Identifying Fraudulently Promoted Online Videos

Vlad Bulakh, Christopher W. Dunn, Minaxi Gupta

April 7, 2014

INDIANA UNIVERSITY

SCHOOL OF INFORMATICS
AND COMPUTING

Motivation

- Online video sharing websites are visited by millions of people every day
- Some people abuse video sharing websites in a number of ways:
  - Video response spam
  - Plagiarizing other's work
  - Uploading videos with questionable content
  - Copyright violations
  - Fraudulent video promotion

Motivation

- Uploaders can earn money on some sharing websites by adding ads to their videos
- Miscreants use fraudulent means of video promotion
- Existing fraud detection mechanisms employed by video sharing websites are not effective in detecting fraudulent video promotion

Approach

- We focus on YouTube due to its popularity, but the approach could be applied to other video sharing websites as well
- We collect a sample of legitimate and fraudulently promoted YouTube videos and profiles
- We analyze the data and come up with a number of differentiating features
- Those features are later used to train machine learning classifiers capable of detecting fraudulent video promotion

Increasing video rank

- The following attributes are believed to influence the ranking of YouTube videos:
  - Number of views, comments, and ratings
  - Popularity of the uploader's channel
  - Quality of views
  - Number of video responses
  - Number of outside links to the video
- A number of paid and free online services specialize in boosting these attributes

Fraudulent video promotion

- There are two ways in which a video can be fraudulently promoted:
  - Premium services, paid
  - Video exchange websites, free
- We chose to study fraudulent video promotion via the exchange sites

Video exchange websites

- Require the installation of a browser plugin or standalone software
- Most of the actions, such as viewing, liking, commenting on, and favoriting a video, are automated
- Most popular video exchange websites:
  - www.vagex.com
  - www.enhanceviews.com
  - www.viewtubetrain.com

- Data collection
- Legitimate vs. fraudulent videos
- Legitimate vs. fraudulent profiles
- Classifier training
- Other insights from data

Fraudulent videos

- Recorded ids of 3,308 fraudulent videos that we ‘watched’ via www.vagex.com and www.viewtubetrain.com plugins
- Collected the information for the fraudulent videos via YouTube API

Challenges in selecting legitimate videos

- YouTube does not provide a way to request random videos
- Randomly selecting a string and then checking if such video exists is impractical:
  - YouTube video id is an 11-character-long string
  - Each character has 64 possible values:
    - a-z, A-Z, 0-9, _, -
- Selecting most popular/viewed/favorited videos would not give one a representative sample of YouTube

Legitimate videos

- Used a frequency dictionary of the top 5,000 most popular English words to collect an unbiased video sample via YouTube API:
  - Selected 1-3 words randomly
  - Issued API queries to search for the selected words
  - Randomly ordered the results
  - Selected random video out of the first 1,000 videos
  - Repeated the steps until we got 4,000 legitimate videos

Fraudulent profiles

- Uploaded three unlisted videos on YouTube
- Used plugins from www.vagex.com and www.viewtubetrain.com to promote them
- YouTube does not provide a way to see which profiles watched a particular video
- Any profile that commented on our videos or subscribed to our channels is fraudulent
- Utilized YouTube API to collect data on 502 fraudulent profiles

Challenges in selecting legitimate profiles

- YouTube does not provide a way to request random user profiles
- Guessing profile ids is out of the question:
  - YouTube profile id can be between 6 and 20 characters long
  - Each character has 62 possible values:
    - a-z, A-Z, 0-9

Legitimate profiles

- Our legitimate video dataset is an unbiased sample of YouTube videos
- We randomly selected 529 profiles that commented on the legitimate videos
- Used YouTube API to collect information on the said 529 legitimate profiles

Data collection challenges

- YouTube API limitations:
  - Limits requests to 1,000 most recent items
  - Profile activity is limited to 150 most recent actions
  - The number of features has diminished over time
  - No access to raw data
- We believe that our classifiers would have been much more accurate if we had video and profile information which was unavailable to us due to these limitations

