Automated Goal Classification and Analysis of Facebook Political Advertisements

Natalie Appel, Lexie Silverman, Erika Franklin Fowler
Wesleyan University, Department of Government, QAC Apprenticeship Summer 2021

Introduction

Over the last decade, the scale and influence of political digital advertising has increased dramatically. There is an overwhelming amount of data available about these ads, especially from large platforms like Google and Facebook, but the sheer volume of ads makes it difficult to analyze. The Wesleyan Media Project has a team of hand coders who have reviewed 3,000 Facebook ads from the 2020 U.S. Election and classified them based on the perceived goal of the ad. This information is extremely useful for our analysis of the sentiment and content of these ads, but hand coding is a time intensive process that cannot keep up with the scale of digital advertising. Our goal is to automate part of this process by training a machine learning classification model on the information we have from the hand labeled data set. In doing so, we can utilize computational analysis to gain insight into campaigns' advertising strategies on social media platforms.

Data

Data on Facebook's political ads is available from their public ad library. 3,000 of these ads were collected from the period leading up to the 2020 U.S. Election and given to hand coders who systematically identify the ad's primary and secondary goal (among other variables). We focused on using primary goal to train models that classify the entire set of ads relating to the 2020 Election. The categories for ad goal include:

Donate	Persuade	Purchase	Info	Other goal
• Asks audience directly to donate to a campaign or provides information on how to do so	• Speaks in a positive or negative way about a candidate or tries to convince audience to vote	• Sells or provides information how one can buy merchandise and/or tickets to a political event	• Encourages users to learn more about a candidate, campaign, or political group	• Anything else an ad may convey, including asking the audience to contact their legislator

Models

We started by testing a series of different classifier models that analyzed the text of each ad, which includes the ad's creative body, title, description, and caption, but not text transcribed from images and/or video.

Random Forest: 75.8% accuracy

Implemented first to get a baseline idea of classifier accuracy. Trained only on the ad text.

	precision	recall	f1	support
DONATE	0.91	0.91	0.91	128
INFO	0.67	0.70	0.68	151
OTHER GOAL	0.50	0.04	0.08	48
PERSUADE	0.74	0.86	0.79	221
PURCHASE	0.85	0.85	0.85	13
accuracy			0.76	561
macro avg	0.73	0.67	0.66	561
weighted avg	0.74	0.76	0.73	561

SVM: 74.5% accuracy

A Support Vectors Machine model that was used as another baseline. Trained only on the ad text.

	precision	recall	f1	support
DONATE	0.88	0.93	0.90	128
INFO	0.65	0.66	0.65	151
OTHER GOAL	0.36	0.27	0.31	48
PERSUADE	0.80	0.78	0.79	221
PURCHASE	0.76	1.00	0.87	13
accuracy			0.75	561
macro avg	0.69	0.73	0.70	561
weighted avg	0.74	0.75	0.74	561

DistilBERT: 76.7% accuracy

A Bidirectional Encoder Representations from Transformers model, initially trained solely on ad text. Later we implemented more versions that analyzed ASR text as well.

	precision	recall	f1
DONATE	0.9206349	0.9062500	0.9133858
INFO	0.6666667	0.6622517	0.6644518
OTHER GOAL	0.3600000	0.3750000	0.3673469
PERSUADE	0.8272727	0.8235294	0.8253968
PURCHASE	0.8666667	1.0000000	0.9285714

Keyword Search: 58% accuracy

Utilized word frequencies and keyword lists to uniquely identify primary ad goals by the ad text.

	precision	recall	f1
DONATE	0.8112948	0.9145963	0.8598540
INFO	0.6414286	0.5923483	0.6159122
OTHER GOAL	0.1154381	0.3458333	0.1730970
PERSUADE	0.7953795	0.4346258	0.5620991
PURCHASE	0.6769231	0.6769231	0.6769231

BERT Variations

From our initial set of models, the distilBERT model appeared to be the most promising approach for higher accuracy. We decided to further explore by testing the BERT model on three variations of our data: the first included only the ad text (*textonly*), the second included ad text and ASR text (text transcribed from video) concatenated together (*combined*), and the third included ad text and then ASR text sequentially (*separate*). Each of these models were applied to the full set of ads so that we could examine the discrepancies between their predictions, which in turn would reveal the impact of including ASR text in the training process.

Overall, the three variations generally produced similar predictions. However, there were several examples in which the predictions disagreed. Any instance in which all three tests did not agree was considered a disagreement. The table to the right shows the percentage of disagreements attributable to each model and each category for ad goal.

