Unlocking New York City Crime Insights using Relational Database Embeddings

Apoorva Nitsure
IBM Watson Research Center
Yorktown Heights, NY 10598
apoorva.nitsure@us.ibm.com

Rajesh Bordawekar
IBM Watson Research Center
Yorktown Heights, NY 10598
bordaw@us.ibm.com

Jose Neves
IBM Systems
Poughkeepsie, NY 12601
jneves@us.ibm.com

arXiv:2005.09617v2 [cs.DB] 20 May 2020

ABSTRACT

This paper demonstrates the use of the AI-Powered Database (AI-DB) in identifying non-obvious patterns in crime data that could serve as an aid to predictive policing measures. AI-DB uses an unsupervised neural network, `db2Vec`, to capture inter and intra-column semantic relationships from a relational table and allows users to exploit such relationships using novel semantic SQL queries. Using the publicly available New York Police Department (NYPD) Crime Complaint Dataset as an example, the paper illustrates how AI-DB can be used to interpret the data and generate useful insights. We demonstrate that AI-DB's database embedding model and semantic queries enable users to identify criminal complaint patterns that are not possible to extract using current crime analysis tools, including NYPD's state-of-the-art *Patternizr* system. We show that the AI-DB system can generate new insights with reduced pre-processing and execution costs (e.g., no labeling, reduced feature engineering, and use of standard SQL queries) with reasonable training performance (i.e., processing and training the 6.5 Million crime complaints in the NYPD Crime Complaint Dataset took less than 4 hours). The SQL-based implementation can be incorporated into any data science pipeline to provide visual representation of the results.

1. INTRODUCTION

Law enforcement work can be broadly classified into reactive policing where officers respond to service calls, make arrests, etc. and pro-active policing where historic policing data is used to understand crime patterns and insights. With such information police departments can strategize long and short term policies to prevent and/or mitigate future significant incidents. Recognizing the importance of pro-active techniques, the National Institute of Justice, an Agency of the U.S. Department of Justice supports several initiatives in this area [12, 7]. One area in particular is *Predictive Policing* where historical crime-data information combined with geo-spatial technologies and evidence-based intervention models are used to reduce crime and improve public safety. Research in this field has been focused on the use of advanced analytics and statistical methods to aid the identification of patterns in crime data. In some cases these methods have been augmented with the use of geo-spatial information as well as dynamic information of citizen activity (travel, shopping, dining, working, etc) within such geo-spatial and temporal data.

Currently, there are multiple systems [17, 6, 15] designed to analyze criminal datasets, including *Patternizr*, which is the state-of-the-art system deployed at the New York Police Department (NYPD) for identifying criminal complaint patterns [4]. All of these approaches use supervised training techniques to classify and predict criminal complaints. Any supervised training approach requires extensive feature engineering as well as data labeling. Furthermore, a supervised model can predict only one attribute that the model has been trained for. As we demonstrate in this paper, our novel approach, AI-Powered database (AI-DB), based on *unsupervised* neural network modeling of relational tables can not only identify obvious patterns as the Patternizr but can also uncover unapparent insights from the NYPD dataset.

In this paper, we apply AI-DB to analyze the publicly available NYPD historic crime complaint dataset [11]. AI-DB is a prototype relational database system currently being developed at IBM Research that uses an *unsupervised* neural network based approach from Natural Language Processing (NLP), called *vector embedding*, to extract *latent* semantic knowledge from a database table or a collection of tables [2, 3, 1]. For a text document, the word embedding model (e.g., Word2Vec [9]) builds a semantic model for the document words, where for each word, the inferred semantic meaning captures contributions from neighborhood words, across all instances of such word. For relational tables, our new vector database embedding model (db2Vec) captures *inter- and intra-column* and *inter-row semantic* relationships between database entities of different types and builds a *meaning* vector for each entity. The vector model is integrated into the existing SQL query processing infrastructure and used to enable a new class of SQL-based semantic analytics queries called *Cognitive Intelligence* (CI) queries. The AI-DB approach enables new capabilities (e.g., inductive reasoning queries), ease-of-use (e.g., no data labeling, no or very limited feature engineering, and standard SQL-based querying), and reasonable pre-processing and training performance.

The remainder of the paper is organized as follows. First, literature in crime prediction is reviewed in Section 2. In Section 3, the main features of AI-DB system are high-

lighted. Section 4 describes the features of an in-house tool developed by the New York Police Department, Patternizr, designed to identify meaningful crime patterns to aid investigators. Section 5 describes the data pre-processing stage in training the database embedding model from the NYPD dataset. Section 6 details different ways of detecting crime patterns and extracting relevant information using various types of CI queries which include similarity, dissimilarity, multi-attribute, analogy, and semantic clustering queries. This information can be visualized in several different forms, such as tables, line and bar charts, and geo-spatial maps. Finally in Section 7, a summary of the results obtained using AI-DB is presented and it concludes with describing future work.

2. RELATED WORK

Analysis of crime data is not a novel concept. There have been multiple approaches taken to realize implicit information embedded in such data. Some examples of these approaches are crime pattern detection and identification, localization and visualization of criminal hot spots organized by crime classification, and forecasting crime types given a set of attributes. Another example is the identification of crime patterns to help identify new crime patterns and aid in making arrests[4]. For this purpose a system, Patternizr, was built using three supervised machine learning models to help identify burglaries, robberies, and grand larcenies. This system was incorporated into the NYPD Domain Awareness System (DAS) and is available to authorized personnel. Starting from the NYPD complaint dataset, multiple crime-to-crime similarity values are combined into a *similarity score* that quantifies the likelihood that a pair of crimes are in a pattern together. Patternizr is designed around a pattern detection algorithm first introduced in [17] called *Series Finder*. Series Finder is a supervised machine learning algorithm that makes use of pattern-specific and pattern-general coefficients. It is a growing algorithm that makes use of certain seed complaints and grows by adding other crimes in the database to achieve a minimum cohesion score. Cambridge Police Department data was used to validate this approach. In [6] the same NYPD crime dataset is used in conjunction with weather and temporal attributes to predict crime using various traditional machine learning algorithms and deep learning.

Another use of multiple datasets is described in [15] where Chicago and Portland crime complaint data is augmented with weather, public transportation and census data to predict crime counts. A variety of neural network architectures including Feed Forward, Convolutional, Recurrent and their combinations are used to predict crime count bins on a spatial region at a daily level. Selection of features for predicting crime counts based on NYPD Felony Crime Data and geo-spatial citizen activity from Foursquare is described in [8]. Data from Foursquare is divided into spatial features representing static information about a venue and human dynamic features representing how citizens interact with a venue. To validate the feature selection algorithms a supervised learning crime count prediction system is described. The system implements different techniques for crime count prediction such as Lasso, Ridge and Random Forest Regressor models while optimizing Root Mean Squared Logarithmic Error. Another interesting dataset that can be used for crime prediction is the 911 emergency calls. In [5], 911 calls from New York City are analyzed to predict local demand of police resources. The initial step is to cluster the calls based on temporal behavior to provide insight to call behavior. This is followed by the development of a predictive model to be used to estimate future resource demands in specific regions. Finally, the system also includes techniques to detect anomalous events allowing the identification of unusual events and allocation of resources accordingly.