- Data collection
- Legitimate vs. fraudulent videos
- Legitimate vs. fraudulent profiles
- Classifier training
- Other insights from data

Video length and number of views

Video length CDF:

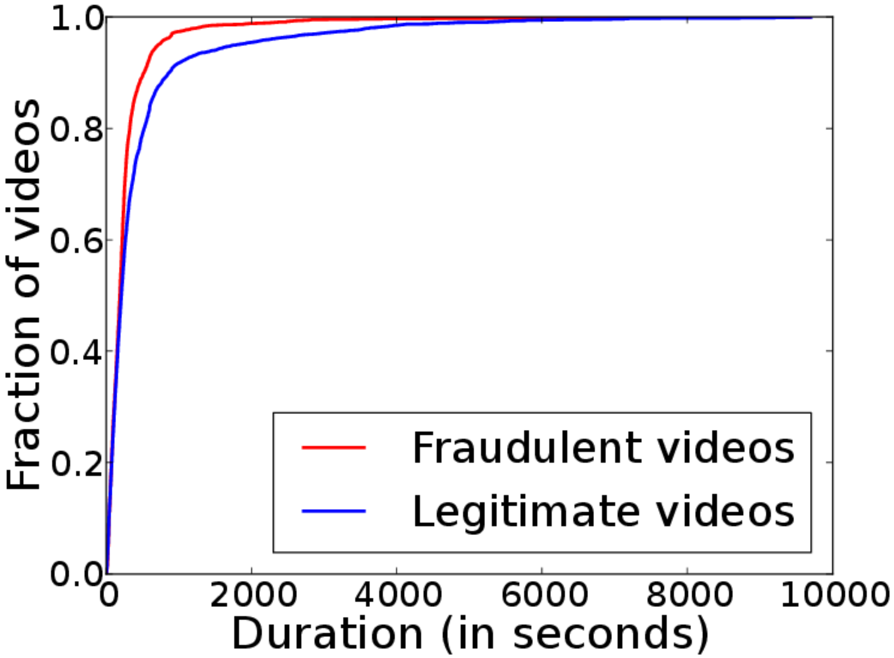

Views per day CDF:

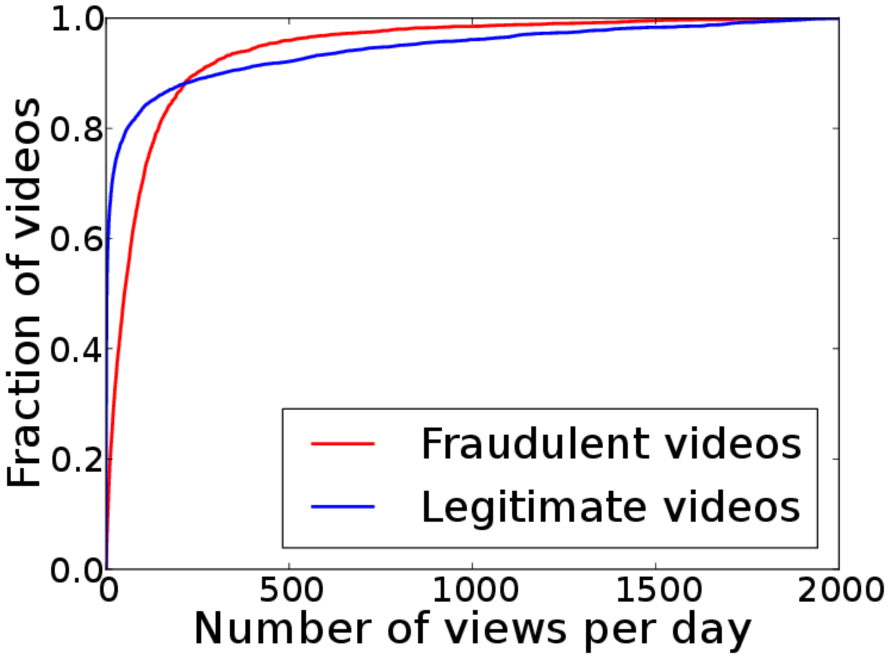

- Fraudulent videos tend to be shorter and have more views per day

Description length and # of URLs

Video description length CDF:

