Mapping Moods: Geo-Mapped Sentiment Analysis During Hurricane Sandy

Cornelia Caragea[1], Anna Squicciarini[2], Sam Stehle[2], Kishore Neppalli[1], Andrea Tapia[2]

[1]University of North Texas,
[2]Pennsylvania State University

ISCRAM 2014

Social Media

- Social media have become part of our daily lives and everyday communication patterns.
- Scholars of disasters see hope in social media, and argue that used around crises can produce accurate results, often in advance of official communications.

- However, responders are generally hesitant to using social media due to insecurity and apprehension concerning the connection between the location of the disaster and those tweeting about the disaster.

Social Media around the Sandy Hurricane

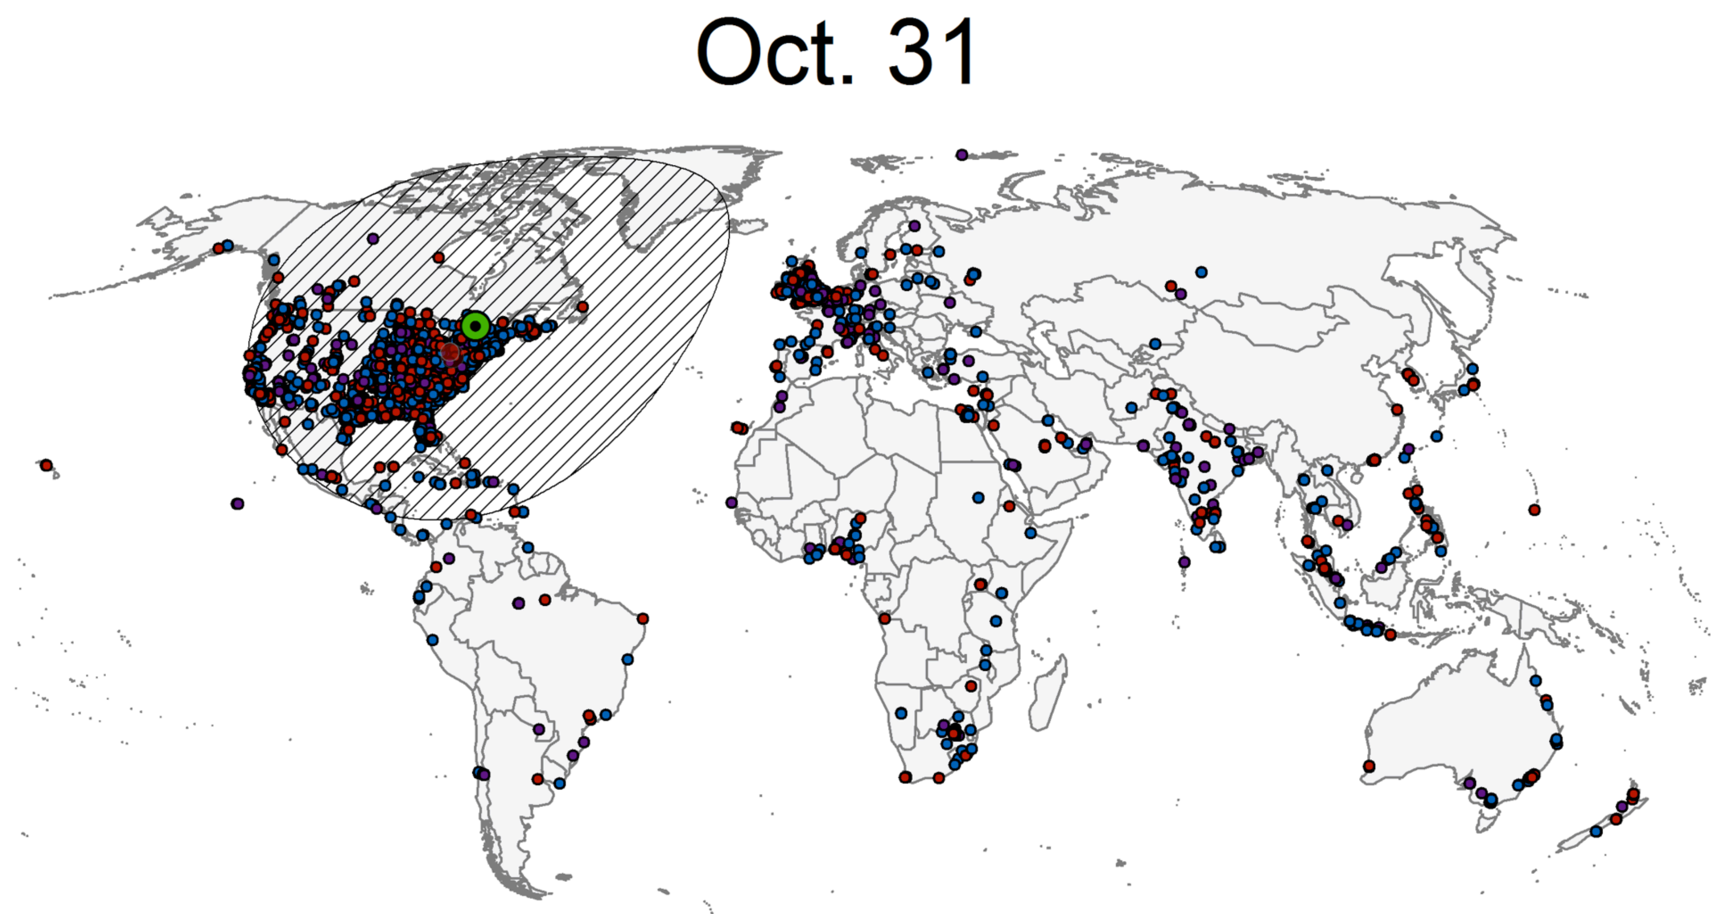

Proof of Concept

- *Using Twitter data from Hurricane Sandy, we identify the sentiment of tweets and then measure the distance of each categorized tweet from the epicenter of the hurricane.*

- *Extracting sentiments during a disaster could help responders develop stronger situational awareness of the disaster zone itself.*

Analyzing Sentiment in Disaster Events

- Can help understand the dynamics of the social network
  - The main users' concerns and panics
  - The emotional impacts of interactions among users.
- Can help obtain a holistic view about the general mood and the situation on the ground.

- Strong value to those experiencing the disaster and those seeking information about the disaster