3. AI-POWERED DATABASE OVERVIEW

AI-Powered Database (AI-DB) is a novel relational database system being developed that uses the power of unsupervised neural networks to enable standard SQL-based semantic queries on relational tables. There are three key phases in the execution flow of an AI-Powered Database [2, 3]. The first *preprocessing phase* takes a relational table with different SQL types as input and returns an unstructured but meaningful text corpus consisting of a set of sentences. This transformation, termed *textification*, allows us to generate a uniform semantic representation of different SQL types. The textification phase processes each relational row separately and for each row, converts data of different SQL data types to equivalent text representation. Once a relational table is converted into a textual training dataset, it is processed by the database embedding (db2Vec) approach to generate semantic vectors for every unique entity in the training dataset (*training phase*). The db2Vec differs from traditional NLP vector embedding models used for language modeling, such as Word2Vec [9] or GloVe [13] in several ways. Key differences include:

- A sentence generated from a relational row is generally not in any natural language such as English. In the db2Vec implementation, every token in the training set has the same influence on the nearby tokens; i.e. we view the generated sentence as *a bag of words*, rather than *an ordered sequence*.
- For relational data, we provide special consideration to primary keys. Traditional word embedding approaches discard less frequent words from computations. In our implementation, by default, every token, irrespective of its frequency, is assigned a vector. For a unique primary key, its vector represents the meaning of the entire row.
- The `db2Vec` model provides special treatment for the entities corresponding to the SQL `NULL` (or equivalent) values. The `NULL` values are processed such that they do not contribute to the meanings of neighboring non-null entities; thus eliminating false similarities.

For each database token in a relational table, the model generates a vector that encodes contextual semantic relationships generated by collective contributions of other tokens within and across rows (each table row is viewed as a sentence). The db2Vec model generates a variety of semantic vectors: (1) each unique primary key is associated with a semantic vector that captures behavior of the entire relational row associated with that key, (2) for all other entities, their semantic vectors capture collective contributions of their neighborhood entities across their occurrences, and (3) table schema types (column names) get their meaning vectors that capture table-wide relationships.

The final *query execution phase* is where the user issues SQL statements to extract information from one or more databases using the trained db2Vec models. Such queries, termed *Cognitive Intelligent (CI)* queries, can support both the traditional value-based as well as the new semantic contextual computations in the same query. Each CI query uses user-defined functions (UDFs) to measure semantic similarity between a pair of sets (or sequences) of tokens associated with the input relational parameters. The core computational operation of a cognitive UDF is to calculate similarity between a pair of tokens by computing the cosine similarity between the corresponding meaning vectors. For two vectors v_1 and v_2, the cosine similarity is computed as $cos(v_1, v_2) = cos(v_1, v_2) = \frac{v_1 \cdot v_2}{\|v_1\| \|v_2\|}$. The cosine similarity value varies from 1.0 (very similar) to -1.0 (very dissimilar). Each CI query uses the UDFs to execute *nearest neighbor* computations using the vectors from the current word-embedding model. Our current implementation supports four types of CI SQL queries: similarity based classification, inductive reasoning, prediction, and cognitive OLAP [2]. Figure 1 presents a simple CI SQL query that given a complaint number (i.e., `229113435`), identifies other complaints with the similar behavior (each table row represents a complaint). The UDF, `ComplaintSimilarityUDF()`, takes relational variables as input, and returns a similarity score computed using corresponding meaning vectors. The SQL CI query then returns those rows whose semantic similarity is higher than a specified bound (0.5), ordered in a decreasing order of semantic similarity.

```
SELECT *,
ComplaintSimilarityUDF(X.CMPLNT_NUM,"229113435")
AS proximityValue
FROM Index_View X
WHERE ComplaintSimilarityUDF(X.CMPLNT_NUM,"229113435") >
0.5
ORDER BY proximityValue DESC
```

Figure 1: Example of a SQL CI similarity query: find complaints similar to a given complaint

One of the key advantages of the vector embedding approach is that one can compute semantic similarity between any pair of entities by computing cosine similarities between the corresponding vectors. Thus, unlike the supervised training model that works only for the target entity it has trained for, a single embedding model can be used to compute different types of similarities between any pair of entities, irrespective of their types, from the associated relational table.

4. *PATTERNIZR*: NYPD'S CRIME ANALYSIS TOOL

Law enforcement agencies have access to a lot of data, however identification of crime patterns is a non-trivial task. The New York Police Department has developed an in-house tool, called Patternizr, to aid investigators in crime analysis and prediction, specifically by identifying meaningful crime patterns[4]. Patternizr has been designed to address the following issues: (1) On encountering a new crime report, an investigator tries to recall crimes with similar characteristics to see if they belong to a pattern. This process is time consuming and memory based. (2) Investigators cannot effectively identify patterns across precincts and patterns linked deeper in the past over longer periods due to the manual process. (3) Use of search engines limits the search to exact categorical matches, unable to consider a broader match on different attributes. Patternizr identifies complaint patterns defined as a group of complaints that are similar and could be committed by the same person. The crimes in a pattern may be nearby in space, similar in times and days of occurrence, and (or) alike in method (also known as the modus operandi, or M.O.)

Patternizr uses supervised machine learning to provide recommendations to crime analysts. The general flow of the process is: first, a seed complaint is chosen, then a comparison of that seed complaint and every other complaint is made to generate a similarity score, and the ranked list of similar complaints is returned for analytic review. The model makes use of the Random Forest algorithm for classifying pairs of crimes as a pattern. A *separate* model is built for three different crime types: burglary, robbery, and grand larceny. Each model is trained over a data corpus of crime complaints recorded over a 10 year period from 2006-2015 for the associated crime type. Each training example is a pair of two crime complaints, with features such as location, date-time, categorical, suspect information, and unstructured text. **Sensitive attributes like race and gender were removed while training the model.**

The Patternizr paper [4] mentions two key shortcomings of the system: first, it compares complaints using only three crime types, but Patternizr users would like compare complaints at a lower granularity, e.g., compare grand larceny against petite larceny complaints. Second, users would like to compare complaints across crime types, e.g., robberies and grand larcenies. The supervised learning approach also restricts the pattern identification to only one aspect of a criminal complaint, crime types; e.g., analysts can not compare crime types against crime location, or crime types against crime times. In addition, the current approach for building the Patternizr model using pairwise computations is computationally expensive (it took **19.4 days** to process all historical pairs on the cloud with **1600 cores** [4]). As discussed in the following sections, the AI-DB approach that uses unsupervised training, can address both the functionality and performance issues, and also provide new capabilities for criminal analysts.

5. DATA PRE-PROCESSING AND MODEL TRAINING IN AI-DB

The NYPD historic crime complaint dataset consists of felony, misdemeanor and violation crimes reported to the New York Police Department from 2006 to present [11]. This dataset has been downloaded from NYC Open Data Portal [10], where publicly available datasets from different city agencies, including the New York City Police Department, can be downloaded. The complaint dataset was made public in 2016 and has had periodic updates since then. The data version that was downloaded for this analysis was last updated in September 2019. It is in CSV format with 2.06 GB in size and contains 6500870 distinct complaints, each containing 35 features. More details on the NYPD schema (i.e., column names and types) can be found in Appendix A.

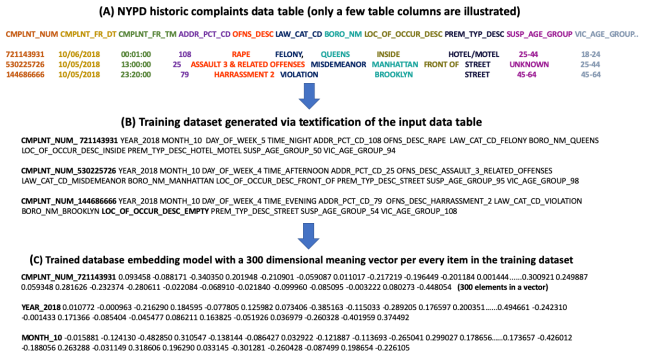

Figure 2: Different steps in generating a database embedding model. (A) Original Table data; (B) Training dataset after textification; (C) Embedding model generated by db2Vec

Figure 2 outlines key steps in building a database embedding model from the raw NYPD crime complaints database. The first step, *textification*, converts the CSV file into a corresponding text document. First, several column attributes from the original dataset was removed. The dropped columns include **sensitive attributes containing suspect and victim gender and race information**, and other columns that presented redundant information or seemed unimportant in adding any value to the analysis (i.e., more than 90% entries were marked "not available" (`NA`)). Appendix A lists the dropped columns along with their descriptions.

