Decoding Visual Sentiment of Political Imagery

Olga Gasparyan [1] Elena Sirotkina [2]

[1]FSU, Political Science [2]UNC Chapel Hill, Political Science

FSU FLORIDA STATE UNIVERSITY

Motivation

- Images direct our attention, shape attitudes, and can reinforce stereotypes.
- How can we identify and/or predict sentiment of such visual content?

How would you label sentiment of this image?

Published by Reuters (Oct 29, 2018). Photo by Reuters/Adrees Latif.

Two known approaches from computer vision literature:

- *treat images holistically*: compile them in databanks with tagged sentiment labels
- *object-centric approach*: focus on specific regions within image and assign labels object-based (eg., smile = positive sentiment; firearms = negative sentiment).

Crucial drawback of these approaches: assumption that the sentiment conveyed by an image is an inherent attribute of the image itself and not of the perceiver's attitudes of this image.

- Especially for politically charged imagery, viewers' interpretations are often subjective and are a function of their beliefs.
- Subjective sentiment should largely depend on human coders' attitudes toward the political issue depicted.

Novel Approach to Visual Sentiment Task

When training a visual sentiment classifier for political images, we incorporate important attitudinal differences that people might have about the depicted topic.

- Identifying an Attitudinal Cleavage: examine whether visual labeling reflects a stable societal gap (eg., age, gender, or political ideology divides) that cannot be neutralized by labelling at scale.
- Creating a Dataset of Sentiment Labels that does not ignore such cleavages by averaging across all the human coders' evaluations but rather assigns separate sentiment labels for different societal groups.
- Training a Multi-task Multi-class Classifier: for prediction task we propose to build a classification model that will incorporate such vector of labels for each image, emphasizing the development of classifiers that more accurately capture sentiments as interpreted by humans.

Empirical Case

- Political imagery that can evoke political divide - we choose visual representation of events related to the topic of immigration.
- Assume political divide as a basis for the attitudinal cleavage
- 816 images acquired from Twitter accounts of U.S. media outlets and Getty Images
- In two survey waves respondents were asked to evaluate sentiment of these images.

Proxy visual sentiment with four questions (on a scale from 1 to 7):

1. Sentiment: Would you say that this image portrays the subject(s) or objects(s) in this picture in a positive or negative light?
2. Subject of harm: In your opinion, the subject(s) who is (are) portrayed in this picture is (are) more likely to be dangerous or harmless?
3. Object of harm: In your opinion, the subject(s) who is (are) portrayed in this picture is (are) more likely to be vulnerable or safe?
4. Accuracy: Do you think that this image is a faulty or accurate representation of the story that actually occurred?

We average across individual evaluation scores for each visual sentiment proxy and produce 3 evaluation scores (AES $\in [1,7]$) for each image: 1. averaged across all respondents; 2. averaged across self-identified Democrats; 3. averaged across self-identified Republicans.

We assign categorical labels to AES in the following way:

$$\text{Categorical Label} = \begin{cases} \text{negative}, & \text{if AES} \leq 3 \\ \text{neutral}, & 3 < \text{if AES} < 5 \\ \text{positive}, & \text{if AES} \geq 5 \end{cases}$$

Image-Level Labels

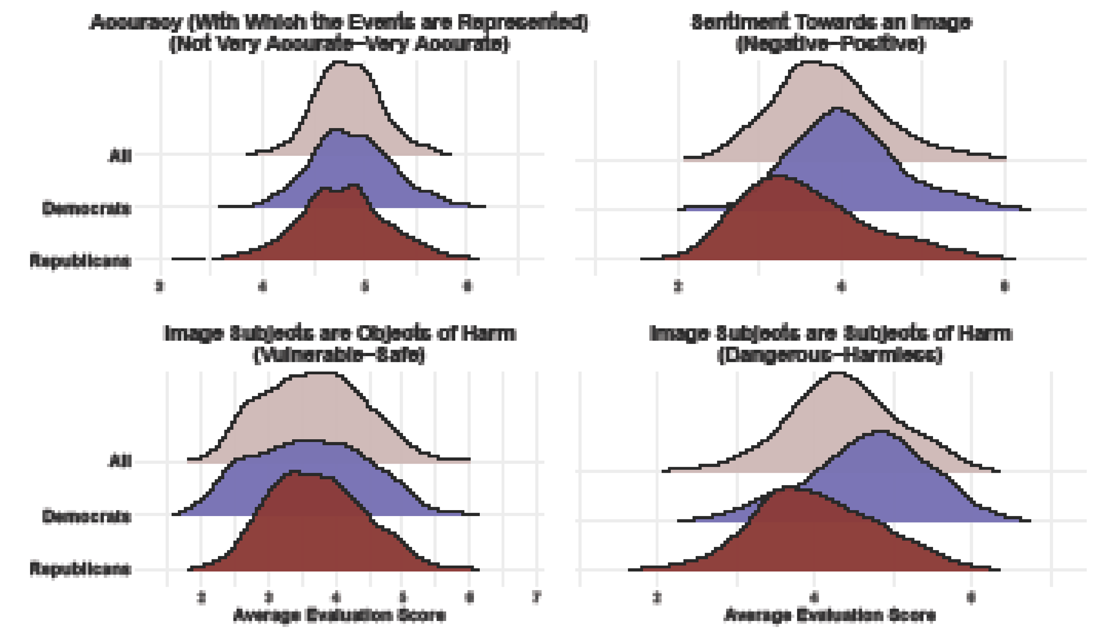

- For 'Sentiment' and 'Subjects of Harm' questions: notable differences in average scores between Democrats and Republicans ...
- ... so here a single averaged label might introduce significant bias
- We need to use separate labels for Democrats (L_D) and Republicans (L_R).

