A Closer Look at Final Statements

Introduction

According to one recent study, **at least** 4.1% of people put to death in the United States are innocent of the crime they have been sentenced for (Pilkington, 2014), meaning that more than 200 people have more than likely been wrongly put to death – a rate much higher than the .027% quoted by Antonin Scalia in 2007. And yet, a 2018 Pew study shows that in our current political and cultural climate, support for the death penalty is rising for the first time in years; up to 54% from 49% in 2016, a four-decade low (Oliphant, 2018). But the fact of the matter is, the question at hand in this study is not whether or not the death penalty is right, or even the guilt or innocence of those who have found themselves on death row. Those questions are beyond the scope of a simple text mining study. The question is much more general: what can last statements tell us about the crime and the convicted?

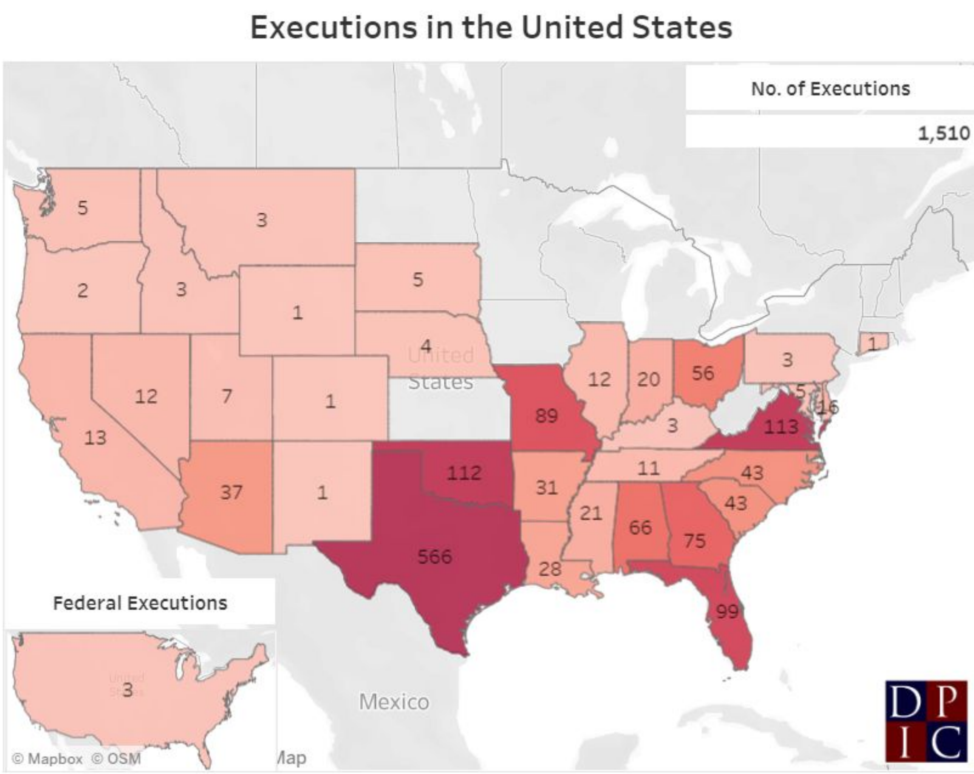

Figure 1 – Executions by state since January 17, 1977

By looking more closely at the final statements of inmates in Texas, the only state to release statements until the recent change in policy which barred the practice (Weber, 2019), trends begin to emerge that can help policy makers, advocates, and on-lookers to get a better understanding not just of the practice, but of its outcome as well as the effects of the criminal justice system. For example, one recent study that examined last statements in Texas from 2002-2017 found that, "executed prisoners in Texas became fewer and older, spent longer on death row and had committed more serious offences. Themes of love and spirituality were constants, but requests for forgiveness declined" (Foley & Kelly, 2018). Another found that, "half of all verbalized last statements contained a religious expression in the early time period (1982–95), but the proportion substantially increased (even as it ebbed and flowed) across the four remaining periods, settling at 66 percent in recent years (2011–16)" (Smith, 2018). And yet another found that, "the fact that a full one-third of the sample spontaneously apologized suggests that apology was important to these offenders. In addition, these apologies were accompanied by indicators of true repentance, such as taking responsibility for their actions, asking for forgiveness, showing empathy, and being sincere" and pointed to how these conclusions, like the ones that might be found in this study, could lead to real world change by saying, "This suggests that more resources should be devoted to finding ways to enable offenders, including those on death row, to apologize directly to their victims or their victims' families (if desired by both parties), either through victim–offender mediation or similar programs. More generally, this research adds to the growing literature indicating that apology and forgiveness are important in the criminal justice system and can have practical advantages for both victims and offenders" (Eaton & Theuer, 2009).

Table 4 Logistic regression predicting apology from last statement variables

Variable	β	Wald test	p	Odds ratio
Remorse	.924	1.316	.251	2.52
Ask for forgiveness	1.285	3.861	.049	3.61
Responsibility	1.909	5.332	.021	6.75
Empathy	1.557	7.874	.005	4.74
Sincerity	3.397	24.969	<.001	29.89

Figure 1 – Logistics regression predicting apology from last statement variables (Eaton & Theuer, 2009)

These outcomes all focus on the offender and situation created by the crime, but other studies suggest that final statements can be used to learn more about victim recovery; showing evidence that rates of remorse and repentance rise with the presence of the victim or victim's family at the execution itself, which may lead to better victim closure (Rice, 2009). The study goes on to "recommend that future research employ interviews with survivors to understand subtle connections between inmate death narratives and survivor transformation" (Rice, 2009).

What is clear is that final statements, alongside victim variables, offender variables, and execution variables, among other things, are rife with opportunities to learn more about the criminal justice system in the United States, the sociology and psychology of inmates, and the true impact on the lives of everyone involved.

Analysis and Models

About the Data

Texas

1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986
0	0	0	0	0	0	1	0	3	6	10
1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997
6	3	4	4	5	12	17	14	19	3	37
1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
20	35	40	17	33	24	23	19	24	26	18
2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	
24	17	13	15	16	10	13	7	7	13	

Figure 3 – Overview of executions in Texas from 1976-2018 (Death Penalty Information Center, 2019)

Gathering the Data

The dataset originally came from kaggle. After discovering many gaps in the kaggle dataset, the researchers decided to compile their own dataset. First, they scraped the **tdcj.texas.gov** for generic data such as last name, first name, TDCJ number, age, date (of execution) race and county (of origin). This initial scrape included links to two supplemental pages, the first containing much of the data that would come to be the y-variables, the second containing the text of the inmate's last statement.

The first link contained a deeper level of data about both the inmate and the nature of the crime committed. The inmate data included data received, age (when received), education level, date of offense, age (at time of offense), county, race, gender, hair color, height (in feet and inches), weight (in pounds), eye color, native county, native state, prior occupation and prior prison record. The crime data included a summary, co-defendants and race and gender of victim, if known. The summary paragraph often included information about weapon used, type of crime committed and additional information about the victims.

Only 186 inmates had the aforementioned information digitized. Three hundred and eighty inmates instead had links to scanned images of printed documents. See examples below.

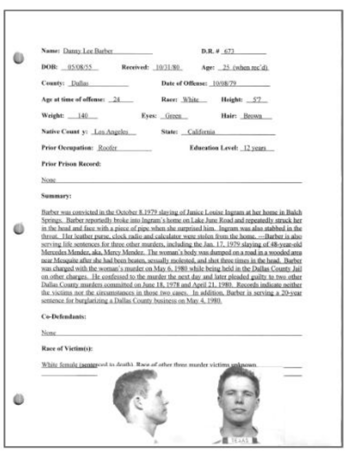

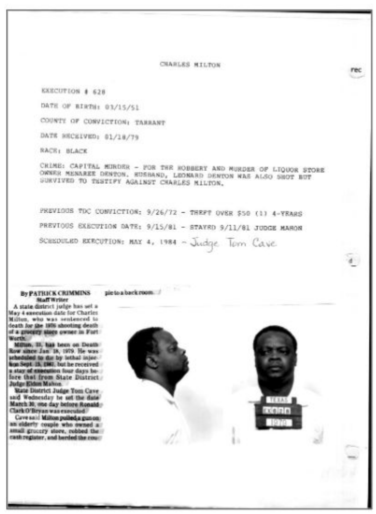

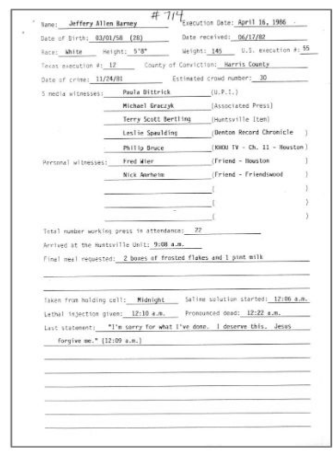

The researchers had to employ Optical Character Recognition (OCR) to these documents to extract the text to get the data to match the other 186 inmates. The result of the OCR was a corpus of non-standardized text documents with various incongruities that posed additional obstacles for the researchers.

```
Name: James Edward smith

DOB: 10 / 19 7 52


Receiver

. Harris

Count;


Age at time of offense: 30


Weight: 158 Eyes: PEOW

Native County: Jefferson

Prior Occupation: retail merchant

D.R.#

uw

84
/ Age: (when rec'c

Date of Offense: > of) oy




5-10


Race Height:



black
Hair: °°

State: Kentucky

Education level: '* Y"ors


Prior prison record:
None
```

```
CHARLES MILTON

EXECUTION # 628

DATE OF BIRTH: 03/15/51
COUNTY OF CONVICTION: TARRANT
DATE RECEIVED: 01/18/79

RACE: BLACK

CRIME: CAPITAL MURDER - FOR THE ROBBERY AND MURDER OF LIQUOR STORE
OWNER MENAREE DENTON. HUSBAND, LEONARD DENTON WAS ALSO SHOT BUT
SURVIVED TO TESTIFY AGAINST CHARLES MILTON

PREVIOUS TDC CONVICTION: 9/26/72 - THEFT OVER $50 (1) 4-YEARS
PREVIOUS EXECUTION DATE: 9/15/81 - STAYED 9/11/81 JUDGE MAHON

SCHEDULED EXECUTION: MAY 4, 1984 ~ Judge Tom Cave

ByPATRICKCRIMMINS —_pletoa backroom.
Staff Writer

A state district judge has set a
May 4 execution date for Charles
Milton, who was sentenced to
death for the 1976 shooting death
of a grocery store owner in Fort
Worth.

Milton, 33, has been on Death
Row sittce Jan. 18, 1979. He was

hedulled to die by lethal injec-

Sept. 15, 1981, but he received
a stay of execution four days be-
fore that from State District
Judge Eldon Mahon.

State District Judge Tom Cave
said Wednesday he set the date '
'March 30, one day before Ronald
Clark O'Bryan was executed.

Cave said Milton pulleda gunon:
an elderly couple who owned a
small grocery store, robbed the
'cash register, and herded the cou-
```

```
Name: John L. Wheat D.R. #99922

DOB: 05/22/44 Received: _03/11/97 Age: 52. (when received)
County:__ Tarrant Date of Offense:_07/30/95

Age at time of offense: _51 Race: white Height: 5-9 _
Weight: 150 Eyes:_blue Mair: _gray

Native County: _Erath State: _Texas

Prior Occupation: welder/mechanic _-_-_~—«-Education Level: 1] years

Prior Prison Record:



'Co-Defendants:

ON is



Race of Victim(s):
```

```
Prior prison record:
None

Summary: Convicted in the February 1992 slaying of his wife and two
children inside the family's home in Progreso. Killed were Leticia Ramos,
age 42, Abigail Ramos, 7, and Jonathan Ramos, 3. All were beaten with a
blunt object and died of skull fractures. Their bodies were found more than
a month later buried beneath the bathroom floor of their home after Mrs.

Ramos' sister reported them missing. Abigail's hands had been bound by tape


and _her mouth gagged prior to her death. Robert Ramos first told relatives
that his family had died in an automobile accident, but couldn't say where

He then claimed to have found his family dead after returning to the home
from a day of job seeking. He later confessed, saying he fled to Arkansas
Co-Defendants: after the murders and discarded the murder weapon, which was
never found. Ramos married another woman three days after
killing his family and moved her into the home where they were
buried.
No co-defendants

@ Race of victim(s): two Hispanic females, one Hispanic male
```

Regex was used to clean this corpus and extract as much usable data as possible. The wide variety of form formats and differing wording made this an excellent use-case for pythons "try" and "except." The cleaning, in english, looked a lot like this -- "try to find the word 'education'" -- "if an exception, return 'no data'". The researchers would then examine the documents that returned "no data" to find that either the OCR had incorrectly scanned "education" or that the form said "schooling" instead of "education." After many iterations of this across all the different data points, the researchers decided that a manual overview was needed.

In summation, there were challenges with both the OCR and the document formats, which lead to a considerable amount of manual work, on top of the automated work, for the researchers. The result of this manual work was a considerably cleaner, more standardized dataset. This dataset is now, to the best of the researcher's knowledge, the most complete dataset for inmates and last statements.

Cleaning

The data set contains 566 rows and 24 columns. Each row represents an executed offender.

	execution	last_name	first_name	age_received	education_level	age_crime	occupation	prior_record	num_of_vic	main_crime
0	566	Hall	Justen	23	9	21	laborer	yes	1	murder
1	565	Sparks	Robert	34	8	33	machine operator	yes	3	murder
2	564	Soliz	Mark	30	8	28	cabinet maker	yes	1	murder, robbery
3	563	Crutsinger	Billy	49	11	48	laborer	yes	2	murder
4	562	Swearingen	Larry	29	11	27	laborer	yes	1	murder, kidnapping

Figure 4 – First 5 rows and 10 columns of the df

vic_kid	vic_male	vic_female	vic_police	inmate_number	age	date_executed	race	county	last_statement
0	0	1	no	999497	38	11/6/2019	White	El Paso	Yeah, I want to address the Roundtree family ...
2	2	1	no	999542	45	9/25/2019	Black	Dallas	Umm, Pamela can you hear me Stephanie, Hardy,...
0	0	1	no	999571	37	9/10/2019	Hispanic	Johnson	It's 6:09 on September 10th, Kayla and David,...
0	0	2	no	999459	64	9/4/2019	White	Tarrant	Hi ladies I wanted to tell ya'll how much I I...
0	0	1	no	999361	48	8/21/2019	White	Montgomery	Lord forgive them. They don't know what they ...