Data from the remainder 18 columns (Appendix A) was then transformed using a Python based framework utilizing key libraries (e.g., pandas, sklearn, and numpy). The input CSV data file is first read into a pandas dataframe in chunks and each chunk operated on using pandas dataframe functions. Entities for the PD_DESC column are processed to create a list of offense descriptions. Entries for the columns `CMPLNT_FR_DT` and `CMPLNT_FR_TM` are converted to datetime datatype. Year, month, day of the week is extracted from `CMPLNT_FR_DT` whereas time is bucketed into intervals from `CMPLNT_FR_TM`. After datetime transformations, various chunks are aggregated to form a single in-memory dataframe. The NA values are replaced by the string `EMPTY`. Using sklearn's preprocessing library, values for the columns `VIC_AGE_GROUP` and `SUSP_AGE_GROUP` are label encoded. After these transformations, all values are converted to strings and each value is prepended with the column name separated by an underscore. This dataframe is then saved as a text document, see Figure 2(B).

For the 2.06 GB NYPD complaints dataset, the entire textification process took around **32 minutes**. The majority of time was spent in Python libraries for feature extraction and label encoding. The textified document is then trained using db2Vec (Section 3) to build a database embedding model. The training step generated a model where each of the 6501958 entries gets a 300 dimensional single-precision vector, Figure 2(C). The training process took **208 minutes (around 3.5 hours)** on a single dual-core x86 system (with 2 Xeon CPU E5-2680 processors, each with 28 cores).

6. INSIGHTS DERIVED FROM AI-DB QUERIES

This section presents the results of evaluating AI-DB capabilities on the NYPD crime complaints dataset. For the evaluation, we used a Spark 2.3 based implementation of the AI-DB system [1]. The original CSV dataset along with the corresponding database embedding model were stored as Spark dataframes. Then SQL CI were invoked on these dataframes. For certain queries, we computed the similarity results using cosine similarity calculations. Our experiments were run on a single dual-core x86 system. For visualization purposes, a Jupyter Notebook was designed combining the results of CI SQL and cosine similarity based similarity queries with map and graphic visualization capabilities. The some of the figures in this section were generated from this Jupyter Notebook.

In this section, we discuss results from executing the following CI queries on the NYPD historic criminal complaints dataset: complaint similarity, single- and multi-attribute similarities, dissimilarity, and inductive reasoning queries.

6.1 Complaint Similarity Queries

In the first set of experiments, we replicate what Patternizr can achieve: given a seed complaint, identify and group related complaints to expose underlying patterns. However, unlike Patternizr, we support similarity queries on complaints of different crime types. For this set of experiments, we use the three seed complaints presented in Figure 3. These three complaints have completely different temporal, spatial, and offense related attributes. For each complaint, we invoke a CI similarity query over two different database embedding models: one trained on data with 18 columns (Section 5), and the other one trained on a different view of the underlying table: it has 17 columns and the precinct information is skipped. For the complaint similarity computations, we compare vectors of the complaint number entries, which are unique primary keys. Therefore, the meaning vector for each complaint number captures the behavior of its entire row.

Complaint Attributes	Complaint A	Complaint B	Complaint C
Complaint Number	715447869	134746083	229113435
Complaint Date	10/06/2018	10/08/2018	10/05/2018
Complaint Begin - End Time	02:35:00 - 02:45:00	04:20:00 – 04:40:00	19:00:00- 19:30:00
Precinct (Borough)	40 (Bronx)	66 (Brooklyn)	17 (Manhattan)
Offense Description	Robbery, Personal Electronic Device	Burglary, Unclassified, Night	Grand Larceny, Open Areas, Unattended
Premise	Street	Inside Synagogue	Inside Residence-APT/HOUSE
Suspect Age Group	25-44	Unknown	Unknown

Figure 3: Seed criminal complaints: A, B, C

In the first scenario, let us consider the crime complaint A in Figure 3 as the seed complaint. Our aim is to find complaints that exhibit *similar behavior* as the current seed complaint. Figure 4 presents the results of executing a similarity query using two different models. In both cases, we present 3 complaints that are most similar to the seed complaint (i.e., the cosine similarity score between the meaning vectors of the complaint entities is the closest to 1.0).

We can see that the top three complaints returned by the similarity query share the precinct, borough, premise, and offense description; however, they differ in date, begin-end times, and suspect age. When the model was built *without* using the precinct information, the top matched complaints

Complaint Number	Date	Begin Time	End Time	Precinct & Borough	Premise	Offense Description	Suspect Age
596459038	07/04/2018	02:30:00	02:45:00	40, Bronx	Street	Robbery, Personal Electronic Device	25-44
175772592	10/09/2016	02:55:00	03:00:00	40, Bronx	Street	Robbery, Personal Electronic Device	Unknown
600101132	10/05/2015	02:20:00	03:00:00	40, Bronx	Street	Robbery, Personal Electronic Device	Unknown
Results from the model generated without using the Precinct information							
243886080	08/11/2018	23:55:00	NA	47,Bronx	Street	Robbery, Personal Electronic Device	25-44
863636601	12/27/2014	17:30:00	NA	44, Bronx	Street	Robbery, Personal Electronic Device	Unknown
736738896	01/25/2014	18:10:00	18:11:00	49, Bronx	Street	Robbery, Personal Electronic Device	Unknown

Figure 4: Similar Robbery Complaints

lie across multiple precincts in the Bronx. This result highlights the key aspect of the AI-DB approach: the inferred meaning of a relational entity (e.g., a complaint number) is based on the relational view used to train the database embedding model. As such the results provide additional information which can be used for deeper insights. Since the results span across multiple years one can use such information to determine if there was a relapse of the crime suspect.

Let's look at the results (Figure 5) for another seed complaint about a burglary (Figure 3:B). The first set of results finds burglary complaints across different precincts in Brooklyn, but with the same premise as the seed complaint: a Synagogue. With the new model, the similar complaints span multiple boroughs with different begin-end times and premises. These results can provide new useful insights since criminals do not stick to a single borough.

Complaint Number	Date	Begin Time	End Time	Precinct & Borough	Premise	Offense Description	Suspect Age
668380136	01/17/2018	04:05:00	04:20:00	90, Brooklyn	Synagogue	Burglary, Unclassified	Unknown
877587424	10/27/2016	03:00:00	03:30:00	66, Brooklyn	Synagogue	Burglary, Unclassified	Unknown
855890575	04/06/2015	04:45:00	05:12:00	61, Brooklyn	Synagogue	Burglary, Unclassified	Unknown
Results from the model generated without using the Precinct information							
214280759	09/22/2016	17:30:00	20:30:00	108, Queens	Street	Burglary, Unclassified, Night	Unknown
298386312	10/18/2018	02:15:00	03:30:00	7,Manhattan	Construction Site	Burglary, Commercial, Night	Unknown
289053420	10/07/2018	04:35:00	04:52:00	79, Brooklyn	Residence – Apt/House	Burglary, Residence, Night	Unknown

Figure 5: Similar Burglary Complaints

One of the shortcomings mentioned in [4] is that Patternizr currently cannot detect patterns across crime types, and a new model needs to be trained for each crime type pair. In the AI-DB approach, not only can you compare across crime types but also without the need to train new models. The same model is used for analyzing complaints across all crime types.

