Fair and Balanced? News Media Bias in the Photographic Coverage of the 2016 U.S. Presidential Election

Markus Neumann The Pennsylvania State University

The 2016 presidential election has renewed discussions about the impartiality of the news media. Scholars have studied this issue extensively, investigating newspapers, television and online news, yet the basic question remains unsettled: Is the media biased? In this paper, I focus on nine popular news websites covering the 2016 presidential election campaign. I apply computer vision techniques to photos of the candidates and their supporters at campaign rallies, automatically identifying their emotions. I expect that co-partisan news outlets would portray members of their favored side as happier and thus more positively. An analysis of this data shows that no such media bias exists. While both Donald Trump himself as well as his supporters consistently appear less happy than Clinton and her followers, there do not appear to be any differences between liberal and conservative news sources in this regard. I argue that these findings are the result of a polarized and profit-driven media environment, where the desire to portray one's side in a positive light is balanced out by the incentive to keep readers apprehensive by playing up the other side's chances of winning.

The question of if and how the news media is biased has featured prominently in political discussions for decades, if not centuries. In the course of the 2016 presidential election and its aftermath, this debate reached new levels of acrimony. The Trump campaign in particular leveled accusations of unfair coverage against the "liberal media", with liberals airing similar grievances against conservative outlets such as Fox News and Breitbart. Accompanied by the proliferation of "fake news" (both the actual concept as well as the now very loosely used term), this development reached a fever pitch with the new president declaring the news media to be an "enemy of the American People."[1]

These events make the following question more relevant than ever: is the news media actually biased? There is no shortage of research on this question, and yet, the evidence remains inconclusive. Numerous studies have provided support for (Friedman, DiMatteo and Mertz, 1980; Waldman and Devitt, 1998; Banning and Coleman, 2009; Moriarty and Garramone, 1986; Moriarty and Popovich, 1991) and against (Larcinese, Puglisi and Snyder, 2011; Budak, Goel and Rao, 2016) this proposition.

In addressing this question, it is critical to make use of a corpus that actually reflects the news consumption habits of a large enough portion of Americans to be relevant for

[1] `https://twitter.com/realDonaldTrump/status/832708293516632065`

the political environment. According to a study conducted by Pew (Mitchell et al., 2016) in 2016, 57% of Americans prefer to get their news through television, 38% rely on online sources, 25% listen to news on the radio, and 20% favor newspapers. The majority of older studies on media bias focused on television or newspapers. Recently, research on news websites has featured much more prominently in scientific publications (Baum and Groeling, 2008; Hehman et al., 2012). Given that online media sources are already preferred by more than a third of Americans and continue to grow at a much higher rate than their competitors, this is hardly surprising.

So far though, studies on online sources - in contrast to research on television - have primarily focused on text. However, in contrast to traditional newspapers, websites (or the web presences of newspapers) place far more emphasis on images. On websites, photos are more numerous and larger, articles are shorter, and most readers only look at the image and the first few lines of text. Even if readers do take the time to read the entire article, pictures still prime the reader, setting the tone for the entire story (Hehman et al., 2012). Consequently, the choice for a picture can convey at least as much editorial bias as the writing itself. In this paper, I analyze the selection of pictures by leading news media sources - such as the New York Times, the Wall Street Journal, or Breitbart News - in the coverage of the 2016 presidential election campaign. In order to ensure comparability, I focus on two specific types of pictures: (1) the faces of the candidates themselves and (2) the crowds of supporters cheering for them at their rallies.

Portraits of the candidates are commonly used in the coverage of the presidential election. The increasing powers associated with the Presidency (Lowi, 1986) make the presidential election an extremely high-stakes event, and the personalization of politics in the U.S. (Bennett, 2012) causes the attention of an entire nation to be focused on two people, and in the case of visual media coverage, their faces. Consequently, existing political science research on media bias in images has largely focused on the faces of the presidential candidates (Moriarty and Garramone, 1986; Moriarty and Popovich, 1991; Waldman and Devitt, 1998; Banning and Coleman, 2009). Campaign rallies channel this cult of personality towards the electoral goals of the candidates, and yet allow them to connect with the people they aim to represent. This was particularly evident in the 2016 presidential election campaign, were Donald Trump (as well as Bernie Sanders in the primaries) tied his fortunes to his “mega-rallies”, citing them as evidence for the success of his campaign and the righteousness of his message. The news media appears to have picked up on this phenomenon, frequently using these photos in articles on the horse race, making them a symbol for the election campaign as a whole. Within these images, I focus on emotions. A large literature (see, for example Brader (2005); Ridout and Searles (2011); Jones, Hoffman and Young (2012); Weber (2013); Huddy, Mason and Aarøe (2015) has shown that emotions displayed by candidates and amplified by their campaigns play a critical role in elections, shaping the behavior of voters. My research question, then, is whether news outlets portray their preferred candidate and their supporters in a more positive manner by showing them in a higher proportion of happy images.

The Achilles's heel of most of the media bias literature is that it largely relies on human-coded news articles and television programs, and thus ultimately rests on the objectivity of researchers. Natural language processing-based techniques have contributed to addressing this problem when it comes to text data (see, for example, Gentzkow and Shapiro (2010); Soroka (2012)), but for measuring media bias in images, manual approaches are still the state of the art. And yet, this is the area in which researcher bias is the most likely and problematic: Hillary Clinton and Donald Trump are both hugely polarizing figures who are ubiquitously known across the U.S. and even the world, so it would be nigh impossible to find coders who are not yet familiar with their faces and do not associate some form of subconscious bias with them. Furthermore, the scene viewing literature in the cognitive sciences has used eye-tracking methods to show that there are a number of factors – such as personality or gender – influencing which part of an image viewers pay attention to (Pan et al., 2004; Risko et al., 2012). Furthermore, Masuda et al. (2012) show that for the specific task of rating the emotions displayed by a person in an image, there are inter-coder differences based on the ethnicity of the raters as well as the content of the image. It is likely that these factors further interact with the political beliefs of the coders. For these reasons, it should be expected that the human coding of images – especially the images of two politicians who are hated and loved by millions – would likely lead to bias.[2]

My approach is to rely on computer vision techniques to parse images of the candidates and their supporters to identify their faces and the emotions they display. Deep neural networks have recently led to rapid progress in several areas of machine learning, especially computer vision. I rely on a number of these models to identify the faces found in the images of the candidates and their supporters and subsequently classify their emotions. The result of this analysis shows little evidence of media bias. While Clinton and the supporters attending her campaign rallies appear to be happier than their Republican counterparts, different media outlets do not seem to differ in this assessment.

The theoretical expectation of the news media portraying their co-partisans as overwhelmingly happy is likely too simplistic. I argue that partisan motives can nevertheless lead to non-partisan news coverage. Both liberal and conservative media outlets have incentives to avoid portraying their side too positively. A prominent feature of Donald

[2] While it is theoretically possible to attempt to select a balanced set of coders through the use of implicit bias tests, doing so would be very complicated and ultimately still not guarantee a complete absence of bias.

Trump's campaign was that he and and his supporters expressed genuine anger over the status quo. It appears plausible that conservative websites attempted to echo that anger. At the same time, portraying Clinton as happy fit the conservative narrative of her as an arrogant person. Furthermore, happiness is associated with the expectation of winning (Huddy, Mason and Aarøe, 2015), so images of happy Clinton supporters would likely spur on the right. Conversely, as both pollsters and pundits predicted a Clinton landslide, liberal news outlets had some incentives to keep expectations down and portray conservatives as happy. A divisive and closely-fought race keeps readers engaged and therefore is directly beneficial for the media's bottom line.

Media Bias

The literature on media bias is extensive, involving a number of disciplines such as political science, communication, sociology, psychology, economics and computational linguistics. The question of what constitutes media bias however depends on the specific line of inquiry. There are two overarching branches of research – selection bias and presentation bias (Groeling, 2013).

The former deals with cases in which bias occurs because editors pick certain stories over others, and thus engage in priming. Measurement of this concept frequently involves the raw number of times an issue or politician gets mentioned by a news source. For example, Larcinese, Puglisi and Snyder (2011) study bias in the coverage of economic news by tracking the volume of stories on unemployment and inflation in U.S. newspapers, and comparing them to their actual level.

Presentation bias on the other hand describes skewed news coverage with regard to how a story frames an issue. For an example on a similar topic, Soroka (2012) conducts an automated content analysis on economic news stories in the New York Times, detecting whether their tone is more favorable to Democrats or Republicans.