Figure 4 – First 5 rows and last 10 columns of the df

Currently, the data frame is not discretized and there are columns that will not serve in the analysis. Execution, inmate number, and date_executed are unique identifiers and therefore were removed from the data set.

```
execution           int64
last_name          object
first_name         object
age_received       object
education_level    object
age_crime          object
occupation         object
prior_record       object
num_of_vic         object
main_crime         object
type_of_crime      object
weapon             object
co_defendants      object
race_vic           object
vic_kid            object
vic_male           object
vic_female         object
vic_police         object
inmate_number       int64
age                 int64
date_executed      object
race               object
county             object
last_statement     object
```

Figure 4 – Data types for each column

The following columns needed to be changed from objects to numeric columns: age_received, age_crime, num_of_vic, vic_kid, vic_male, and vic_female. An issue that arose in this process was the fact that not all of the values in the columns were in fact numeric. There were unknowns. The unknowns were changed from unknown to an empty value. Once the columns were converted to numeric, the entries that originally had "unknown" now displayed nan. To rectify the missing values, the average for each column was found and inserted in the entries with missing values.

```
The number of missing values in age_received is 2
Now the number of missing values in age_received is 0
The number of missing values in age_crime is 2
Now the number of missing values in age_crime is 0
The number of missing values in num_of_vic is 1
Now the number of missing values in num_of_vic is 0
The number of missing values in vic_kid is 1
Now the number of missing values in vic_kid is 0
The number of missing values in vic_male is 2
Now the number of missing values in vic_male is 0
The number of missing values in vic_female is 2
Now the number of missing values in vic_female is 0
The number of missing values in co_defendants is 1
Now the number of missing values in co_defendants is 0
```

Figure 4 – Number of missing values for each column

The following columns were changed from object to category (factor): occupation, main_crime, type_of_Crime, weapon, race, race_vic, county, last_name, first_name, prior_record, and vic_police. A new column, time_on_death_row was aggregated by taking the age the prisoner was executed and subtracting the age the prisoner received the death row sentence.

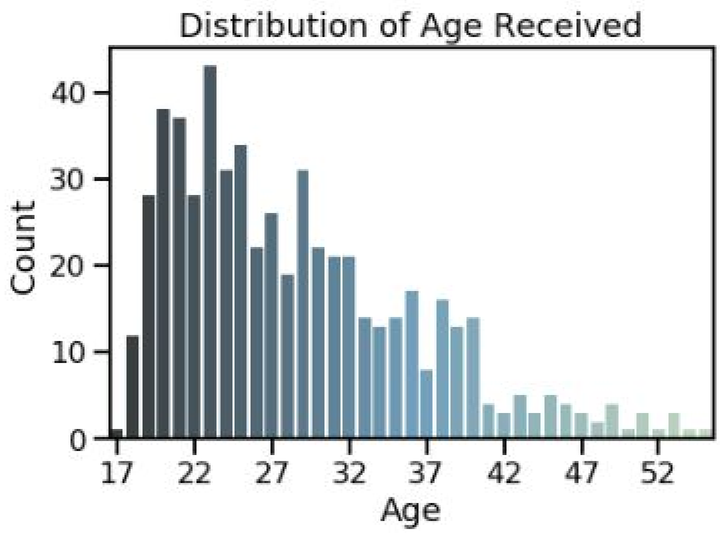

Figure 4 – Distribution of Age Received Prior to Discretization

There is an issue with the way age_received is currently stored. Every column in the data frame, with the exception of last statement will serve as a potential label. Currently, the way age_received is broken down there is not sufficient information to run a prediction model. Therefore, the age_received column was discretized into 3 different categories: teens, twenties, and thirty+. The breakdown of labels is self-explanatory.

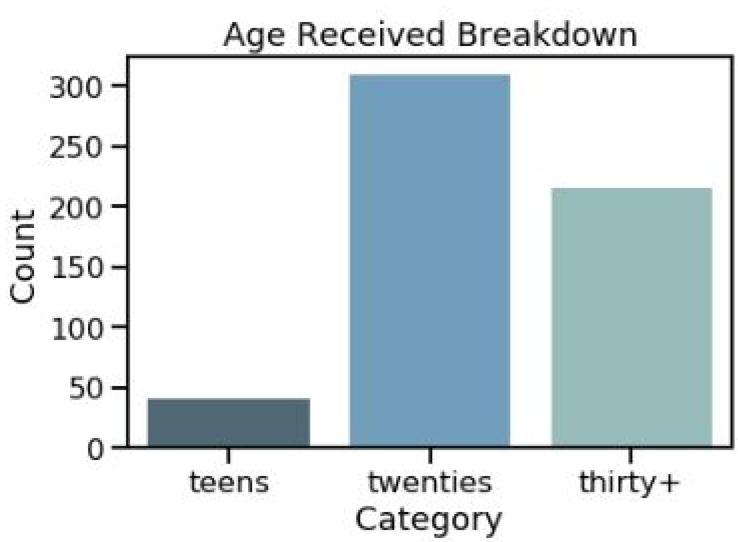

	count	category
0	309	twenties
1	216	thirty+
2	41	teens

Figure 4 – Distribution of Age Received

After discretization, there were 41 inmates who received the death penalty in their teens, 309 in their twenties, and 216 who were 30+. For the prediction, it might be prudent to combine the teens with those in their twenties, as the teens sample is rather small.

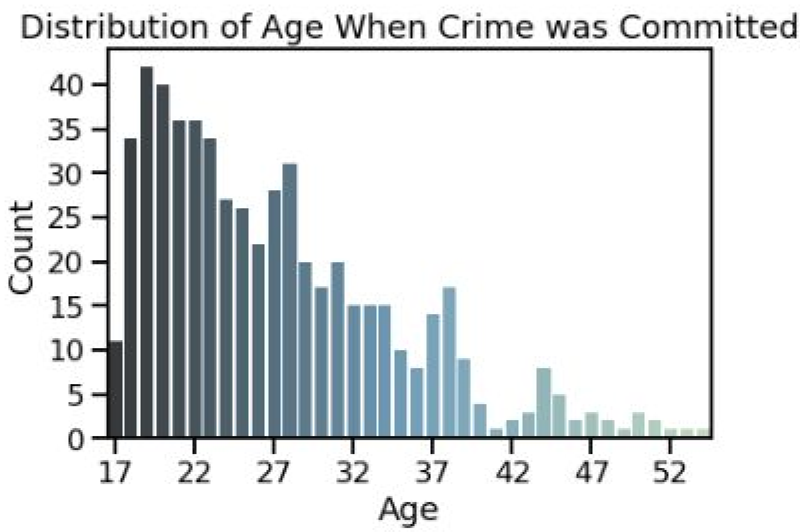

Figure 4 – Distribution of Age when Crime was Committed Prior to Discretization

There is an issue with the way age_crime is currently stored. Currently, the way age_crime is broken down there is not sufficient information to run a prediction model. Therefore, the age_crime column was discretized into 3 different categories: teens, twenties, and thirty+. The breakdown of labels is self-explanatory.

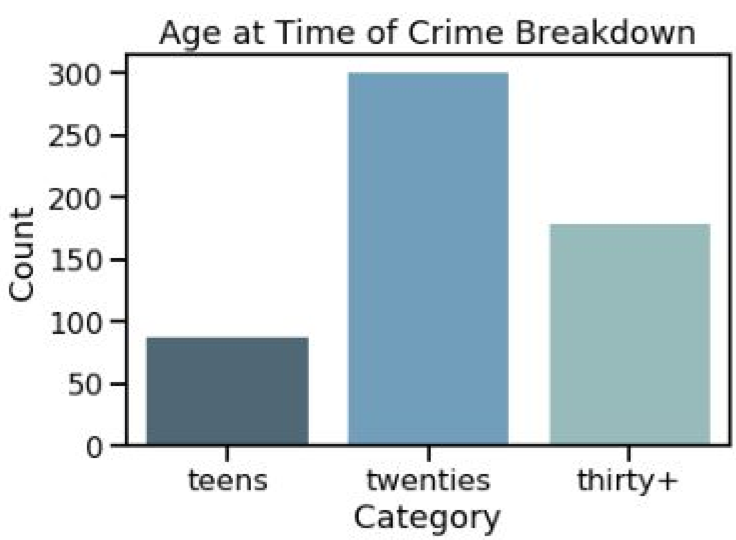

	count	category
0	300	twenties
1	179	thirty+
2	87	teens

Figure 4 – Distribution of Age at Time of Crime

After discretization, there were 87 inmates who received the committed the crime that led to the death penalty in their teens, 300 in their twenties, and 179 who were 30+. For the prediction, it might be prudent to combine the teens with those in their twenties, as the teens sample is rather small.

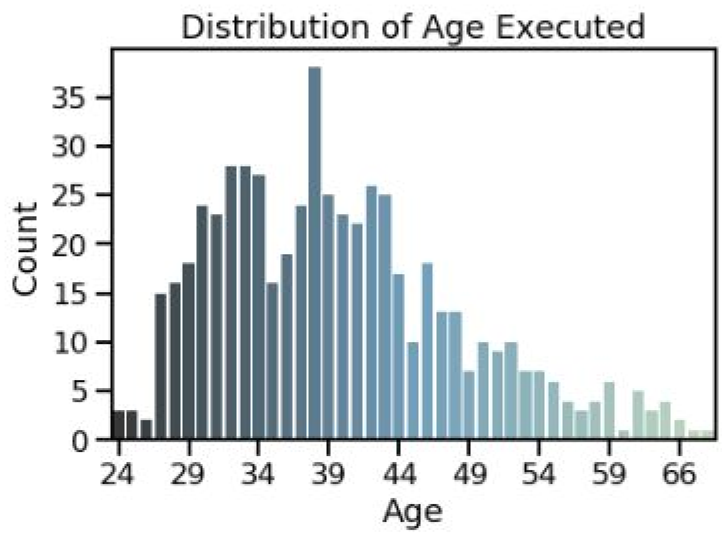

Figure 4 – Distribution of Age Received Prior to Discretization

There is an issue with the way age is currently stored. Every column in the data frame, with the exception of last statement will serve as a potential label. Currently, the way age is broken down there is not sufficient information to run a prediction model. Therefore, the age column was discretized into 3 different categories: 18-34, 35-45, and 45+. The breakdown of labels is self-explanatory.

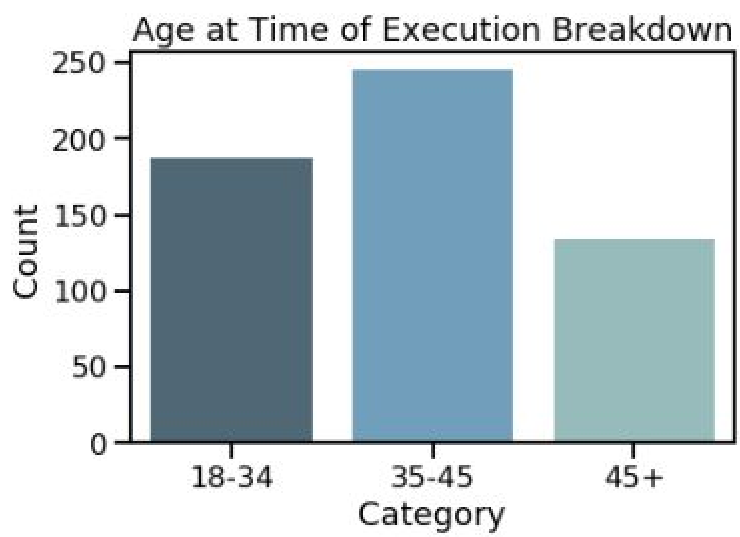

	count	category
0	245	35-45
1	187	18-34
2	134	45+

Figure 4 – Distribution of Age Executed

After discretization, there were 187 inmates who were executed between 18 to 34 years old, 245 between 35-45, and 134 who were 45 and above.

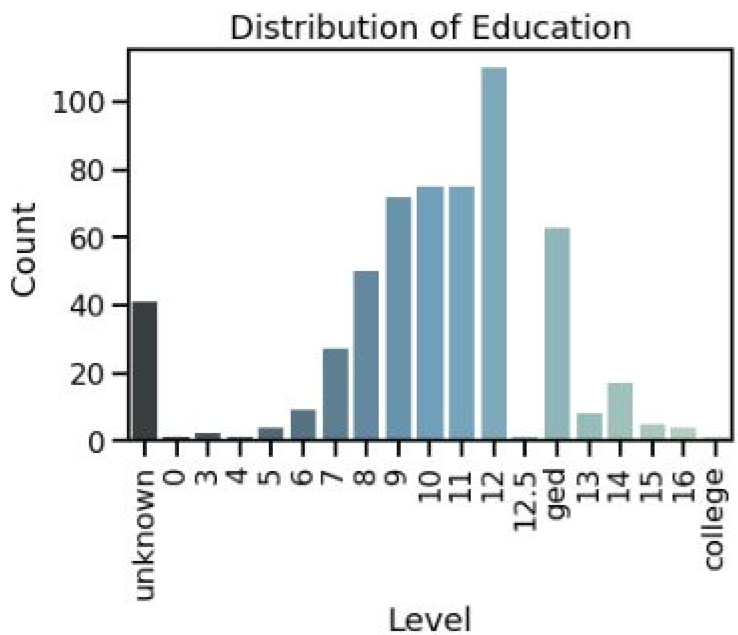

Figure 4 – Distribution of Education Prior to Discretization

There is an issue with the way education level is currently stored. Currently, the way education level is broken down there is not sufficient information to run a prediction model. Therefore, the education level was discretized into 5 different categories: unknown, no_highschool, some_highschool, highschool, and college. No highschool is comprised of people who did not reach 9th grade. Some highschool is comprised of prisoners who attended highschool but did not graduate. Highschool is comprised of prisoners who either graduated or attained their ged. College is comprised of people who had education after highschool.