Let us look at a grand larceny complaint from 2018 (Figure 3:C) and find petit larceny complaints similar to it in the same year. Note that the same model is used on all queries, independent of crime type such as robbery, grand larceny, or burglary. The differences and commonalities between these complaints can be seen from the output of the query highlighted in Figure 6 which depicts the important temporal, spatial and other crime related information of the similar complaints.

Complaint Number	Date	Begin Time	End Time	Precinct & Borough	Premise	Offense Description	Suspect Age
333288538	10/07/2018	18:40:00	18:51:00	17, Manhattan	Chain Store	Petit Larceny, From Store Shop Lifting	Unknown
966407047	10/12/2018	18:40:00	18:50:00	13, Manhattan	ATM	Petit Larceny, By False Promise	Unknown
451648666	07/04/2018	18:51:00	18:55:00	17, Manhattan	Chain Store	Petit Larceny, From Store Shop Lifting	Unknown
Results from the model generated without using the Precinct information							
393835975	04/22/2018	08:30:00	NA	33, Manhattan	Residence-APT/ House (Inside)	Petit Larceny, From Building	Unknown
755614165	05/03/2018	16:15:00	16:20:00	18, Manhattan	Drug Store (Inside)	Petit Larceny, From Store Shop Lifting	Unknown
120150545	10/19/2018	19:00:00	06:00:00	75, Brooklyn	Street	Petit Larceny, From Auto	Unknown

Figure 6: Similar Petit Larceny Complaints

From these experimental results, one can note that CI queries generate vastly diverse results based on what view of the data is used for model training. As we can see from the attributes of similar crimes, they all occur in similar time intervals and at the same premise; however the precincts are different. This provides a view for investigators to expand their realm in terms of identifying patterns out of their assigned precincts as it is very likely that a suspect may commit crimes in different precincts. It also provides information that can be used to visualize crimes by given premises and time windows.

6.2 Single-attribute Similarity Queries

In the previous section, we demonstrated similarity queries over complaint number entries that are unique primary keys of the table. In this section, we demonstrate AI-DB similarity queries over entities that are not unique primary keys, and their meaning vectors capture collective contributions of other neighborhood entities. Also, since AI-DB generates meaning vectors for all unique entities in a relational tables, it is possible to compare entities irrespective of their relational schema type. For these experiments, we used a CI SQL query similar to shown in Figure 1 using the full model (with 18 columns).

In the first experiment, we find similar premises based on their *collective* criminal behavior inferred by using all attributes of the dataset. In the first case, we get the sorted list of premises similar to **Synagogue** as follows: (1) Other House of Worship, (2) Mosque, (3) Public School, (4) Church, and (5) Jewelry, which belong mostly to places of worship. In the second experiment, we aim to find premises similar to **Street**, CI similarity query returns the following sorted list of premises: (1) Other, (2) Bus Stop, (3) Bus (NYC

Transit), (4) Taxi (Livery Licensed), (5) Public Building, and (6) Gas Station. All these choices are related to Street: bus, bus stops, taxis and gas station are all on streets, and public building also makes sense as crimes might occur in the buildings on prominent crime ladden, busy streets.

The next query finds most similar borough to a bucketed time interval. As discussed in Section 5, before training the database embedding model, the time value for each complaint is bucketed into 4 intervals: **Morning** (6 am-noon), **Afternoon** (noon-6 pm), **Evening** (6 pm-midnight), and **Night** (midnight-6 am). Figure 7 presents the results of the similarity query as a heatmap. The result matches the anecdotal evidence that in the morning most complaints arise from commuting boroughs like Queens, Brooklyn, and Staten Island; Manhattan gets most complaints in the working hours; Staten Island is the safest borough in the evening and night intervals; and The Bronx and Brooklyn generate most complaints at night.

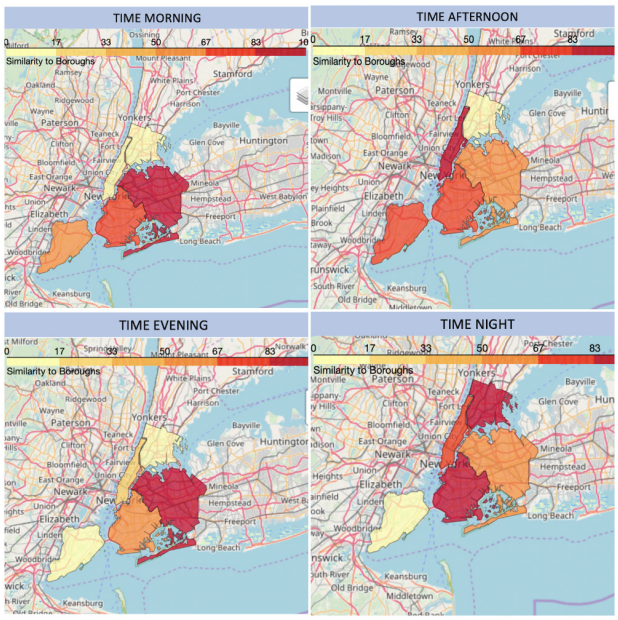

Figure 7: Boroughs Most Similar to Time Intervals

Similar to the query which compares time and precinct, one can also explore meaningful relationship between time and premise. It would be advantageous for the police to determine which premise there is likely to be a crime at certain time period in a day. The results from the query of finding most similar premise to a bucketed time interval can be seen in Figure 8. We observe that the results depict some of the usual hot spots of premises suitable for that time (e.g., schools in the morning, stores in the afternoon, etc.) If we look at time evening, we see that the top 2 results are parking areas. This is quite natural as people would go to pick up their cars after work in the parking lots, increasing the chances of crime. Another interesting premise is bar and night clubs which perfectly matches with the time people visit bars. Finally, at night most crimes are reported from public transportation system (e.g., buses and other facilities), and public building.

6.3 Dissimilarity Queries

TIME MORNING(6AM – 12PM)	TIME AFTERNOON(12PM – 6PM)
PUBLIC SCHOOL	CHURCH
DOCTOR DENTIST OFFICE	SMALL MERCHANT
CHURCH	GYM FITNESS FACILITY
COMMERCIAL BUILDING	FOOD SUPERMARKET
PAROCHIAL SCHOOL	STORE UNCLASSIFIED
TIME EVENING(6PM – 12AM)	**TIME NIGHT(12AM – 6AM)**
PARKING LOT GARAGE (PRIVATE)	TRANSIT FACILITY (OTHER)
PARKING LOT GARAGE (PUBLIC)	OTHER
BAR AND NIGHT CLUB	OTHER HOUSE OF WORSHIP
OPEN AREAS (OPEN LOTS)	PUBLIC BUILDING
TAXI (LIVERY LICENSED)	BUS (NYC TRANSIT)

Figure 8: Sorted List of Premises Most Similar to Time Intervals

Another novelty of the AI-DB approach is the ability to find entities dissimilar to a given entity. Dissimilarity is an interesting concept; because just like similarity one value being dissimilar to another is a holistic notion which encompasses the influence of multiple attributes. A dissimilarity CI query is a very minor modification of the similarity query shown in Figure 10; instead of checking for similarity score closer to 1.0, we check for similarity score closer to -1.0, and return the results in the ascending order. Statistical dissimilarity compares different entities based on their frequency and purely using statistical measures. As such, one can define dissimilarity only for specific attributes; it can not capture collective meanings based on the entirety of attributes.