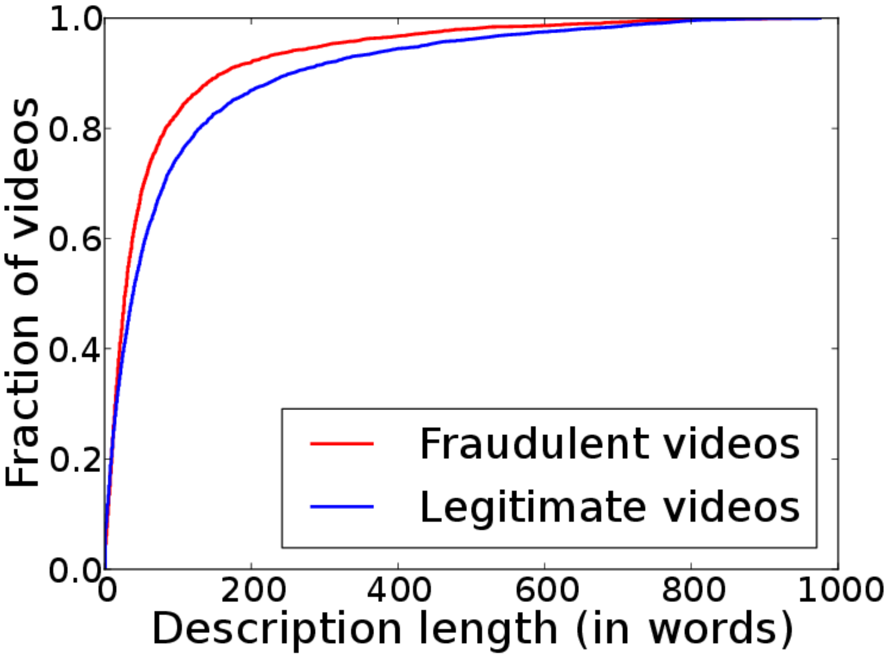

Number of URLs in descriptions CDF:

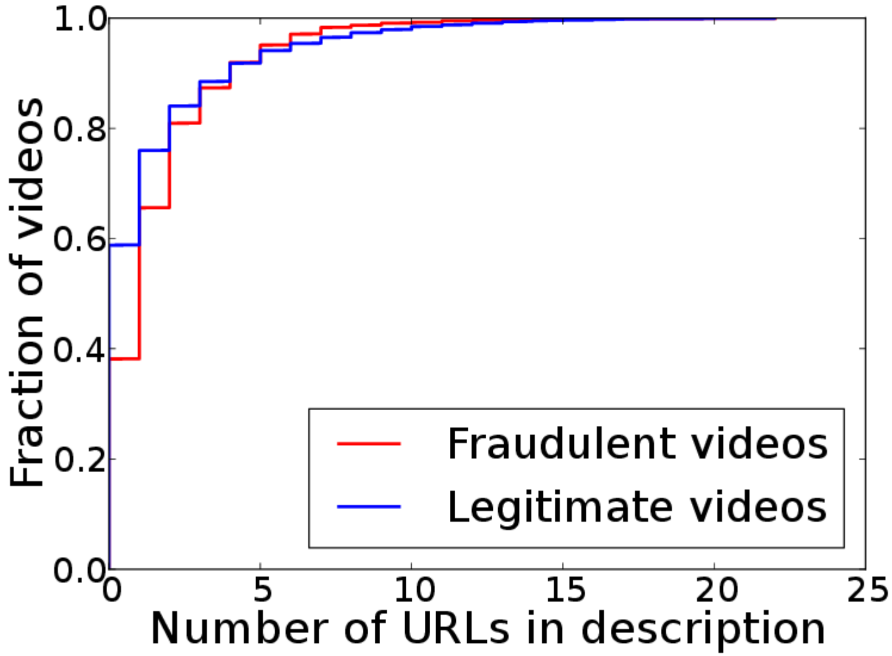

- Legitimate videos have longer descriptions
- Fraudulent videos have more URLs in the descriptions