Another important question regarding the definition of bias is whether it refers to a "systematically [...] distorted" "portrayal of reality" (Groeling, 2013), or as an inevitable consequence of limited human information processing, which cannot be avoided (Guerra et al., 2011). The former suggests that there is an objective reality of what happened, with bias being the media's deviation from it. Under the latter, this question is inconsequential - even if there is a ground truth, humans are incapable of detecting it. Bias in a political context then merely refers to the ordering of attitudes and opinions, without a baseline. In this paper, I follow the latter approach because there is no way to establish the actual level of happiness and anger experienced by the subjects of the images. Ergo, my results describe the degree of bias media outlets exhibit *in relation to each other*. As a baseline for the assumed partisan leaning of the news outlets covered in this paper, I rely on the Pew Political Polarization & Media Habits study (Mitchell et al., 2014). Here, news outlets are classified according to the position of their readers on the ideological spectrum. This allows me to formulate expectations about their preferred candidate.

Scholars have studied bias in a variety of news sources, the classical example being newspapers. Frequently conducted as content analyses in which research assistants are tasked with coding the partisan slant of stories, this type of study can take on many forms. A fairly conventional example, Barrett and Peake (2007) analyze local newspaper coverage of presidential travel, relying on manual content coding. The authors show that the partisan leaning of a newspaper affects both the amount as well as the tone of coverage. In addition to traditional approaches like this newspapers can also be utilized in a more innovative fashion. Butler and Schofield (2010) conduct a randomized field experiment in which they sent ideologically slanted letters to newspapers to determine whether editors would be more likely to print letters that conform with their paper's ideological position. Surprisingly, the opposite turned out to be the case – newspapers appear to be encouraging the spread of dissenting opinions.

With the increasing ubiquitousness of the internet, analyses of bias in the written word have increasingly turned to online sources. The digital versions of traditional newspapers continue to be the go-to source for researchers, but online-only outlets such as the Huffington Post or Townhall, as well as the political blogs of partisans are starting to see frequent use in the study of media bias (Baum and Groeling, 2008; Hehman et al., 2012). Lin, Bagrow and Lazer (2011) show that bias is generally more pronounced and polarized on blogs compared to traditional news sources – in either direction. One advantage of relying on online data is its sheer volume – Larcinese, Puglisi and Snyder (2011) study a total of 140 newspapers (via automatic sentiment analysis) while Budak, Goel and Rao (2016) rely on crowdsourcing to crawl through a trove of over 10,000 news articles.

No matter the medium, content analyses of articles always suffer from one problem: coder subjectivity. The need for, as well as the difficulty of being objective, varies by research design – coding ideological bias presents different challenges to coding whether a story represents reality in an accurate manner – but ultimately it always comes down to the same problem: Humans are inherently political (Hatemi and McDermott, 2011), so the coders' own biases have the potential to distort their evaluations of the media's. Machine learning-based approaches attempt to solve this problem by putting the burden of decision-making on an algorithm, but ultimately, those algorithms still rely on human-coded training data (Caliskan-islam, Bryson and Narayanan, 2016; Caliskan-Islam, Bryson and Narayanan, 2017).

Researchers have also studied bias of television news, with Friedman, DiMatteo and

Mertz (1980) presenting one of the first accounts. The authors analyzed video footage of newscasts during the 1976 presidential election campaign. Rather than taking the obvious route of analyzing the show's content, Friedman et al. coded the emotional reactions of newscasters, operationalized through their facial expressions as they are saying the names of the candidates. This approach, though one of the oldest, mirrors my own the most closely. It also happens to be one of the surprisingly few cases in which researchers actually do uncover a liberal bias in the media.

Mullen et al. (1986) build on this study by replicating it for the 1984 presidential election and combining it with a telephone survey in which they poll respondents on their vote choice, as well as their TV viewing habits. The results show that people who habitually watch a show in which the newscaster displays a particular kind of partisan bias, are significantly more likely to vote for that party. In doing so, Mullen et al. (attempt to) establish a causal effect of media bias on voting behavior. Unfortunately, the authors simply dismiss the alternative explanation of an echo chamber effect without presenting any evidence against it.

Banning and Coleman (2009) present a more recent account of television news, featuring still images from the 2000 presidential election. The authors analyze emotional content of candidates, rating for favorability of expression, appearance, nonverbal behavior, etc. The results point to a slightly more favorable coverage for Republicans.

Aside from television, print and online media have also been used as a source for studies of media bias on photographs. Moriarty and Garramone (1986) conduct a content analysis of images of presidential candidates in 1984, featured in U.S. News and World Report, Time and Newsweek, with Reagan receiving more favorable coverage than Mondale. A similar study on the 1988 presidential election produces comparable findings (Moriarty and Popovich, 1991). By contrast, Waldman and Devitt (1998) show that in 1996, Clinton received slightly more favorable coverage, although the horse-race polling at any given moment served as a better predictor of flattering photographs. This, in combination with the studies by Moriarty and Popovich also suggests an incumbency advantage. Hehman et al. (2012) presents one of the most recent analyses, rating photos of George W. Bush and Barack Obama on online news websites for features such as warmth, competence, or dominance. The results suggest that ideologically aligned news sources frequently feature more complimentary images.

In addition to detailing where media bias is originating from, it is also worth noting what it is aimed at. A large portion of studies detail media bias with regard to presidential candidates (Friedman, DiMatteo and Mertz, 1980; Mullen et al., 1986; Moriarty and Garramone, 1986; Moriarty and Popovich, 1991; Banning and Coleman, 2009). Actual presidents also feature as the object of studies, albeit less frequently (Barrett and Peake, 2007; Hehman et al., 2012). Congress has not received the same kind of scholarly attention as the presidency, presumably because of the equally lower media attention (Gentzkow and Shapiro, 2010; Lin, Bagrow and Lazer, 2011). As far as actual political issues are concerned, the accuracy of reporting is a frequent topic (Larcinese, Puglisi and Snyder, 2011; Soroka, 2012; Parks, 2016). The wars in Afghanistan and Iraq have also been covered (Aday, 2010; Glazier and Boydstun, 2012) and share one important quality with my own study: the object of the media is inherently subjective and the ground truth is unknown to the researcher. Overall political ideology also features as the object in a number of studies (Budak, Goel and Rao, 2016).

Ultimately, the most important question however is: Is the media actually biased? Evidence for the vaunted liberal news media is certainly more rare than expected, but can be found in some studies (Friedman, DiMatteo and Mertz, 1980; Waldman and Devitt, 1998). However, bias in favor of Republicans occurs just as much (Banning and Coleman, 2009; Moriarty and Garramone, 1986; Moriarty and Popovich, 1991). Many studies report no bias (Larcinese, Puglisi and Snyder, 2011; Budak, Goel and Rao, 2016), or bias towards the side a particular outlet is leaning to (Barrett and Peake, 2007; Hehman et al., 2012).

The fact that emotions have become a frequent object in the study of media bias (Friedman, DiMatteo and Mertz, 1980; Banning and Coleman, 2009; Hehman et al., 2012) is owed to the central role they play in political campaigns. Scholars have uncovered the effect of emotions on participation (voting, donating, volunteering) (Jerit, 2004; Kiss and Hobolt, 2011; Huddy, Mason and Aarøe, 2015), the retention of information on candidate platforms (Civettini and Redlawsk, 2009), as well as the psychosocial functioning of partisans (Westen et al., 2006; Vigil, 2010). Emotional content (and even specifically emotional images (Burton et al., 2005)) aids both cognitive processing and memory retention, thus influencing how political stimuli are perceived (Fazio, 2001; Spezio and Adolphs, 2006). Enthusiasm and anger have received a particularly high degree of attention. Enthusiasm among supporters is both a response to positive appeals made by politicians (Brader, 2005; Ridout and Searles, 2011; Jones, Hoffman and Young, 2012; Weber, 2013), as well as the belief that their side is winning, (Huddy, Mason and Aarøe, 2015), an effect that is amplified among the strongest partisans. Similarly, anger is the product of candidates with a negative message, as well as the expectation to lose (Weber, 2013; Huddy, Mason and Aarøe, 2015).