  - However, there has not been much uptake of message data by large-scale, disaster response organizations.

Problem Statement

- We seek to find mechanisms to *automatically* classify the sentiment of users' Twitter posts during the Hurricane Sandy.
  - Supervised classification problem:
    - Classify a post (or tweet) into one of the following classes: *positive*, *negative* or *neutral*, based on the polarity of the emotion expressed in the tweet.

Tweet	Sentiment
1. *"RT @User1: During this hurricane we are all going to reunite on Xbox like the good ole days."*	Positive
2. *"RT @User2: It doesnt look like a hurricane is coming."*	Neutral
3. *"User3: I got a feeling that #Sandy is about to screw up my work schedule for the week :(smh"*	Negative

Challenges

- Sentiment classification of tweets faces many challenges:
  - Dealing with very short texts, e.g., a tweet is at most 140 characters in length
  - Dealing with unstructured text and noisy user input, e.g., tweets contain many misspellings, *“ole”* instead of *“old”*, or acronyms, *“smh”*

Problem Statement

- To detect sentiment of tweets, we propose to use a combination of *bag of word* and *sentiment features* such as emoticons, acronyms, and polarity clues.
- To understand the general mood during the Hurricane Sandy, we perform a geo-mapped sentiment analysis:
  - We first identify all geo-tagged tweets in our collection and label each of these tweets using our trained classifiers.
  - We then associate the sentiment of tweets with their geo-locations.