Video likes, dislikes, and ratings

Video likes/dislikes CDF:

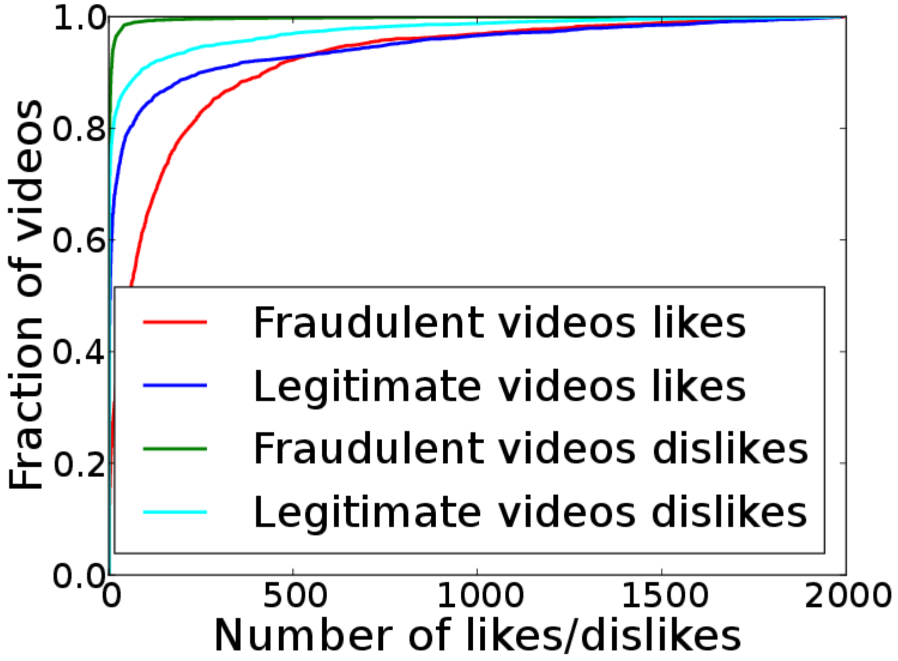

Video rating on a 5-point scale CDF:

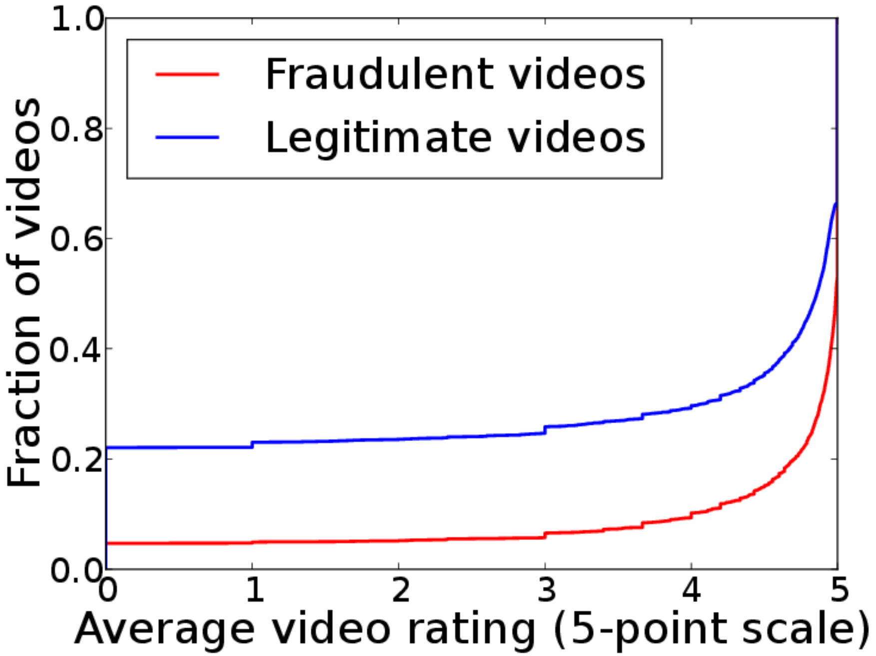

- Fraudulent videos have very few dislikes and are rated higher than legitimate videos