Given the central function these emotions perform, media bias likely plays a role in the way they are portrayed. Both liberal and conservative news outlets have incentives to frame their favored campaign as enthusiastic, and their opponents as angry: One, both sides, despite the cynicism with which they conduct themselves at times, still believe in the constructive role their cause has to play for the good of the country. Consequently it makes sense to portray co-partisans as having a positive message (enthusiasm), whereas opponents only channel obstructionism and negativity (anger). Two, due to the existence of the bandwagon effect (McAllister and Studlar, 1991), there is a strategic advantage to be gained by casting an opponent as the losing side (anger), and co-partisans as winning (enthusiasm). My hypothesis then, is simple: Liberal media outlets are expected to portray Clinton and her supporters as enthusiastic and Trump and his supporters as angry, with the reverse being true for conservative websites.

Data & Methods

To measure media bias in photos of the candidates and their supporters, I build two datasets. Henceforth, I refer to the former, depicting only Clinton and Trump, as the "candidate dataset", and the latter, showing crowds at their campaign rallies, as the "rally dataset".

Scraping

The first step towards building these datasets consists of acquiring the images themselves. To this end, I scrape pictures from nine different online media sources. The selection of news outlets is based on the precedents set in the literature (Larcinese, Puglisi and Snyder, 2011; Hehman et al., 2012; Budak, Goel and Rao, 2016). These studies have chosen their sources based on two main criteria: One, maintaining a mix of traditional newspapers (New York Times, Wall Street Journal, USA Today, Chicago Tribune), TV networks (CNN, Fox News, MSNBC), and online only (Huffington Post, Breitbart) outlets. Two, ensuring that both sides of the political spectrum are equally well-represented.[3] With the exception of Breitbart, which I added because of the considerable attention it received during and after the 2016 presidential election campaign, all of these websites have featured in the studies cited above.

The goal of this first step is to build a database as large as possible, prioritizing volume over accuracy. This means that I prefer including false positives to omitting false negatives. Practically, the scraping runs entirely through Google Images. Since Google indexes the entire known web[4], all news websites can be searched inside Google.[5] The advantage

[3]The expected ideological positions of news outlets are given by the Pew study on Political Polarization & Media Habits (Mitchell et al., 2014). In my sample, USA Today has the median ideology, with the New York Times, Huffington Post, MSNBC and CNN to its left, and the Wall Street Journal, the Chicago Tribune, Fox News, Breitbart to its right.

[4]https://www.google.com/insidesearch/howsearchworks/crawling-indexing.html

[5]For the campaign rallies, the scraping was carried out between February 16-19, 2017. For the candidate images, the scraping was done on February 10, 2019, using the date range 1-1-2016 to 11-8-2016 (election day)

of this approach is that a) images are ordered consistently between websites (instead of using each website's own search algorithm, which might differ drastically from that of another) and b) I don't have to adapt my scraping program[6] to the layout of each website individually.

The specific search terms used are "Trump rally crowd"/"Clinton rally crowd"[7] for the images of supporters, and "Donald Trump"/"Hillary Clinton" for the candidates. I have experimented with different terms, as well as combinations of the results of several terms, but found the above to lead to the highest percentage of usable images. Even so, the proportion of pictures actually portraying campaign rallies gets progressively lower as I go further down the list of search results. Consequently I only retain the first 500 (an arbitrarily chosen number) hits from each search. This means that for both datasets, I start out with 500 images from each campaign for each website.[8] In the next step, these (and other) false positives are filtered out.

Filtering

Not all images in this pool actually fit the search parameters. In some cases, photos of rallies outside the U.S. are included in the results. Similarly, images depicting the other candidate occasionally turn up in the wrong place. The reason for this is simple – both candidates' names generally appear in any one article on the election, even if it focuses on one of the two specifically. Consequently I filter out undesirable images by hand. I do so because while deep learning-based models for face recognition are quite advanced at this point, there are several reasons to assume that they would not perform at 100% accuracy on this dataset: Some of the images feature impersonators of the candidates, others portray them at a different age and some are of such bad quality that while the candidate may be identifiable, emotion recognition is not possible. Similar issues exist with the campaign rally dataset.

For the candidates dataset, I filter the images after the face detection model has been applied and the images have been cropped to only feature the faces (meaning that one original image can be turned into multiple portraits). For the followers dataset, filtering is done before face recognition, as I additionally need to decide whether the image actually

in Google's image search to ensure that only pictures from the appropriate timeframe would be selected.

[6]I use a webdriver-controlled browser (Firefox), implemented with the Selenium package in Python, to circumvent Google's anti-scraping measures.

[7]For example, the following search term would yield pictures from Trump rallies, covered by the New York Times: "site:www.nytimes.com Trump rally crowd" (without quotation marks)

[8]For the candidate dataset, some websites, such as the New York Times, returned less than 500 results.

depicts a campaign rally (of the respective candidate).[9,10] I also omit all photos (for the rally dataset only) on which a facial detection algorithm cannot find at least three faces. Images with too few faces would a) not actually capture the concept of a rally crowd and b) be inefficient to use. Finally, for both datasets, images below 36x36 pixels are removed, as they do not contain sufficient detail to reliably perform emotion recognition.

At the end of this process, the original 9000 images in the rally dataset have been winnowed down to 1,158, which, in the next step, produce 12,825 faces. For the candidates dataset, I am left with a total of 2665 faces.

Computer vision

Four computer vision techniques - face, age, gender and emotion detection - are applied to the images (age and gender estimation are only used in the rally dataset). These are separate methods, but they all share the same underlying process - convolutional neural networks (CNNs). Since this approach is still relatively new in political science[11], I first I explain it in principle, and then move on to detailing which steps where applied to the two datasets in what way.

The architecture of CNNs consists of several types of layers. The most important component is a convolution layer. Color images are three-dimensional arrays - consisting of height, width and depth (depth refers to the color channels - red, green and blue). Convolution involves matching a filter (which can be thought of as a kind of feature, for example a small, prototypical representation of an edge or a curve) of a smaller height and width, but the same depth - against parts of the image. By applying this filter to an entire image through a sliding window, each part of the image can be assigned a numerical value, indicating how closely it matches that filter - thus creating a feature map (also referred to as an activation map). In a convolutional layer, a number of these maps - created from different features - are then stacked depth-wise. Figure 14 in the appendix illustrates how these filters look like in the emotion detection model applied to the candidate dataset. Figure 15 shows their outputs, detailing which part of a face the

[9]For the rally dataset, another potentially problematic type of image is one that primarily focuses on family members, co-partisans and staff, rather than an actual crowd of supporters. Programming a computer vision algorithm to specifically find these photos would be quite labor-intensive, as there is currently no labeled training data available.

[10]Besides manual filtering and computer vision, another possibility for the removal of false positives would be to rely on the "alt" attribute (describing the content of an image in case it fails to load), that websites are supposed to (according to W3 specifications `https://www.w3schools.com/tags/att_img_alt.asp`) associate with images on HTML pages. Unfortunately, not all of the media sources in this study are sufficiently diligent about following this guideline, so this would drastically reduce and potentially bias the sample.

[11]See Anastasopoulos et al. (2016); Casas and Williams (2017) for some exceptions.

model pays attention to.

Most CNNs also involve a ReLU (Rectified Linear Unit) layer, which normalizes the feature maps, as well as a pooling layer, which essentially downsamples the image - retaining the same basic information, but at a lower resolution and higher level of abstraction. Several convolutional, ReLU and pooling layers are then stacked, yielding more high-level features further down the line. Deeper neural networks generally also lead to better performance, but require considerably more processing power. This process is also where the term "deep learning" comes from.

Finally, the last pooling layer forms the input for the fully connected layer. By this point, features have reached the highest level of abstraction, corresponding to, for example the eyes, the mouth or the nose. The fully connected layer takes these features and turns them into probabilities associated with output classes, for example female/male, happy/angry, etc. The category with the highest value is then chosen as the observed class.

Training requires a dataset that consists of a large number of images (in this case depicting human faces), each of which is labeled (based on hand-coding) on the class of interest, for example gender. The key to training such a network is the process of backpropagation. Its first step is the forward pass, where a training image goes through the neural network, leading to a set of probabilities in the fully connected layer. In the first attempt, those probabilities will likely be completely naive, for example [0.5,0.5] (the first probability for female, the second for male) for gender classification. Since the actual image is labeled, for example with [0,1], this result can then be passed through a loss function, determining how far off the neural network's prediction was. In the backward pass, the weights responsible for this result are determined, and subsequently updated. Then, additional rounds of forward pass, loss calculation, backward pass and weight update are repeated, slowly "learning" how to perform this type of classification through gradient descent. Once a neural network is trained, it can be used to classify unlabeled images.

