike was built. I have no idea. Uh, but now that's where the country club is, that's where Steadmans built their new fancy apartments down and built everything up so that they're above floodplain, rentals that make a thousand dollars a month, and now all of a sudden, the rest of the city got to deal with a flood problem.

I: Mmhmm.
R2: I'm not saying if those people caused it.
I: Mmhmm.
R2: Because I actually think Steadman's did very –– who (open?) themselves, I love 'em. They're great people –– but all of that stuff is part of that history that…
I: Right, right. So the, so the Tongue River essentially was contained or moved because it, if I'm understanding it right, because it braided through the city, uhm, but it's been kind of –– what I'm hearing you say is almost erased. They leveled the bowl or that pla –– the football field, is that right?
R2: Mmhm.
I: Uhm, and so maybe it erased from site and maybe a little bit from memory?
R2: Uh, good possibility.
I: Mmhmm.
R2: And then we have the jokesters in town. I mean, you know, the––I love the Feds always pull out this boat float down Main Street. Have you ever seen the picture?
I: I haven't seen that photo. I'll have to find it. A boat floating on Main Street?
R2: [Laughs/coughs] Yeah, a boat floating on Main Street. Two guys sitting in a flat bottom boat, floatin' on Main Street. It was like six inches of water on Main Street after a good, heavy rain.
I: Uh-huh.
R2: And yes, the (indistinguishable) may have backed up a little bit, no one's gonna question that.
I: Mmhmm.
R2: But, here's proof that Miles City floods. That's not proof. These are two old drunks who got out of a bar and sat in a boat and took pictures for a joke.
I: Mmhmm, mmhmm.
R2: And, almost everybody in town will tell you that's what that is, but the Feds, this is proof that we flood. No, it's not. You know.
I: So that's interesting. So there's, like, yeah, what constitutes evidence of flooding, right? And so you can –– they, they're pointing to something that local knowledge is, like, no, that didn't –– that's a joke, and that's not flooding, and they're saying, "No, that constitutes flooding."
R2: Well it floods bad enough that the Indian history is that's where they always camped, right down at the Tongue, confluent down there, between the Tongue and Fort (Tio's?) where they camped. The Indians don't get flooded.
I: Mmhmm, mmhmm. And so the other sort of general question is like you––so we kind of get an idea about the value, but what about––do–– and I think you may have answered some of this, uhm, in terms of what are the problems with these rivers and maybe part of what you answered is, well, they, at least the Tongue is braided through, but that's problem's taken care of, and so––I don't' know what, the '40s when that dike was built?

R2: I have no honest idea when it was built.
I: I think that's what Sam said, was the '40s, uhm … But, are there any other kinds of problems related to the river that might be even present today… uhm, other that flooding? I mean, we're gonna get into flooding obviously, but——or maybe that's not even a problem.   Uhm … Like, for example, someone said, you know, I think a big problem is that there's a lot of trash that gets collected, so we're trying to organize this community to do trash pickup, so that's what one person said.   I thought that was interesting.
R2: Okay. Now, now we're gonna get into some history that I don't know if I can even validate if I had to try. I was a member of the Miles City (JCs?) years and year ago.
I: Mmhmm, mmhmm.
R2: Miles City (JCs?) was located at a (federal grants) for indoor/ outdoor recreation.
I: Mmhmm.
R2: This is forty-five years ago? [Laughs/coughs] And one of the grants was that we actually take the Slough and clean it out and make it a running water, kind of like the San Antonio walkway ——
I: Mmhmm.
R2: And, use it as a draw to the whole north side 'cause it's a beautiful areas, and all we had to do was get a bunch of backhoes and clean it up and tapper the sides and plant grass and trees and dah-dah-dah-dah-dah. (Indistinguishable) reclamation had outdoor grants for that kind of stuff, and the (JCs?) went to research and do some of that and put the grant together and we found out, one, the Slough has questionable ownership itself, and some people who claim ownership actually have titles and deeds that it is theirs.
I: Ooh.
R2: So they fill it in, or have filled parts of it in, so the Slough no longer functions as it was designed to function and it, frankly, I don't know if it could, but its design was, and it takes water from upriver, channels it across through town, and dumps it out downriver, after that low spot and after we're away from that that flooding area.
I: Mmhmm.
R2: So if that number of gallons of water was actually improved and nice banks and parks, and we, you know, it'll beautify the areas, uh, it would be a pretty nice place to be. Well it also would increase land values, and it would make a nice, beautiful homes along the beautiful walkway, and (powers?) would be and a bunch of little guys who were nothings doing (JCs?) [chuckles/coughs], a bucket of heads with country commissioners and other people who says, "But our country club's on the other end of town" and it died in the process of trying to get done. So was there a vision? Yes. Is there a vision? I think so. People could still have a wonderful vision.
I: Mmhmm.
R2: But, a lot of things are gonna have to be straightened out as to how the Slough was originally laid out, planned, or whatever, you know——and you——I don't know if you can ever do that.
I: So part of that problem then really is about——I meant that's

complex, right? That's about, there's property rights issues in there, there's also sort of, uhm, political powers, and what people want Miles City to be.
R2: Correct, there's all of that stuff. Every cotton-pickin' bit of it.
I: Mmhmm, which is so interesting because we put––this is a side note here, but we, I think, we put a lot of money into the science of, like, hydrology and, uhm, how these rivers function and flooding, and very little to understand these problems that you're describing of, uhm, you know, property and sort of the politics of decisions that sound like they're backed by science, but they're really political. It's tough. It's really complex.
R2: (Indistinguishable)
I: Yep. So, I kind of know the answer to my next question, but I want to hear you say it again 'cause I stopped you in the middle of it. Uhm. What is you experience with river flooding events here? Or elsewhere, but here is what I'm also interested in.
R2: We had one flood since we've lived here, and water backed up near … school over there, is that Roosevelt?
R1: I think it's Roosevelt School.
R2: Roosevelt School?
I: And it backed up from the Tongue?
R2: It backed up from the Yellowstone dike having been violated, some people said. Other people said, no, it wasn't violated, it had just washed out. Uh, we had a lumberyard guy down there who went down to the river, he collected sand and rock and gravel and stuff and (sole standing?) rock. Some people blamed him. Some people said it was the Slough didn't function the way it was supposed to because the gate was frozen so it washed the, (yeah?). I've read everything in the world about it.
I: When –– and when was that?
R2: Uhm.
R1: I'd say '76, '77, somewhere in there.
R2: When did we move here?
R1: We moved here in '75.
R2: So it'd be '78 or '79.
I: Okay.
R2: David Bryce lived down there on that (edge wood?), and his house was flooded with a wooden basement. Uh …
I: And that was actually –– so I haven't heard about this, this is the first time I've heard about this flood –– so this is from the Yellowstone, not the –– I've heard about the Tongue backing up, but you're saying that the dike, it breached the dike ––
R2: Yeah, this (undistinguishable) –– well ––
I: –– or it seeped up behind it? Or both? I don't know …
R2: Like I say, I've heard everything in the world about it. I don't know. I could drive you down there and show you where it was at, where the flood was. Like, that's not a problem. The head start is right (next?) to the school. And the head start school had one layer, six inches of sandbags, and water never got into the school. I mean,