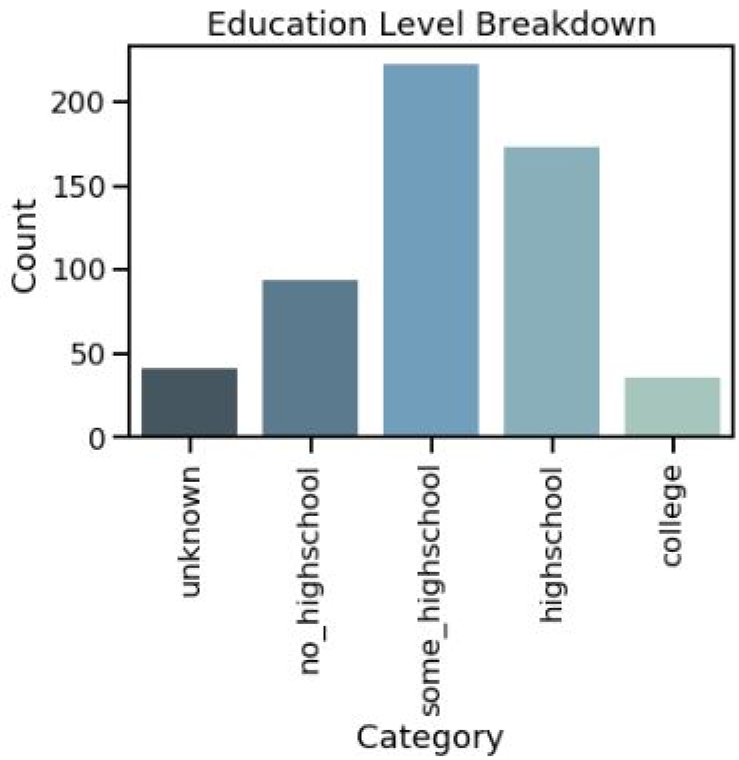

count	category
222	some_highschool
173	highschool
94	no_highschool
41	unknown
36	college

Figure 4 – Distribution of Education

After discretization, there were 94 inmates who had no highschool, 222 with some highschool, 173 who graduated, and only 36 who attended college. For the prediction, it might be prudent to remove the unknown prisoners and to combine the highschool and college inmates together, as the college sample is rather small.

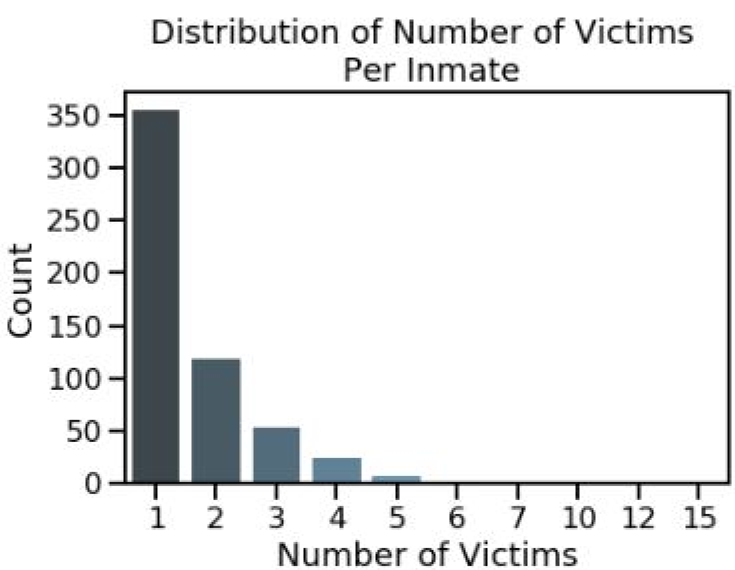

Figure 4 – Distribution of Number of Victims Per Prisoner Prior to Discretization

Most prisoners are on death row for killing one person, however there are 212 prisoners on death row that have killed multiple people. There is an issue with the way num_of_vic is currently stored. Currently, the way num_of_vic is broken down there is not sufficient information to run a prediction model. Therefore, the num_of_vic was discretized into 2 different categories: one, and two+. The breakdown of labels is self-explanatory.

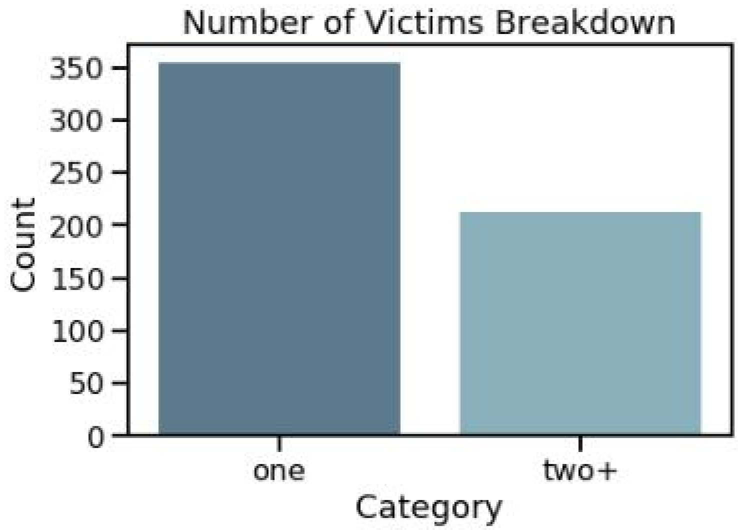

	count	category
0	354	one
1	212	two+

Figure 4 – Distribution of Number of Victims Per Prisoner

After discretization, there were 354 inmates who one victim and 212 prisoners with two or more victims. A victim, in this case, is described as a person that was murdered by the prisoner.

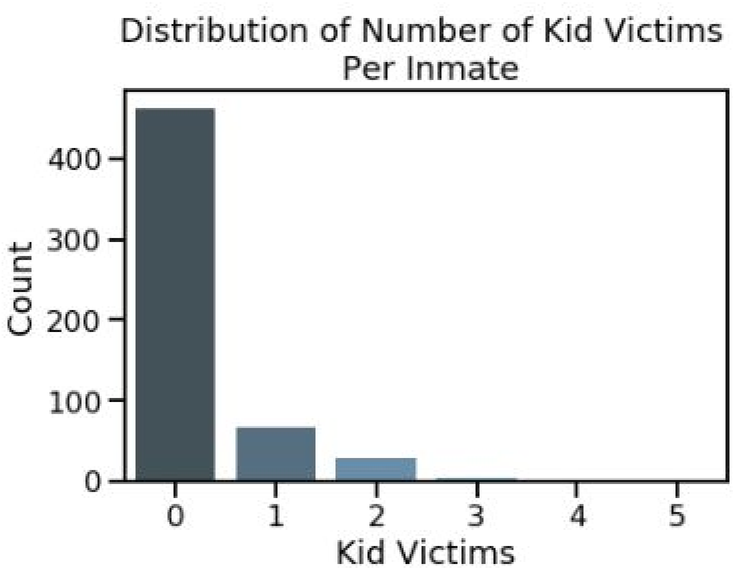

Figure 4 – Distribution of Number of Kid Victims Per Prisoner Prior to Discretization

There were a total of 154 children who were harmed or worse by people on death row. Most prisoner who have been executed did not murder or harm a child. There is an issue with the way vic_kid is currently stored. Currently, the way vic_kid is broken down there is not sufficient information to run a prediction model. Therefore, the vic_kid was discretized into 2 different categories: yes and no. The breakdown of labels is self-explanatory.

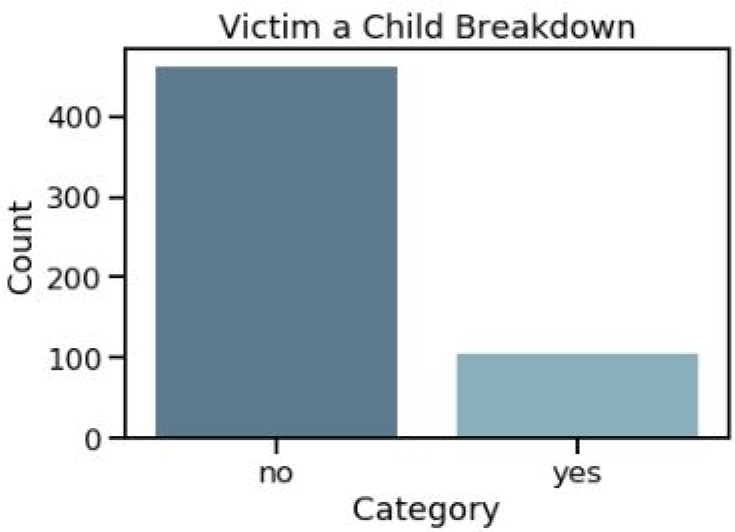

	count	category
0	461	no
1	105	yes

Figure 4 – Distribution of Kid Victims

After discretization, there were 461 inmates who did not have a child victim and 105 prisoners with a child victim. A victim, in this case, is described as a person that was murdered, violated, and/or injured by the prisoner.

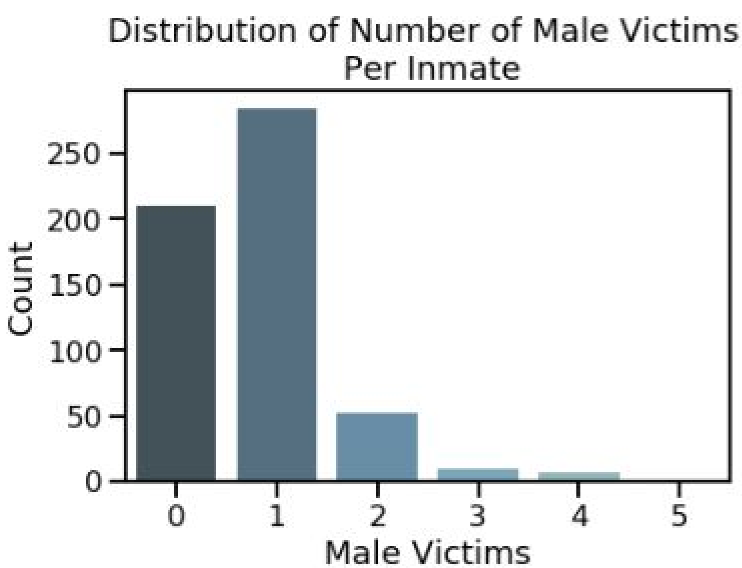

Figure 4 – Distribution of Number of Male Victims Per Prisoner Prior to Discretization

There were a total of 458 males who were harmed or worse by people on death row. There is an issue with the way vic_male is currently stored. Currently, the way vic_male is broken down there is not sufficient information to run a prediction model. Therefore, the vic_male was discretized into 2 different categories: yes and no. The breakdown of labels is self-explanatory.

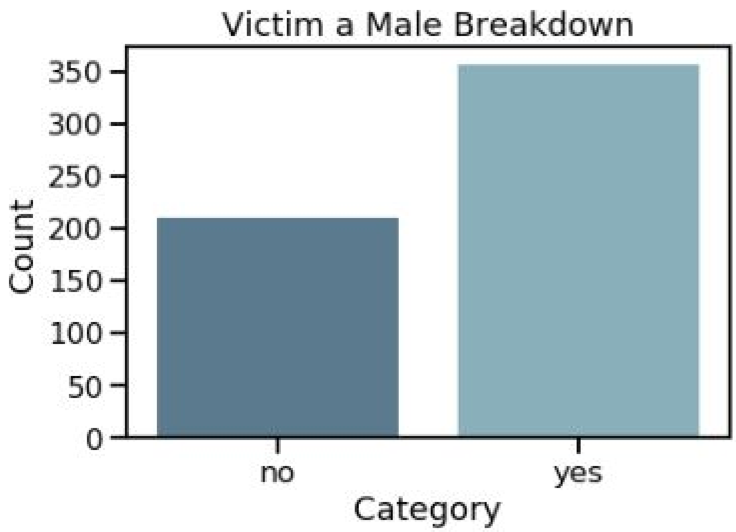

	count	category
0	356	yes
1	210	no

Figure 4 – Male Victim

After discretization, there were 461 inmates who did not have a child victim and 105 prisoners with a child victim. A victim, in this case, is described as a person that was murdered, violated, and/or injured by the prisoner.

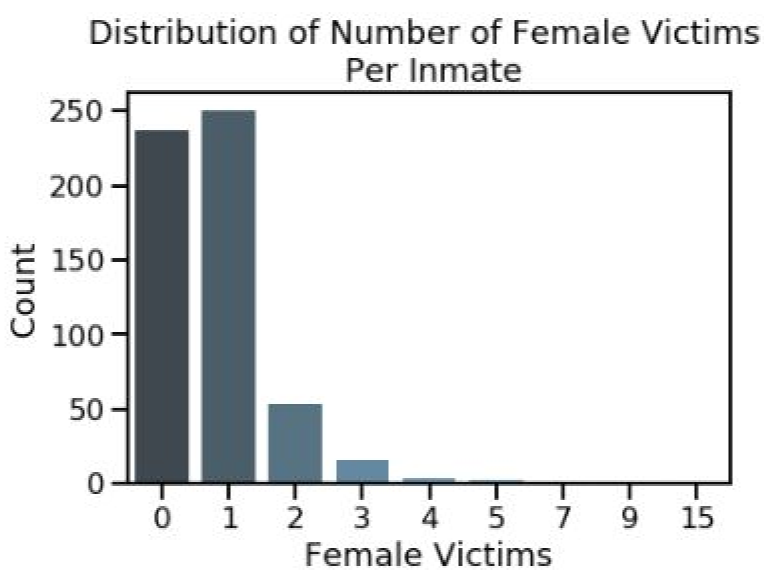

Figure 4 – Distribution of Number of Female Victims Per Prisoner Prior to Discretization

There were a total of 466 females who were harmed or worse by people on death row. There is an issue with the way vic_female is currently stored. Currently, the way vic_female is broken down there is not sufficient information to run a prediction model. Therefore, the vic_female was discretized into 2 different categories: yes and no. The breakdown of labels is self-explanatory.