To illustrate this capability, let's look at an example query to find dissimilar precincts. According to a raw count of complaints, precinct 75 has the most number of crime complaints. From Figure 9, one can see the red-colored precinct is 75 which serves the easternmost portion of Brooklyn and envelopes East New York and Cypress Hills. Using the AI-DB CI query, the top 10 most dissimilar precincts to precinct 75 can be seen from Figure 11.

Precinct 22, which covers Central Park, is identified to be the most dissimilar precinct. It is followed by other precincts which are primarily in Manhattan, Staten Island and Queens. If we take a look at the precincts which have the least number of complaints (Figure 11), we can see that our result does more than just listing the dissimilarity in order of frequency; it can capture impact of individual attributes (e.g., crime types, offenses, etc.) on the precinct meaning.

6.4 Multi-attribute Similarity Queries

The multi-attribute query is an extension of the single inter-column (attribute) similarity query. In a multi-attribute query, given a value of an attribute, we aim to find values of the other relevant attributes that are most similar (dissimilar) to the input attribute value. Since such a query *cannot*

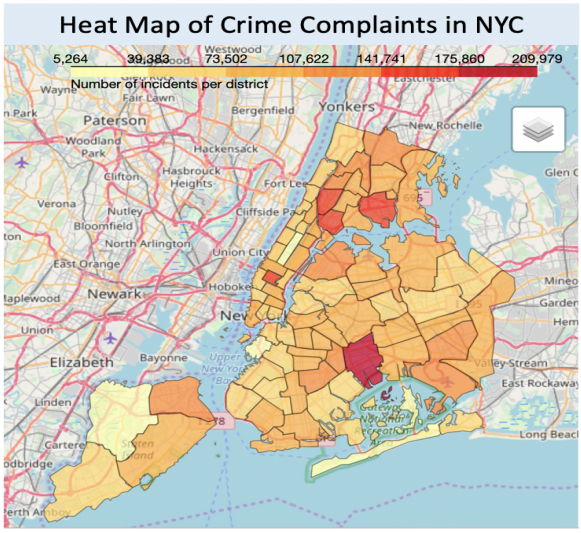

Figure 9: HeatMap of total crime complaints in NYC till '18

```
val result_premdis_df = spark.sql("""SELECT DISTINCT
X.ADDR_PCT_CD, precinctSimilarityUDF("75", X.ADDR_PCT_CD)
AS proximityValue
FROM Index_View X
WHERE precinctSimilarityUDF("75", X.ADDR_PCT_CD) > -1.0
ORDER BY proximityValue ASC """)
```

Figure 10: Example of a SQL CI dissimilarity query

be easily expressed in SQL, we use simple cosine-similarity based scores to find similar (dissimilar) entities.

As an example, consider a query that for **a given offense key,** KY_CD_235, finds most similar values for the target set of attributes: {TIME_NIGHT, ADDR_PCT_CD_25, BORO_NM_MANHATTAN, PREM_TYP_DESC_PARK_PLAYGROUND, OFNS_DESC_DANGEROUS_DRUGS, MONTH_6 }.

If we explore the output, it paints a believable story. The offense key 235 corresponds to the offense description of dangerous drugs. Precinct 25 encompasses the northern portion of East Harlem which has the Marcus Garvey Park, Harlem Art Park, and the 125th Street Metro-North Station. The most similar premise in the output is park or playground which relates to the precinct description. Month 6 corresponds to June where the weather starts getting hotter in summer and days are longer and temperatures start to cool in the night time which seems like a favorable climate to commit such a crime. The most similar time is night (closer to 2 am) which seems a likely time to commit such a crime. If we look at Figure 12, we can see that many people tweet that precinct 25 has drug related issues.

Another example of a multi-attribute query is as follows: **given the evening time**, find most similar values to the target set of attributes. We observe the following result: {BORO_NM_MANHATTAN, DAY_OF_WEEK_4, ADDR_PCT_CD_10, LAW_CAT_CD_FELONY, PREM_TYP_DESC_PARKING_LOT_GARAGE_(PRIVATE), OFNS_DESC_ROBBERY, MONTH_1}.

This result tries to tell us that, according to the past data in Manhattan, near Hudson yards and Chelsea area

Precincts ranked in the order of least number of Complaints			AI-DB Answer	
Precinct Number	**Precinct Description**	**Number of Complaints**	**Precinct Number**	**Precinct Description**
22	**Central Park**	5264	**22**	**Central Park**
100	Southern portion of Queens on the Rockaway Peninsula	33898	**17**	**East side of Midtown, Manhattan**
			123	**Portion of the South Shore of Staten Island**
121	**Northwestern shore of Staten Island**	36110	1	A square mile on the southernmost tip of Manhattan
76	South Brooklyn	36695		
111	Northeastern portion of Queens	40987	**121**	**Northwestern shore of Staten Island**
123	**Portion of the South Shore of Staten Island**	41694	14	Southern portion of Midtown, Manhattan
26	Upper West Side of the Manhattan	43189	**112**	**Centrally located portion of Queens**
			100	**Southern portion of Queens on the Rockaway Peninsula**
17	**East side of Midtown, Manhattan**	43461		
94	Northernmost portion of Brooklyn	44705	6	Southwestern Manhattan
112	**Centrally located portion of Queens**	46358	122	Portion of the South Shore of Staten Island

Figure 11: Top 10 precincts dissimilar to Precinct 75

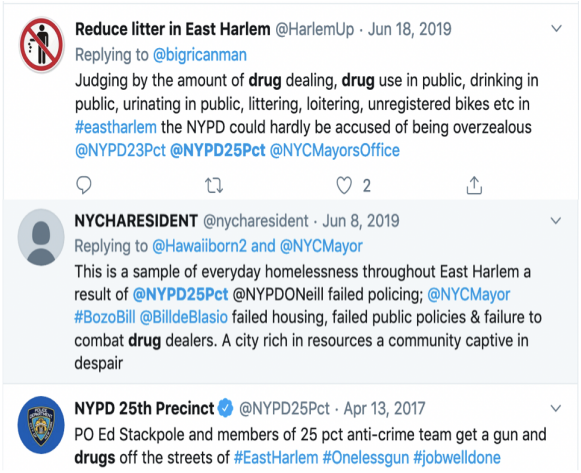

Figure 12: Twitter feed for precinct 25

(Precinct 10) on a Friday near parking lots in January, robberies are likely committed in the evening.

6.5 Inductive Reasoning Queries

A unique feature of embedding models such as word and database embeddings is their capability to answer inductive reasoning queries that enable an individual to reason from *part* to *whole* or from *particular* to *general* [16, 14]. Solutions to inductive reasoning queries exploit latent semantic structure in the trained model via algebraic operations on the corresponding vectors. AI-DB encapsulates these operations in UDFs to enable different types of inductive reasoning operations on relational entities using CI queries. In this section, we demonstrate application of two inductive reasoning queries, semantic clustering and analogy, to extract novel insights from the NYPD complaints dataset.

6.5.1 Semantic Clustering Queries

Given a set of input entities, {X, Y, Z, ..}, the semantic clustering process identifies a set of output entities, {W, ..}, that share the most dominant trait with the input entities. Figure 13 presents the corresponding CI query that returns a set of complaints that share the dominant traits with the input complaint set, {387152338, 259754425, 752017184}. Results of this query are presented in Figure 14. The `semanticCluste` first computes a centroid vector from the meaning vectors of the input entity set, and then finds those entities whose cosine similarity with the centroid vector is greater than some bound (0.5).

```
val result_boro_cluster = spark.sql("""SELECT
X.CMPLNT_NUM,
semanticClusterUDF("387152338", "259754425",
"752017184", X.CMPLNT_NUM)
AS proximityValue
FROM Index_View X
WHERE semanticClusterUDF("387152338", "259754425",
"752017184", X.CMPLNT_NUM) > 0.5
ORDER BY proximityValue DESC """)
```

Figure 13: Example of a SQL CI semantic clustering query: Find a complaint number showing the same traits as input complaint numbers

From Figure 14, we can see that all of the source complaints have the offense descriptions related to financial crimes involving a compromise in bank or credit card information. This dominant trait in the chosen seed complaints is present in all the resultant complaints. Additionally the boroughs and time periods are also in a similar range to that of the seed complaints.