that's how far it came up.
I: Okay, okay.
R2: Uh, there's a house down further north of it that just has berm built up like a basement house and it never go flooded.
I: Mmhmm.
R2: So (some of us?) had flooded basement, some didn't. You know, it's –– (indistinguishable) construction and preparedness. There was, uh, there was ranch land down there and I know the guy. I always depended on moving his cattle every spring, 'cause it's gonna happen.
I: Yeah.
R2: He did it every year. Now I don't even know if he'd down there still farming or not. I have no idea. Because this is like, you know, (forty?) years ago. Then, uh, 2010? '09? We had the wettest year we've had. One hundred and thirty-five percent more rain, snow, moisture, runoff, I mean, just –– we had springs out in the mountains that came alive that hadn't been alive in twenty years, and the Yellowstone held. Everybody was watching, everybody was holding their breath. Uh, is it gonna hold? Is it gonna hold? Well the Tongue's putting a lot of weight on it, too. Is it gonna hold? Is it gonna hold? And everything held fine. They say we had a five-hundred-year flood in a hundred-year time, and everything held fine. In fact, some of those springs were still running last summer, which surprised me, I mean, because we were going back through a drought cycle [chuckles/coughs].
I: Right, right.
R2: And yet some of those springs had sprung up during that heavy, heavy rain and heavy snow, were still coming out of the mountainsides. So, kind of a, Oh, how come all of this crap? Well then they come out with, we were gonna reassess us, we made us into a floodplain. So the initial paperwork I got from FEMA was telling how great FEMA in terms that they paid, I'm going to say $800,000 back to the community from 1940 to 1959 or whatever it was time period, that uh, they had taken, you know, and given back to the county for flood insurance, how good flood insurance is. Well, I sat down and just did a little bit of math and I figured, you know, if just half of the people on the north side had flood insurance, that would account for around five hundred homes, and if those five hundred were paying a little (undistinguishable) about $500 a year during those same time frame, that's like $4.8 million, and they only gave us back $800,000? That's not a very damn good deal.
I: Mmhmm.
R2: Now somebody's lying to me when they tell me this is a good deal. I don't want a damn thing to do with it [chuckles/coughs].
I: And ––
R2: Well and then it gets –– you know, we had a (indistinguishable), we have all this dictated, you don't have a choice. If you have a loan, you're gonna have to have it. Well, how come the banks became the collectors all of a sudden? How can you change my contract in the middle of a contract?
I: So you had (American?) ––

R2: I had a real vicious nastiness of that, I really do.
I: Yeah, so, well I can understand that frustration. You had a mortgage and you –– and then, and this is what I'm understanding you saying, I'm just making sure I'm understanding –– and then, uh, FEMA came in and said that with, it was a new map, right, and that, so this house, is that right? Or maybe some of your rentals?
R1: One block over.
I: Okay.
R2: Well, it's this one, too, but both of 'em, yeah.
I: But came –– were in their new maps, meant that you were in a higher risk of flood, and so you had to pay the premium flood insurance price.
R2: Yeah, that's correct.
I: Okay. So your mortgage probably doubled or whatever. There's a chuck of change on top or your mortgage.
R2: Well, no it really didn't. I just went out and bought insurance 'cause I had to.
I: Uh-huh.
R2: Because they said if you don't, we're gonna want full payment right now.
I: Yeah.
R2: I don't think anybody has that right. I have a contract that's different than that.
I: Mmhmm.
R2: And then when I read it, and it says that I can't get paid anything unless that house moves off the blocks … it will not pay a silly thing unless my house physically moves off the foundation.
I: So you have to get, so if I'm understanding you, you have to get insurance, and yet the insurance only pays if your house shifts off the blocks.
R2: Correct.
I: Not if there's damage down in the basement or in the wood or ...
R2: That's correct.
I: Hmm.
R2: You wanna see the policy? [Chuckles]
I: No, I believe ya. I'm just trying, I'm just making sure I'm following ya.
R1: I have one of these (indistinguishable) policies (are?) just statements.
I: Right.
R2: You know, that just is, that's just wrong.
I: Yes.
R2: And then we had one company –– it was an out-of-state company that had a mortgage with us, and that company did not want to accept the insurance that I bought locally, so they were always buying a policy for me and charging it onto my (Esco?) account, which the policy doesn't have an (Esco?) account, 'cause I don't set 'em up with (Esco?) accounts. I take care of my own taxes and insurance and so forth. But, now I have an insurance company writing policies for me because of a mortgage company who says they're not getting the

insurance papers. Well I sent them to them, so then I've got three bills coming out [laughs/coughs], and I'm fightin' with insurance companies, I'm fightin' with the mortgage company, I'm fightin' with FEMA, (saying?) all of this stuff is garbage. Well if you want to survey your property and get it exempt, it's gonna run you about $350 to survey. We don't have $350 extra, especially if it's the three properties we have problems with. And then you have to—if—if the survey takes it out, then you have to pay another $800, $900 to get a certified letter to get certified out of the floodplain, and if you don't do it, we're going to triple the cost of your insurance next year or you gotta keep payin'.

I: So the problem with flooding isn't water. The problem with flooding are entities ... federal, the bank, the insurance companies ...

R2: And, at the same time, we have a government who's trying to shut off some of the stuff that our last twelve years of presidency or eight years of presidency shoved at us, who needs more money, 'cause they've got Sandy Hook, they've got ... uh, New Orleans and they've got [chuckles/coughs] and they've got and they've got and FEMA's broke so where do they get more money? We go out west and we find out who we can screw over.

I: So you — so part of the harm of floods is that places like Miles City and you with the homes that you have hear are being squeezed to support outside —

R2: Even their own data said that where they said they gave us $800,000 but on my estimates, and I don't have all figures by not means but just guessing — we didn't even get about a third of what we gave 'em.

I: Mmhmm, mmhmm.

R2: And during that time frame, too, I haven't, uh, well I have friends that live in Fargo, North Dakota —Fargo, North Dakota's been flooded and flooded and flooded and flooded and flooded and rebuilt and rebuilt and rebuilt and rebuilt —

I: The Red River, right.

R2: After eleven years of rebuilding Fargo they finally took all that area over and made it into a park.

I: (JCs?) are on to something.

R2: [Chuckles] Uhm, I mean it's just, this is what we wanted to do to begin with, folks.

I: Mmhmm.

R2: Yeah, we go through all this stuff with the (flooding stuff?) and they come out and they last thing that really just irritated me beyond all ends ... The Yellowstone River dike is okay. FEMA'll certify it. It's been certified. It's the Tongue River dike that they won't certify.

I: Huhh.

R2: Because it was done privately. And it was added to privately and it was kept in condition privately, and yet the only dike that ever rushed out that I can remember —

I: Was the Yellowstone.

