Predicting Fine-Grained Crime Types using Social Media

Deepak Venugopal
Assistant professor
Department of Computer Science
University of Memphis

Partying in downtown!

Depressed about grades

Demonstration on 5th street

I Hate XXXxxx

School party

Afraid to go to school

Noisy neighbors can't concentrate!!

Partying in downtown!

Depressed about grades

Demonstration on 5th street

I Hate XXXxxx

School party

Afraid to go to school

Noisy neighbors can't concentrate!!

Signals of potential crime events

Partying in downtown!

Depressed about grades

Demonstration on 5th street

I Hate XXXxxx

School party

Afraid to go to school

Noisy neighbors can’t concentrate!!

Signals of potential crime events

Research Challenge: Can we predict fine-grained, potential criminal activity in real-time from tweets?

Twitter-based Model

- Collect a tweet corpus offline and annotate it with the following categories
  - Violent Crime
  - Racism
  - Fraud
  - Narcotics

Twitter-based Model

- Collect a tweet corpus offline and annotate it with the following categories
  - Violent Crime
  - Racism
  - Fraud
  - Narcotics
- Pre-process tweet text (lots of noise)

Twitter-based Model

- Collect a tweet corpus offline and annotate it with the following categories
  - Violent Crime
  - Racism
  - Fraud
  - Narcotics
- Pre-process tweet text (lots of noise)
- Learn a Machine learning model from features obtained from processed tweets to distinguish between crime categories
- Capture live tweets and predict crime categories based on learned model

Implementation

- Implemented on top of Amazon cloud services
- Spark streaming to process thousands of tweets quickly in a cluster environment

Implementation

- Implemented on top of Amazon cloud services
- Spark streaming to process thousands of tweets quickly in a cluster environment
- Use a filter to look only for tweets that suggest potential criminal activity
- Detailed visualization dashboards to explain prediction results in real-time

Challenges

- Twitter data is very noisy

Challenges

- Twitter data is very noisy
- Twitter allows users access to a miniscule percent of their tweets

Challenges

- Twitter data is very noisy
- Twitter allows users access to a miniscule percent of their tweets
- Among these tweets, those relevant to crime are even fewer

Challenges

- Twitter data is very noisy
- Twitter allows users access to a miniscule percent of their tweets
- Among these tweets, those relevant to crime are even fewer
- Among the crime-relevant tweets, very few are specific to a geo-location
  - Geo-tags are generally empty in most tweets
- In short, very hard to get reliable signals for specific geo-locations from twitter content alone

Can we use more reliable data?

- Utilize Memphis specific data from Bluecrush
  - Incidents recorded along with type of incident and date

Can we use more reliable data?

- Utilize Memphis specific data from Bluecrush
    - Incidents recorded along with type of incident and date
- Developed a Hidden Markov Model
    - Latent crime states
    - Observable crime events (7 types of crime)

Can we use more reliable data?

- Utilize Memphis specific data from Bluecrush
  - Incidents recorded along with type of incident and date
- Developed a Hidden Markov Model
  - Latent crime states
  - Observable crime events (7 types of crime)
- Predict future incidents from the HMM
  - Different granuarity levels (by zipcode, precinct zone, etc.)

HMM-based predictions

- Can be tailored to specific geo-locations
  - More reliable signals

HMM-based predictions

- Can be tailored to specific geo-locations
  - More reliable signals
- Not a real-time model
  - Incidents must be recorded and published
  - Need to update model with new incidents
- Augment the Twitter model
  - Combine the models systematically (ongoing work)

Results

- 5000 tweet training corpus collected and annotated
- Several Machine learning methods evaluated
  - Support Vector Machines with Linear Kernels work best
  - Around 88% weighted F1-score using 5-fold cross validation
- Scalable to process thousands of tweets in real-time
- Real-time visualizations of predictions