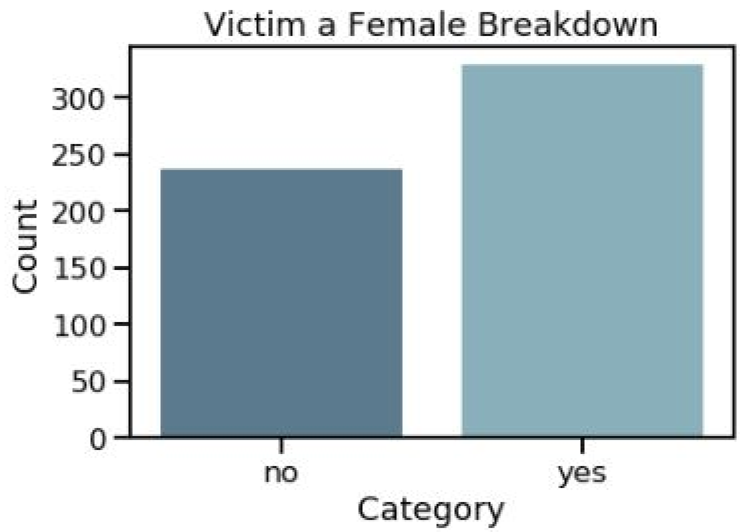

	count	category
0	329	yes
1	237	no

Figure 4 – Female Victim

After discretization, there were 237 inmates who did not have a female victim and 329 prisoners with a female victim. A victim, in this case, is described as a person that was murdered, violated, and/or injured by the prisoner.

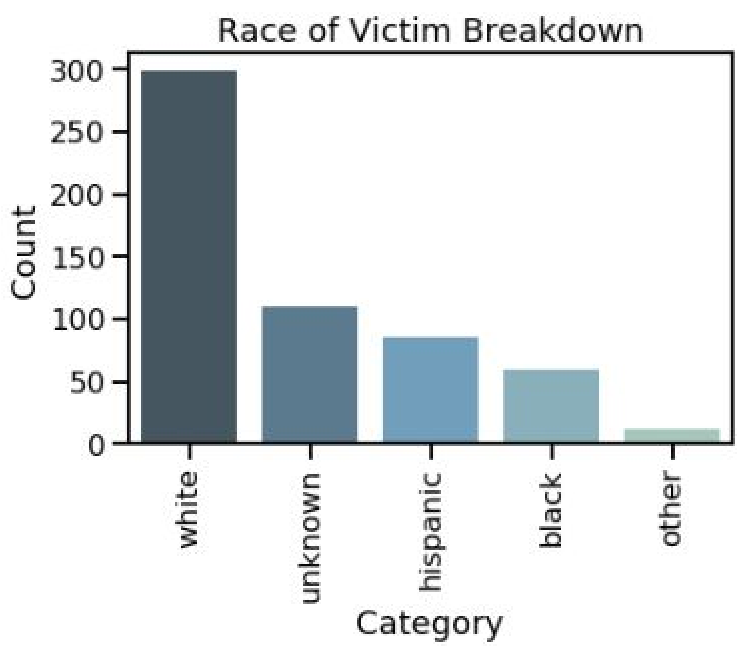

	count	category
0	299	white
1	110	unknown
2	86	hispanic
3	59	black
4	12	other

Figure 4 – Race of Victim

299 of the victims were white, 110 unknown, 86 hispanic, 59 black, and 12 identified as other. A victim, in this case, is described as a person that was murdered. For the prediction, it might be prudent to remove the unknown and other victims.

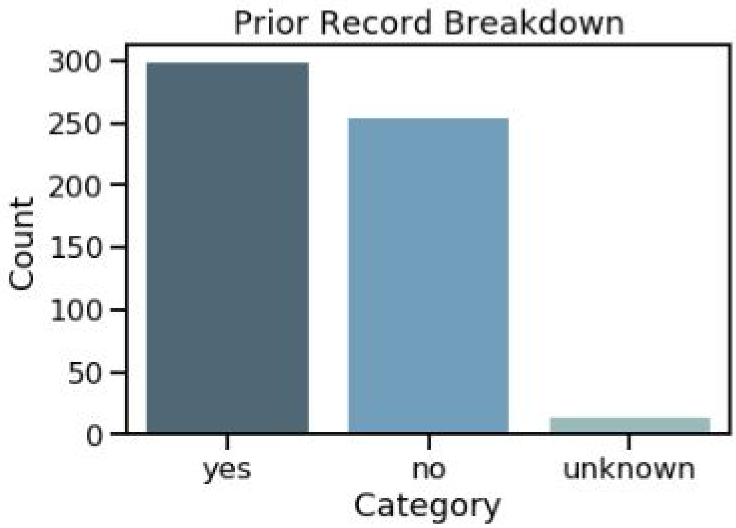

	count	category
0	298	yes
1	254	no
2	14	unknown

298 of the executed inmates did have a prior record, and 254 did not. For the prediction, it might be prudent to remove the unknown entries.

Results

Data Plotting

The 566 rows of data all had a column with the 113 counties in which the crime was committed. This allowed for the use of Plotly's Cloropleth maps which uses the county boundaries along with heatmaps to better represent the data, in this case the locations of the crimes. The county columns were all given as names of the counties, however, the data that is needed to create the plot is the Federal Information Processing Systems (FIPS) county code. The FIPS is a 5-digit code that is assigned to each county based on the state. The first two digits represent the state, and the last 3 digits represents one of the 254 counties. For example, the county code for Texas, Dallas County is 48113, and Texas, Harris County is 48201. The 48 in both of those examples represents the state of Texas, and the 113 and 201 respectively identify Dallas and Harris counties. A new data frame that contained all of the Texas FIPS codes and county names was imported as a .csv, and merged with the original data frame based on the county names. This built a new FIPS column which was used for the following figure.

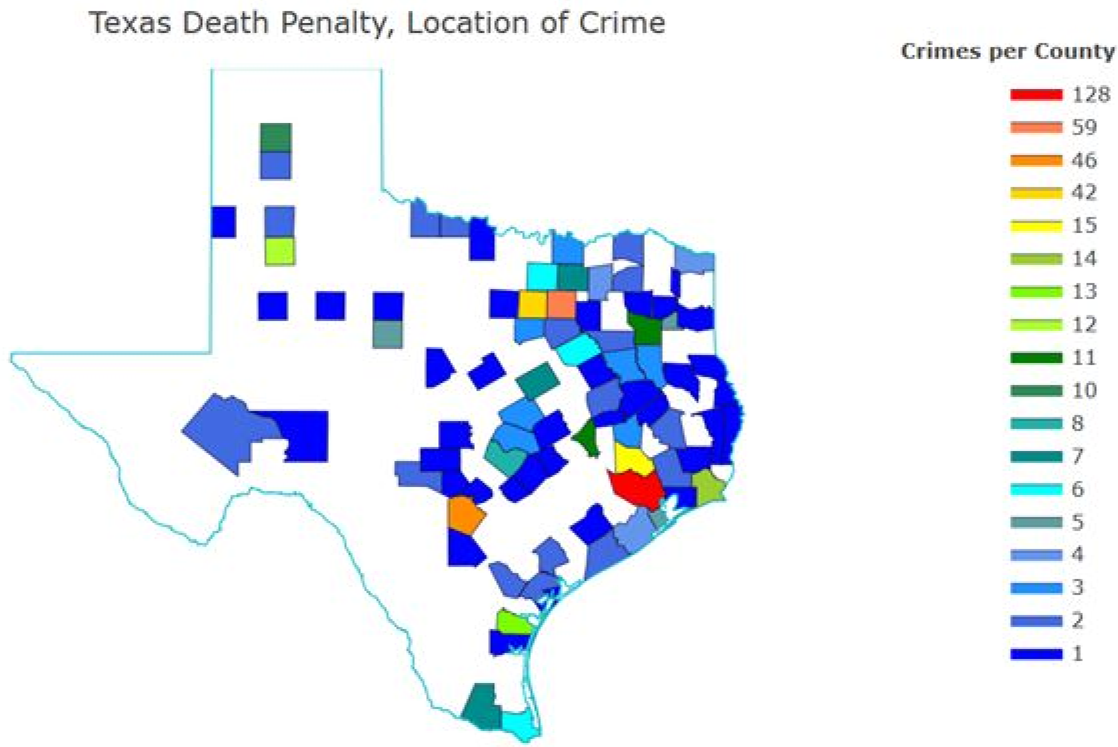

Figure 17 – Texas county plots

This image shows the counts of the criminal acts that led to death row for each of the counties. The highest county which is shown in bright red is Harris County which is Houston's county with 128 crimes that warranted death row. The next highest is Dallas County with 59, with Bexar County which is San Antonio with 46, and closely followed by Tarrant County which is Fort Worth with 42.

Topic Modeling

To give a high-level look at the statements that make up this study, and in an effort to better understand the distribution of topics down the road, the following bar chart was made.

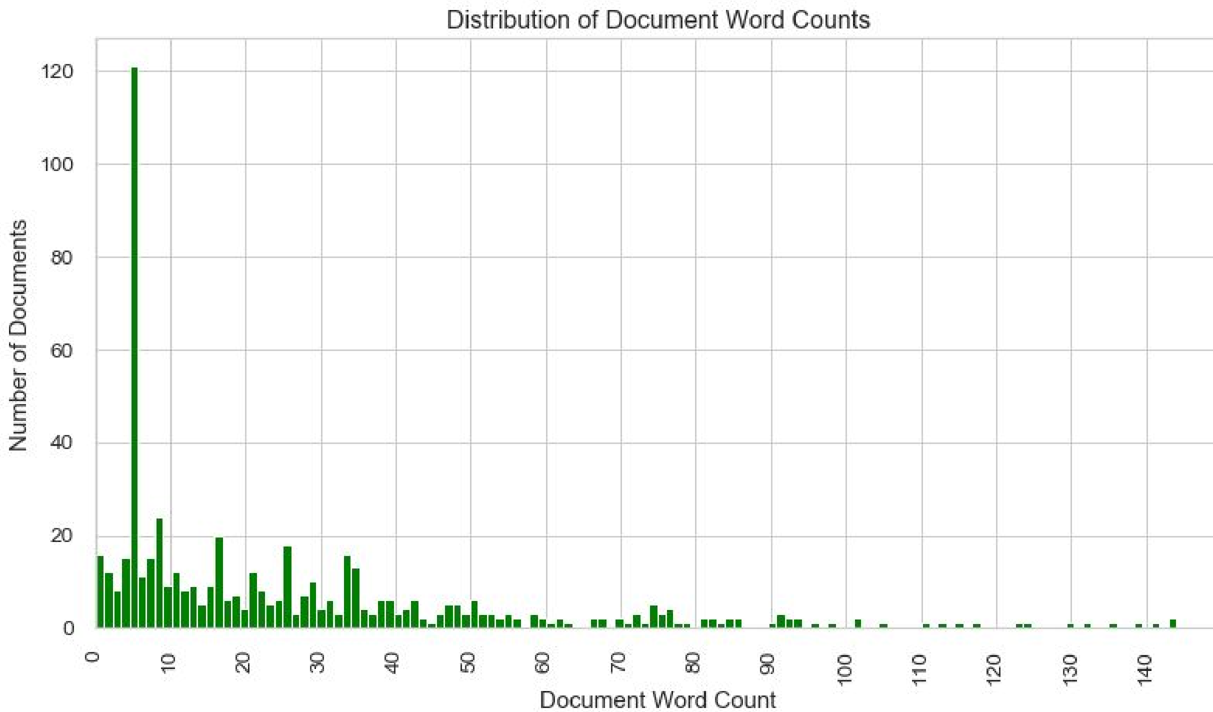

Figure 11 – Distribution of document word counts

Although this graph doesn't give insight to the topics which they cover, it does give a clear indication of why normalizing the data, and the normalization that is built into the LDA model, is so critical. This step will ensure that that the importance of each word within the individual documents, but also within the topics as a whole.

Another exploratory measure would be to look at term frequency of the pre-processed text for early context clues.

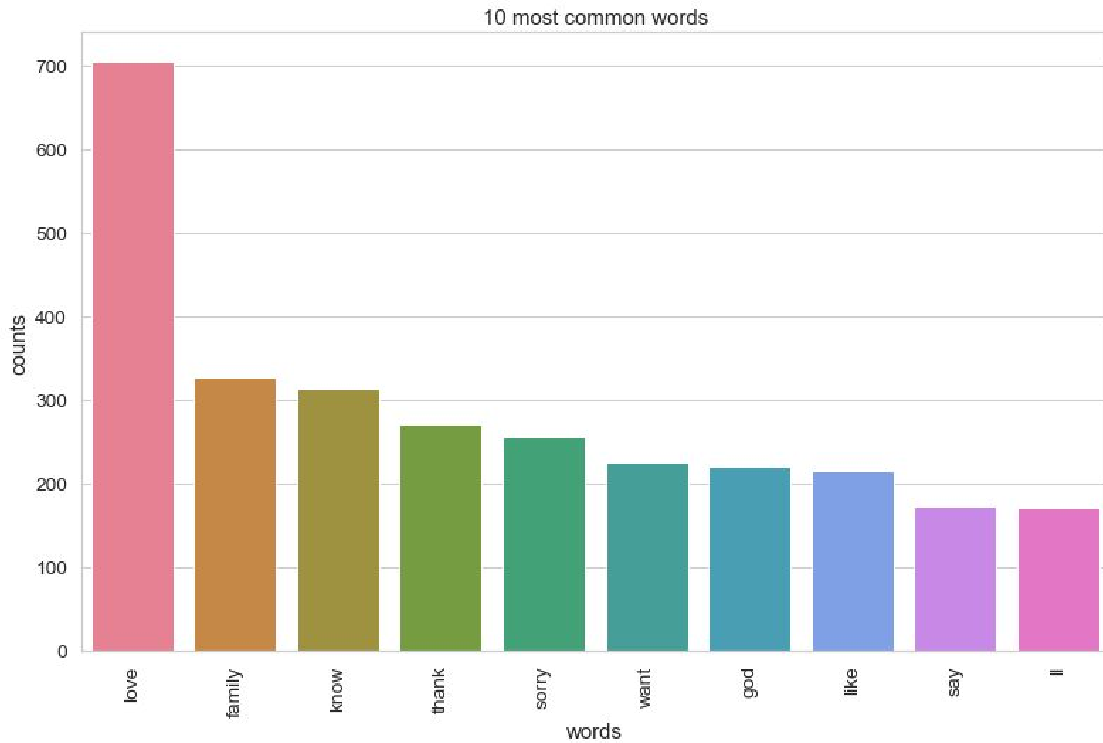

Figure 11 – Top 10 most commons words, by frequency, in the pre-processed corpora

After modeling the book descriptions using the LDA model outlined above, seven topics were created, and the top ten words were printed alongside their weighting.

```
[(0,
  '0.064*"love" + 0.037*"want" + 0.032*"everybody" + 0.018*"take" + '
  '0.017*"keep" + 0.017*"hold" + 0.016*"always" + 0.016*"strong" + '
  '0.014*"tell" + 0.012*"wife"'),
 (1,
  '0.025*"life" + 0.022*"people" + 0.020*"want" + 0.016*"justice" + '
  '0.015*"live" + 0.015*"know" + 0.014*"anyone" + 0.012*"innocent" +
'
  '0.011*"call" + 0.011*"bitterness"'),
 (2,
  '0.030*"forgive" + 0.030*"will" + 0.028*"family" + 0.027*"thank" +
'
  '0.022*"love" + 0.022*"hope" + 0.022*"know" + 0.018*"sorry" +
0.014*"done" + '
  '0.012*"death"'),
 (3,
  '0.102*"statement" + 0.098*"make" + 0.095*"last" + 0.084*"offender"
+ '
  '0.084*"declined" + 0.026*"allah" + 0.017*"father" + 0.012*"fear" +
'
  '0.009*"evil" + 0.009*"power"')]
```

Figure 11 – 4 topics, top 10 words and their corresponding weighting

To begin to better understand how word count and topic overlap, those clustered documents can be mapped alongside one another.