The Supervised Learning Problem

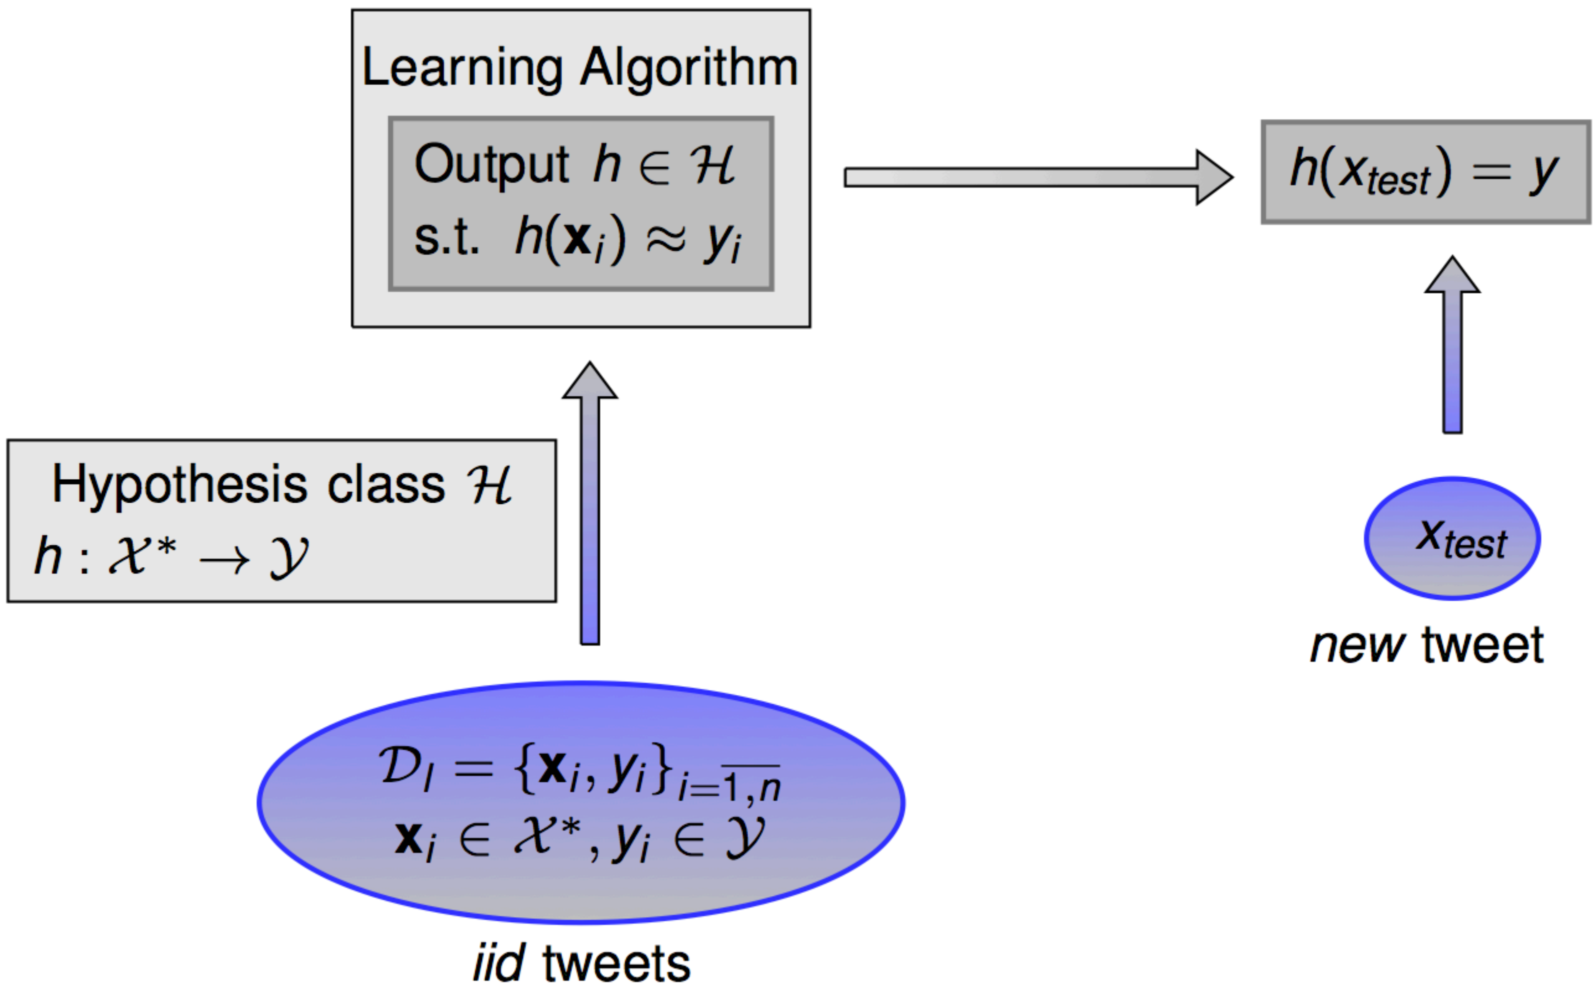

Features Used

- Unigrams: occurrences of words in a tweet
- Polarity Clues:
  - PosDensity: pos. words / num. of words in a tweet
  - NegDensity: neg. words / num. of words in a tweet
  - PosVsNegDensity: (PosDensity+1)/(NegDensity+1)
- Emoticons: ☺ ☹
- Internet Acronyms: *lol = laughing out loudly*
- Punctuation: “I hate this!” and “I hate this!!!!!!”
- SentiStrength positive and negative scores.

Sandy Twitter data

