Emotion Classification and Data Mining on Social Networking Sites

Justin Prieto (jp393), Sujith Vidanapathirana (srv3), Stephen Purpura (sp559), Claire Cardie (ctc9), John Pollak (jpp9), Geri Gay (gkg1)
Cornell University, College of Engineering

Introduction

Over the past few years, social networking services like **Twitter** and **Facebook** have become pervasive throughout the web. Efficient and accurate analysis of the emotions present in the user-generated content on these services could prove very valuable in electronic commerce, marketing, and web search.

We have performed a comparative analysis of multiple machine learning techniques to assess their applicability to **detecting emotion** in Twitter posts (Tweets). Due to the 140 character limit imposed on Tweets, this is much more difficult than the usual text classification task. Using workers provided by Amazon Mechanical Turk to provide training data for our classifiers, we built an emotion detection application for Tweets.

We also examined the feasibility of harvesting data from Facebook while adhering to their terms of service and respecting the users' **expectations of privacy**. We extended our emotion detector to provide information about the emotional content of users' posts to the users themselves.

Data

- Our data consisted of **5,000 Tweets** labeled by experts as belonging to one of four classes: no emotion, positive emotion, negative emotion, and uncertain emotion.
- We experimented with preprocessing (e.g., removal of stop words and Porter Stemming) and different feature representations.

Example Classification

Tweet	Classification
W00t, just got a new job!	Positive Emotion
So much homework :(	Negative Emotion
Just saw harry potter	No Emotion
Can you believe she's like 60?????@$$	Uncertain Emotion

Results

- **Feature selection matters more than algorithm type -** Like other sentiment applications, binary term presence features in a variety of algorithms perform better than other methods.
- **Community language reduces reliability** - Use of jargon and sociolinguistic differences affect the quality of classification.
- **Mechanical Turk users can provide valuable accurate training data for classifiers -** Untrained users have provided labels for Tweets via Amazon that are more accurate than those provided by Cornell students.

System Architecture

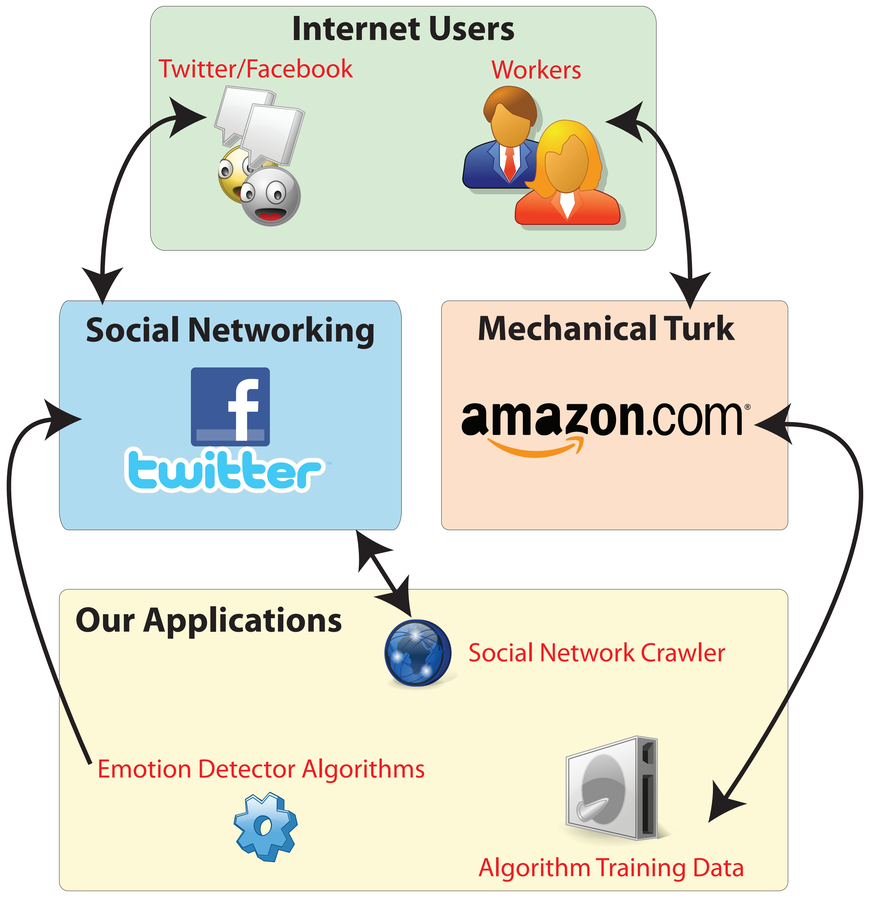

Future Work

- **Experiment with Mechanical Turk worker payment** - Examine the effect different price points have on the quality and volume of work obtained from MTurk workers. Implement pay for performance -- workers who do a better job get paid more.
- **Build a search engine that analyzes emotional content to deliver tailored results** - A search engine can analyze the data on the emotions of a user and provide tailored results to that user based on their opinions or feelings about a product or idea.
- **Extend emotion detector Facebook application** - Include features to compare users' emotions with their friends' and calculate statistics on their mood over time.
- **Modify parameters of ML methods and try new ones** - Changing the selection of features used to represent Tweets as well as trying other methods such as Hidden Markov Models with various tweaks could yield better results.

Where is the BOOM?

Among other applications, NIH funded projects are beginning to examine the use of social networking sites to interact with medical patients. Using natural language processing and machine learning techniques to analyze user posts on social networking sites for emotional display is complicated. This work has extended our knowledge of some of the problem areas.