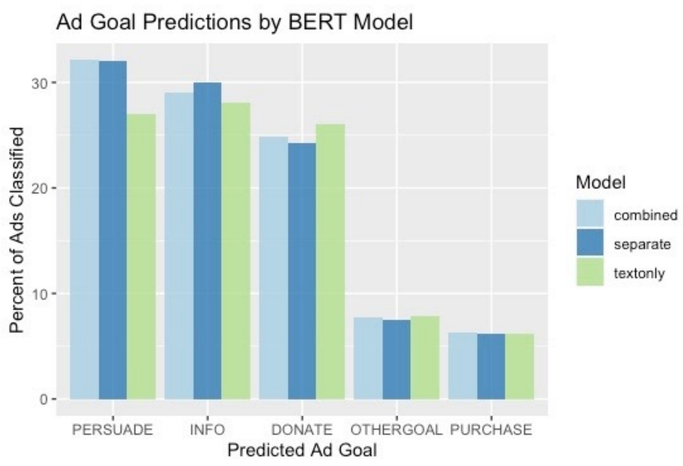

BERT disagreements

	textonly	combined	separate	Total
DONATE	50.68	13.49	5.75	33.02
INFO	24.05	26.07	58.19	32.10
OTHER GOAL	9.02	10.99	1.44	7.73
PERSUADE	15.95	38.45	34.60	24.73
PURCHASE	0.30	11.00	0.01	2.43
Total	57.16	20.46	22.38	100.00

The most noticeable pattern is that over half of disagreements occur when the *textonly* model disagrees with the two models that include ASR text. Additionally, all three models are much more likely to disagree on Donate, Info, and Persuade; this is especially evident from the bar chart below, which illustrates how often each category was the ad goal detected by the model that disagreed.

Conclusions

Our research provides insight into the content of political ads and their main use in election campaigns. The figure below combines the results of all the models we have discussed and illustrates the distribution of predicted ad goals determined by each classifier. It is clear the majority of political Facebook ads from the 2020 U.S. Election fall into the Donate, Info, and Persuade categories, which implies that campaigns are utilizing digital advertising to gain supporters and promote their ideology. On the other hand, fewer ads than we anticipated were aimed at actions within the Purchase or Other goal categories, such as buying campaign merchandise or contacting a representative. Again, this suggests that Facebook ads are primarily trying to gain support for a candidate, issue, or cause by informing their audience.

In addition, it appears that the most variation occurred within the Other goal and Persuade groups. This is likely in large part due to the Keyword Search and Random Forest classifiers as is evident in the figure below. Conversely, the three variations of the distilBERT model seem to produce the most consistent results, which confirms our belief that the BERT classifier approach will result in the highest level of accuracy.

Moving forward we plan to continue improving the accuracy of our models and expanding them to create more classifiers. We hope this research can serve as a foundation and be used to predict other ad characteristics, such as the ideology of ad sponsors and whether they are tied to individual candidates or groups.

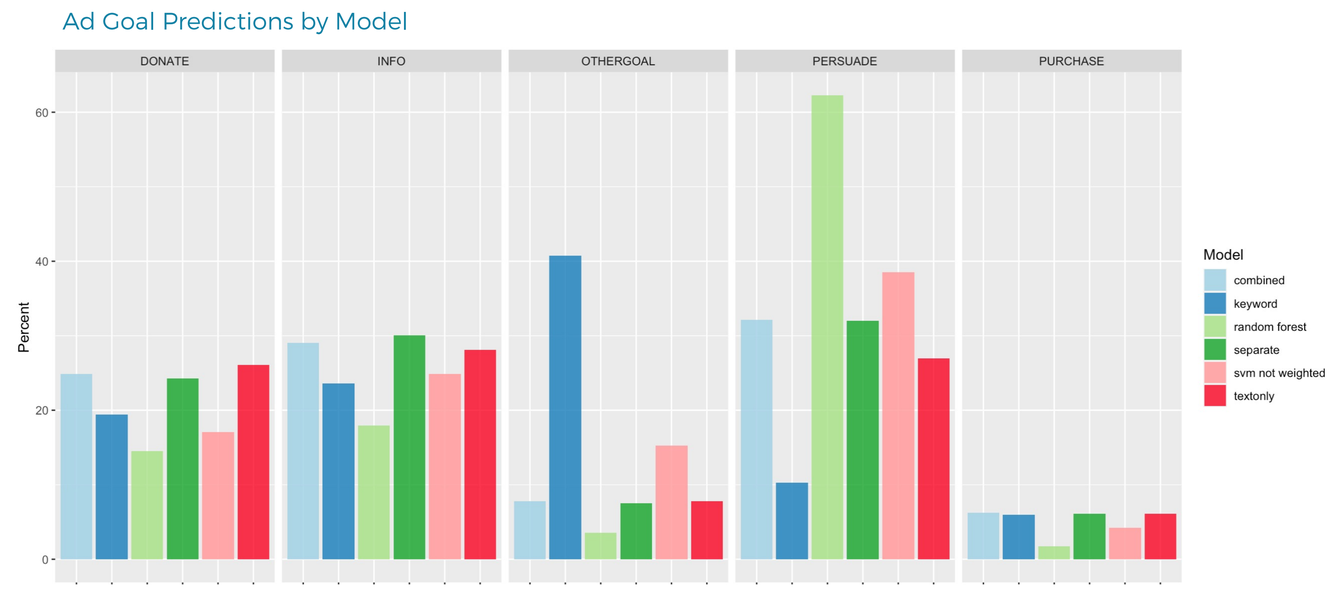

Acknowledgements

We are extremely thankful to the entire Delta Lab Team, especially:
Erika Franklin Fowler, Laura Baum, Pavel Oleinikov, Markus Neumann, Jielu Yao, Spencer Dean, and Sam Feuer.