- We crawled 12,933,053 tweets between 10-26-2012 and 11-12-2012.
- Among these tweets, 4,818,318 have links to external sources, 6,095,524 are retweets and 622,664 contain emoticons.
- We manually labeled a subset of tweets (≈600) from the crawled data as positive, negative and neutral.

Experimental Design

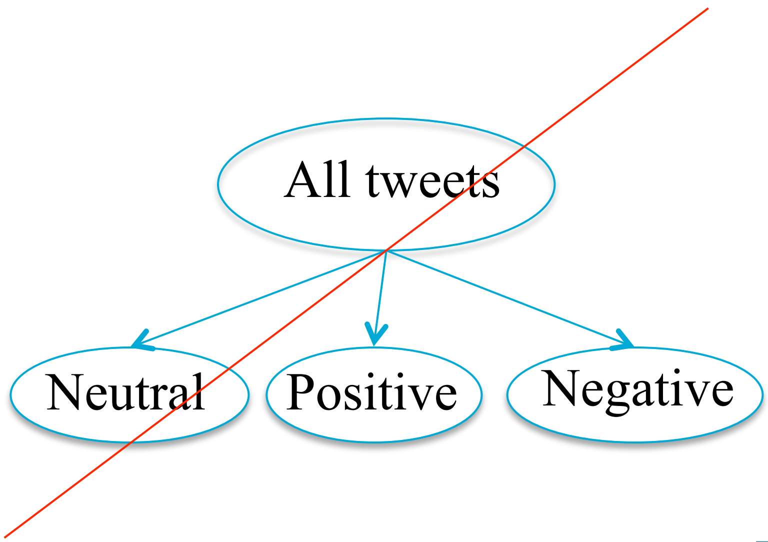

Not so good results!

SentiStrength for Neutral vs. Polar

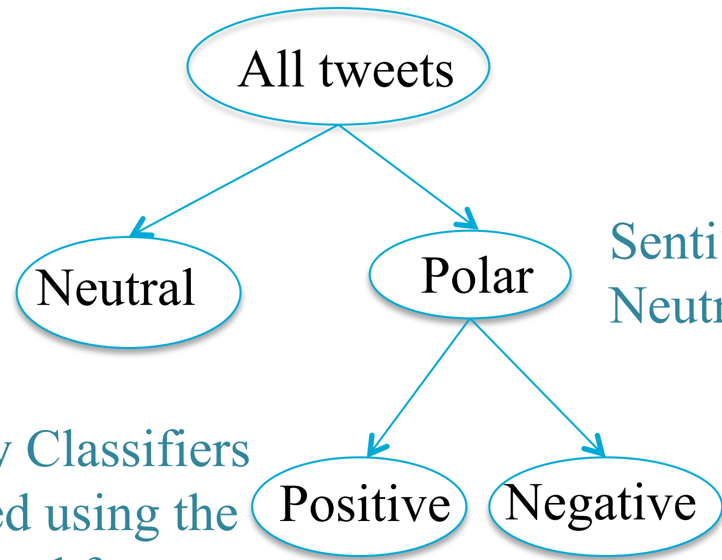

Binary Classifiers Trained using the Proposed features

Good results!