Complaint Number	Precinct & Borough	Offense Description	Date	Time	Premise
Source Complaints					
387152338	106, QUEENS	LARCENY,GRAND BY ACQUIRING LOST CREDIT CARD	09/16/2018	19:00:00	RESIDENCE-HOUSE
259754425	14, MANHATTAN	LARCENY,GRAND BY OPEN CREDIT CARD (NEW ACCT)	09/22/2018	08:00:00	CHAIN STORE
752017184	110, QUEENS	LARCENY,GRAND BY BANK ACCT COMPROMISE-TELLER	07/09/2018	08:00:00	RESIDENCE - APT. HOUSE
Semantic Clustering Result					
431328491	110, QUEENS	LARCENY,GRAND BY BANK ACCT COMPROMISE-ATM TRANSACTION	09/07/2018	08:00:00	RESIDENCE - APT. HOUSE
506201381	106, QUEENS	LARCENY,GRAND BY OPEN CREDIT CARD (NEW ACCT)	04/28/2018	08:00:00	RESIDENCE-HOUSE
918826290	115, QUEENS	LARCENY,GRAND BY ACQUIRING LOST CREDIT CARD	04/15/2018	08:00:00	RESIDENCE-HOUSE
482110137	14, MANHATTAN	LARCENY,GRAND BY OPEN CREDIT CARD (NEW ACCT)	01/23/2018	08:00:00	RESIDENCE - APT. HOUSE
327843091	110, QUEENS	LARCENY,GRAND BY OPEN CREDIT CARD (NEW ACCT)	07/03/2018	09:00:00	RESIDENCE - APT. HOUSE

Figure 14: Semantic Clustering query over a set of complaints

As a second experiment, we used a query similar to the query shown in Figure 13 for complaints related to vehicle larceny offenses, we get the result shown in Figure 15. We can observe that all source complaints have the offense descriptions involving some kind of vehicle larceny or unauthorized use. Along with that, all the seed crime complaints occur on the street. These are the dominant traits in those complaints. Similar to the results in Figure 14, all the result complaints share the offense description: unauthorized use of vehicle, specifically bicycle or motorcycle. The premise also matches with the seed complaints. Additionally the boroughs are also in a similar range to that of the seed complaints.

Complaint Number	Precinct & Borough	Offense Description	Date	Time	Premise
Source Complaints					
717362228	48, BRONX	LARCENY,GRAND FROM VEHICLE/MOTORCYCLE	09/28/2018	14:08:00	STREET
666090055	84, BROOKLYN	LARCENY,GRAND OF BICYCLE	10/06/2018	20:30:00	STREET
955211258	115, QUEENS	UNAUTHORIZED USE VEHICLE	10/06/2018	09:30:00	STREET
Semantic Clustering Result					
953756172	84, BROOKLYN	LARCENY,GRAND FROM VEHICLE/MOTORCYCLE	07/19/2018	20:15:00	STREET
985597988	94, BROOKLYN	LARCENY,GRAND FROM VEHICLE/MOTORCYCLE	10/06/2018	20:30:00	STREET
144886198	84, BROOKLYN	LARCENY,GRAND OF BICYCLE	10/06/2018	10:30:00	STREET
610038721	109, QUEENS	LARCENY,GRAND FROM VEHICLE/MOTORCYCLE	10/06/2018	18:00:00	STREET
182793476	115, QUEENS	LARCENY,GRAND FROM VEHICLE/MOTORCYCLE	10/01/2018	22:30:00	STREET

Figure 15: Semantic Clustering query over a set of complaints

6.5.2 *Analogy Queries*

Wikipedia defines analogy as a process of transferring information or meaning from one subject to another. A common way of expressing an analogy is to use relationship between a pair of entities, `source_1` and `target_1`, to reason about a possible target entity, `target_2`, associated with another known source entity, `source_2`. The CI SQL query for analogies is similar to the semantic clustering query (Figure 13), where the UDF executes vector operations to extract the possible answers.

For the NYPD dataset, we can encode the following relationship, **Manhattan:14 :: Bronx:?**, as a CI query (i.e., identify a precinct in the Bronx that has the same behavior as precinct 14 in Manhattan). The top analogous precinct returned by the CI query is precinct 43. It shares certain commonalities with the precinct 14 which are presented in Figure 16. Specifically, both precincts have the largest number of complaints in the borough; they share the most common offense and premise (petit larceny and street), and crimes are reported mostly at evening and night. Precinct 14 is the southern precinct of Mid-town and serves commercial areas like Grand Central, Times Square, Penn Station, Madison Square Garden, Koreatown section, and the Manhattan Mall Plaza. Whereas Precinct 43 is also in the southeast section of the Bronx. It has four primary commercial strips: Westchester Avenue, Castle Hill Avenue, White Plains Road, and Parkchester. We can see that both are commercial areas in similar locations within each borough.

Since the meaning vector of a complaint captures the holistic meaning of the crime complaint based on contributions of different attributes in that row, an analogy between a borough and a complaint would provide much utility. If an analyst knows about a certain crime complaint that happened in the borough of Queens, and wants to find another complaint that has the same relation in the Bronx, the analogy would look like this: **Queens:586406852 :: Bronx:?**

Precinct	Number of Complaints	Most Common Offense	Most Common Premise	Top 2 times of Max Complaints
14	Max in Manhattan	Petit Larceny	Street	EVENING NIGHT
43	Max in Bronx	Petit Larceny	Street	NIGHT EVENING

Figure 16: Analogy of precincts in different boroughs of NYC

Complaint Number	Complaint Date & Time	Precinct & Borough	Offense Description	Premise	Suspect Age Group	Victim Age Group
586406852	04/11/2018 09:25:00	103, Queens	Robbery, Personal Electronic Device	Street	18-24	18-24
692937181	01/23/2018 10:00:00	42, Bronx	Robbery, Personal Electronic Device	Street	18-24	18-24

Figure 17: Analogy for finding similar complaints in different boroughs of NYC

The top analogous complaint number returned by the CI query is: 692937181. From Figure 17, we can see that complaint number 692937181 shares a lot of characteristics with complaint number 586406852. These complaints both have the same offense, occur during similar time periods and at the same premise. In addition to that the age groups of the suspect and victim both have the same range.

7. CONCLUSIONS AND NEXT STEPS

Automated crime analysis is garnering huge importance in recent times to undertake crime prevention strategies for ensuring public safety. AI-Powered Database (AI-DB) serves as a unique solution to enhance efficient policing and aid in decision making process of crime analysts. The application of AI-DB to a real crime dataset generates insights which are either not easily made available or difficult and time consuming to implement by traditional machine learning, SQL analytics or statistical approaches. The different queries demonstrate capabilities of identifying patterns among the same type of crimes (Patternizr can accomplish), patterns across crime types(shortcomings of Patternizr), finding similarity among varied data types and other more semantic queries like analogies, clustering and multi-attribute similarity. The end user is not bothered with subtleties associated with exploiting the neural network models: the only interface to the database embedding model is via standard SQL.