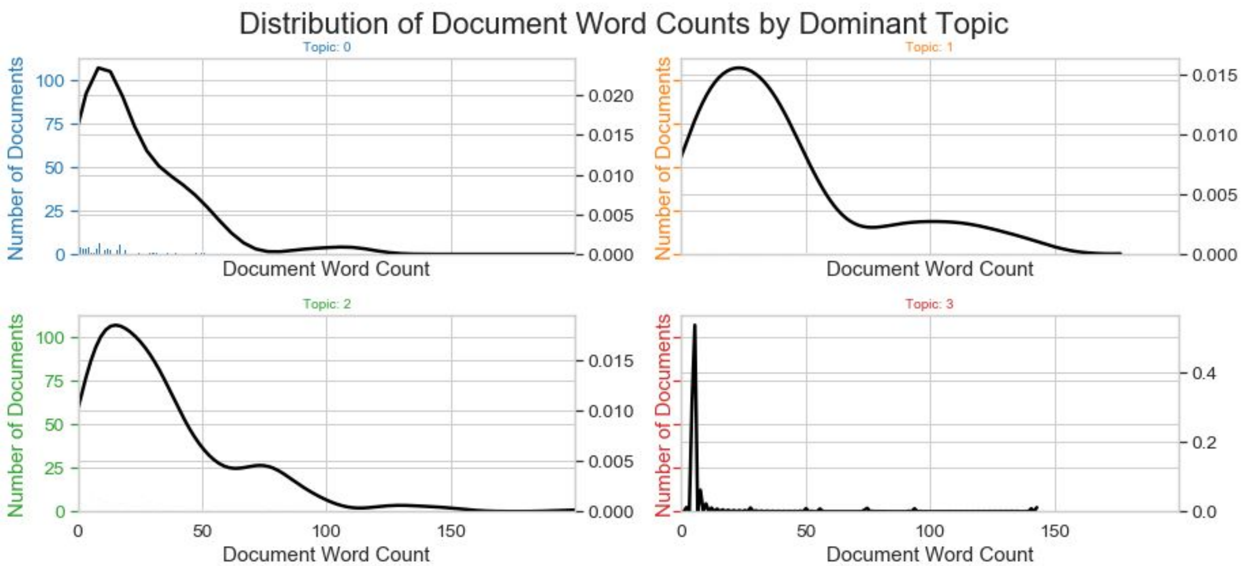

Figure 11 – Distribution of document word count by dominant topic

To get a better understanding of what those dominant topics truly reflect, in order to build a more efficient predicting or recommendation tool by developing true labels, word clouds were created for each dominant topic.

Figure 11 – Word clouds by topic cluster

This small look into the corpus exhibit a clear skew where those who did not wish to make final statements and final statements that include Allah are concerned and may point to a need for a larger dataset. Because these represent such a small portion of the sample, they populate the same topic model, despite not being related necessarily.

Further insight can be gleaned by breaking down the word count and corresponding importance of keywords within each topic.

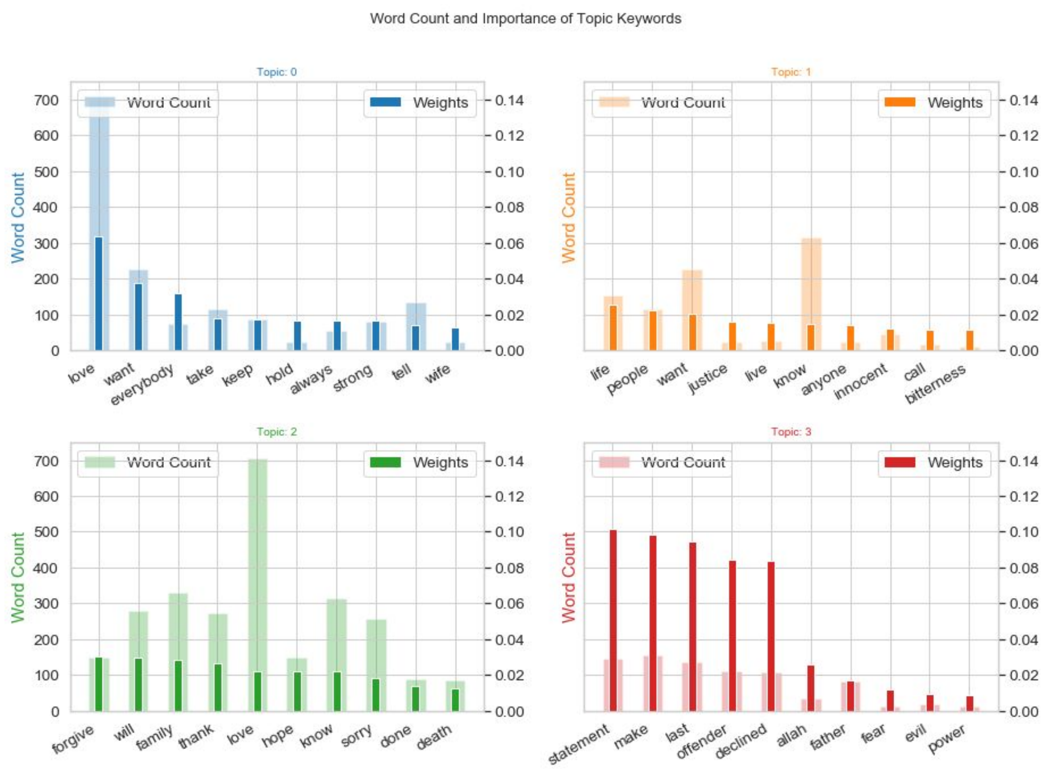

Figure 11 – Word count and importance of keywords by topic

Another way of getting a better understanding of how topics are distributed and how each document might fall within those clusters, the processed vocabulary of individual sentences can be color coded by topic.

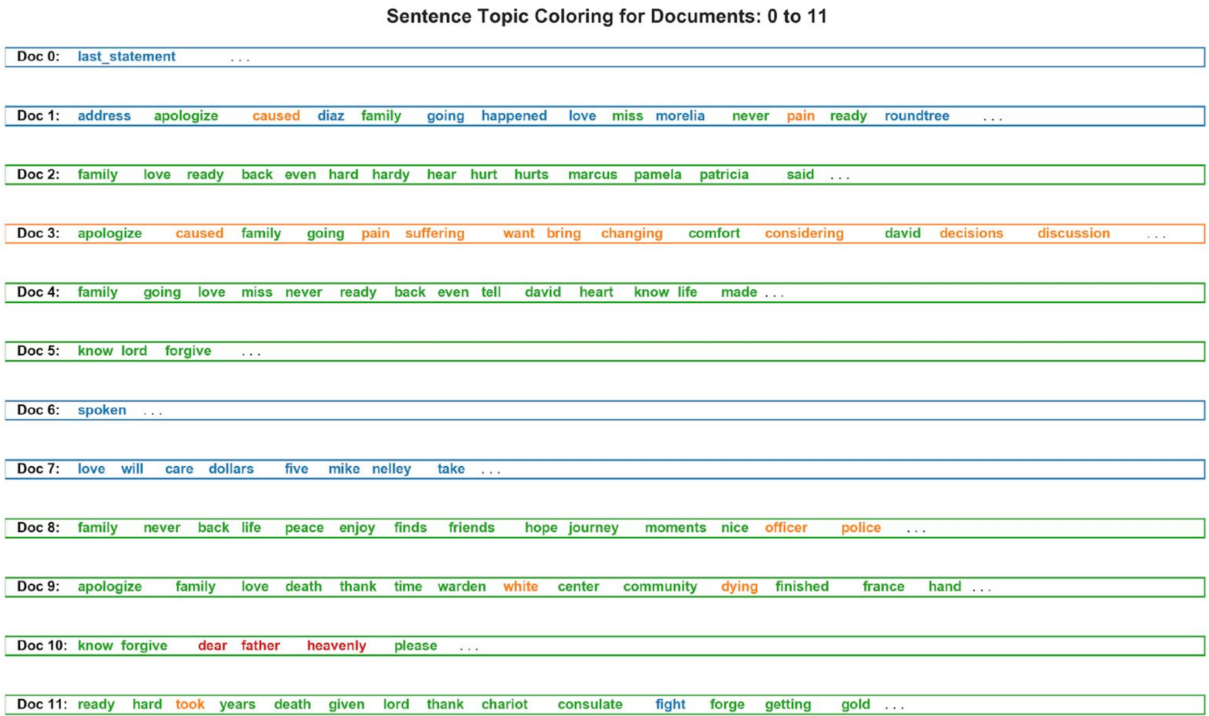

Figure 11 – A sampling of sentences color-coded by topic

The most important visualizations, however, are the clusters themselves to help not only give a high-level view of how the popular each topic is (by how many documents are in each cluster) but also to ensure the correct number of topics has been selected for the model. Although the t-SNE cluster graph for this model shows some minor bleeding between clusters, for the most part they are clearly defined and separated from the other clusters, implying good measure selection.

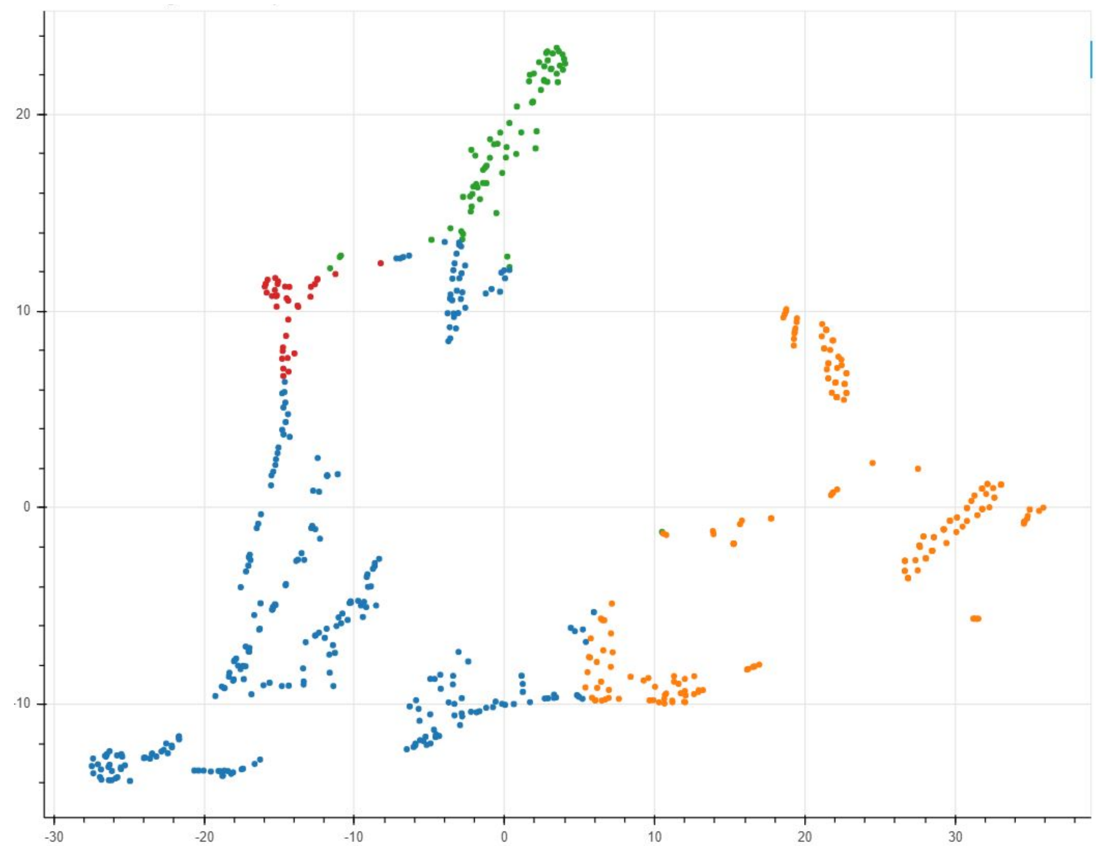

Figure 11 – Topic clusters

This is further confirmed by a bokeh visualization of the clusters which shows overlap between all of the topics but via the Intertopic Distance Map but shows clear topic delineations. It also illustrates the top 30 salient terms for each topic, alongside their overall term frequency, and estimated term frequency within each document. In many of the topics, what the salient terms graph shows is that the terms are highly associated to their topic; in these cases, the bars for each term within the topic are nearly or totally the same as the bars for each term's total use, suggesting the topics are well defined. The exception being the word "love," which appears in three of four topics, although even then it falls heavily into Topic 1, as compared to 2 and 3.