Video comment length

Comment length in words CDF:

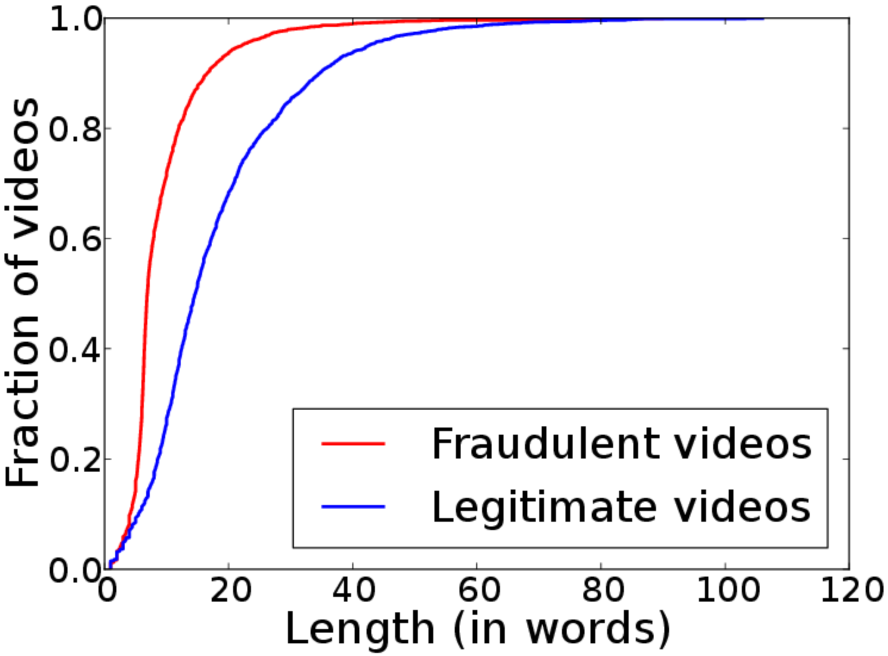

Comment length in characters CDF:

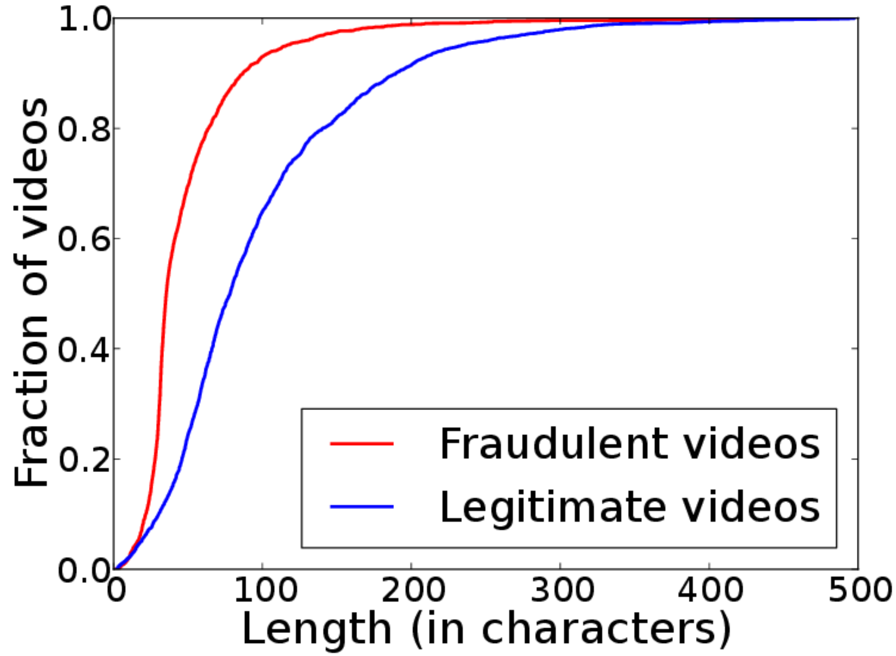

- Fraudulent videos tend to have much shorter comments compared to the legitimate videos