Deep Learning Approach

- Our approach: Transfer learning of known convolutional neural networks (CNN) — ResNet50V2, DenseNet121, and DenseNet169 — with adapted dual-task multi-class classification in the last softmax layer.
- Fine-tuning with 80/20 train/validation sets split, lr=0.0001, bs=32 with added Batch Normalization and Dropout layers on 50 epochs with early stopping.
- Example of CNN architecture adapted for dual-task classification:

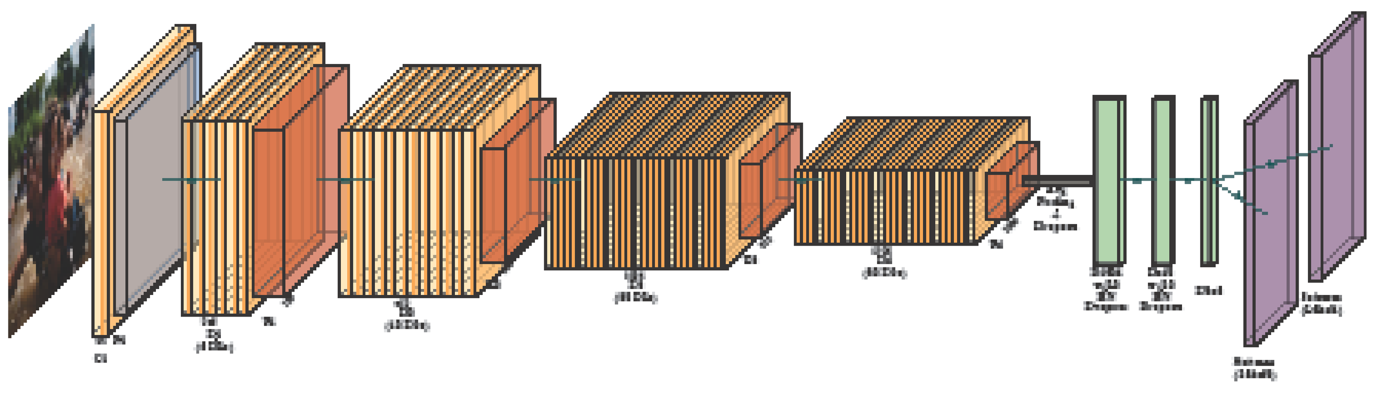

Model Results on Validation Set:

- *Sentiment:*
  - Dem: F1: 0.7 and Val Accuracy: 83%
  - Rep: F1: 0.62 and Val Accuracy: 64%
- *Subject of Harm:*
  - Dem: F1: 0.7 and Val Accuracy: 73%
  - Rep: F1: 0.72 and Val Accuracy: 76%

Model Testing

"Sentiment" Dem: [Neutral], Rep: [Negative]
"Subject of Harm" Dem: [Positive], Rep: [Neutral]

Conventional Labels:
"Sentiment": [Neutral], "Subject of Harm": [Neutral]

"Sentiment" Dem: [Neutral], Rep: [Negative]
"Subject of Harm" Dem: [Positive], Rep: [Neutral]

Conventional Labels:
"Sentiment": [Neutral], "Subject of Harm": [Positive]

Note: "Sentiment" and "Subject of Harm" are used to proxy a broader concept of visual sentiment. If for "Sentiment" outcome negative, neutral, and positive labels are quite straightforward, for "Subject of Harm": negative label means that subjects were perceived as more dangerous and positive label means that subjects were perceived as more harmless. Conventional labels are based on average calculations across all the respondents without taking into account partisan divide in sentiment perceptions.

Implications and Contributions

- Visual sentiment should be understood as an interplay between the image content itself and individual attitudes towards such content.
- Importance of recognizing and addressing substantial and systematic influence of coders' social characteristics (such as political beliefs) on the labeling of visual sentiment.
- Improving interpretation of prediction results:
  - *'For group X (e.g., Democrats/Republicans), our model predicted a positive/negative visual sentiment label for this image.'*
- Implications for politically polarizing content:
  - researchers can use our model to determine whether an image's sentiment is influenced by partisan differences;
  - our trained classification model can forecast the sentiments that images elicit in different partisan groups;
  - that allows constructing an index measuring the polarizing impact of an image: $\|L_D - L_R\|$, where L_D and L_R - sentiment labels from Democrats and Republicans.