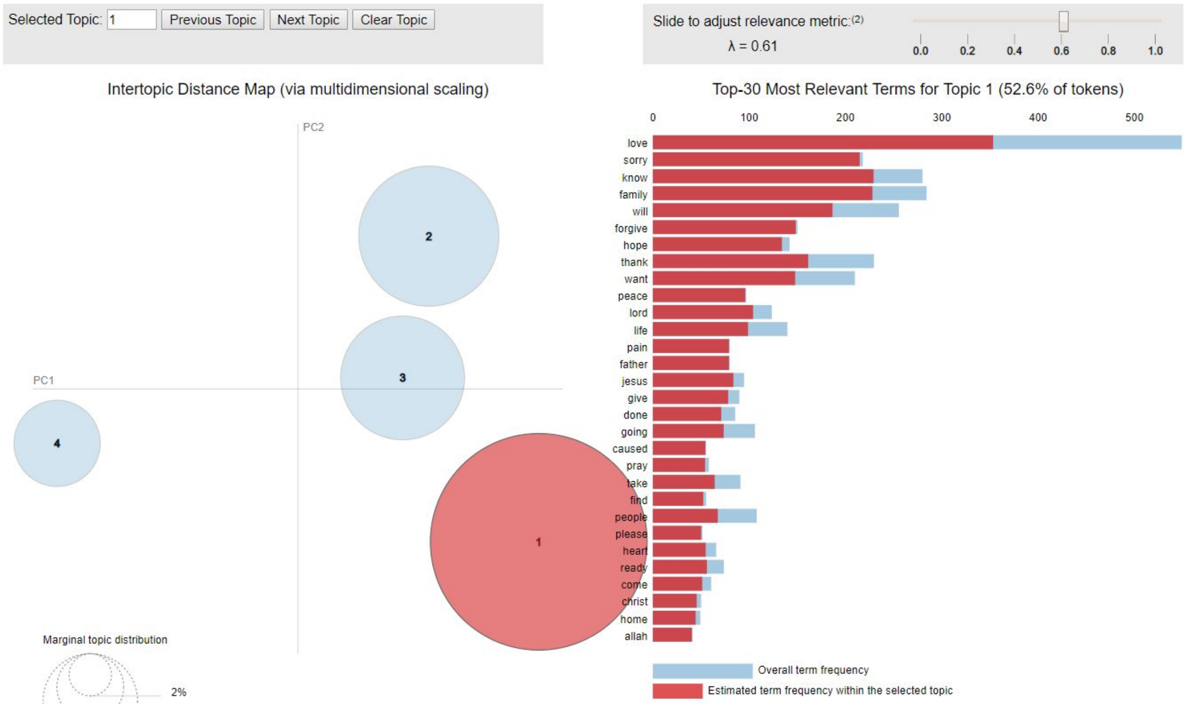

Figure 11 – Topic cluster visualization for Topic 1, including salient terms

Unlike Topic 1, above, Topic 4 shows less strong of an association and some muddying with other topics. This might suggest the topic numbers analyzed needs to be changed, but the dataset might also be too small or improperly processed. The frequency within the selected topic as compared to overall frequency for Topic 4, however, also does suggest that those terms are highly related to one another.

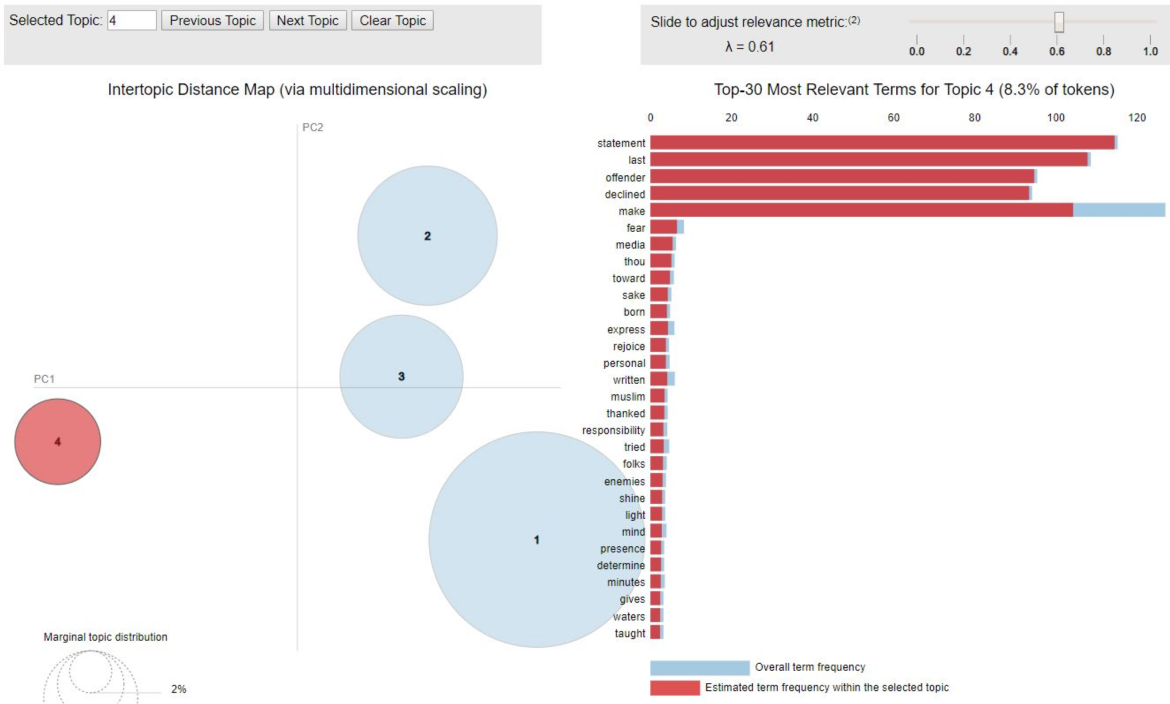

Figure 11 – Topic cluster visualization for Topic 4, including salient terms

General Topic Results

Based on this small initial study, three of the four labels for the resulting topics could be confidently chosen. Those include:

- Topic 1 – Love
- Topic 2 – Life/Justice
- Topic 3 – Forgiveness/Family

The remaining topic is too unclear to make conjecture or seem so topic specific (such as the possibility that Topic 4 is about lack of statements only) that it would be irrelevant as the basis of model building. Before those topics or the ones guessed above can be assigned, a larger dataset should be procured and analyzed to ensure the topic designations remain consistent with the content.

Prediction Modeling

After processing the text, what is left are nine dataframes which were split and used to train and test three classifiers in order to compare their accuracy. Each dataframe was created with an eye for creating variety in an effort to optimize the classification results.

Those dataframes include:

Dataframe	Parameters
DFOne	Drops stop words, words less than or equal to three and any words that contain digits

DFTwo	Drops any words that contain digits
DFThree	Drops stop words, words less than or equal to three and any words that contain digits
DFFour	Drops words less than three and any words that contain digits
DFFive	None – all tokens left in
DFSix	Drops stop words, words less than or equal to three and any words that contain digits. Data normalized using TFIDF.
DFSeven	None – all tokens left in. Data normalized using TFIDF.
DFEight	Drops stop words, words less than or equal to three characters long, any words that contain digits, stemmed words, and words that occur less than three times
DFNine	Drops stop words, words less than or equal to three characters long, any words that contain digits, stemmed words, and words that occur less than three times. Data normalized using TFIDF.
DFTen	Limited to bigrams with a minimum frequency of 3 and a maximum feature set of 5,000.
DFEleven	Limited to bigrams with a minimum frequency of 3 and a maximum feature set of 5,000 represented in binary format.
DFTwelve	Drops stop words, words that occur less than three times, a maximum feature set of 5,000, and any words with digits.

Figure 11 – Dataframes and their parameters

Those nine dataframes, after processing as outlined above, contained the following word count:

Dataframe	**Word Count**
DFOne	2,370
DFTwo	2,666
DFThree	2,370
DFFour	2,619
DFFive	2,700

DFSix	2,700
DFSeven	2,370
DFEight	835
DFNine	835

Figure 12 – Dataframes and their word counts

Education

Multinomial Naïve Bayes Classifier

The first attempt at prediction focused on the education of the inmates – specifically if the inmate only attended school through junior high or high school. Each of the above dataframes were run once for initial accuracy and then run through a 10-fold cross validation to judge consistency in accuracy.

Dataframe	Initial Accuracy	10-fold Accuracy
DFOne	76.79%:	73.09%
DFTwo	76.37%	72.91%
DFThree	75.95%	74.23%
DFFour	71.31%	75.11%
DFFive	73%	74.95%
DFSix	76.37%	74.27%
DFSeven	77.22%	79.33%
DFEight	68.63%	68.63%
DFNine	70%	72.24%

For the sake of concision, only the top performing models will be investigated more closely and the rest will be discussed generally.

The runs showed DFSeven dataframe had the highest initial at 77.22% as well as the highest 10-fold accuracy at 79.33%. Interestingly, both labels had an f1-score of 77%.

```
MultinomialNB 10-Cross Validation Score: 0.7932792207792209
Accuracy Score :  0.7721518987341772
Report :              precision    recall  f1-score   support

  Highschool       0.74      0.81      0.77       113
     JuniorH       0.81      0.74      0.77       124

    accuracy                           0.77       237
   macro avg       0.77      0.77      0.77       237
weighted avg       0.77      0.77      0.77       237
```

Figure 13 – Accuracy report for DFSeven

Bernoulli Naïve Bayes Classifier

Only one dataframes (DFThree) included binary operators as the matrix's vocabulary indicators, a requirement to be run through the Bernoulli Naïve Bayes Classifier, which means only one dataframe were tested.

```
Bernoulli 10-Cross Validation Score: 0.731492303992304
Accuracy Score :  0.6624472573839663
Report :              precision    recall  f1-score   support

  Highschool       0.81      0.49      0.61       128
     JuniorH       0.59      0.86      0.70       109

    accuracy                           0.66       237
   macro avg       0.70      0.68      0.66       237
weighted avg       0.71      0.66      0.65       237
```

Figure 15 – Accuracy report for DFThree

The run showed that the DFThree dataframe as having a fairly similar result as compared with the classifier outlined above. DFThree held an initial accuracy of 66.24% and a 10-fold accuracy of 73.15%.

For the sake of posterity, another dataframe with a differing vocabulary set should be run through the Bernoulli classifier for comparison. It was not done in this case because the initial test fell so short of the other classifiers.

SVM Classifier

Each model was run once for initial accuracy and then run through a 10-fold cross validation to judge consistency in accuracy. The model was run with a C score of 1 and max iterations of 10,000.

Dataframe	Initial Accuracy	10-fold Accuracy
DFOne	77.64%:	81.82%
DFTwo	79.75%	76.03%
DFThree	81.86%	78.22%
DFFour	79.75%	76.77%
DFFive	81.43%	77.12%
DFSix	84.81%	83.17%
DFSeven	85.65%	82.25%
DFEight	78.48%	74.78%
DFNine	80.17%	79.67%

The runs showed the DFSeven dataframe had the highest initial accuracy at 85.65%. The model also exhibits a close spread between the f1-score between labels, just 1%.

```
SVM 10-Cross Validation Score: 0.8224915824915824
Accuracy Score :  0.8565400843881856
Report :              precision    recall  f1-score   support

  Highschool       0.82      0.91      0.86       118
     JuniorH       0.90      0.81      0.85       119

    accuracy                           0.86       237
   macro avg       0.86      0.86      0.86       237
weighted avg       0.86      0.86      0.86       237
```

Figure 17 – Accuracy report for DFSeven

The runs showed the DFSic dataframe had the highest 10-fold accuracy at 83.17%, which also had an f1-score spread of just 1%.

```
SVM 10-Cross Validation Score: 0.8317195767195769
Accuracy Score :  0.8481012658227848
Report :               precision    recall  f1-score   support

  Highschool       0.79      0.91      0.84       108
     JuniorH       0.91      0.80      0.85       129

    accuracy                           0.85       237
   macro avg       0.85      0.85      0.85       237
weighted avg       0.86      0.85      0.85       237
```

Figure 17 – Accuracy report for DFSix

Priors

Multinomial Naïve Bayes Classifier

The second attempt at prediction focused on whether or not the inmate had any priors. Again, each of the above dataframes were run once for initial accuracy and then run through a 10-fold cross validation to judge consistency in accuracy.

Dataframe	Initial Accuracy	10-fold Accuracy
DFOne	58.13%:	46.45%
DFTwo	55.63%	49.89%
DFThree	56.25%	53.35%
DFFour	48.75%	49.35%
DFFive	50%	48.83%
DFSix	50.63%	50.36%
DFSeven	50%	50.65%
DFEight	46.86%	49.59%
DFNine	47.5%	48.2%

The runs showed DFOne dataframe had the highest initial at 58.13% and DFThree had the highest 10-fold accuracy at 53.35%. Both far below the outcome of the education prediction.

```
MultinomialNB 10-Cross Validation Score: 0.46450924608819344
Accuracy Score :  0.58125
Report :               precision    recall  f1-score   support

          no       0.56      0.41      0.47        73
         yes       0.59      0.72      0.65        87

    accuracy                           0.58       160
   macro avg       0.57      0.57      0.56       160
weighted avg       0.58      0.58      0.57       160
```

Figure 13 – Accuracy report for DFOne

```
MultinomialNB 10-Cross Validation Score: 0.5335150940414098
Accuracy Score :  0.5625
Report :               precision    recall  f1-score   support

          no       0.61      0.42      0.50        83
         yes       0.53      0.71      0.61        77

    accuracy                           0.56       160
   macro avg       0.57      0.57      0.56       160
weighted avg       0.58      0.56      0.55       160
```

Figure 13 – Accuracy report for DFOne

Bernoulli Naïve Bayes Classifier

Only one dataframes (DFThree) included binary operators as the matrix's vocabulary indicators, a requirement to be run through the Bernoulli Naïve Bayes Classifier, which means only one dataframe were tested.

```
Bernoulli 10-Cross Validation Score: 0.5307254623044095
Accuracy Score :  0.4875
Report :               precision    recall  f1-score   support

          no       0.39      0.14      0.21        76
         yes       0.51      0.80      0.62        84

    accuracy                           0.49       160
   macro avg       0.45      0.47      0.42       160
weighted avg       0.45      0.49      0.43       160
```

Figure 15 – Accuracy report for DFThree

Again, the run showed that the DFThree dataframe as having a fairly similar result as compared with the classifier outlined above. DFThree held an initial accuracy of 48.75% and a 10-fold accuracy of 53.07%. Interesting to note, however, is how low the f1-scores are – with the "no" label at just 21%.

SVM Classifier

Each model was run once for initial accuracy and then run through a 10-fold cross validation to judge consistency in accuracy. The model was run with a C score of 1 and max iterations of 10,000.

