Deep & Machine Learning Approaches to Analyzing Gender Representations in Journalism

Maggie Davis, Stephanie Campa, Daniela Gonzalez

Motivation

Journalism plays a key role in modern society and is widely regarded as a crucial component of a functioning democracy. However, despite increasing social pressure in support of equality of the sexes, men and women continue to be represented differently in the news. Not only are men more frequently the subject of journalistic articles than women, but the way men and women are talked about in the news differs. Our goal is to employ natural language processing and deep learning to delve more deeply into this phenomenon.

Objective

- Given an input headline from a news article, our model aims to classify the gender of the subject (not the author).
- There are two output classes: 'Man' and 'Woman'.
- Although several researchers have attempted to classify the gender of writers/authors with NLP and deep learning (Mukherjee & Liu), few have tried to classify the gender of the subject.
- We hope to capitalize on the biases that exist in NLP to learn more about the gender-based biases that are most prevalent in journalism.

Dataset & Features

- Dataset is comprised of approximately 1,800 article headlines from eleven popular news sources that regularly report about people.
- We scraped, labeled, and parsed the data ourselves.
- We removed words that were obvious indicators of gender, including gender pronouns and names of key public figures, from the input data.
- All of our features were text-based, ranging from bag-of-words to more complex NLP features like word classes, word2vec, and GloVe.

Headline	Label
"Trump Wants to Start a Fourth of July Parade Tradition"	Man
"Women Sue Yale Over a Fraternity Culture They Say Enables Harassment"	Woman
"Senate Passes a Sweeping Land Conservation Bill"	Omitted from dataset

Experiments

BASELINE: *Naive Bayes*

- Multinomial Naive Bayes classifier.
- Use word counts and the NB assumption to calculate the probabilities that words (and therefore articles) are about men vs. women.

INTERMEDIATE: *Support Vector Machines*

- Implemented classification with support vector machines (SVM) with bag of words features, leveraging the SVC class of sklearn.SVM.
- Experimented with both linear and Gaussian kernels.

ADVANCED: *Convolutional Neural Networks*

- Feed-forward convolutional neural network (CNN).
- First implemented this using bag-of-words embeddings and later experimented with pre-trained word2vec, pre-trained GloVe, word classes (factor analysis), and part of speech (F-measure).
- After the embeddings, this model has a 2-D convolution layer with filter sizes 3, 4, and 5, followed by a max-pooling layer and, finally, a fully-connected layer with dropout, using softmax as the loss function.
- Also experimented with L2 regularization to minimize overfitting.

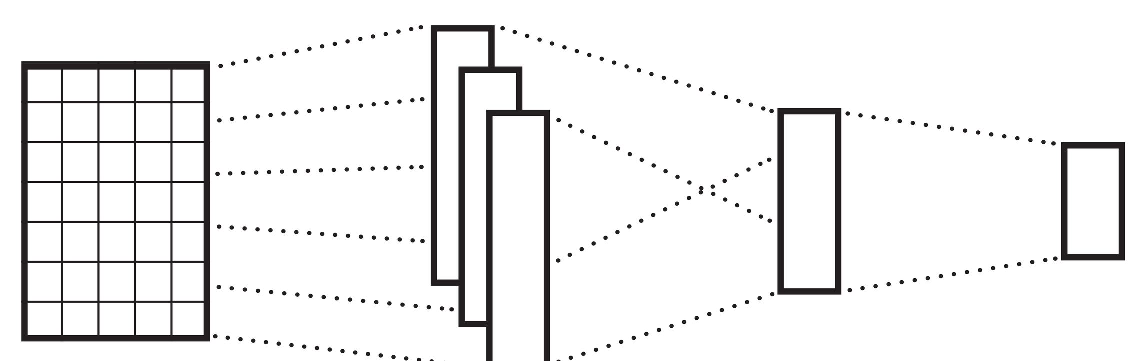

Results

- CNNs were more effective than Naive Bayes and SVM, as expected.
- CNN with GloVe embeddings and word classes (factor analysis) was most successful, with accuracy of **86.7%**.

	Predicted: Man	Predicted: Woman
Actual: Man	114	2
Actual: Woman	24	40

Model Performance

Model	Accuracy
Naive Bayes	73.9
SVM (Gaussian Kernel)	66.2
SVM (Linear Kernel)	72.5
CNN w/BOW	82.2
CNN w/word2vec	83.3
CNN w/GloVe	85.6
CNN w/Factor Analysis + BOW	75.0
CNN w/Factor Analysis, BOW, F-measure	76.1
CNN w/Factor Analysis + word2vec	82.8
CNN w/Factor Analysis, GloVe, F-measure	84.4
CNN w/Factor Analysis + GloVe	**86.7**

Indicative Words

- **Man:** 'guilty,' 'summit,' 'emergency,' 'congress,' 'report,' 'do,' 'wall'
- **Woman:** 'dress,' 'israel,' 'mom,' 'look,' 'daughter,' 'ex,' 'paris,' 'breaks,' 'husband'

Hyperparameter Tuning

Embedding Dimension*	Num Filters	Dropout Probability	L2 Reg Lambda	BOW Accuracy	FA + GloVe Accuracy
128	128	0.5	0.0	71.6	**86.7**
64	128	0.5	0.0	76.7	86.7
256	128	0.5	0.0	78.9	86.7
128	64	0.5	0.0	77.2	85.0
128	256	0.5	0.0	75.0	86.1
128	128	0.2	0.0	78.3	85.0
128	128	0.3	0.0	81.1	85.0
128	128	0.4	0.0	74.4	86.1
128	128	0.5	1.0	**82.2**	86.1
128	128	0.5	5.0	80.0	85.6

**Note that the embedding dimenson does not effect the FA + GloVe column, since the GloVe word embeddings are pre-trained to a set dimension.*

Analysis & Conclusion

- We anticipated CNNs would perform well on our task because they have hierarchical architectures and are good at identifying patterns.
- Model seems to assume that headlines regarding crime or politics are about men, whereas headlines regarding family are about women.
- High model accuracy is indicative of both the gender bias in journalism and the discrepancies in gender representation in society as a whole.

References

- Arjun Mukherjee and Bing Liu. Improving Gender Classification of Blog Authors. *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, 2010.
- Yoon Kim. Convolutional Neural Networks for Sentence Classification. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, 2014.