Candidates Dataset

To detect the faces in the images, I rely on the method developed by Zhang et al. (2017).[12] The goal of this approach is a face detector that is scale-invariant, meaning that it is robust to very small images. Given that this dataset has a large degree of variance in this regard, it is very important for the face detection model to be able to spot even the smallest faces. It accomplishes this through an anchor-based approach, which essentially overlays

[12] Face detection is applied using the following implementation of Zhang et al. (2017): `https://github.com/clcarwin/SFD_pytorch`

(a) *RGB representation* **(b)** *LBP representation*

Figure 1: *On the left, an example image in its RGB (red, green and blue) representation. On the right, the same image in its LBP (Local Binary Pattern) representation. This is the input to the neural network and provides better accuracy than its RGB counterpart.*

different grids on the image and then checks each tile for the presence of a face. The model is trained on the WIDER FACE dataset[13] and then subsequently applied to all of the images depicting Clinton and Trump.[14]

I use the method described by Levi and Hassner (2015) to measure the emotions displayed in the images of the two candidates.[15] Here, Local Binary Patterns (LBP), which describe the texture of an image, are used in lieu of its RGB (i.e. red, green and blue) representation. Originally popularized by Ojala, Pietikäinen and Mäenpää (2002), the advantage of this approach is the construction of features that are more robust to scale, rotation and light conditions. This is important in an applied case, because the photos of politicians and their followers are not always shot under perfect conditions. LBP encodings of images are created by thresholding the surrounding 8 pixels to the pixel in their center so that pixels that exceed its RGB value are coded as 0, and those that are lower as 1. These sets of eight unordered LBP codes is then mapped back into an image using multidimensional scaling. Figure 1 illustrates this with an image of Donald Trump in its RGB and LBP representation.

[13] I use the pre-trained weights provided by the model's authors.

[14] In their benchmarks, the authors demonstrate that the model performs with high accuracy, achieving an average precision score (i.e. a weighted mean of precisions at different levels of recall) of 99.85 on the AFW dataset, 98.49 on the PASCAL dataset, and 93.7 on the WIDER FACE dataset.

[15] I use the implementation provided by the authors for replication: `https://github.com/GilLevi/AgeGenderDeepLearning`

Then, a convolutional neural network using the `VGG_S` architecture (Chatfield et al., 2014) is trained on the 891 training images of the 2015 EmotiW challenge.[16] The model is trained to recognize seven emotions: Anger, disgust, fear, happiness, neutral, sadness and surprise. Of these, my substantive interest leads me to focus on happiness and anger, which, according to the authors, the models also performs best on (Levi and Hassner, 2015).[17] The trained classifier is then applied to the 2665 images of Hillary Clinton and Donald Trump. The results from the classification are the outputs of the fully connected layer - probabilities associated with each emotion that sum to 1. That means the best way of interpreting this result is a categorical variable, corresponding to whichever emotion has the highest value for a face.

Rally Dataset

For the rally dataset, I implement face, age, gender and emotion detection through Microsoft's Cognitive Services API[18] (using R and the httr package) - a set of pre-trained machine learning tools based on deep neural networks. The choice of using a pre-trained black box method (meaning that I have no information about the training dataset or the hyperparameters[19]) does have some downsides - mainly the fact that replicability may decrease over time as Microsoft improves its product. However, the rally dataset is considerably more noisy than the candidates dataset, meaning that the images of supporters are often of very poor quality. So far, Microsoft's API has given me better performance in terms of error rates than the alternatives.[20,21] For gender classification, the output of the API is a simple binary value. For age, an estimation is given. The output of the emotion recognition classifier are the same as described above.

Data Overview

Tables 1 (candidate dataset) and 2 (rally dataset) show how many images remain in the datasets after the filtering is done.[22] The total number of (usable) images varies quite

[16] I use the pre-trained weights provided by the model's authors.

[17] In the authors' benchmarks (on the EmotiW benchmark dataset), 75.79% of happiness emotions are correctly identified as such. Given that this is a 7-class problem on a dataset which contains many hard cases, this is a good score, and more reliable than neutral (62.69%) and anger (54.55%).

[18] `https://www.microsoft.com/cognitive-services`

[19] The framework is based on the company's own open source "Microsoft Cognitive Toolkit": `https://github.com/microsoft/cntk`

[20] I have tested the API and alternatives on the widely-used JAFFE emotion dataset and the IMDB-WIKI dataset for age and gender.

[21] In a future version of this paper, I hope to use the same open-source implementation for both datasets.

[22] While the same number of images was scraped for all news outlets, a much smaller proportion of them actually pertained to Clinton, Trump and their rallies. Instead, these searches returned, in addition

heavily depending on the source, as well as the type of image. In the candidate dataset, the Wall Street Journal has the most (569) images, but ranks towards the bottom when it comes to pictures of the campaign rallies (62). The opposite is true for the New York Times, which has large (189) number of rally images, but the lowest number of candidate images (38). Breitbart and CNN have a large quantity of images in both categories.

While these differences do not amount to media bias per se, they nevertheless reveal something about the preferences of these outlets. It appears that for example, the Wall Street Journal was not nearly as interested in Donald Trump's 'mega-rallies' as some of its competitors, and therefore provided less images on them. Outlets that are either online only (Breitbart) or primarily television-based (CNN, MSNBC) have larger number of images, although the Huffington Post[23] and FOX News do not necessarily fit that pattern. Given that images have always been the central component of television, and are becoming more and more prominent on websites, this makes sense.

Outlet	Total	Clinton	Trump	Clinton %	Trump %
Breitbart	433	308	125	71.13	28.87
Chicago Tribune	94	62	32	65.96	34.04
CNN	462	257	205	55.63	44.37
FOX News	459	215	244	46.84	53.16
Huffington Post	60	41	19	68.33	31.67
MSNBC	482	287	195	59.54	40.46
New York Times	38	18	20	47.37	52.63
USA Today	68	39	29	57.35	42.65
Wall Street Journal	569	306	263	53.78	46.22
All	2665	1533	1132		

Table 1: *Candidate dataset – number of images (and faces), per candidate, by media outlet*

Clear differences also exist with regard to coverage of the two campaigns – contingent on the image category. For the candidate dataset, all outlets but Fox News and the New York Times have more images on Clinton. By contrast, the exact opposite is true for the rallies. With the exception of CNN, all websites host more images of supporters at Trump rallies (table 2, last column). Breitbart, the most conservative outlet in the sample provides a particularly high share of such pictures, with fellow conservatives Chicago Tribune (71%) and Fox News (70%) not far behind. On the other side, the more left-leaning websites,

to images that were topical but unusable, a surprising amount of photographs from rallies, protests and riots from all around the world. While Breitbart did produce a much larger corpus, its unsuitable images share this characteristic. It appears that this is mostly a right-wing phenomenon, perhaps arising from the conservative tendency to portray the world outside the United States as hostile (Jost et al., 2003).

[23] A possible explanation for the low number of images returned by Google for the candidate search on the Huffington Post is the fact that the website has changed its domain name since the election, although old URLs are still accessible. Similarly, the New York Times has a very large website with multiple subdomains, some of which Google may have missed.

Outlet	Total	Clinton	Trump	Clinton %	Trump %
Breitbart	166	28	138	16.87	83.13
Chicago Tribune	121	35	86	28.93	71.07
CNN	188	96	92	51.06	48.94
FOX News	43	13	30	30.23	69.77
Huffington Post	108	29	79	26.85	73.15
MSNBC	147	70	77	47.62	52.38
New York Times	189	75	114	39.68	60.32
USA Today	134	41	93	30.6	69.4
Wall Street Journal	62	27	35	43.55	56.45
All	1158	414	744		

Table 2: *Rally dataset – number of images, by news outlet and candidate rally. Note that the number of faces, as opposed to images is much higher, as each image contains at least three faces. See figure 18 in the appendix for an overview.*

CNN (49%), MSNBC (52%) and the New York Times (60%) show less of a bias towards Trump. That being said, the Huffington Post (73%) and the Wall Street Journal (56%) defy this partisan trend. So what does this entail? Table 1, viewed on its own, might give off the impression of liberal media bias, whereas table 2 points in the opposite direction. The most plausible explanation is that since Donald Trump placed such great emphasis on his rallies, the media followed his lead and provided plenty of coverage of these events. Meanwhile, Clinton ran a more traditional campaign, and as a result, appeared in a higher number of images that featured only herself. This is supported by the fact that the total number of images – both rallies and the candidates themselves – is fairly even, 1947 for Clinton and 1876 for Trump.