Dataframe	Initial Accuracy	10-fold Accuracy
DFOne	45.63%:	50.42%
DFTwo	56.25%	46.36%
DFThree	58.13%	48.97%
DFFour	45.3%	45.3%
DFFive	50%	49.29%
DFSix	50.63%	44.83%
DFSeven	58.13%	45.56%
DFEight	46.25%	48.8%
DFNine	48.13%	50.18%

The runs showed that DFThree and DFSeven had the same (highest) initial accuracy, but DFThree also had the highest 10-fold accuracy at 58.13% and 48.97%, respectively.

```
SVM 10-Cross Validation Score: 0.4897423739529003
Accuracy Score :  0.58125
Report :               precision    recall  f1-score   support

          no       0.56      0.51      0.53        75
         yes       0.60      0.65      0.62        85

    accuracy                           0.58       160
   macro avg       0.58      0.58      0.58       160
weighted avg       0.58      0.58      0.58       160
```

Figure 17 – Accuracy report for DFThree

Victim's Race

Multinomial Naïve Bayes Classifier

The final attempt at prediction focused on the race of the victim's, which was narrowed down to only White, Black, and Hispanic occurrences because the others listed (Asian, Samoan, and Unknown) were so small as to be (a) insignificant and (b) unlikely to be correctly predicted by the models. Again, each of the dataframes were run once for initial accuracy and then run through a 10-fold cross validation to judge consistency in accuracy.

Dataframe	Initial Accuracy	10-fold Accuracy
DFOne	52.25%:	59.60%
DFTwo	61.26%	62.7%
DFThree	57.66%	62.66%
DFFour	67.57%	55.6%
DFFive	52.25%	65.43%
DFSix	64.86%	65.42%
DFSeven	67.57%	64.21%
DFEight	49.55%	55.77%
DFNine	69.37%	63.48%

The runs showed DFSeven dataframe had the highest initial at 69.37%.

```
 0.6348376068376069
Accuracy Score :  0.6936936936936937
Report :              precision    recall  f1-score   support

       black       0.00      0.00      0.00        16
    hispanic       0.00      0.00      0.00        18
       white       0.69      1.00      0.82        77

    accuracy                           0.69       111
   macro avg       0.23      0.33      0.27       111
weighted avg       0.48      0.69      0.57       111
```

Figure 13 – Accuracy report for DFNine

The runs showed DFFive dataframe had the highest initial at 65.43%. So while this initial batch performed better than predicting priors, it still fell far short of the scores associated with predicting education.

```
MultinomialNB 10-Cross Validation Score: 0.6542977207977207
Accuracy Score :  0.5225225225225225
Report :               precision    recall  f1-score   support

       black       0.00      0.00      0.00        18
    hispanic       0.12      0.04      0.06        25
       white       0.58      0.84      0.69        68

    accuracy                           0.52       111
   macro avg       0.24      0.29      0.25       111
weighted avg       0.38      0.52      0.43       111
```

Figure 13 – Accuracy report for DFFive

Bernoulli Naïve Bayes Classifier

Only one dataframes (DFThree) included binary operators as the matrix's vocabulary indicators, a requirement to be run through the Bernoulli Naïve Bayes Classifier, which means only one dataframe were tested.

```
Bernoulli 10-Cross Validation Score: 0.6187407407407408
Accuracy Score :  0.6306306306306306
Report :               precision    recall  f1-score   support

       black       0.00      0.00      0.00        17
    hispanic       0.20      0.05      0.08        19
       white       0.68      0.92      0.78        75

    accuracy                           0.63       111
   macro avg       0.29      0.32      0.29       111
weighted avg       0.49      0.63      0.54       111
```

Figure 15 – Accuracy report for DFThree

SVM Classifier

Each model was run once for initial accuracy and then run through a 10-fold cross validation to judge consistency in accuracy. The model was run with a C score of 1 and max iterations of 10,000.

Dataframe	Initial Accuracy	10-fold Accuracy
DFOne	55.86%:	44.36%

DFTwo	58.55%	47%
DFThree	49.55%	56.47%
DFFour	51.35%	48.89%
DFFive	49.55%	47.9%
DFSix	62.2%	63.87%
DFSeven	59.46%	66.67%
DFEight	47.75%	51.8%
DFNine	63%	56.1%

The runs showed the DFSix dataframe had the highest initial accuracy at 62.16% and the highest 10-fold accuracy of 63.87%.

```
SVM 10-Cross Validation Score: 0.6386794871794871
Accuracy Score :  0.6216216216216216
Report :              precision    recall  f1-score   support

       black       0.33      0.07      0.12        14
    hispanic       0.33      0.04      0.07        26
       white       0.64      0.94      0.76        71

    accuracy                           0.62       111
   macro avg       0.43      0.35      0.32       111
weighted avg       0.53      0.62      0.52       111
```

Figure 17 – Accuracy report for DFSix

The most important thing to note in the predictions for race is that, unlike education, because the data wasn't run through over or under sampling methods, the data was skewed far too heavily for the model to accurately predict. To be more accurate, this test should be run again with those methods taken into account.

Sentiment Analysis

Some extra data preparation was necessary to optimize the data frame for sentiment analysis. There were initially 566 rows of data. Many of the last statements were actually missing data and were labeled as none. These were removed and decreased the total number of data points to 461. The last statements were all treated the same and removed any numbers, and eliminated any special characters. The words were all made into lowercase. The common first-person pronouns such as I, me, mine, etc.

were replaced as first-person pronouns while other pronouns such as she, he, it, they, etc. were replaced by the word pronoun. Using TextBlob, a common python library for text data, the sentiment polarity was calculated from the last statements. The range for the sentiment values was between -1 to +1, with -1 being the most negative sentiment, and +1 being the most positive sentiment. The score of zero lacked sentiment in the last statement. The distribution of the last statements is shown below.

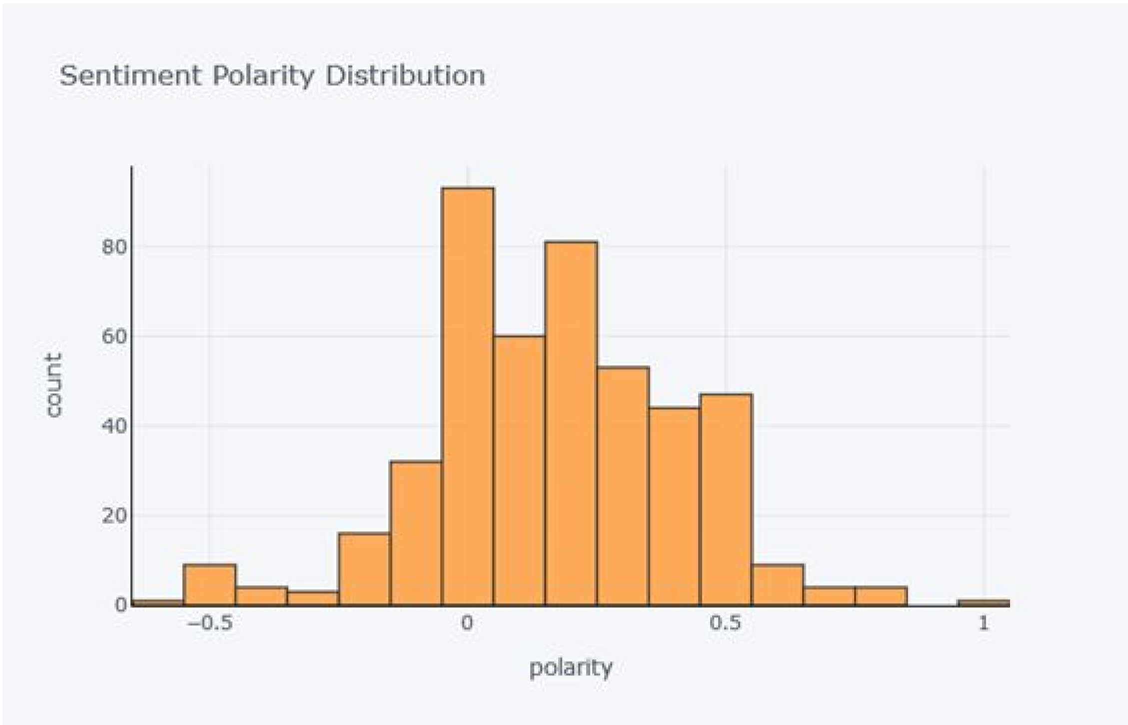

Figure 17 – Sentiment Polarity Distribution

The plot illustrates that there were very few that were marked as negative sentiment. The majority of the last statements were either neutral or positive in sentiment. There was one last statement that was scored as a +1 for sentiment, and this is the last statement of that person.

```
first_person_pronoun want to say god forgives as first_person_pronoun forgive god is the greatest thank pronoun
```

Figure 17 – Most Positive Sentiment

Some of the other highly rated positive sentiments that were rated higher than 0.75 are shown below.

```
i would like to tell first_person_pronoun family first_person_pronoun love pronoun first_person_pronoun attorneys did
their best all of first_person_pronoun brothers on death row those who died and those who are still there to hang in
there and thats all first_person_pronoun have to say
i would like to say  first_person_pronoun just hope ms fielder is happy now first_person_pronoun would like to thank
first_person_pronoun lawyer nancy for pronoun help on first_person_pronoun case and for being with first_person_prono
un now
for the pain first_person_pronoun have caused pronoun first_person_pronoun am ashamed to even look at pronoun faces p
ronoun are great people to first_person_pronoun brothers on death row mexico mexicospanish
 first_person_pronoun want to say god forgives as first_person_pronoun forgive god is the greatest thank pronoun
i want to let all of first_person_pronoun people know and everybody who is here and supported first_person_pronoun th
at first_person_pronoun love pronoun and wish pronoun all the best
```

Figure 17 – Highly Positive Sentiment

The most negative sentiments were not as negative as would be expected. The most common values for the most negative sentiments were only -0.5 whereas the most positive were closer to >= +0.75.

```
god please forgive first_person_pronoun of first_person_pronoun sins look after first_person_pronoun people bless and
protect all people first_person_pronoun am sorry for first_person_pronoun sins lord take first_person_pronoun home wi
th pronoun amen a couple of sentences garbled
i would like to thank all of pronoun for coming first_person_pronoun am sorry for all of the pain first_person_pronou
n have caused both families  first_person_pronoun family and yours
i hereby declare robert steven everett and nicholas velasquez guilty of crimes against first_person_pronoun douglas a
lan feldman either by fact or by proxy first_person_pronoun find pronoun both guilty first_person_pronoun hereby sent
ence both of pronoun to death which first_person_pronoun carried out in august  as of that time the state of texas ha
s been holding first_person_pronoun illegally in confinement and by force for  years first_person_pronoun hereby prot
est first_person_pronoun pending execution and demand immediate relief
i just ask everybody first_person_pronoun ever hurt or done anything wrong to to just forgive first_person_pronoun fo
r whatever wrongs first_person_pronoun done to them
i would like to tell the victims families that first_person_pronoun am sorry very sorry
```

Figure 17 – Most Negative Sentiment

Using these sentiment polarity scores, and the data for every column, boxplots for every combination could be made for each of these. Some of these boxplots were so similar that they were not useful in determining anything, while some were fairly indicative of a sentiment trend based on the discretized attribute. The first nine plots shown did not have any significant differences between the attribute and the sentiments. The first two are the age of the crime, and the age at which they received their sentencing. In these cases, there was minimal difference between the them regardless if they were in their teens, twenties, or thirties and older.

Figure 17 – Age Boxplots

The occupation when discretized into laborer or other did not show any difference in the sentiment. The type of crime, which was discretized into with gun or other did not show any difference between the median scored sentiment either.

Figure 17 – Occupation and Type of Crime Plots

When the crime committed had codefendants, the sentiment polarity values were still very similar. The amount spent on death row did not have any effect on their last statement sentiment either.

Figure 17 – Codefendants and Time on Death Row

Lastly, the victim of the crime did not matter either. these last 3 plots were tested with female, male, or children as victims and the resulting sentiment polarities were on average the same.

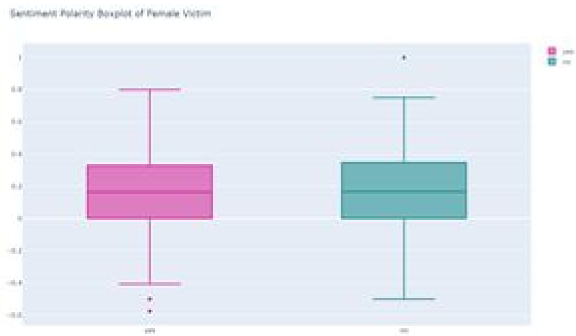

Figure 17 – Victim Types

The sentiment polarity with blox plots for a lot of these was difficult to see a drastic change in due to the fact that most of the sentiments were normally distributed between 0 and 0.5. However, there were some cases when the discretized data was showing a difference between the groups using the sentiment polarity. This first case is the educational level of the offender. The average sentiment of those with high school was lower than those with some or no high school education. The college or unknown also showed a lower sentiment polarity score.

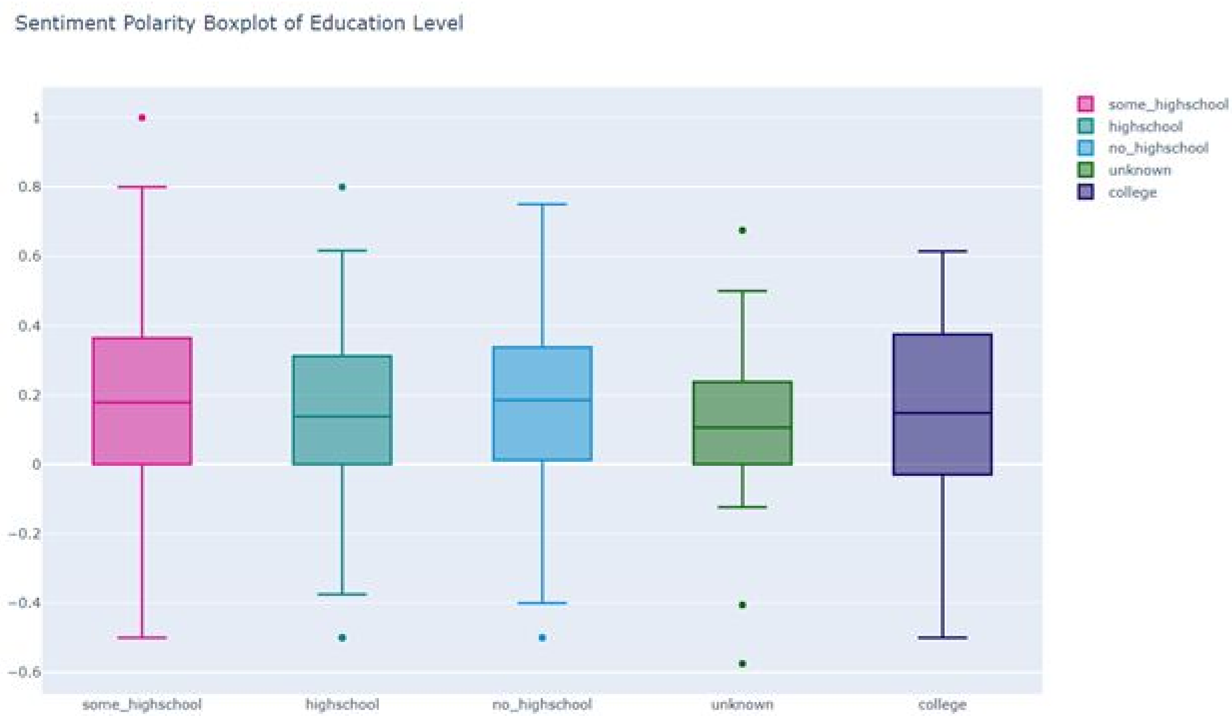

Figure 17 – Education Level Boxplot

The sentiment score for those offenders that had a past criminal background was on average slightly lower than those without a past criminal background. Those that lacked criminal background history showed the lowest average sentiment score, but had the tightest range with the fewest sentiments that were actually negative.