R2: —— was the Yellowstone [Chuckles/coughs].
I: And so because —— so you're saying that because the community or because the city of Miles City built those Yellowstone dikes but this was a private, uhm, entities —— are they land owners that built the dy —— do I have my directions right? The Tongue ...
R2: Well, the Tongue's over here, but yeah.
I: Sorry, the Tongue.
R2: Again, I don’t know, I don't know. The Yellowstone River was built by the Army Corps of Engineers.
I: Huh, okay. I need to write that down. Okay.
R2: Uh, the Army Corps of Engineers had to have input on the Tongue River, but I don't know how much, I don't know what year it was built. Butch Grenze would be an excellent man for you to talk to.
I: Who is that?
R2: Butch Grenze, a former mayor.
I: G-R-I? ——
R1: G-R-E-N-Z-E.
I: ——E-N-Z—— I would never have gotten that one ——G-R-E-N-Z-E.
R2: Stop in at 600 and have a cup of coffee or something and say you're looking for Butch.
R1: And ask for Butch.
R2: He and his wife own the place. And he's a past mayor who fought it and he's still fighting with 'em, trying to get somebody to say they will indeed certify it. Because Butch said in one of the last meetings that I went to that he was at that even if we get certified, FEMA's already said they won't accept a certification. It has to be privately certified by some entity.
I: By an engineering company?
R2: Well, you're gonna pay somebody some mega bucks to do that, too. The Army Corps' gotta take responsibility but (the Army Corps' gonna make you do it?)
I: Mmhmm. So there's both, like, individual harm in terms of or burden in terms of the cost of this insurance but also the community in terms of baring any co —— potential costs for getting these certified then, is that ... which is you, too, right?
R2: No, it is.
I: It's a collective you.
R2: It's all of us, it's all of us. We're all (taken into shorts?)
I: Huh.
R2: Now the county——this house now has paid——and that's the other part of my thing with finances. I had to actually go out and finally refinance things to get the insurance policy straight 'cause I cannot deal with the bureaucrats in companies that live out of state and have nothin' to do with us but keep billing me.
I: Mmhmm.
R2: There is now a class action suit that I'm involved in to get my money back.
I: With other residents here in Miles City?
R2: Uhm, Miles City and the country.
I: Oh. Alright.

R2: 'Cause evidently this (unclear) company was doing it elsewhere.
I: Okay.
R2: And the suit is actually going out of Spokane. Now twenty-five years from now, who knows where that's gonna end up 'cause that's how they play.
I: Right. Delay, wear out.
R2: But, in the meantime, you know, here we sit in our 70s and still trying to check out the bucks and now if I want to get certified, I have to pay, I think it's almost over $1000 just to get surveyed now and, uhm, it's gone up. This survey company come in and when it's all started at $350 and certify me or not but no guarantee, so I just had to pay the $350. Now they'll come in and certify you and I think their rates are now, like, $700 and a seven month waiting list and then you still have to buy a certification letter at $1000. So the two pieces of property that I currently––or three, I guess, that I currently own that could all be exempt are gonna cost me four grand to get out of. Maybe.
I: Right, and a question mark at that. So there's a question mark about getting houses exempted. There's a question mark whether the dikes would be certified, so a lot of the problem is uncertainty.
R2: And, then a question mark as to I can't sell a home unless I get a survey to certify it.
I: Oh, right. 'Cause no one will –– can buy it 'cause no one can get a mortgage. I think they call that being backed into a corner.
R2: I think they call it crooked business.
I: Alright, there you go. Mmhmm.
R2: I mean, I mean if you or I or (inaudible) public citizen did what we are being exposed to from our government, it would be a Ponzi scheme.
I: Mmhmm.
R2: The last guy out turn off the lights, 'cause there's nobody left [chuckles/coughs].
I: Mmhmm.
R2: And I really believe that our politicians have been paying attention to politics rather than ––Washington, D.C. politics –– our senators or representatives (usually end up in the Bat Force?) but not a one of 'em went to (Bat Force?) 'cause they were arguing over whether or not we could get healthcare, whether or not we could get (chips?), whether or we could get –– in the meantime, they sneak this under the (channel?) and it's part of being able to pay for everything else, let's get more money.
I: Mmhmm. And so this came as a really big shock to the system in Miles City in, fairly recently, like when they redrew that map, is that right in whenever? 2010?
R2: 2012, I think, is when that last map came out.
I: Mkay, not sure.
R2: But, it –– it's –– it crept in slowly with rumors and innuendos and, well we're gonna get the dike fixed, we're gonna this, we're gonna that, we're gonna this, we're gonna that and nobody can make up (their?) mind, and then when it came down it was, like, oh crap.

I: Mmhmm. I can't even imagine.
R2: Like when I got letter from the banks it says pay off or have an insurance policy tomorrow.
I: Yeah. That's a big shock, it is. And so, do you —— uhm, to what extent do you think that you, given all this, that I'm here and, uhm —— to what extent do you think you're actually at risk of flooding? ... yeah, zero.
R2: Zero.
R1: Zero, because you're (taping?).
[Collective laughter]
I: It's not a video tape.
R2: Not a video tape.
[Collective laughter]
R2: Uh, well, I g —— I own a house on 7th Avenue and Tatro. You go up 7th across the bridge up to the hill there (unclear) in.
I: Mmhmm. Okay, yep.
R1: And, the Slough runs down the backside of our property.
R2: And, the Slough runs down the backside of our property, and I've seen it flow full and over the top, and my property sits there with a basement, and I've never been wet.
I: Mmhmm. So —— yep. So your experience here, your other two houses ...
R2: And yeah, I can't dig a basement in this house if I want to.
I: Mmhmm.
R2: Because I can't no longer build a basement in a home. Unless it's certified out in this, that, and the other and jumps through all the hoops.
I: Wow, yeah. So, that's pretty intense. So zero risk and yet, uhm, a lot of economic pain and, and uhm, and some real anger, really.
R2: I don't know if you can truly call it anger. God's blessed us with so much all we, all we have to do is be proper stewards of what he's given us.
I: Mmhmm.
R2: And we're tryin' our utmost to do that. Frustration, yeah. Anger, no. I'm not angry about it. I feel sorry for who's going to straighten it out, 'cause we're 70s and we're gonna live in this house 'til we get put in a coffee can.
I: [Chuckles]
R1: We're not going in a casket. We drink too much coffee, it's gonna be a Maxwell House coffee can.
I: Right [laughs].
R2: And our poor children will have to deal with it then, I guess, but, you know, we ... we (only?) thank God for the life we've had here. I mean, it's been a good life. It's put our children through school, it's done what (all?) it needs to do, we're in our senior citizen, and ... (maybe like?) a little miserable at time, yes, but uh, but ... it's only fun (when you're) keeping track. I'd like to see a change and go back to what it were.
I: Mmhmm. And so, what it was before really was that you were trusted to be able to make decisions.

R2: Yeah.
I: For yourself and what you wanted to do. Is that about right?
R2: Yeah. You want to get off on another tangent?
R1: [Laughs]
R2: If you figure the amount of money that FEMA put into New Orleans, and go find those figures 'cause they're available, and you find a population of New Orleans at the time of the great hurricane, do a division work. Every man and woman and child should have had $1.7 million dollars.
I: Mmhmm.
R2: And $1.7 million, they could have moved wherever they wanted or they could have rebuilt, they could have whatever. Now we have homes down there that still haven't been rebuilt, we still have problems in New Orleans, we have, you know, on and on and on it goes. Where does all this money gone? FEMA doesn't do its job.
I: Mmhmm.
R2: They put trailers up for people to live it, they kill 'em. Now we're full of formaldehyde, and FEMA gets away with it. If I, as a homeowner, have lead-based paint, I have to pay to (paint?). If have a home that has asbestos, I have to pay six contractors to come in and remove my asbestos. It's sittin' there, it ain't doing nothing. But, I gotta remove it because I'm a homeowner and I rent that house, but FEMA can do that kind of crap and nobody says anything? Our government has lost its credibility to such a degree that it's just unreal.
I: Mmhmm. So holding standards for us, the citizens, and not themselves to those same standards?
R2: Oh, and now you're holding about healthcare. They have a better healthcare than we have because we're on national healthcare and they're not. They're on their own special system that we get to pay for, too?
I: Mmhmm.
R2: I mean, I can go down this tangent or sixteen or seventeen different things that are all --
I: So this is part of the same pattern, is what you're saying.
R2: There you go. Is it not a pattern?
I: With flooding, you know, uhm, so with healthcare, with ... mmhmm.
R2: I'm sorry that's where my --
I: No, no! Don't -- never apologize for that. So, but this is interesting. So th -- so your story about flooding is really the story about the relationship with and what government is doing.
R2: Is it not? I mean, is it not?
I: Mmhmm.
R2: I don't know how many people think the way I think, but the things are all tied together.
I: Yeah.
R2: You cannot deny that they're not tied together.
I: Mmhmm, yep.
R2: I mean, they call that profiling, but we're not supposed to do that, are we? But facts are facts, and if they lead to certain