An alternative to analyzing the quantity of pictures is to look at their quality. Figures 2 and 3 show the average number of pixels,[24,25] broken down by news outlet and campaign. For the candidate portraits, pictures of Clinton tend to be larger – except for the Chicago Tribune, Huffington Post and Fox News. For the rallies, Trump once again appears to be clearly favored, and depending on the website, quite heavily so. Breitbart, the Huffington Post and USA Today (all of which hosted more than 70% Trump rally images) show a particularly large gap in terms of image quality.[26]. Strikingly, not a single website appears feature more high-quality images of Clinton rallies. The reason for this divergence is likely economic. High-quality images are more expensive to shoot as well as to host, so will likely only be used if they can drive pageviews and thus increase revenue for the outlet. The greater commercial appeal of covering the Trump campaign (as noted by the

[24] Analyzing file size instead of number of pixels leads to nearly identical results.

[25] This analysis was conducted on the original images, before they were resized as described above.

[26] There is also a very large gap for USA Today among the candidate images, but this is likely due to the small sample size

candidate himself) means that Trump's rallies received favorable coverage, but Clinton was somewhat compensated through the higher quality images of herself.

Figure 2: *Candidate dataset – average number of pixels per picture, by news outlet and candidate. For most outlets, images of Clinton are larger.*

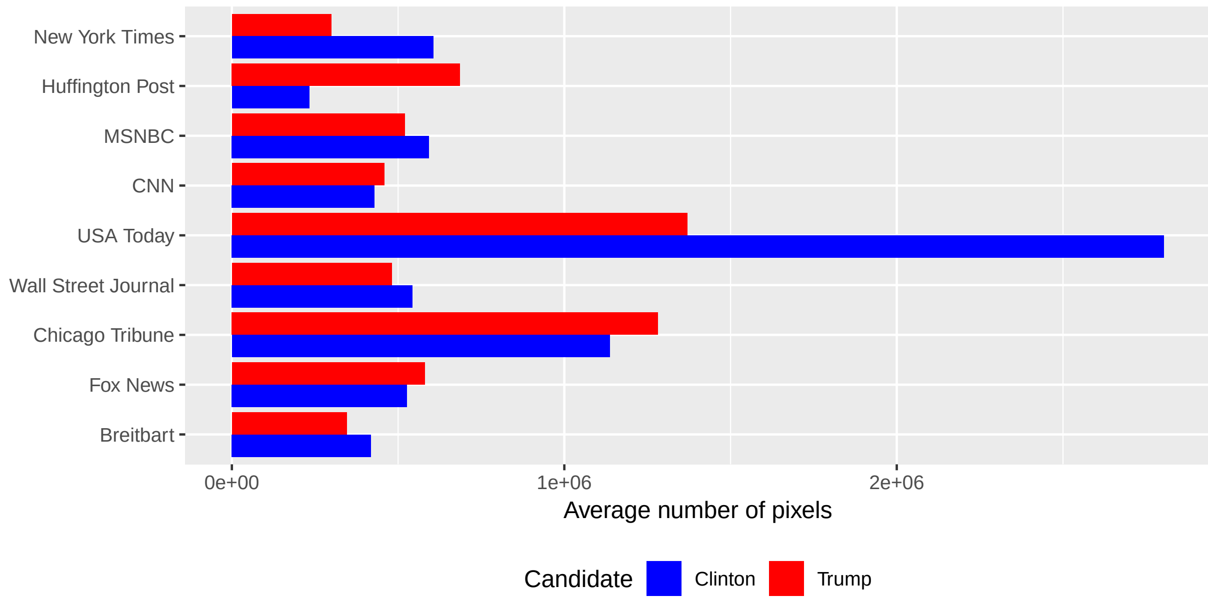

Figure 3: *Rally dataset – average number of pixels per picture, by news outlet and candidate. For every news outlet, images of Trump are, on average, larger.*

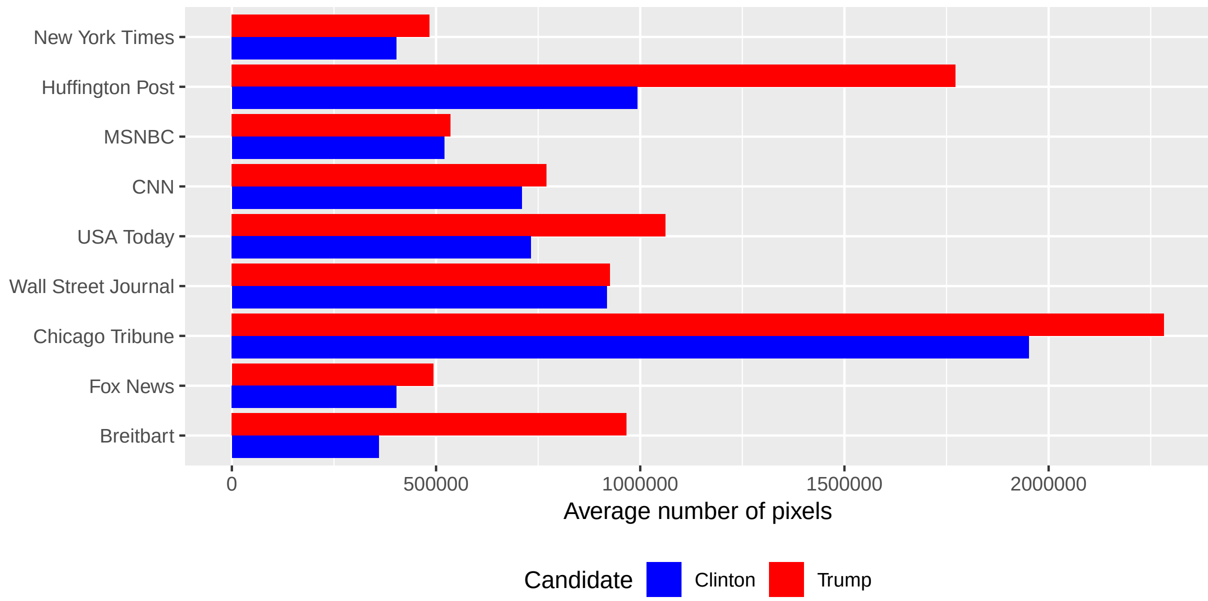

Results

In order to assess whether the news media portrays emotions in photos of candidates and campaign rallies in a biased manner, I employ a multinomial logit model, in which individual faces are the unit of observation. As noted above, the output of the fully connected layer of the neural network is a set of probabilities denominating the likelihood of each emotion being displayed. This result is turned into the categorical dependent variable of the model, where the expressed emotion is the one with the highest probability. The independent variables I am substantively interested in are (1) the candidate whose rally a particular supporter is attending and (2) the news outlet the corresponding photo is appearing on. Specifically for assessing media bias, an interaction term between the two measures the effect of partisanship on emotion, contingent on the media source.[27]

Candidates Dataset

For the candidates dataset, figure 4 shows the predicted probability of the respective candidate exhibiting happiness on an image, depending on which news outlet is using said image (the order of the news websites in the plot is based on the partisan composition of their readership, according to Mitchell et al. (2014)). Two observations can be made here. One, while the point estimates generally predict Hillary Clinton to be slightly more happy, the 95% confidence intervals overlap with those of Donald Trump, meaning that there is no statistically significant difference between them. Two, while the point estimates for Hillary Clinton for Breitbart and Fox News, ostensibly the two most conservative news outlets in the sample, are lower than for most of the other websites, they are not the lowest, and the confidence intervals once again overlap. This means that these two news sources don't publish less happy images of Clinton than the rest of the sample. The results for Trump are similar. Overall, this means that there does not appear to be any media bias in the selection of the candidate images, as measured by the computer vision model.

There is a second, simpler way to assess this hypothesis. Instead of detecting the emotion displayed in each image automatically, I aggregate all of the images belonging to a candidate, for each news source. When an image is converted to greyscale, each pixel is represented by a value between 0 (black) and 255 (white) (as opposed to three values for red, green and blue). Consequently the result is a matrix with the same dimensions as the image. An average image can be generated by coercing all images to the same dimensions (this will result in some stretching), calculating the average value for each pixel. This

[27] Tables 3 (candidate dataset) and 4 (rally dataset) in the appendix show the full regression tables.

