

Meeting notes

- Meeting with Sandra 1-28-2020
- March 22-2019
- April 20-2019

Meeting with Sandra 1-28-2020

Sandra wants to work on getting baselines running for all the languages we're examining

- Still running into difficulty getting Petya access to the system since Brandi left

She thinks that using YASS for Arabic may be useful

- It may not be the most linguistically sound way to do things but it will be consistent and not be too aggressive as a root identification system would (if we reduce the words to only their roots, maybe we lose too much information)
- Maybe for Arabic and other languages it makes sense to use YASS to do splitting rather than stemming (e.g. keep the suffix that gets stripped)
-

March 22-2019

absolute 2 way ig is not working well. The negative class always has higher absolute values.

Should probably stop running experiments soon, the deadline is in a month or so.

Is the difference between German and English due to differences in the dataset size? What if we tried to see if the classifiers are cluing into topics and not malicious language.

Look at the English data and get a sense of what we think people are bitching about.

Look at the ig features and see what winds up in there. to see if particular features are showing up.

Send danny an email with the number of features used.

Yue has high 70% for accuracy. She will update us once it it converges.

April 20-2019

LDA in sklearn running over words only (unigrams). Cutoff of 3 2 topics pick number of top words it spits out.

Compare to XGBoost and SVM baselines. E.g. How often does the output from the SVM correspond to a particular topic from LDA? How often does the output from the XGBoost correspond to a particular topic from LDA?

Check out cases where topic modeler is not confident e.g. probabilities are close to 50. Maybe ones that are missed are on the borderline in the LDA model.

Could also run svm and get the coefficients for different word features.

Ken is running sampling experiments again. We're both running with 70,000 ig features 30,000 features and then 10,000 features

Ken is having some issues with cluster centroids. It has been running for 2 days and still hasn't finished downsampling

It appears to be using a sparse to dense conversion because it's using up to 50 gb of ram