things, that's not profiling. That is what it is.
I: Mmhmm.
R2: And, we are supposed to do that.
I: Mmhmm.
R2: But, back to my story. God's been good to us, and I want to be a good steward for what he's given me and make decision over things that he's given me, not the government making bad decisions.
I: Mmhmm, yeah. I think that, that's beautifully put. So, so this is –– well, and maybe, I don't want put words in your mouth, but I'm going to ask you if this makes sense. So the next question is, like, how do you prepare? How do you prepare for flooding events and how does your community prepare ... uhm, for flooding events. So I think part of how you prepare is really to be at peace with some of these relationships, particularly with God for you, uhm, but are there other ways they people prepare, maybe you or in the community, for flooding?
R2: I think our community does a really good job with the county sheriff who monitors the river during, you know, ice breaks (unclear). I think there's a lot of activity about preparedness. Sandbags are made available to those that need 'em. Uhm, I think there's, you know, in the isolated times that we've had stuff, they've really done quite well. Uhm ...
I: So they ––
R2: (Am I?) in my 70s ready for any of it? No [coughs]. My simplicity is truly that, just pray. Lord, bless the property you've given and we try to do our best. Be kind.
I: Yep, yep.
R2: And I know that I've seen hail storms go around Miles City when it should have wiped us out because of prayer.
I: Mmhmm, mmhmm.
R2: Or, whatever, but I'll give it prayer every time. Uh, but no, I think the city is, and the county do a really good job in preparedness. And I think (Sam's?) done a great job trying to get this thing understood by people, and it's still not, I don't think, understood. It's, it's ...
I: Well it's certainly complex, and it is hard to, to, uh, to relay that, but you're, yeah, I think that, uhm, I think that there is a lot of effort in trying to talk. I mean, Miles City, you guys are the ones who responded, right? "Yeah, I'll come talk to you," so it's obvious that there's some, there's real interest here.
R2: Well I really think for your benefit you gotta talk to Butch Grenze.
I: Mkay, I will. I'll give him a ring.
R1: I think one way to prepare is every dollar that every family has to put in to flood insurance was put in to flood insurance but put in to Miles City to where Miles City could hire the contractor to do whatever is necessary, if it's the dike, the Yellowstone or the Tongue, get that taken care of.
R2: Butch had an agreement with FEMA to try and do that and FEMA wouldn't do it.
I: Mmhmm. So part of pre––a good idea for preparation would be to

have those insurance agencies be local so that those moneys can be kept in-house?
R1: Mmhmm.
I: Mmhmm, mmhmm.
R1: And set aside specifically for repairing the dikes.
R2: Well, I'll look at California's situation right now with that dam.
I: Right, Oroville dam, mmhmm.
R2: They purposefully altered funds from maintaining that spillway because it had never been needed, (the dam had?) never been used, and gave it to their, let's help these illegal aliens get full benefits. Well now they have to use it and they found out, we should have been repairing it 'cause this thing has been here. Oh well, now that it's broken down, we'll just ask the Feds to put more money in here, FEMA will come through. Oh, by the way, if FEMA doesn't come through, it's gonna wash out more and cause a bigger problem.
I: Mmhmm.
R2: So now they have blackmail. I (haven't never?) thought about what Vivian just said, but that's a good idea. So much of——eighty percent of what we put into FEMA should be in this community. Twenty percent can go to the federal. And then they can disperse it throughout the nation. Or twenty-five percent stay here and seventy-five throughout the nation.
I: It's still (twenty percent?) more, yeah.
R2: What the bottom line is, we're not getting a third of it back now.
I: Gotcha, mmhmm.
R2: And they tell us it's a good deal. That's not a good deal. I buy house insurance for fire and if my house burns down, I do get to rebuild.
I: Mmhmm.
R2: But, I can't rebuild in this area right now because I can't put a basement back under the house that had one for the last hundred years. Why am I buying insurance then, even for fire and flood?
I: Mmhmm, right.
R2: I mean, it's, to me it's like, stupidity.
I: Mmhmm, mmhmm ... We, so I wonder when was, so when was, uhm, Butch the past mayor when he was working with FEMA on this? How long ago was that?
R2: Uh, Butch was mayor for two terms and then he was out for a term and then went back in for two.
R1: I don't remember the time.
R2: Butch has got a lot of history of him.
I: Okay.
R2: And he's only been out since this last election cycle.
R1: Yeah, 'cause with this last election cycle, Hollowell got it.
R2: Yeah.
I: Uh-huh, right.
R2: And Hollowell's a neat, young guy.
I: Good, I don't know him. Well that is very interesting in terms

of, uhm, how the city can be prepared or how the community can be prepared for floods. That's an excellent idea. Uhm, are there other people in, you know, that you know who prepare for floods in ways that may have a different level or perception of their own risk? Do other people prepare in ways?
R1: 'Cause there hasn't been a flood, not too many people ever prepared for it. I mean, yes, they prepared for their home insurance and other things, but not for flood.
I: Yep. So was there, yeah, there's no reference for it, so why prepare for it? Since there hasn't, it hasn't happened, then ... then there's no need to really prepare for it because ...
R2: Well I think the –– there's people around town, this guy over here on this corner, since the flood insurance thing came in, he had someone come in and contour his property up onto the wall of his house, put in a brick retaining wall. Now I don't know if that was just for him to get high enough to say he's out of the floodplain or if he's actually doing it in preparedness or why, but there's been several people who have done things ––
I: Yeah, okay.
R2: –– individually. Now I don't know any of them personally. I can see work done around town.
I: Gotcha.
R2: And if, you know, if that's what it takes, fine. I've done a GPS of my house, and it's, it's a certified GPS from a government institution, but it's, it can be used for surveying purposes. I'm six inches higher than people who are out of the ground, so under theory, I should be able to pay the $1800, $1600, whatever it costs and get out of it.
I: Mmhmm.
R2: Well when we refinanced this last time, everything paid for except one house, so I only have one now with an insurance policy, and I'm not gonna have a problem with any of 'em until I have to sell 'em.
R1: That's flood insurance policy [laughs].
R2: Flood insurance. Flood insurance, yes. And, that's because of banks involved.
I: Mmhmm.
R2: Uh, but other than that, we got rid of the other two that I, that I had flood insurance for that I no longer need, 'cause they're mine. I own 'em. They're free and clear. But if I go to rebuild or try to sell I'm still not really free and clear.
I: Yeah, and that's –– yeah, I can see how that can be trouble in one –– that's, that's your, that's your nest egg. Or, I don't know what I'm saying ––
R2: Well there's ––
I: –– you know, could be.
R1: Mmhmm.
R2: We've been, uh, we've bought rentals and we've been fixing rentals and we tried always to make them not a dump, but ... that's what we live on, is our income.
I: Yep.