Key observations

- Not only legitimate videos generally run longer than their fraudulent counterparts, but they also have longer comments
- Although fraudulent videos tend to have shorter descriptions, they are more likely to contain links to outside websites
- Fraudulent videos have very few dislikes

- Data collection
- Legitimate vs. fraudulent videos
- Legitimate vs. fraudulent profiles
- Classifier training
- Other insights from data

Age and activity

Profile activity per day CDF:

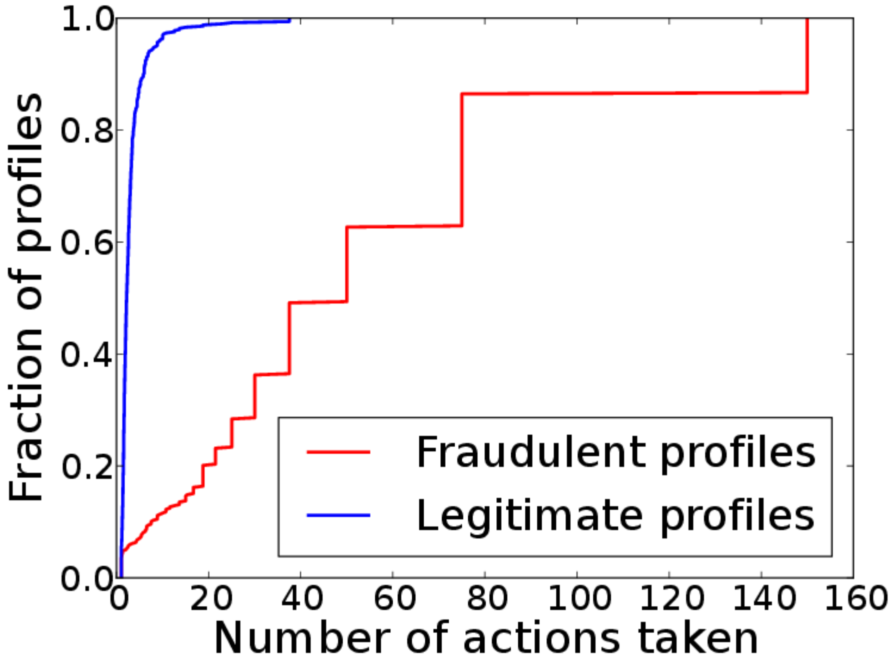

Number of days on YouTube CDF:

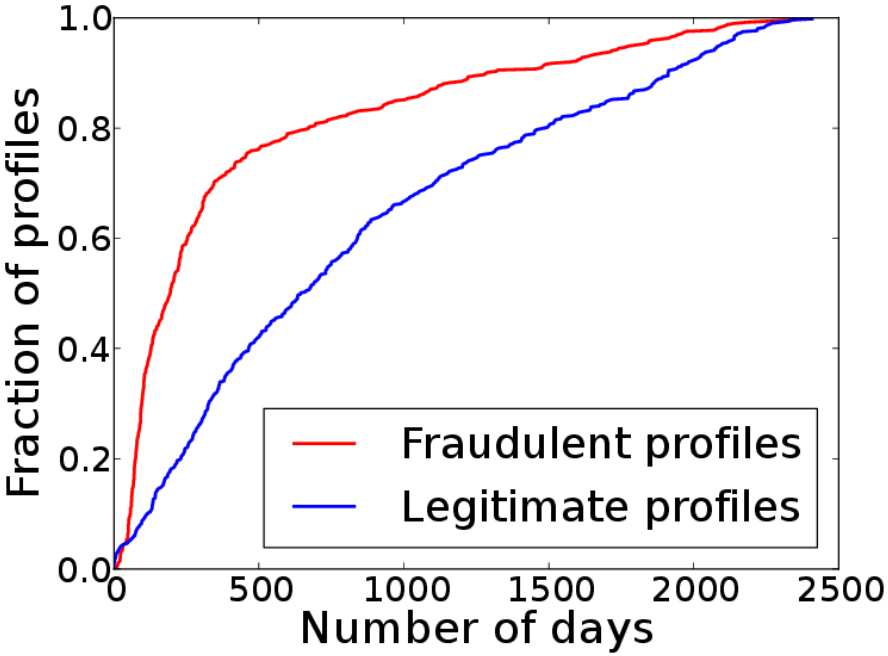

- Although fraudulent profiles are relatively new to the system, they are much more active than the legitimate profiles