Figure 4: *Candidate dataset – predicted probabilitiy of exhibiting happiness, by news outlet and candidate. The point estimates are consistently higher for Clinton, but the confidence intervals mostly overlap, so no statistically significant difference between the candidates exists. There are no systematic differences between left- (top) and right-leaning (bottom) media outlets, hence no media bias.*

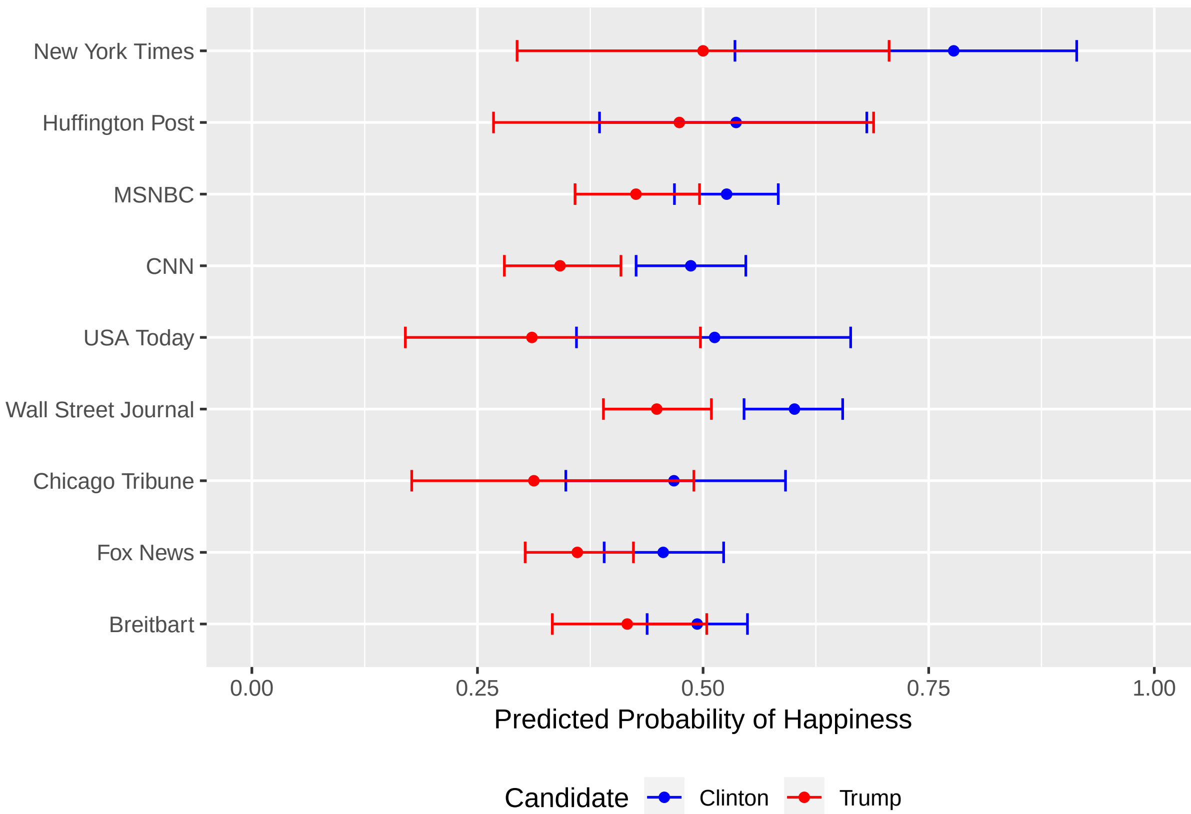

picture then represents how the subject(s) of the image look(s) on average. Figure 5 shows this kind of plot for Hillary Clinton for Breitbart (left) and MSNBC (right). While smaller details (nose, eyebrows, etc.) have vanished due to variation in the sample, schematic forms of the eyes and mouth remain. The average MSNBC image of Clinton looks happier, as the corners of the mouth are raised and the eyes are rotated. See figures 12 and 13 for the complete set of comparisons. Overall, the average Breitbart face of Clinton appears to be the least happy, and her average Fox News face looks similarly glum. Hence, the average face approach does offer some evidence of media bias, albeit at the detriment of once again relying on the reader's own judgment. That being said, the aggregation process obfuscates the face to a degree where an automatic partisan reaction to the identity of the candidate might not be as strong, making it easier to remain impartial about it.

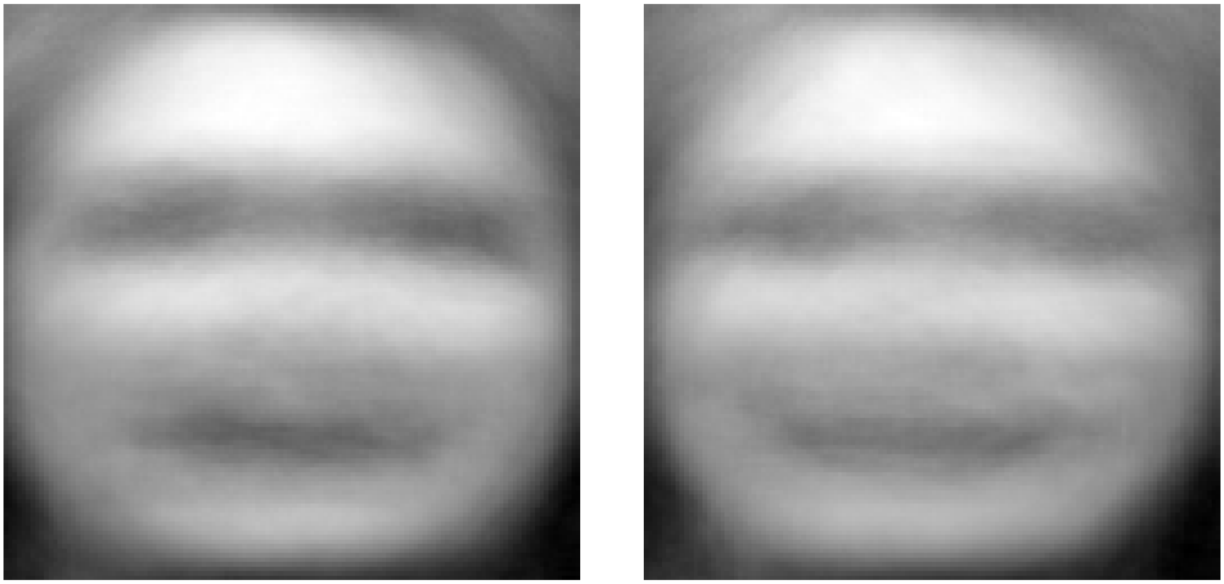

(a) *Breitbart* **(b)** *MSNBC*

Figure 5: *Average faces of Clinton for Breitbart and MSNBC, constructed by coercing all images to the same size and averaging the grayscale value of each pixel. The average MSNBC image looks friendlier than the one from Breitbart, indicating some small degree of media bias.*

Rally Dataset

In addition to the two main independent variables discussed above (candidate and news outlet), the model predicting emotion for rally attendants also features controls for age and gender of the face, as well as the number of days until the election, from when the image was published (since happiness has been found to be more common earlier, and anger more prolific later in a campaign (Ridout and Searles, 2011)). To model how the prospects of victory (and thus the optimism of supporters) at a given point in time factor into the equation, I use the distance to the other candidate in the horse race poll average of FiveThirtyEight on that day[28,29].

Figure 6 shows the results of this model with regard to happiness. The predicted probability of being happy for Clinton voters generally appears to be considerably higher than for attendants of Trump rallies - the Huffington Post is the only news outlet where no difference can be observed. Given Clinton's lead in the polls and the fact that Democrats were the incumbent party and therefore have reasons to be satisfied with the status quo,

[28]If Hillary Clinton is 3 percentage points ahead of Donald Trump, the value would be 3 for supporters at a Clinton rally, and -3 for supporters at a Trump rally.

[29]Since this data is only available for the general election campaign, this restricts the sample to the time between June 8, and November 8, 2016. As a result, the number of pictures in the sample is reduced by more than half. I have run the same analysis without the polling data variable, for the full sample. With regard to media bias, nothing changes.