R2: It is a limited income. Now, I can increase rent to everybody because that's being increased to me, but frankly, that's not something that's good for the single lady that's livin' there that's on a fixed income worse than mine.
I: Mmhmm.
R2: Uhm, the single mom who can't afford –– I mean, we're some of the cheaper places in town, our (unclear) rentals.
R1: And they're not dumps, either.
R2: And they're not dumps.
I: Mmhmm. So they appreciate that, I'm sure.
R2: Right. I think we just put $4000 in this year to one.
R1: Five for the roof alone.
I: Mmhmm.
R2: $9000 (or two?)
[Collective laughter]
I: Alright, but (inaudible)
R2: Well, and that's the one in the floodplain, too [laughs/coughs].
I: Hmm. Uhm, so in terms of information, uhm, when you need information, whether it's about flooding itself or information, uhm, you know, about these different, uhm, governmental processes, where do you go for that kind of information?
R2: Well, I've got to a couple of the college programs that Sam's sponsored, put on.
I: Uh-huh.
R2: And, uh, most of it I get from, at (Dan Mayor?) 'cause I see him and know and we talk.
I: Mmhmm, mmhmm.
R2: So I get most of it right from, from (unclear). We don't subscribe to newspapers, so unfortunately we don't get it there.
I: Mmhmm.
R2: So if somebody says something to us, we'll pick it up, but no, we don't get information unless it's ...
I: So from that, from really the city officials who are working in this every day is where you get your information, and then neighbors or whoever you end up talking to per chance?
R2: Yeah.
I: Yeah.
R2: That's right.
I: Uhm, and so what kinds of information? (Unclear) where you get your information, but what kinds of information?
R2: I don't know. What do you want to count? Some of it I guess could be called gossip just because it's from the neighbors and stuff, but then when I talk with the mayor and the mayor says, this, this, and this and this, that's pretty factual and pretty close to the surface. Go to the college meetings, we get to meetings that, uh, engineering companies and stuff, you know, passing out.
I: Mmhmm.
R2: A lot of that I still question for the same reason I question this. You know, we haven't had a flood, we have never been flooded. How can you do this?

I: Mmhmm.
R2: Well that railroad track over there and this railroad track over here makes this a natural floodplain. No, it doesn't. 'Cause right now with the cotton-pickin' Tongue River's on the other side of that railroad [laughs/coughs].
I: So --
R2: I don't know.
I: Right. So the kinds of information, though, or what, you know, it could be what, about, uhm, how do you -- that conversation that you just said. So that's, you know, how do you determine if this is floodplain. Another kind of information is, you know, what's happening with the river right now, like, where's the ice? Is it melting? Are there channels? So, so there are different kinds of information, I think, out there that are --
R2: Well I think the Army Corps of Engineers has done computer models, and they've taken six different models and take the very, very worst one and called that our floodplain, when in fact, such things as the dam up on Decker, you know, it's way up on the Tongue River ...
I: Oh, right, right.
R2: That dam is a earthen dam. It has been getting widened and improved and the State of Montana's been suing Wyoming to allow us to refill it, because they drained it in order to do dam repair. Now the dam has been repaired, the dam has been raised higher than it has been and now we're suing to get the water back into it because Wyoming says they don't have to give us their water, and finally Montana wins -- I mean, this -- you talk about complicated. Finally within the last month or so, Montana wins the thing, so Wyoming has to give up the water and allow us to fill the dam back up to the height and above that it used to be, but it's now improved. This has been a twenty-year process. How much of that twenty year process has been included in with the Army Corps of Engineers is saying, this is what's gonna happen? You take out of here on the TY ditch -- we talked about that. TY runs up on the high side of the Yellowstone valley and then runs back into the Yellowstone. A lot of farmer, ranchers have contoured their property, so that the Tongue River dam, or Tongue River ditch feeds downhill into the farmland.
I: Using gravity.
R2: Using gravity. So now at the low end of that, they might be seven, eight, ten inches below , quote-on-quote, floodplain because they've contoured it to be (unclear) floodplain to utilize the water in an overpass. How much of that is currently even into -- and the only reason I come up with that is, the one company that built a big building out there had to fill the one end of their field in seven feet in order to build their plant. This end's out of the floodplain but this end wasn't, so they had to build everything up.
I: Oh, wow.
R2: So here you got this major, uh, chemical plant out there that had to redo all of its blueprints because of floodplain got changed, and they did, but some of our land may actually be floodplain for a reason, like Tongue River is supposed to carry through the Tongue

Rive, TY ditch and dump into the Yellowstone through farmland, and the farmland is damaged by a flood, but it actually is healed, too.
I: So benefits.
R2: I live, I lived, was born and raised in Arizona, so we had big, huge cannels.
I: Mmhmm.
R2: Maybe our little TY ditch needs to a TY cannel and carry more water if needed. So can you take more water and run it the whole way.
I: It's like (a hand of the pulse?), right? Mmhmm.
R2: But, that's never been discussed that I've ever heard of. That's not part of Bureau of Reclamation or is it or isn't it?
I: Yeah. So that could be a solution.
R2: There could be all kinds of solutions with people thinking outside (hats?).
I: Mmhmm.
R2: But, they've got this little hat that says the Army Corps of Engineers or somebody else has to do it. No. But, the Army Corps comes in here with their set of plan, which is the worst scenario (committee?), and I don't think the worst ever happens.
I: Right, right. So planning on the worst, the worst that could happen, uhm, is unavailable ——
R2: (Unclear) It's unrealistic. It's making it look worse than it is.
I: Yep, yep.
R2: More like a reality, and it's not.
I: Yep, mmhmm. Gotcha, gotcha. Well do you wish that you had any other kinds of information?
R2: Do we need other information?
R1: Truth.
R2: Truth [laughter].
I: Alright, so it's not necessarily, like, yeah, it's not like a different like hydrological engineering information. It is just truth. (Unclear).
R2: And every engineer's gonna design a different floodplain.
I: Mmhmm.
R2: You know, (unclear) I think has already designed two or three different ones and showed 'em different ways, and then they go back, the Army Corps of Engineers say, "Well will you please accept ours." Well, there's only one way that can happen.
I: One way what can happen?
R2: Water goes downhill.
I: Ahh.
R2: And it spreads as it goes.
I: And it spreads, yeah.
R2: And if TY ditch is gonna get overflooded, maybe they do need to address TY ditch and carry it further away from the town.
I: Mmhmm.
R2: If the dike down here is really that big of a problem and yet it survives the worst of the worst that's ever been seen [laughter] in recorded history, (everything's?) pretty good.

I: Mmhmm.
R2: But, now in order to get it certified, Butch was saying to me, that we have to tear the entire dike down to its roots and start with a new foundation, new base, 'cause everything that's there hasn't been certified so they can't guarantee it ... It's worked, it's good.
I: Mmhmm.
R2: You can drive pylons through anything and reinforce anything, but to pay rip-out, new rip-in, (unclear), to me that just sounds like somebody in the Union's getting' a hell of a lot of money for something that doesn't need to be done. Or the Army Corps of Engineers does it and charges us for it again [laughs/coughs].
I: Part of that Ponzi scheme that you were talking about, I think, is what you're saying ...
R2: I don't know.
I: Yeah. Nah, that's interesting, it's really interesting. So, but ––
R2: Bet you didn't expect this kind of –– [laughter]
I: No! I didn't, but it's awesome. So, but I –– but I just want to make sure I'm understanding, with that, with that dike out here, at one point I thought you said that FEMA would've certified that but now they won't ... or?
R2: The Yellowstone they have. The Yellowstone's not the problem.
I: Oh, okay. It's the Tongue.
R2: It's the Tongue River that's the problem.
I: Gotcha, gotcha. Alright.
R2: And yet, they're always saying it's the Yellowstone that's gonna overflow.
I: Gotcha.
R2: See, it's ...
I: Yep, yep, I was –– my brain was on the Yellowstone, but you're talking about the Tongue just now.
R2: And it's on the Yellowstone River, that's the other thing. See, it's the Yellowstone River, but it's gonna flood because Decker breaks under certain circumstances. The highest rainfall ever recorded in the month of June, the highest snowfall for the winter 'til it's there, and the slow melt-off that doesn't melt until we have a super high rain, [telephone rings] and then we have the warmest June that we've ever had in the world's history –– we have (freak?) catastrophes [telephone rings] that have all occurred because the Decker Dam to get overloaded to flood to break the dam that will come downriver through the Tongue River, flood everything from the Wyoming border at Sheridan all the way down to here to where the Tongue joins, if you have that much water and you have all these freak catastrophes could possibly occur at the same time.
I: Mmhmm.
R2: You think something along the way might go out like, uh, part of the river they have counted on going out and it'd flood all that plain.
I: Mmhmm
R2: Or [laughs]