As the next step, we would like to perform extensive user studies to validate our insights. We also plan to enhance AI-DB capabilities for the crime analysis by utilizing external data sources and generating vector embeddings for them as well. We also would like to evaluate methods of explainability of the results based on certain statistical information of database tokens.

8. REFERENCES

[1] R. Bordawekar. Cognitive Database: An Apache Spark-Based AI-Enabled Relational Database System. Spark+AI Summit 2018, June 2018.

[2] R. Bordawekar, B. Bandyopadhyay, and O. Shmueli. Cognitive database: A step towards endowing relational databases with artificial intelligence capabilities. *CoRR*, abs/1712.07199, December 2017.

[3] R. Bordawekar and O. Shmueli. Using word embedding to enable semantic queries in relational databases. In *Proceedings of the 1st Workshop on Data Management for End-to-End Machine Learning*, DEEM'17, pages 5:1–5:4, New York, NY, USA, 2017. ACM.

[4] A. Chohlas-Wood and E. S. Levine. A recommendation engine to aid in identifying crime patterns. *INFORMS Journal on Applied Analytics*, 49(2):154–166, 2019.

[5] A. Chohlas-Wood, A. Merali, W. Reed, and T. Damoulas. Mining 911 calls in new york city: Temporal patterns, detection, and forecasting. In *AAAI Workshop: AI for Cities*, 2015.

[6] L. Elluri, V. Mandalapu, and N. Roy. Developing machine learning based predictive models for smart policing. In *2019 IEEE International Conference on Smart Computing (SMARTCOMP)*, pages 198–204, June 2019.

[7] P. A. Haskins. Research will shape the future of proactive policing. *NIJ Journal 281*, November 2019.

[8] C. Kadar, J. Iria, and I. Pletikosa. Exploring foursquare-derived features for crime prediction in new york city. In *The 5th international workshop on urban computing (UrbComp 2016)*, 2016.

[9] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781, 2013.

[10] New York City. New york city open data portal. https://opendata.cityofnewyork.us/, 2020.

[11] New York City/NYPD. New york police department complaint data historic. https://data.cityofnewyork.us/Public-Safety/NYPD-Complaint-Data-Historic/qgea-i56i, 2016.

[12] N. I. of Justice:nij.ojp.gov. Overview of predictive policing, June 9 2014.

[13] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014*, pages 1532–1543, 2014.

[14] D. E. Rumelhart and A. A. Abrahamson. A model for analogical reasoning. *Cognitive Psychology*, 5(1):1 – 28, 1973.

[15] A. Stec and D. Klabjan. Forecasting crime with deep learning, 2018.

[16] R. J. Sternberg and M. K. Gardner. Unities in inductive reasoning. Technical Report Technical rept. no. 18, 1 Jul-30 Sep 79, Yale University, 1979.

[17] T. Wang, C. Rudin, D. Wagner, and R. Sevieri. Learning to detect patterns of crime. *Machine Learning and Knowledge Discovery in Databases. ECML PKDD 2013.*, August 2013.

APPENDIX

A. NYPD CRIME COMPLAINT DATA OVERVIEW

The raw data has a total of 35 columns and their detailed description is as follows:

Column Name	(Type) Description	Column Name	(Type) Description
CMPLNT_NUM	(Number) Randomly generated persistent ID for each complaint	LAW_CAT_CD	(Plain Text) Level of offense: felony, misdemeanor, violation
ADDR_PCT_CD	(Number) The precinct in which the incident occurred	LOC_OF_OCCUR_DESC	(Plain Text) Specific location of occurrence in or around the premises; inside, opposite of, front of, rear of
BORO_NM	(Plain Text) The name of the borough in which the incident occurred	OFNS_DESC	(Plain Text) Description of offense corresponding with key code
CMPLNT_FR_DT	(Date & Time) Exact date of occurrence for the reported event (or starting date of occurrence, if CMPLNT_TO_DT exists)	PARKS_NM	(Plain Text) Name of NYC park, playground or greenspace of occurrence, if applicable (state parks are not included)
CMPLNT_FR_TM	(Plain Text) Exact time of occurrence for the reported event (or starting time of occurrence, if CMPLNT_TO_TM exists)	PATROL_BORO	(Plain Text) The name of the patrol borough in which the incident occurred
CMPLNT_TO_DT	(Date & Time) Ending date of occurrence for the reported event, if exact time of occurrence is unknown	PD_CD	(Number) Three digit internal classification code (more granular than Key Code)
CMPLNT_TO_TM	(Plain Text) Ending time of occurrence for the reported event, if exact time of occurrence is unknown	PD_DESC	(Plain Text) Description of internal classification corresponding with PD code (more granular than Offense Description)
CRM_ATPT_CPTD_CD	(Plain Text) Indicator of whether crime was successfully completed or attempted, but failed or was interrupted prematurely	PREM_TYP_DESC	Specific description of premises; grocery store, residence, street, etc.
HADEVELOPT	(Plain Text) Name of NYCHA housing development of occurrence, if applicable	RPT_DT	(Date & Time) Date event was reported to police
HOUSING_PSA	(Plain Text) Development Level Code	STATION_NAME	(Plain Text) Transit station name
JURISDICTION_CODE	(Number) Jurisdiction responsible for incident. Either internal, like Police(0), Transit(1), and Housing(2); or external(3), like Correction, Port Authority, etc.	SUSP_AGE_GROUP	(Plain Text) Suspect's Age Group
JURIS_DESC	(Plain Text) Description of the jurisdiction code	SUSP_RACE	(Plain Text) Suspect's Race Description
TRANSIT_DISTRICT	(Number) Transit district in which the offense occurred	SUSP_SEX	(Plain Text) Suspect's Sex Description
VIC_AGE_GROUP	(Plain Text) Victim's Age Group	X_COORD_CD	(Number) X-coordinate for New York State Plane Coordinate System
VIC_AGE	(Plain Text) Victim's Race Description	Y_COORD_CD	(Number) Y-coordinate for New York State Plane Coordinate System
VIC_SEX	(Plain Text) Victim's Sex Description	LATITUDE	(Number) Latitude coordinate for Global Coordinate System
KY_CD	(Number) Three digit offense classification code	LONGITUDE	(Number) Longitude coordinate for Global Co- ordinate System
		LAT_LON	(Location) Geospatial location point (latitude and longitude combined)

Figure 18: NYPD Historic Complaints Dataset Schema

Among the 35 columns, the following columns were dropped from the dataset before the model was trained:

- SUSP_RACE, SUSP_SEX, VIC_RACE, VIC_SEX: Sensitive attributes containing Suspect and Victim gender and race information were removed
- CMPLNT_TO_TM, CMPLNT_TO_DT : We do not need the ending time and date of the crime
- RPT_DT : The date when the crime was reported would not be helpful in finding similar crimes or aid in clustering
- PD_CD : Encodes the PD_DESC in the form of a numerical code. Redundant as PD_DESC is kept
- PARKS_NM : This column provides information if the location was a park. Park names are not important for our analysis and this column has 99.67% values as NA
- HADEVELOPT: Name of the housing complex would not add any value in analysis as this column has 95.02% values as NA
- HOUSING_PSA: This column like HADEVELOPT provides information about development codes and has 92.26% NA values making it not that useful for crime analysis
- JURISDICTION_CODE: The jurisdiction code provides a numerical value for the values in the column JURIS_DESC, hence was dropped
- TRANSIT_DISTRICT: This attribute states if there was a transit district when the crime occurred and there being a few cases like this has 97.78% of NA values
- STATION_NAME: Like the above columns, this column also has a large percentage of NA values (97.78%) hence is dropped as it adds no value
- X_COORD_CD, Y_COORD_CD, Latitude, Longitude, Lat_Lon : These columns are repetitive and redundant given that we have the spatial information captured by precincts