Channel statistics

Number of video uploads CDF:

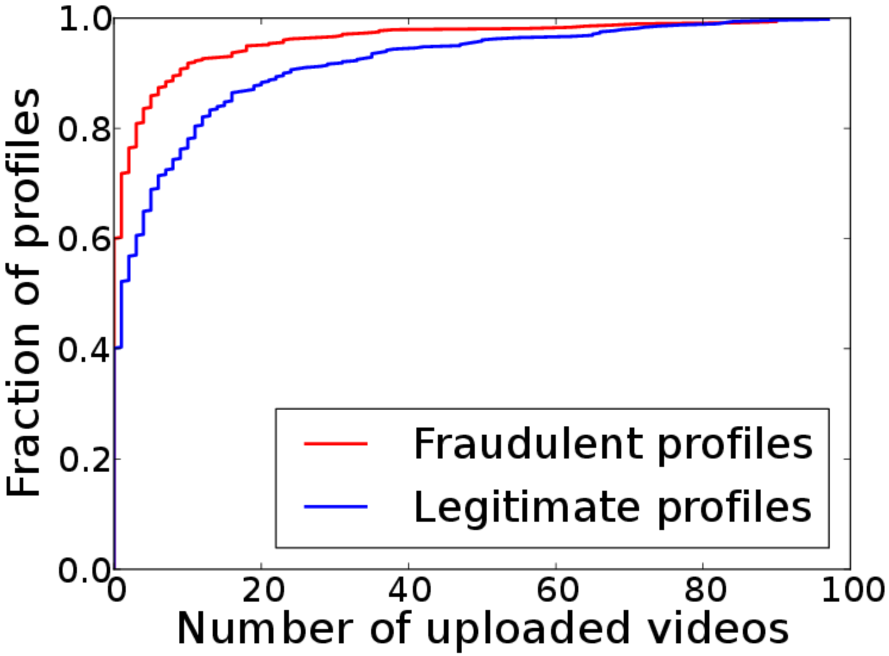

Number of channel views CDF:

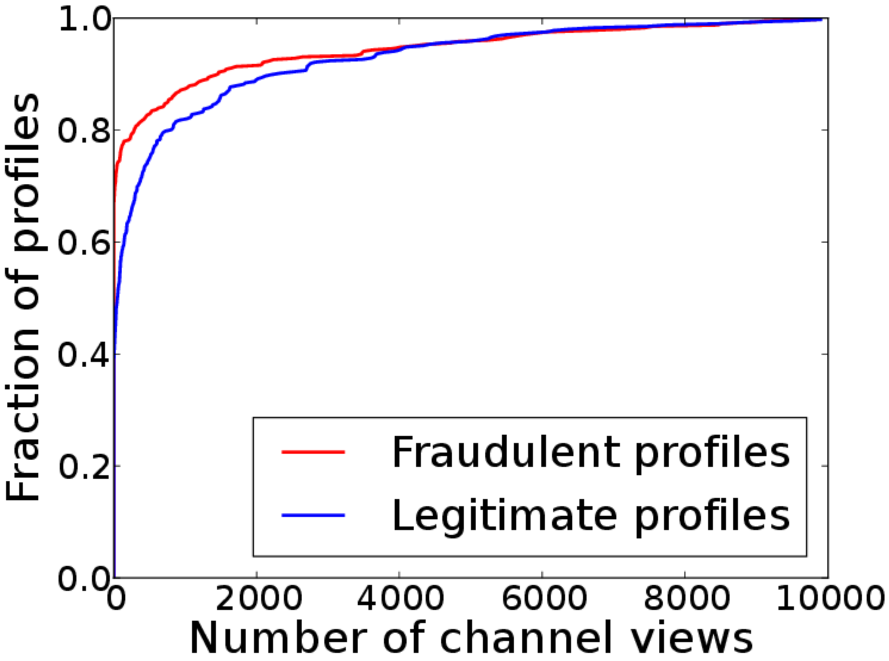

- Fraudulent profiles upload fewer videos, and, as a result, have fewer channel views