Figure 6: *Rally dataset – predicted probabilitiy of exhibiting happiness, by news outlet and candidate rally. Clinton supporters are clearly happier than Trump supporters. There are no systematic differences between left- (top) and right-leaning (bottom) media outlets, hence no media bias.*

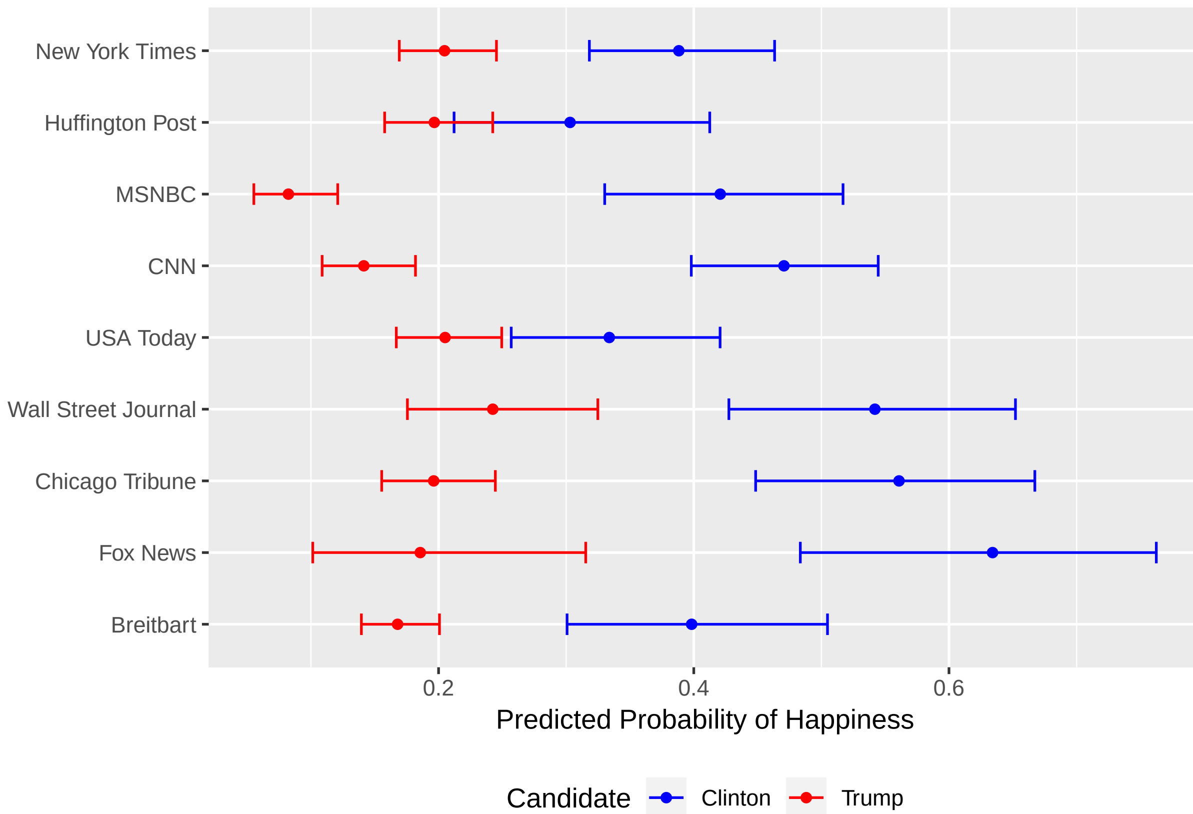

this makes sense. However, there is no clear evidence of media bias: For both Clinton and Trump rallies, the confidence intervals of the different media sources largely overlap. Even in the few cases in which they do not, no systematic bias consistent with the hypothesis of partisan cheerleading is detectable. If anything, conservative websites, compared to liberal ones, actually show Clinton supporters as more happy.

Average faces of supporters, comparable to the method described above, were also constructed and can be found in figure 11. Unlike for the faces of Clinton herself, no systematic differences between news outlets are apparent, although the faces of Clinton supporters clearly appear more happy (and more female) compared to Trump's followers.

Overall, partisan supporters do show different emotions, but the media does not appear to portray these emotions in a biased manner.

Conclusion

So why, contrary to my expectations, do I not observe partisan media bias? For one, the selection of images by media outlets is likely driven, at least to some degree, by purely practical reasons. Images that convey the desired message might not always be available, for example because no photo of smiling Democrats has been taken in the days preceding the news article. Similarly, factors such as image size and content may be playing a role in meeting format requirements. In many cases, thumbnails (i.e. smaller versions) of an article's image are shown on a website's frontpage. Not all images resize equally well however - a photo showing a huge crowd looks considerably worse when downsized, compared to a picture featuring only the candidate. Furthermore, assets like in-house photographers, rather than images from the Associated Press or Getty Images may be prioritized. And even if a media outlet is relying on stock photography, it may simply choose whichever option is cheapest. Ergo, even if writers intended to use photos in a way that is concurrent with their ideology, there is no guarantee they would always be able to do so.

Furthermore, there is a possibility that the multitude of authors employed by a news source may be diluting the message. It seems probable that a website's staff occupies different positions on the ideological spectrum, which means that an outlet's election coverage might not be representative of its ideological position as a whole. I also do not have any information on who selects a picture - it might be the person writing the article, a photo editor, or someone else entirely - a factor which likely also varies between the media sources covered here.

Most important however, is the argument that the ideological incentives of media outlets are likely not as straightforward as they seem. On the face of it, any news outlet would likely want to portray its side as the happiest, both to prove the positivity of its message, as well as the success of its campaign. However, part of the conservative, and particularly Donald Trump's message in 2016, has been outrage over the status quo. If conservative news outlets did want to engage in partisan cheerleading, they may well have been trying to portray that anger. Concurrently, happy Clinton photos may serve the purpose of portraying her as arrogant, another part of the conservative narrative in 2016. Furthermore, showing the opposing candidate and his supporters as happy might serve a purpose: In a polarized media environment, the message that the "other side" is winning motivates partisans and binds them even closer to their preferred news outlet. Lack of overt media bias can thus be explained with purely economic reasons.

Finally, it should also be noted that in the field of media bias, null results are a fairly common finding. Just because the flawed and ideologically-colored perceptions of humans lead us to believe that media bias exists, does not make it so.

In addition to these substantive considerations, the methodological aspects of this paper also merit further discussion. At least for the rally dataset, the greatest limitation lies in the quality of the pictures, and more specifically, that of the faces. On many photos, the crowd is in the background, out of focus, and frequently faces are in profile rather than visible from the front. Deep neural networks cope with these issues better than other computer vision techniques, but even they are ultimately only as good as the data they rely on.

Computer vision has not been widely used in political science, but it has started to receive some attention recently. For example, Torres (2018) uses local key point detection in order to build a bag of visual words for images of political protest, identifying components of images that are politically salient. Notably, this approach does not require the use of neural networks and relies on more traditional (and arguably intuitive) features instead. Casas and Williams (2017) also apply computer vision techniques to images of protests as well as the emotions they evoke, relying on the popular AlexNet architecture. Object recognition, which has developed tremendously in recent years, has also seen its first use in political science (Anastasopoulos et al., 2016). Other social sciences have also applied computer vision techniques. For example, satellite images have been employed to estimate population size (Sutton et al., 2001) and wealth (Sutton, Elvidge and Ghosh, 2007) through nighttime lights, an idea that could easily be transferred to other socially relevant concepts. The media is an area of research for which machine learning also holds great promise: Americans get their news primarily through television, and yet, political scientists have largely focused on the content of print media so far, an oversight that might very well be rectified in the coming years. Neural networks are able to make use of high-dimensional data to a much greater degree than conventional methods. The news media provides such data: whether it is through images, text or even audio, deep learning has much in store for the study of the media and the way it shapes the political environment.

References

Aday, Sean. 2010. "Chasing the bad news: An analysis of 2005 Iraq and Afghanistan war coverage on NBC and Fox News channel." *Journal of Communication* 60(1):144–164.

Anastasopoulos, L. Jason, Dhruvil Badani, Crystal Lee, Shiry Ginosar and Jake Williams. 2016. "Photographic home styles in Congress: a computer vision approach.".

Banning, Stephen A. and Renita Coleman. 2009. "Louder than Words: a content analysis of presidential candidate's televised nonverbal communication." *Visual Communication* 16:4–17.

Barrett, a. W. and J. S. Peake. 2007. "When the President Comes to Town: Examining Local Newspaper Coverage of Domestic Presidential Travel." *American Politics Research* 35(1):3–31.