Polar vs. Neutral Classification

- SentiStrength returns two sentiment scores for a given short text:
  - A positive score ranging from 1 to 5
  - A negative score ranging from -5 to -1
  - A tweet with +1 or -1 scores is labeled as neutral; otherwise, it is labeled as polar.

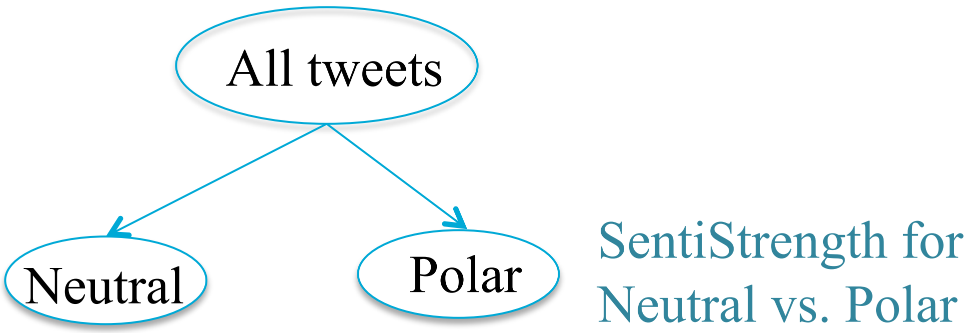

Positive vs. Negative Classification

- Two machine-learning classifiers:
  - Naïve Bayes (NB)
  - Support Vector Machine (SVM)
- We report the average classification accuracy obtained in 10-fold cross-validation experiments.

Binary Classifiers Trained using the Proposed features

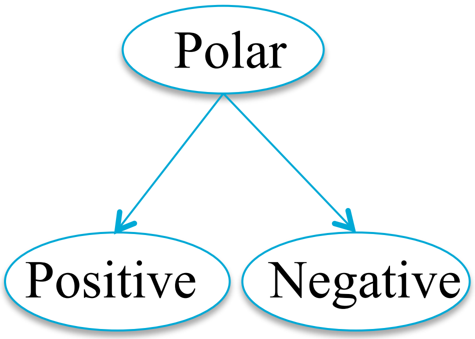

Results

Feature type	Naïve Bayes	SVM
Sentiment-based	68.60	67.95
Unigrams	71.82	72.25
Combination	73.33	**75.91**

Accuracy of Naïve Bayes and SVM for sentiment classification using various feature types.

- A SentiStrengths baseline that classifies tweets as positive vs. negative gives an accuracy of 59.13%.
- Using SVM on the combination of features, we made predictions for all unlabeled tweets with geo-location.