I: Right, right, right.
R2: You know, Cherry Creek would open up and end up being Cherry Creek again and dumping in on the other side of the city and we don't know. There's a lot of what-ifs and all that's th ––
(indistinguishable) Creek might open up again [laughs/coughs]. Who knows?
I: That's awesome. So it's, part of it is, like, the, ha, there's so much un ––
R2: (Unclear) they do?
R1: Are we picking up girls tonight? Are you going?
R2: I had planned on not going, just because.
R1: I'll call her back.
I: Are you talking grandkids?
R1: No.
R2: You can't get her on that phone, can you?
R1: (House?) phone.
R2: Oh ... No, I just, I don't think I'm going. I have basically been bedridden for three weeks.
I: Oh, jeeze!
R2: House-stuck anyway, not bedridden. But, I had a, I have COPD and I think (you can hear?) my breathing change.
I: Yeah.
R2: [Coughs]. And, I've gotten some kind of infection and it just laid me out.
I: Yeah, that's tough.
R2: Today's the first day I've been off oxygen 100%.
I: Oh, wow. Well congratulations, and jeeze, I hope that maybe with some warmer weather and sunshine you can, uh, oh ...
R2: (Yep?). It's gonna occur from time to time.
I: Yeah. Wow. Well is there anything else that you can think of that, uhm, that we haven't covered, that I've missed?
R2: Oh, darlin', you only wanna go down that list.
I: With flooding.
[Collective laughter]
R2: No, I think they're trying to do what they can do, and I really, you know, I think it's high time that several communities get together and say no.
I: Mmhmm. So some kind of collective action with, uh, several communities ...
R2: Well ...
I: To, uhm, to say no to ––
R2: I think if all those communities had gotten together at the beginning and told our senators and representatives to get on the stick and say no to FEMA ...
I: Mmhmm.
R2: … uh, that 2012 thing that was passed through Congress was passed through ... a whole bunch of crap.
I: Mmhmm.
R2: Other things around it, surrounding it and it got stuck in there.
I: Mmhmm.

R2: I don't think anybody knew how bad it was gonna be 'cause it's one of those things, well you can tell until you read the bill.
I: Mmhmm.
R2: Well that's garbage.
I: Right, right.
R2: And then, you know, like I told Vivian, for us to have to change the finance companies and fight the way we fought to try and get them from raping us [coughs].
I: Well, and you're on a collective action (unclear) ––
R2: And yeah, and now that's gone finally.
I: Yeah, mmhmm.
R2: But, I've been arguing that since, well, five years?
I: Yeah.
R2: And, it just came through that I got that notice that's it, uh, part of a class action.
I: Mmhmm. Wow. So the preparedness is really, tell us about the river and more about communities preparing for, uhm, trying to make things right for individuals and communities.
R1: Oh, we have mills levies for school. If they did mill levies for flood and that money stayed locally, other than paying FEMA or whoever for the flood insurance, that's one way of being prepared.
I: Yep.
R1: And set aside and designate it, "This is for preparing for flood."
R2: They're trying to do that with this current thing. They haven't been able to get anybody to agree.
I: Who needs to agree? What do you mean anybody?
R2: Uh, who's gonna certify it? Who's gonna build it? Will the Army Corps come behind it? Will they, they FEMA say okay to altering our insurance policies to them to go to rebuilding, to preventing?
I: Mmhmm, mmhmm.
R2: If you slapped a, uh, another tax on me and says it's just for rebuilding that and FEMA's not gonna give me some kind of way out, now I get double taxation until I can fix the dike?
I: Right.
R2: But, nobody is saying we can fix it, certify it, and guarantee that we can get out of FEMA? Ah, to my way of thinking, FEMA is a bunch of bureaucrats who are put in to power by a const –– by a Congress movement. They can damn well be taken out of by a Congressional movement.
I: Mmhmm.
R2: And if our Congressional people do not get their crap together, then we need to stop electing the same people and sending the same people back to Washington to do the same crap over and over and over again.
I: Yep.
R2: And this goes back to, how long did Max Baucus sit there in that office and do very little? How long did his predecessor sit in that box and do so little? How much money have they donated to your university and built buildings there because they're the great (mucky–

whats?) of Montana? And we can honestly say they were good people, but they did not do their job taking care of some of these issues that (we are now facing?), that my grandchildren will face. And maybe my great-grandchildren will face.
I: That's a problem that's been passed forward.
R2: And if they keep doin' it, the people of Montana really need to sit down and say no.
I: Mmhmm.
R2: And if it takes a collective action of three, four, five communities -- look at the red map and the blue map and see who got what this time and where are we at with power (and?) the structure. And no matter what goes on there's still that ... you know, but ...
I: Mmhmm ... Well you guys have come up with some really interesting ideas. Uhm ...
R2: (You?) could have run out of tape. I don't know how long your tape's good for.
I: I think I got sixteen hours. I don't know --
R2: (Indistinguishable). That is a good one.
I: Yeah, it's great. It has a little USB, so I just plug it in. It's awesome. Uhm, so we are definitely going to, uhm, look up Butch, is it Grenze? Gren --
R1: Grenze.
R2: Grenze.
I: Grenze. I just want to make sure I'm pronouncing his name right. We're coming back through -- uhm, we're gonna to Glendive in two weeks and so I'm, they're a couple other people that we're trying to catch here, so I'll have a couple weeks to catch him. But maybe, do you think I could, like, go to the 600 Café and say, "Hey ... "?
R2: Oh, if Butch is there, he would in a heartbeat.
I: Really?
R2: Yeah.
I: Okay, I might just try that, too. [Laughs]
R1: But, they closed at two o'clock, so you're not gonna get to go there today.
I: Well tomorrow, yeah.
R2: Go in there for breakfast even.
I: Mmhmm. What time do they open, I wonder.
R2: 5:30?
R1: 5:00 or 5:30 in the morning.
I: That's probably a good time. I would imagine it's a little slow at that hour?
R2: Well, he may not even get there until 8:00, I don't know.
I: Oh, okay. Hmm, I could try, though. I really appreciate it. I'm gonna turn this thing off. Uhm --

APPENDIX C

TERM FREQUENCY CALCULATIONS

This is a sample of all the term frequency calculations we provided for narrative construction. The full spreadsheet can be viewed at `https://github.com/hawk132/Flood-Text-Analysis/blob/master/kahuna.pdf` The terminology for this appendix is defined as:

1. Term - a word that appeared in one of the corpora
2. Term.Count- the occurrences of a term in a corpus (Hero, Victim, Villain)
3. RF - Relative frequency
4. TRF - Taking the Coded language by interview, this is the relative frequency of a term to the size of the corpus, where the term count has been transformed by the square root. The documents in these corpora are the coded language of a label(Hero, Victim, Villain) by interview.
5. TF-IDF - Term Frequency Inverse Document Frequency.
6. DFc - The number of documents, or coded narratives, the term appeared in divided by the total number of documents in the corpus.
7. IFc - Interview Frequency. The number of citizen interviews the term appeared in.
8. IRF - the relative frequency of a term in the interviews corpus (ignoring narrative label), where each document was all the coded language from that interview. Transformed by the square root.
9. HRF-VicRF, HTRF-VicTRF, and so on. - The difference between the relative frequencies of a term between corpora. Positive values indicate the term was used more in the first corpus than the second.
10. HRF / Hero DFc and so on. - The relative frequency of a term divided by its document frequency.
11. HRF * Hero DFc and so on - The relative frequency of a term multiplied by its document frequency.