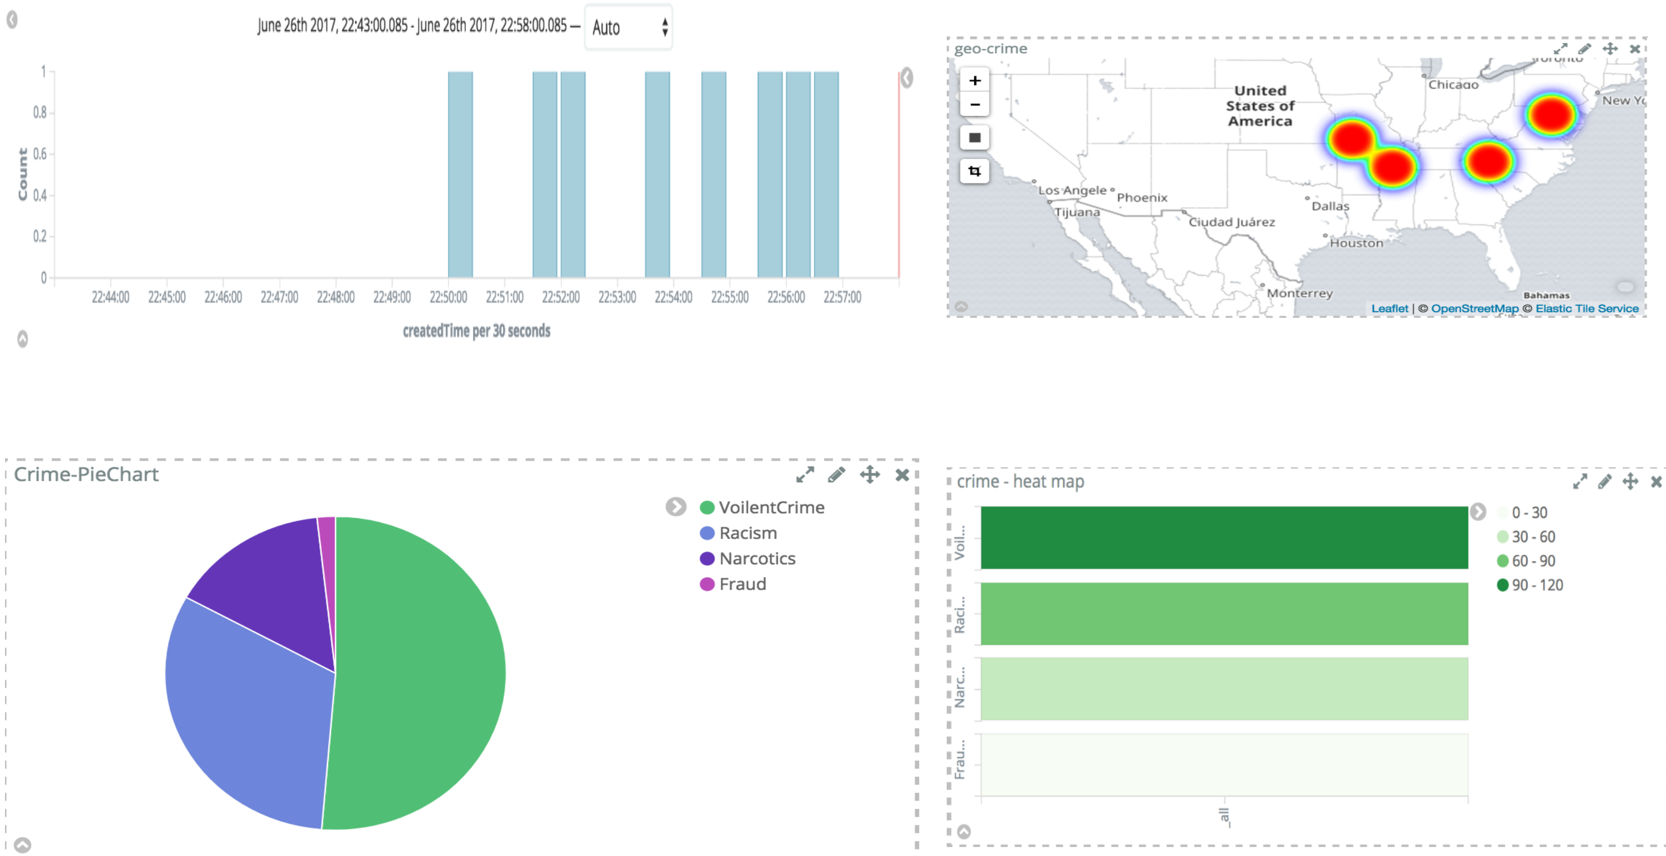

Example visualization dashboards (updated in real-time)

HMM Model

- Collected data recorded over a two month period
- Predict specific crime types for the last 10 days in the collected period
- Error rate of around 12 – 15% on average
  - Shows promise in modeling based on recorded instances
- Granularity affects performance
  - Merge/Split incidents based on zip-codes
  - Split by wards/precinct-zones, etc.

Future work

- Combine the HMM and Twitter model
  - The twitter model may not be particularly sensitive to geo-location
  - The HMM model is much more specific
- Advanced linguistic features from neural embeddings using non-annotated data
  - The duck tacos were delicious and the chocolate fondant with jalapeños and lime cream **to die for**
  - I made a **killing** from the Apple stocks today!!!
  - Whoever said head **wounds bleed** the most never **skinned** his shin while shaving!!
- Look for more events in specific tweets
  - Follow specific local individuals or organizations

Acknowledgements

- Several students
  - Saichand Upputuri (MS Summer 2017)
  - Lindsey Warren (BS Fall 2017)
  - Chris Kent (ongoing)