Key observations

- Legitimate profiles upload more videos and have more channel views and subscribers
- Fraudulent profiles are relatively new to the system and tend to be more active in terms of viewing, commenting, rating, and favoriting videos whilst uploading very few videos themselves
- This suggests that most such profiles are only being used for fraudulent video promotion

- Data collection
- Legitimate vs. fraudulent videos
- Legitimate vs. fraudulent profiles
- Classifier training
- Other insights from data

Training video and profile classifiers

- We utilized RapidMiner data mining package to train and test our classifiers
- The classifiers used include:
  - SVM
  - Decision tree
  - Decision tree ensemble
- 10-fold cross-validation with stratified sampling was used to predict the accuracy

Feature importance

- Information gain and chi-squared tests were used to find the most differentiating features
- Most differentiating user profiles features:
  - Average number of comments left per day
  - Recent activity
  - Comment length
- Most differentiating video features:
  - Video age
  - Average number of views per day
  - Average rating

Video classifiers

- The data set used for training and testing contained 6,000 videos (3k legitimate and 3k fraudulent)
- Ensemble of decision trees:
  - 91.3% accuracy
  - 8.4% False positive rate
  - 9.1% False negative rate
- A single decision tree classifier took the second place with an accuracy of 88.5%

User profile classifiers

- The data set used for training and testing contained 1,000 profiles (500 legitimate and 500 fraudulent)
- SVM classifier with Anova kernel:
  - 99.2% accuracy
  - 0.2% False positive rate
  - 1.5% False negative rate
- A single decision tree classifier took the second place with an accuracy of 98.7%

Video vs. profile classifier accuracy

- Our profile classifier performs much better than its video counterpart:
  - Most fraudulent profiles are created for the purposes of promoting fraudulent videos, and, thus, take very few actions (if any) on the legitimate videos
  - The majority of the fraudulent videos attract organic views and actions from the legitimate profiles (which is the main goal of the fraudsters)
  - The differences between fraudulent and legitimate profiles are more pronounced than those between fraudulent and legitimate videos

- Data collection
- Legitimate vs. fraudulent videos
- Legitimate vs. fraudulent profiles
- Classifier training
- Other insights from data

Do fraudulent videos have organic views?

- In total, 57,548 profiles commented on 3,308 fraudulent videos in our data set
- We ran our user profile classifier on them
- Out of 57,548 profiles, 30,000 were marked as legitimate by the classifier
- This shows that *fraudulent videos attract organic views and actions from legitimate profiles*, which is arguably the main goal of the miscreants

Does YouTube remove fraudulent profiles and videos?

- User profiles:
  - Only 16% of the fraudulent profiles in our data set were closed or suspended over the course of one year
  - 6% of legitimate profiles were closed in the same time frame
- Videos:
  - Out of 3,308 fraudulent videos, 30% have been removed over the course of a year
  - On the other hand, out of 4,000 legitimate videos, 27% have been removed during the same time period
- *YouTube failed to remove the majority of fraudulent videos and profiles*

Can YouTube detect all kinds of fraudulent video promotion?

- Starting in December of 2012, YouTube began removing fake views from videos
- Despite this, only one fraudulently promoted video in our dataset had a net loss of views
- This suggests that *current fraud detection methods employed by YouTube are not adequate*

Conclusion

- Fraudulent video promotion is thriving
- Current methods used by YouTube to detect fraudulent videos and profiles are not very effective
- Although our results could have been better had YouTube made more data available, our approach may still be useful as a first line of defense against fraudulent video promotion on YouTube