Term	H.Term.Count	Vic.Term.Count	Vil.Term.Count	Hero.RF	Victim.RF	Villain.RF	HTRF	VicTRF	VilTRF
help	79.00	30.00	51.00	51.41	14.66	17.13	27.97	9.15	10.45
river	200.00	241.00	446.00	130.14	117.75	149.84	46.27	31.38	35.78
dike	158.00	170.00	255.00	102.81	83.06	85.67	41.50	27.83	26.84
done	71.00	65.00	101.00	46.20	31.76	33.93	26.67	15.47	13.75
good	69.00	42.00	70.00	44.90	20.52	23.52	24.51	13.70	12.85
think	143.00	149.00	242.00	93.05	72.80	81.30	41.60	31.37	28.07
city	93.00	118.00	158.00	60.52	57.65	53.08	31.42	21.38	19.08
built	91.00	92.00	120.00	59.21	44.95	40.31	26.69	17.00	16.74
job	25.00	14.00	20.00	16.27	6.84	6.72	12.58	3.39	4.18
miles	64.00	71.00	91.00	41.64	34.69	30.57	23.92	15.16	13.17
interview	1.00	0.00	0.00	0.65	0.00	0.00	26.14	18.46	13.67
pretty	42.00	33.00	51.00	27.33	16.12	17.13	18.25	10.81	10.08
together	20.00	3.00	8.00	13.01	1.47	2.69	8.67	1.46	2.39
levee	44.00	40.00	42.00	28.63	19.54	14.11	14.50	7.52	5.45
helped	20.00	1.00	5.00	13.01	0.49	1.68	7.38	0.49	1.75
yellowstone	43.00	32.00	61.00	27.98	15.63	20.49	15.78	8.95	11.65
people	182.00	241.00	293.00	118.43	117.75	98.43	45.35	38.70	30.91
water	125.00	144.00	272.00	81.34	70.36	91.38	35.05	28.47	29.81
engineers	48.00	47.00	63.00	31.23	22.96	21.17	16.18	9.67	10.01
big	73.00	62.00	102.00	47.50	30.29	34.27	25.01	18.53	17.33
community	44.00	58.00	59.00	28.63	28.34	19.82	21.07	14.70	10.31
corps	70.00	79.00	113.00	45.55	38.60	37.96	20.90	14.54	14.65
really	100.00	109.00	172.00	65.07	53.26	57.78	32.68	26.36	23.67
care	27.00	16.00	30.00	17.57	7.82	10.08	12.60	6.30	7.24
keep	35.00	29.00	43.00	22.77	14.17	14.45	15.84	9.70	9.57
dirt	36.00	22.00	31.00	23.43	10.75	10.41	9.99	3.96	4.53
ice	68.00	86.00	151.00	44.25	42.02	50.73	20.74	14.72	17.80
things	49.00	58.00	79.00	31.88	28.34	26.54	21.53	15.73	14.66
stuff	56.00	63.00	81.00	36.44	30.78	27.21	20.75	15.02	13.51
took	41.00	38.00	59.00	26.68	18.57	19.82	15.44	9.80	10.82

Term	Hero.TFIDF	Victim.TFIDF	Villain.TFIDF	Hero.DFc	Victim.DFc	Villain.DFc	IFc	IRF	HRF-VicRF
help	9.59	5.88	7.45	0.10	0.03	0.03	30.00	6.63	36.75
river	12.20	21.64	28.77	0.24	0.17	0.22	38.00	15.96	12.39
dike	18.54	21.17	26.56	0.21	0.16	0.17	31.00	13.20	19.75
done	7.49	8.70	12.02	0.12	0.07	0.07	32.00	7.99	14.44
good	9.44	6.11	8.10	0.11	0.05	0.06	32.00	7.18	24.38
think	12.61	19.81	23.62	0.20	0.16	0.18	38.00	13.61	20.25
city	11.46	11.67	13.87	0.15	0.11	0.12	29.00	9.90	2.86
built	10.98	12.38	14.16	0.11	0.08	0.09	31.00	9.00	14.26
job	4.06	3.29	3.67	0.04	0.01	0.02	16.00	3.14	9.43
miles	7.04	9.20	10.82	0.11	0.08	0.08	28.00	7.63	6.95
interview	0.15	0.00	0.00	0.00	0.00	0.00	39.00	6.23	0.65
pretty	5.05	7.44	7.91	0.08	0.04	0.05	29.00	5.83	11.21
together	3.91	0.90	2.01	0.04	0.00	0.01	14.00	2.32	11.55
levee	8.17	6.17	6.48	0.06	0.03	0.03	14.00	4.29	9.09
helped	3.75	0.05	0.90	0.03	0.00	0.01	11.00	1.81	12.53
yellowstone	3.35	5.00	6.61	0.06	0.03	0.05	24.00	5.84	12.35
people	13.82	19.26	21.90	0.26	0.22	0.20	39.00	16.03	0.68
water	10.77	14.99	19.95	0.16	0.12	0.15	38.00	13.53	10.98
engineers	5.74	7.89	11.20	0.07	0.05	0.05	22.00	5.58	8.27
big	6.22	7.77	9.92	0.13	0.07	0.09	36.00	8.60	17.21
community	8.32	7.48	7.05	0.08	0.05	0.05	29.00	6.94	0.29
corps	7.52	12.83	18.09	0.10	0.08	0.09	23.00	7.48	6.95
really	9.70	13.89	16.72	0.17	0.12	0.14	37.00	11.61	11.81
care	3.75	3.62	7.61	0.05	0.02	0.03	23.00	4.21	9.75
keep	5.49	5.32	6.77	0.06	0.03	0.04	27.00	5.47	8.61
dirt	4.60	3.57	4.67	0.05	0.02	0.02	11.00	3.11	12.68
ice	7.23	13.71	19.81	0.08	0.07	0.09	28.00	8.27	2.23
things	5.29	9.44	11.36	0.09	0.06	0.07	31.00	7.73	3.55
stuff	5.97	7.88	10.42	0.10	0.07	0.08	33.00	7.46	5.66
took	5.57	5.52	6.56	0.06	0.04	0.05	27.00	5.73	8.11

