

DREU Summer 2013 Final Report

Analysis of Forum Interactions to Determine and Predict User Deviance

Ruyan Chen

Abstract

Popular forums can result in hundreds and thousands of posts a day. It can be difficult for administrators to manually monitor and enforce rules on their website as the user base grows. To achieve this, we graph the interactions of users in threads to create a bipartite network. Users with a majority of posts in a thread that are negative have the edge marked between the user and thread as so. Infractioned users—users that have been reprimanded in the past—are also marked as so with a color. By giving administrators this visual information, we can aid them in determining which users are a negative influence on the community and should be monitored more.

Introduction

Before social network services such as Facebook and Google+ became popular, most social websites are forum, or bulletin board, style. These forums often have subsections where users can post about topics in a specific category determined by the section. Users can also join user groups if they share a common interest. One major flaw in forums is the lack of tools to assist the enforcement of rules. Usually, assigned administrators must read through posts to determine if they are deviant and harmful to the community. The intention of this research is to assist administrators in identifying deviant users through showing their interactions with other users. Users who argue with each other would be both seen as deviant, and

users can also respond to negative comments in a deviant manner. Interactions such as these show that deviant users are not always isolated cases.

Analysis

The analysis involves three metrics. The weights for these metrics can be adjusted through the GUI and is determined by the administrator rather than fixed in the program. These three metrics are: sentiment, word ratio, and content relevance. We focused on the first two of these metrics.

Sentiment is determined to be positive, negative, or neutral. We use Alchemy API for this determination. All user posts in a thread are consolidated and analyzed for sentiment. This process takes into account the positive, negative, or neutral connotations of all the words used by the user in the thread and returns the overall sentiment. The sentiment is then weighted with the following values: 1 for positive, 0.5 for neutral, and 0 for negative.

The ratio of harmful words to the total number of words in a post is calculated by removing all the punctuation from a post and counting the number of matches to a list of negative words acquired from a previous thesis paper. It is determined that at least 10% of the words must be harmful for the post to be considered negative. This is determined by determining the ratios of all posts marked as infracted and averaging the value.

The intention of the content relevance is to determine if the user's posts are on-topic. Since off-topic posts are often against the rules of a forum, it can indicate the user's unwillingness to follow rules.

Graphical User Interface

These images show a blank interface and an interface with the graph. The graph is a collapsed version of the bipartite network which only shows the user nodes. The graph is constructed in C++, exported as a GML file, and then visually constructed using R.

Conclusion

This research provides a tool for forum administrators to better maintain their forums. However, it is limited in that the administrator must analyze the data and come to a conclusion using the tool, rather than being provided predictions by the tool. I hope further endeavors of this research will result in a development of prediction tools.