Figure 17 – Prior Record Boxplot

The number of victims on average had a higher sentiment polarity score if they had 2 or more victims than those with only one victim. These values were not too different, but the overall range of those with two or more showed fewer negative scores.

Figure 17 – Number of Victims Boxplot

The main crime was discretized into 5 different bins, either murder, murder rape, murder rap robbery, murder robbery, or murder other. The sentiment polarity scores of those that were strictly murder rape showed to have the lowest average score compared to the other 4 categories. This may be because the other acts may have been the main criminal act, leaving a murder as a secondary last resort in which these offenders regret having done that. The murder rapists also have the lowest peak sentiment score from the others.

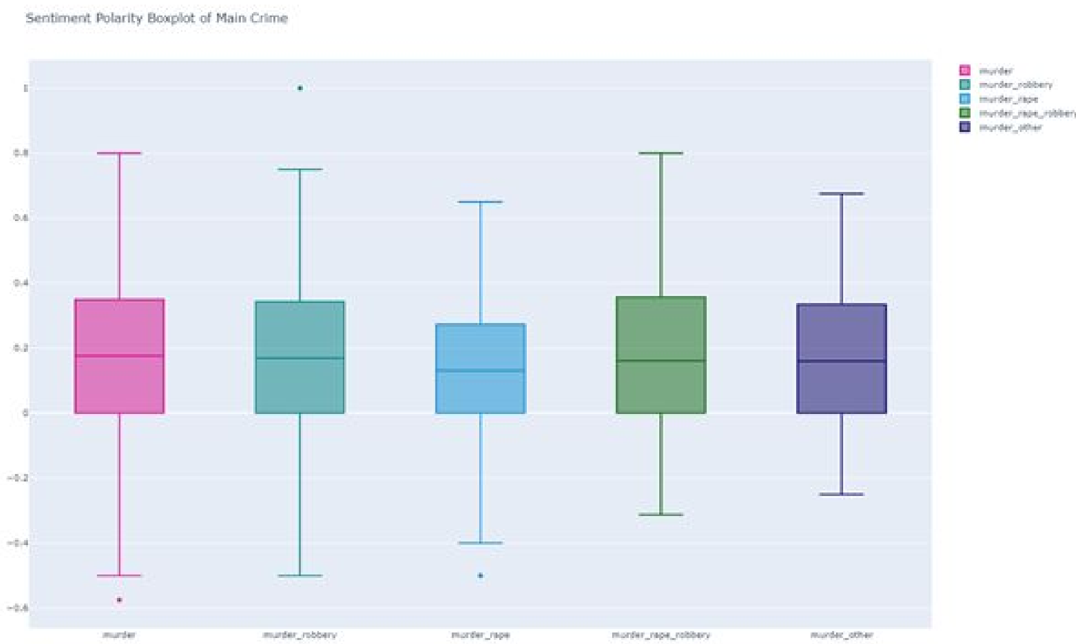

Figure 17 – Main Crime Boxplot

The type of weapon used did show a difference in the sentiment polarity score. Those that used a knife showed lower polarity scores on average to those that used a gun or another form of weapon. The knife sentiment score was almost half that of the other two categories.

Figure 17 – Weapon Type Boxplot

The sentiment score of the victim when they are a police officer or not slightly shows a difference in the sentiment score. Those that were offenders towards police showed a slightly average lower polarity score than those whose victims were not police officers. Overall, the offenders who killed police officers seem to have a more neutral final statement whereas those that did not have a police officer as the victim felt both much more negative sentiment, and positive sentiment towards their victims.

Figure 17 – Police Victim Boxplot

The final statement of those with victims of a specific race did show some differences in their sentiment. While white, hispanic, and unknown victim races were all similar in average sentiment score, those with black victims were much more likely to have a positive sentiment score than others. Also, those that had victims of other races had less negative sentiment, but overall their sentiment average score was the lowest.

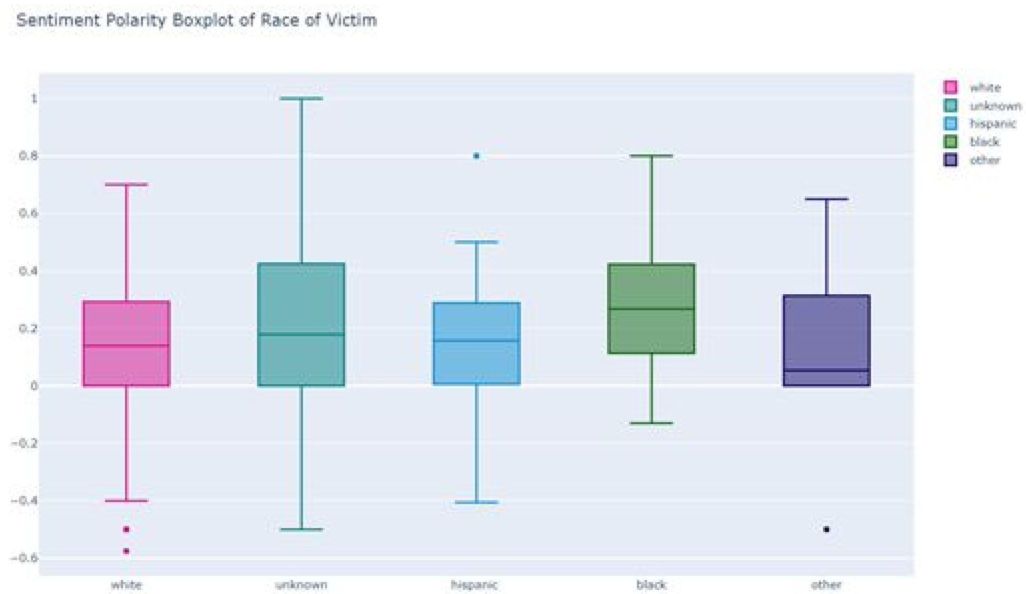

Figure 17 – Victim Race Boxplot

Continuing with the analysis of the race, the last statement sentiment data showed that those offendered of black or hispanic race had twice the sentiment score over those that were white. The values were close to 0.1 and 0.2. While the other's boxplot looked higher, there was not enough data to truly use these values for the analysis.

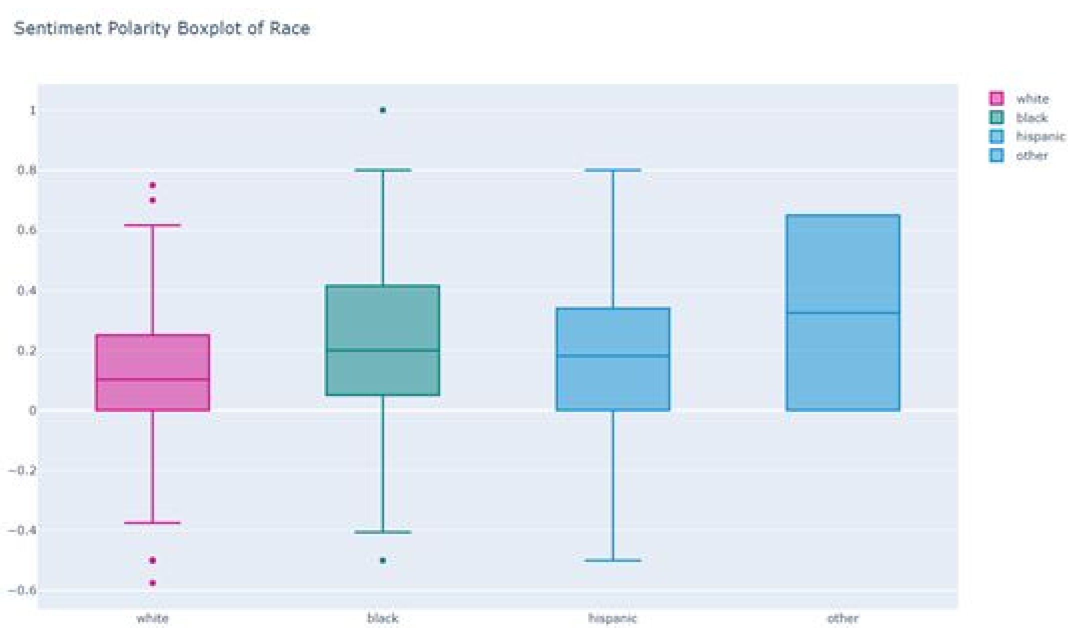

Figure 17 – Offender Race Boxplot

While the sentiment analysis did show some differences in the data and the sentiment polarity analysis, there is not enough data to truly evaluate these methods. Sentiment may not be the optimal way to look at the last statements. A better way would be to try to understand if the offender was remorseful or not, rather than trying to understand their sentiment. While the words that they speak may have a very remorseful tone, the words may also have little or no sentiment. They may actually have a negative sentiment. One of the most negative sentiment scores that was created with TextBlob actually read fairly remorseful, but was graded as a very negative sentiment.

Conclusions

As of July 1, 2019, there are 2,656 inmates on death row in the United States and 1,500 inmates have been put to death since 1979 (Fins, 2019). Much like the inmates in Texas whose last statements were made available to the public, the population of death row skews White or Black, male, and arrested in their 20s. Although that population might appear narrow at first glance, the final statements of these inmates offers a unique and intimate opportunity to learn more about everything from the consequence of long-term imprisonment, how religion plays a part in rehabilitation/outlook, and even the psychological implications of regret and repentance of these communities.

This kind of insight doesn't just provide soft insights, it has the potential to provide a closer look into what brings inmates to death row, the demographics most affected by the practice, the criminal mindsets, and even recidivism efforts – in fact, less than 10% of inmates on death row had a prior homicide conviction so changes in language from conviction to execution could point to its effectiveness as a deterrent, which also points to its effectiveness as a punishment culturally (ACLU, 2012). Maybe most importantly, "nationally, at least one person is exonerated for every 10 that are executed" (ACLU, 2012) and what this means is that, properly tuned, these studies and models have the potential to save lives. At the end of the day, what this small-scale analysis shows is that final statements may be a fruitful source of data.

By dissecting the topics, sentiment, and exploring the data more in depth, researchers of all disciples can learn more about an age-old practice and the people it affects.

References

ACLU. (2012). *The Case Against the Death Penalty*. Retrieved from ACLU: https://www.aclu.org/other/case-against-death-penalty

Alharthi, H., Inkpen, D., & Szpakowicz, S. (20017). Unsupervised Topic Modelling in a Book Recommender System for New Users. *SIGIR 2017 eCom.* Tokyo.

Death Penalty Information Center. (2019). *Executions by State and Year*. Retrieved from Death Penalty Information Center: https://deathpenaltyinfo.org/executions/executions-overview/executions-by-state-and-year

Eaton, J., & Theuer, A. (2009). Apology and Remorse in the Last Statements of Death Row Prisoners. *Justice Quarterly*, 327-347.

Fins, D. (2019). *Death Row U.S.A. : Summer 2019.* NAACP Legal Defense and Educational Fund, Inc. .

Foley, S., & Kelly, B. (2018). Forgiveness, spirituality and love: thematic analysis of last statements from Death Row, Texas (2002–17). *QJM: An International Journal of Medicine*, 399–403.

McAllister, A., Naydenova, I., & Quang. (n.d.). *Humboldt-Universitat zu Berlin*. Retrieved from Building a LDA-based Book Recommender System: https://humboldt-wi.github.io/blog/research/information_systems_1819/is_lda_final/#conclusion

Mimno, D., & McCallum, A. (2007). Organizing the OCA: Learning Faceted Subjects from a Library of Digital Books. *ACM/IEEE Joint Conference on Digital Libraries.* Vancouver, BC.

Oliphant, J. B. (2018, June). *Public support for the death penalty ticks up*. Retrieved from Pew Research: https://www.pewresearch.org/fact-tank/2018/06/11/us-support-for-death-penalty-ticks-up-2018/

Pilkington, E. (2014, April). *US death row study: 4% of defendants sentenced to die are innocent*. Retrieved from The Guardian: https://www.theguardian.com/world/2014/apr/28/death-penalty-study-4-percent-defendants-innocent

Rice, S. K. (2009). Of Guilt, Defiance, and Repentance: Evidence from the Texas Death Chamber. *Justice Quarterly*.

Rowe, A. (2018, July 22). *Estimated U.S. Book Publisher Revenue Was North Of $26 Billion In 2017.* Retrieved from Forbes: https://www.forbes.com/sites/adamrowe1/2018/07/22/estimated-u-s-book-publisher-revenue-was-north-of-26-billion-in-2017/#7ff83f2d3196

Smith, R. A. (2018). The Functional Use of Religion When Faced with Imminent Death: An Analysis of Death Row Inmates' Last Statements. *The Sociological Quarterly*.

Weber, P. J. (2019, April). *Texas stops sharing death row inmate's final statements*. Retrieved from AP News: https://apnews.com/774e30a3ea8a4f718d78c0b133711751