The columns used for model training include:

- CMPLNT_NUM: This column provides the an ID associated with each crime complaint. In the model building process the vector for this columns the overall behavior of a crime complaint.
- ADDR_PCT_CD: This column has 77 unique values corresponding to the different precincts in NYC.
- BORO_NM: This column has the names of the boroughs in which the crime occurred. The five borough names are: Manhattan, Bronx, Brooklyn, Queens and Staten Island
- CMPLNT_FR_DT: This column provides the date of when the crime took place. Other attributes like: Year, Month, Day of the week are extracted from the values in this column
- CMPLNT_FR_TM: This column provides the time of occurrence of the crime. The time is converted into hours and then bucketed into larger bins.
- CRM_ATPT_CPTD_CD: This column provides information about whether a crime was completed or attempted.
- JURIS_DESC: Provides the description of the responsible jurisdiction with 25 unique names and no missing values
- KY_CD: This column provides 74 offense classification codes with no missing values
- LAW_CAT_CD: This column provides the high level classification of crime as either Violation, Misdemeanor or Felony
- LOC_OF_OCCUR_DESC: This column provides more spatial information of where in the premise a crime occurred
- OFNS_DESC: This column provides the offense description and has 71 unique values
- PATROL_BORO: This column states the patrol borough in which the incident occurred and has 9 unique values. Manhattan, Queens and Brooklyn has 2 patrol boroughs for the North and South section respectively whereas Bronx and Staten Island have a single patrol borough
- PD_DESC: It provides a more granular description of the offense and is more close to natural language sentence
- PREM_TYP_DESC: This column provides information about the premise of the incident and has 72 different types of premises
- SUSP_AGE_GROUP, VIC_AGE_GROUP: This column represents binned age groups in the range: UNKNOWN, ¡18, 18-24, 25-44, 45-64, 65+

B. VISUAL CHARACTERIZATION OF THE NYPD CRIMINAL COMPLAINTS DATA

To gain better understanding of the dataset we decided to use data visualization techniques and spot some interesting trends.

B.1 Has NYC become safer?- Crime rates by year

Figure 19 shows the trend of total number of complaints filed across 12 years.

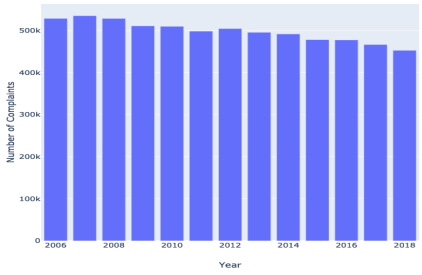

Figure 19: Number of crime complaints vs Year

We find that the number of complaints filed has seems to decrease overtime. However, when looking from a statistical standpoint, see Figure 20, complaints have been some-what steady over the years. Though there is a decreasing trend in the actual number of complaints over the years, there doesn't seem to be a drastic difference. It is seen as the mean and the rolling mean of the number of complaints each year is almost the same. Also the standard deviation(green line)is uniformly close to zero.

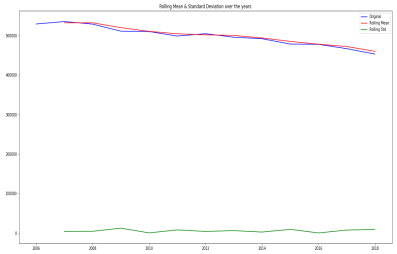

Figure 20: Rolling Mean and Standard Deviation of the Number of Crime Complaints over the years

B.2 What are the most common offenses among Felonies?

Figure 21 shows the most common offenses categorized as felonies. Grand Larcenies contain almost 43 different types of crimes which are termed as grand larcenies. The domain includes finance frauds, vehicle larcenies, electronic issues, shoplifting and more.

B.3 Are particular premises unsafe?

Figure 22 shows that almost the majority of complaints are recorded to have happened on the Street followed by residences(which includes apartment/house/public housing).

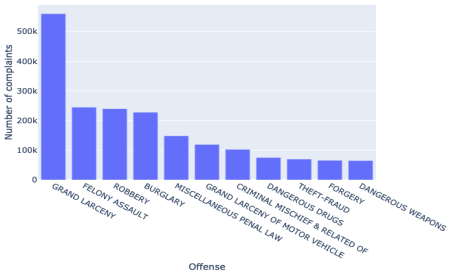

Figure 21: Most frequent offenses within Felony Crime Category

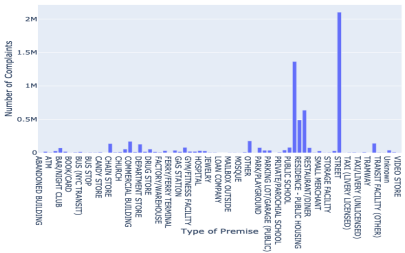

Figure 22: Number of complaints in different premises

B.4 Is a particular day of the week unsafe?

Figure 23 indicates that most crimes historically occur on a Friday.

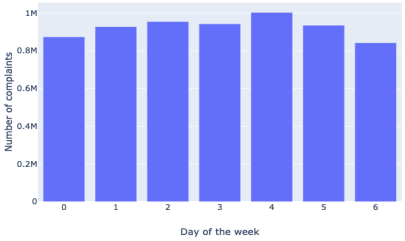

Figure 23: Number of crime complaints vs Day of the Week

B.5 Is a particular crime type more prevalent at a specific time?

From Figure 24 we can see that the most frequent time for misdemeanors is 6 pm, felony is 1 am and violation is 3 pm. Overall majority of historical crimes seem to have occurred in the evening and night times (after 4 pm to 2 am).

B.6 Are crime types similarly distributed in all boroughs?

Figure 25 indicates that the distribution of major crime types is similar in the 5 different boroughs. Misdemeanors

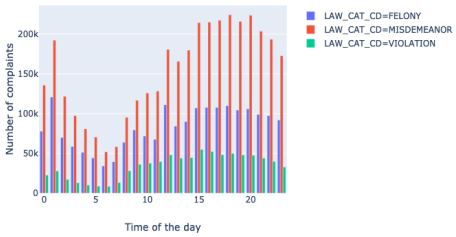

Figure 24: Number of crime complaints vs Time of the Day

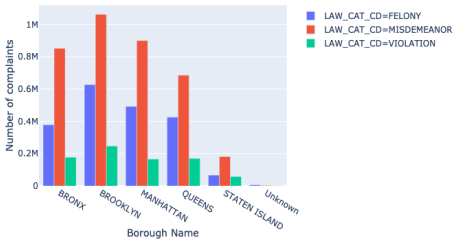

Figure 25: Distribution of crime types in different boroughs

always the most frequent, followed by Felonies and then Violations. Brooklyn has the highest number of crime complaints, followed by Manhattan then Queens, Bronx and the least being in Staten Island.

B.7 Query Execution

Once the vectors were trained, we formulate SQL queries in Scala and execute them in the Spark environment using the spark-shell.

The following Figure 26 is a screen-shot of the query execution to find premises similar to a given premise (street).

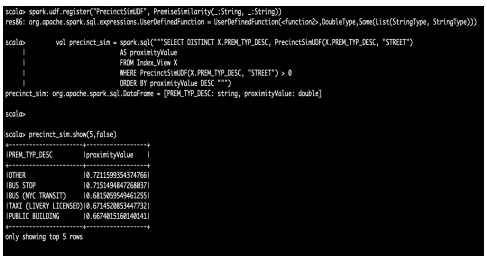

Figure 26: Execution of a Spark SQL Query