Term	HRF-VilRF	VicRF-VilRF	VicRF-HRF	VilRF-VicRF	VilRF-HRF	HTRF-VicTRF	HTRF-VilTRF	VicTRF-VilTRF
help	34.27	-2.48	-36.75	2.48	-34.27	18.82	17.52	-1.30
river	-19.69	-32.08	-12.39	32.08	19.69	14.89	10.49	-4.40
dike	17.14	-2.61	-19.75	2.61	-17.14	13.68	14.67	0.99
done	12.27	-2.17	-14.44	2.17	-12.27	11.20	12.92	1.72
good	21.38	-3.00	-24.38	3.00	-21.38	10.81	11.66	0.84
think	11.75	-8.50	-20.25	8.50	-11.75	10.23	13.53	3.30
city	7.43	4.57	-2.86	-4.57	-7.43	10.04	12.34	2.30
built	18.90	4.64	-14.26	-4.64	-18.90	9.69	9.95	0.26
job	9.55	0.12	-9.43	-0.12	-9.55	9.19	8.40	-0.80
miles	11.07	4.12	-6.95	-4.12	-11.07	8.76	10.74	1.98
interview	0.65	0.00	-0.65	0.00	-0.65	7.68	12.47	4.79
pretty	10.20	-1.01	-11.21	1.01	-10.20	7.44	8.17	0.73
together	10.33	-1.22	-11.55	1.22	-10.33	7.21	6.28	-0.94
levee	14.52	5.43	-9.09	-5.43	-14.52	6.98	9.05	2.06
helped	11.33	-1.19	-12.53	1.19	-11.33	6.89	5.62	-1.27
yellowstone	7.49	-4.86	-12.35	4.86	-7.49	6.82	4.13	-2.69
people	19.99	19.32	-0.68	-19.32	-19.99	6.65	14.44	7.79
water	-10.04	-21.02	-10.98	21.02	10.04	6.58	5.24	-1.34
engineers	10.07	1.80	-8.27	-1.80	-10.07	6.50	6.16	-0.34
big	13.23	-3.97	-17.21	3.97	-13.23	6.48	7.68	1.20
community	8.81	8.52	-0.29	-8.52	-8.81	6.37	10.76	4.39
corps	7.59	0.64	-6.95	-0.64	-7.59	6.36	6.25	-0.11
really	7.29	-4.53	-11.81	4.53	-7.29	6.32	9.00	2.69
care	7.49	-2.26	-9.75	2.26	-7.49	6.30	5.36	-0.94
keep	8.33	-0.28	-8.61	0.28	-8.33	6.14	6.27	0.12
dirt	13.01	0.33	-12.68	-0.33	-13.01	6.04	5.47	-0.57
ice	-6.48	-8.71	-2.23	8.71	6.48	6.02	2.95	-3.08
things	5.34	1.80	-3.55	-1.80	-5.34	5.80	6.88	1.08
stuff	9.23	3.57	-5.66	-3.57	-9.23	5.73	7.24	1.50
took	6.86	-1.25	-8.11	1.25	-6.86	5.64	4.61	-1.03

Term	VicTRF-HTRF	VilTRF-VicTRF	VilTRF-HTRF	HRF / Hero.DFc	VicRF / Vic.DFc	VilRF / Vil.DFc	HRF-VicRF / HDFc
help	-18.82	1.30	-17.52	495.17	547.47	507.58	353.98
river	-14.89	4.40	-10.49	553.39	698.09	679.64	52.69
dike	-13.68	-0.99	-14.67	500.28	525.82	507.58	96.11
done	-11.20	-1.72	-12.92	396.48	465.17	466.19	123.93
good	-10.81	-0.84	-11.66	392.45	393.05	365.47	213.08
think	-10.23	-3.30	-13.53	457.50	468.81	445.51	99.56
city	-10.04	-2.30	-12.34	396.71	518.88	457.46	18.76
built	-9.69	-0.26	-9.95	517.57	532.98	471.83	124.67
job	-9.19	0.80	-8.40	383.91	510.97	424.65	222.48
miles	-8.76	-1.98	-10.74	385.42	446.78	397.01	64.37
interview	-7.68	-4.79	-12.47	307.13	0.00	0.00	307.13
pretty	-7.44	-0.73	-8.17	348.64	430.15	360.95	142.95
together	-7.21	0.94	-6.28	341.26	364.98	318.48	302.82
levee	-6.98	-2.06	-9.05	450.46	583.96	495.42	142.97
helped	-6.89	1.27	-5.62	409.51	364.98	318.48	394.13
yellowstone	-6.82	2.69	-4.13	489.14	467.17	413.35	215.81
people	-6.65	-7.79	-14.44	461.97	546.33	496.36	2.64
water	-6.58	1.34	-5.24	511.89	590.53	597.43	69.11
engineers	-6.50	0.34	-6.16	421.21	476.50	418.01	111.53
big	-6.48	-1.20	-7.68	380.01	411.43	382.18	137.67
community	-6.37	-4.39	-10.76	365.24	516.31	436.99	3.73
corps	-6.36	0.11	-6.25	457.43	488.70	444.31	69.80
really	-6.32	-2.69	-9.00	388.77	437.17	411.87	70.58
care	-6.30	0.94	-5.36	331.70	364.98	341.23	184.11
keep	-6.14	-0.12	-6.27	358.32	441.01	391.28	135.39
dirt	-6.04	0.57	-5.47	502.58	573.54	493.65	271.96
ice	-6.02	3.08	-2.95	580.14	615.45	559.20	29.22
things	-5.80	-1.08	-6.88	358.32	470.42	399.37	39.85
stuff	-5.73	-1.50	-7.24	365.94	418.07	358.29	56.82
took	-5.64	1.03	-4.61	419.75	478.25	417.57	127.63

Term	HRF- VilRF / HDFc	VicRF- HRF / VicDFc	VicRF - VilRF / VicDFc	VilRF -HRF / VilDFc	VilRF -VicRF / VilDFc
help	330.13	-1372.53	-92.47	-1015.30	73.35
river	-83.75	-73.46	-190.22	89.33	145.53
dike	83.42	-125.03	-16.51	-101.57	15.45
done	105.29	-211.52	-31.83	-168.56	29.85
good	186.89	-466.93	-57.38	-332.29	46.56
think	57.77	-130.41	-54.74	-64.39	46.58
city	48.74	-25.75	41.16	-64.07	-39.41
built	165.20	-169.12	54.97	-221.19	-54.26
job	225.34	-704.22	9.05	-603.46	-7.66
miles	102.48	-89.58	53.04	-143.80	-53.48
interview	307.13	0.00	0.00	0.00	0.00
pretty	130.07	-298.96	-26.95	-214.79	21.28
together	270.78	-2875.52	-304.24	-1223.68	144.79
levee	228.46	-271.53	162.36	-509.84	-190.78
helped	356.65	-9356.52	-889.81	-2148.98	225.85
yellowstone	130.88	-368.88	-145.16	-151.02	97.99
people	77.99	-3.14	89.62	-100.82	-97.40
water	-63.20	-92.16	-176.45	65.65	137.44
engineers	135.78	-171.60	37.32	-198.86	-35.52
big	105.87	-233.72	-53.98	-147.60	44.33
community	112.38	-5.33	155.18	-194.22	-187.77
corps	76.19	-88.00	8.05	-88.79	-7.44
really	43.53	-96.98	-37.17	-51.93	32.27
care	141.42	-455.27	-105.57	-253.60	76.56
keep	131.04	-267.84	-8.62	-225.59	7.50
dirt	279.14	-676.37	17.84	-616.71	-15.85
ice	-84.98	-32.65	-127.58	71.44	96.01
things	60.06	-58.87	29.85	-80.42	-27.05
stuff	92.66	-76.85	48.47	-121.49	-46.99
took	107.89	-208.96	-32.32	-144.47	26.43