Geo-Tagged Tweets Sentiment Analysis

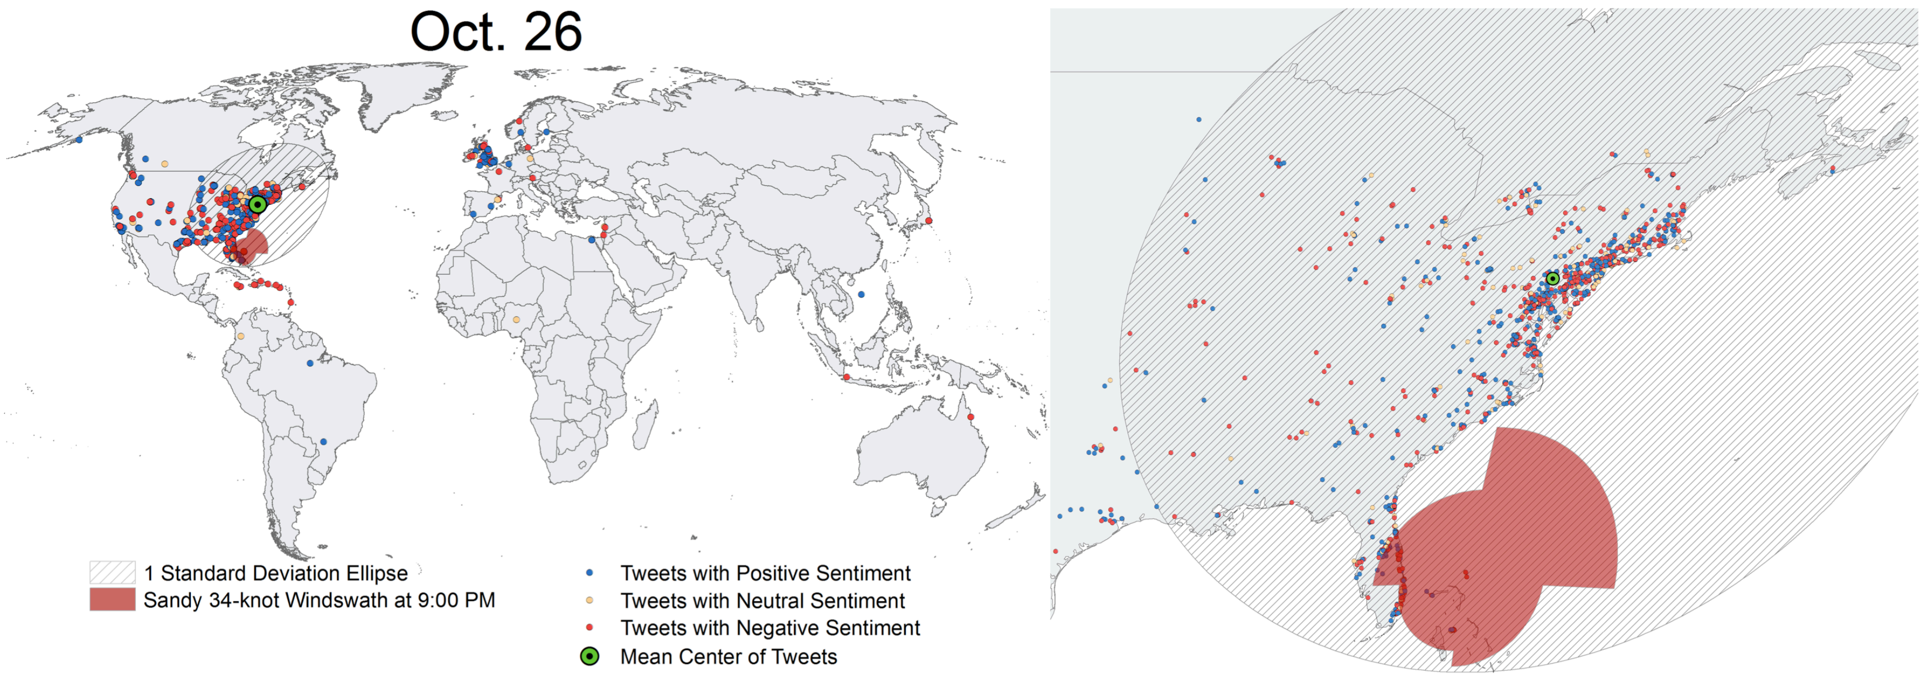

Geo-Tagged Tweets Sentiment Analysis

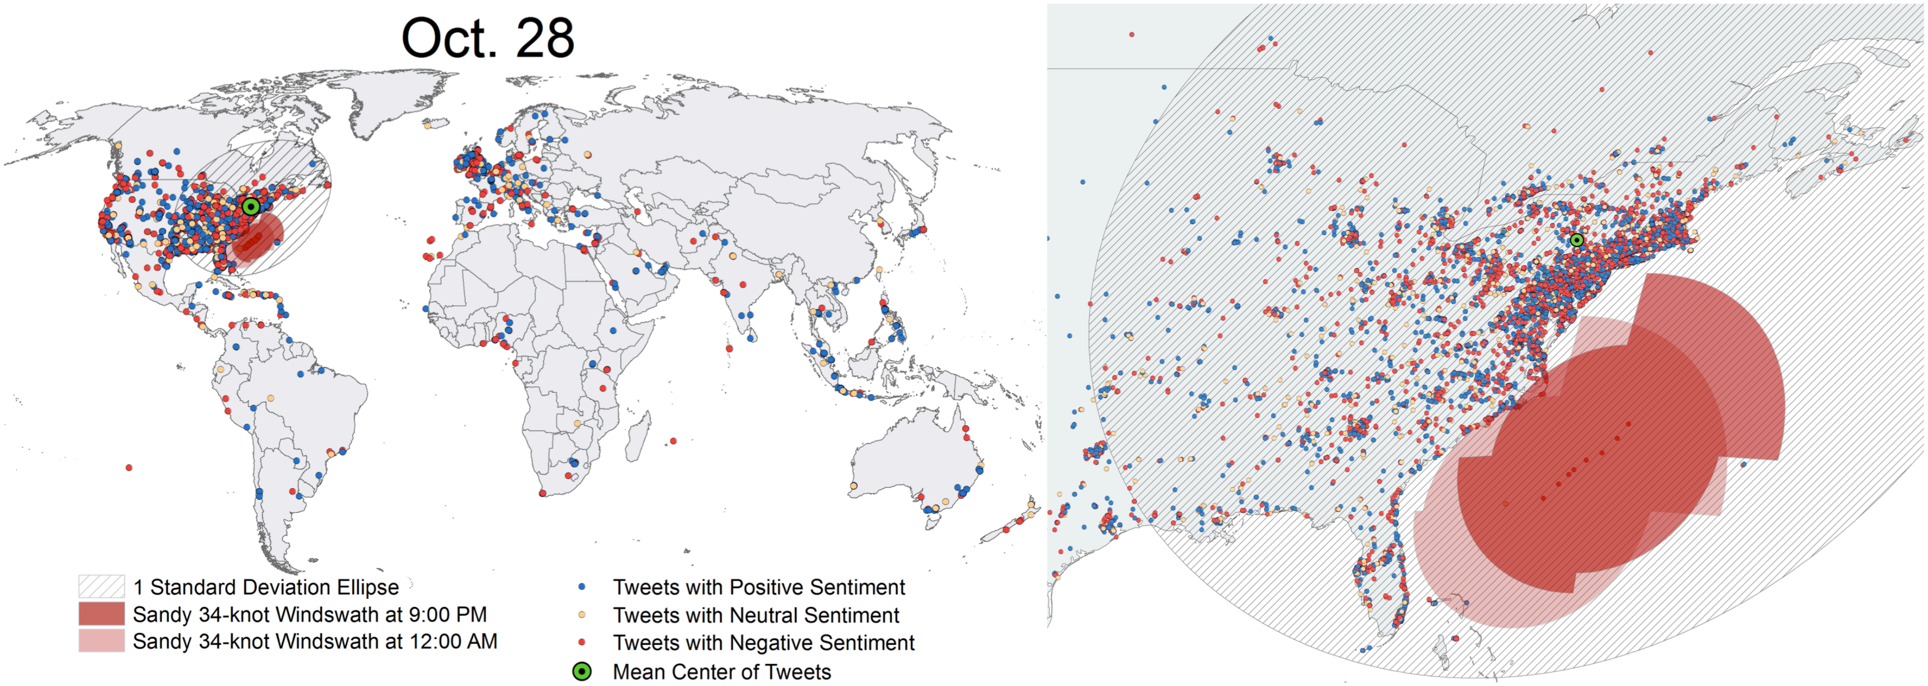

Geo-Tagged Tweets Sentiment Analysis

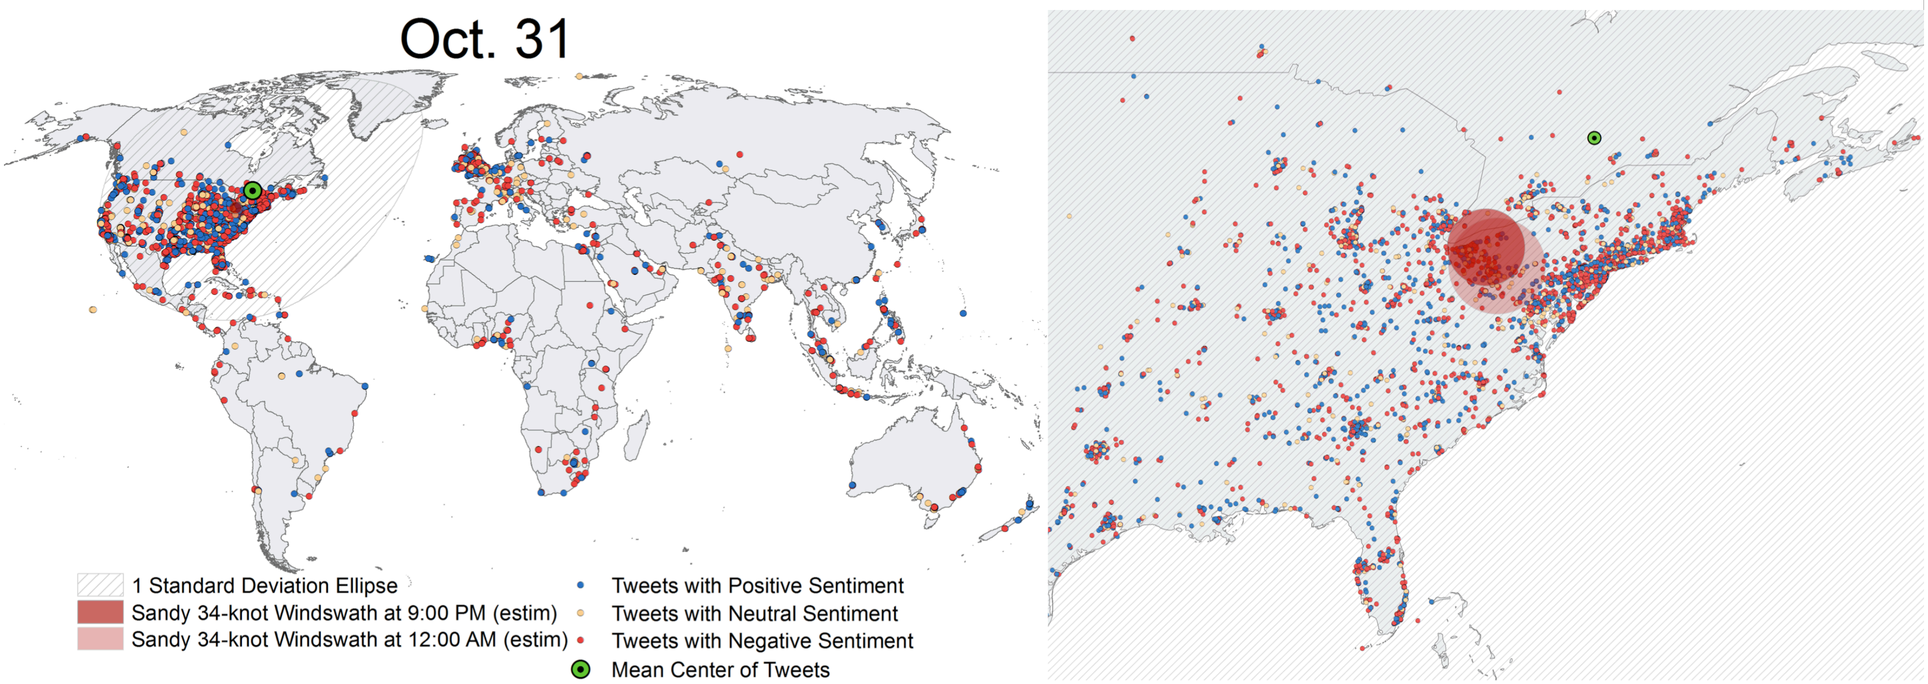

Summary and Conclusion

- We performed sentiment classification of user posts in Twitter during Hurricane Sandy and visualize sentiments on a geographical map centered around the hurricane.
- Our model can be integrated into systems that can help response organizations to have a real time map, which displays both the physical disaster and the spikes of intense emotional activity in proximity to the disaster.

- Future directions:
  - Automatically infer tweets geo-location
  - Automatically identifying trustworthy information spread around disaster events.

Thank you!

National Science Foundation
WHERE DISCOVERIES BEGIN


Cornelia Caragea


Anna Squicciarini


Andrea H. Tapia


Kishore Neppalli


Sam Stehle