Baum, M and T Groeling. 2008. "New Media and the Polarization of American Political Discourse." *Political Communication* 25(4):345–365.

Bennett, W. Lance. 2012. "The Personalization of Politics: Political Identity, Social Media, and Changing Patterns of Participation." *Annals of the American Academy of Political and Social Science* 644(1):20–39.

Brader, Ted. 2005. "Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions." *American Journal of Political Science* 49(2):388–405.

Budak, Ceren, Sharad Goel and Justin M Rao. 2016. "Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis." *Public Opinion Quarterly* 80(S1):250–271.

Burton, Leslie A., Laura Rabin, Gwinne Wyatt, Jonathan Frohlich, Susan Bernstein Vardy and Diana Dimitri. 2005. "Priming effects for affective vs. neutral faces." *Brain and Cognition* 59(3):322–329.

Butler, D. M. and E. Schofield. 2010. "Were Newspapers More Interested in Pro-Obama Letters to the Editor in 2008? Evidence From a Field Experiment." *American Politics Research* 38(2):356–371.

Caliskan-Islam, Aylin, Joanna Bryson and Arvind Narayanan. 2017. "A Story of Discrimination and Unfairness: Implicit Bias Embedded in Language Models.".

Caliskan-islam, Aylin, Joanna J Bryson and Arvind Narayanan. 2016. "Semantics derived automatically from language corpora necessarily contain human biases." *arXiv* pp. 1–14.

Casas, Andreu and Nora Webb Williams. 2017. "Computer Vision for Political Science Research: A Study of Online Protest Images.".

Chatfield, Ken, Karen Simonyan, Andrea Vedaldi and Andrew Zisserman. 2014. "Return of the Devil in the Details: Delving Deep into Convolutional Nets." *BMVC* .
URL: *http://arxiv.org/abs/1405.3531*

Civettini, Andrew J W and David P Redlawsk. 2009. "Voters, Emotions, and Memory." *Political Psychology* 30(1):125–151.

Fazio, Russell H. 2001. "On the automatic activation of associated evaluations: An overview." *Cognition and Emotion* 15(2):229–238.

Friedman, Howard S., M. Robin DiMatteo and Timothy I. Mertz. 1980. "Nonverbal Communication on Television News: The Facial Expressions of Broadcasters during Coverage of a Presidential Election Campaign." *Personality and Social Psychology Bulletin* 6(3):427–435.

Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence From U.S. Daily Newspapers." *Econometrica* 78(1):35–71.

Glazier, Rebecca a. and Amber E. Boydstun. 2012. "The President, the Press, and the War: A Tale of Two Framing Agendas." *Political Communication* 29(4):428–446.

Groeling, Tim. 2013. "Media Bias by the Numbers: Challenges and Opportunities in the Empirical Study of Partisan News." *Annual Review of Political Science* 16(1):129–151.

Guerra, Pedro H Calais, Adriano Veloso, Wagner Meira Jr and Virgílio Almeida. 2011. "From Bias to Opinion: A Transfer-Learning Approach to Real-Time Sentiment Analysis." *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining* pp. 150–158.

Hatemi, Peter K. and Rose McDermott, eds. 2011. *Man Is by Nature a Political Animal: Evolution, Biology, and Politics*. Chicago: University of Chicago Press.

Hehman, Eric, Elana C Graber, Lindsay H Hoffman and Samuel L Gaertner. 2012. "Warmth and Competence : A Content Analysis of Photographs Depicting American Presidents." *Psychology of Popular Media Culture* 1(1):46–52.

Huddy, Leonie, Lilliana Mason and Lene Aarøe. 2015. "Expressive Partisanship: Campaign Involvement, Political Emotion, and Partisan Identity." *American Political Science Review* 109(1):1–17.

Jerit, Jennifer. 2004. "Survival of the Fittest: Rhetoric during the Course of an Election Campaign." *Political Psychology* 25(4):563–575.

Jones, P. E., L. H. Hoffman and D. G. Young. 2012. "Online emotional appeals and political participation: The effect of candidate affect on mass behavior." *New Media & Society* pp. 1132–1150.

Jost, John T., Jack Glaser, Arie W. Kruglanski and Frank J. Sulloway. 2003. "Political conservatism as motivated social cognition." *Psychological Bulletin* 129(3):339–375.

Kiss, Zsolt and Sara B. Hobolt. 2011. The Emotional Voter. An Experimental Study of the Moderating Effect of Emotions on Partisan Behavior. In *APSA 2011 Annual Meeting Paper*. pp. 1–35.

Larcinese, Valentino, Riccardo Puglisi and James M. Snyder. 2011. "Partisan bias in economic news: Evidence on the agenda-setting behavior of U.S. newspapers." *Journal of Public Economics* 95(9-10):1178–1189.

Levi, Gil and Tal Hassner. 2015. "Emotion Recognition in the Wild via Convolutional Neural Networks and Mapped Binary Patterns." *Proceedings of the 2015 ACM on international conference on multimodal interaction.* pp. 503–510.

Lin, Y-R, J P Bagrow and D Lazer. 2011. "More Voices than Ever? Quantifying Bias in Social and Mainstream Media." *arXiv preprint arXiv* 1111(1227).

Lowi, Theodore J. 1986. *The Personal President: Power Invested, Promise Unfulfilled*.

Masuda, Takahiko, Huaitang Wang, Keiko Ishii and Kenichi Ito. 2012. "Do surrounding figures' emotions affect judgment of the target figure's emotion? Comparing the eye-movement patterns of European Canadians, Asian Canadians, Asian international students, and Japanese." *Frontiers in Integrative Neuroscience* 6(September):1–9.

McAllister, Ian and Donley T. Studlar. 1991. "Bandwagon, Underdog, or Projection? Opinion Polls and Electoral Choice in Britain, 1979-1987." *The Journal of Politics* 53(3):720–741.

Mitchell, Amy, Jeffrey Gottfried, Jocelyn Kiley and Katerina Eva Matsa. 2014. "Political Polarization & Media Habits.".
URL: *http://www.journalism.org/interactives/media-polarization/*

Mitchell, Amy, Jeffrey Gottfried, Michael Barthel and Elisa Shearer. 2016. "How Americans get their news.".
URL: *http://www.journalism.org/2016/07/07/pathways-to-news/*

Moriarty, S. E. and M. N. Popovich. 1991. "Newsmagazine visuals and the 1988 Presidential election." *Journalism & Mass Communication Quarterly* 68(3):371–380.

Moriarty, Sandra E. and Gina M. Garramone. 1986. "A Study of Newsmagazine Photographs Of the 1984 Presidential Campaign." *Journalism Quarterly* 63(4):728–734.

Mullen, Brian, David Futrell, Debbie Stairs, Dianne M Tice, Kathryn E Dawspn, Catherine A Riordan, John G Kennedy, Roy F Baumeister, Christine E Radloffand, George R Goethals and Paul Rosenfeld. 1986. "Newscasters' Facial Expressions and Voting Behavior of Viewers: Can a Smile Elect a President?" *Journal of Personality and Social Psychology* 51(2):291–295.

Ojala, Timo, Matti Pietikäinen and Topi Mäenpää. 2002. "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24(7):971–987.

Pan, Bing, Geri K Gay, Helene A Hembrooke, Laura A Granka, Matthew K Feusner and Jill K Newman. 2004. "The Determinants of Web Page Viewing Behavior: An Eye-Tracking Study." *ETRA '04 Proceedings of the 2004 symposium on Eye tracking research & applications* 1(212):147–154.

Parks, Amanda Jo. 2016. The Competitive Communications Environment: How the News Media Report and Distort Economic News PhD thesis.

Ridout, Travis N. and Kathleen Searles. 2011. "It's My Campaign I'll Cry if I Want to: How and When Campaigns Use Emotional Appeals." *Political Psychology* 32(3):439–458.

Risko, Evan F., Nicola C. Anderson, Sophie Lanthier and Alan Kingstone. 2012. "Curious eyes: Individual differences in personality predict eye movement behavior in scene-viewing." *Cognition* 122(1):86–90.
URL: *http://dx.doi.org/10.1016/j.cognition.2011.08.014*

Sandoval, Greg. 2005. "Breitbart.com has Drudge to thank for its success.".
URL: *https://www.cnet.com/news/breitbart-com-has-drudge-to-thank-for-its-success/*

Soroka, Stuart N. 2012. "The Gatekeeping Function: Distributions of Information in Media and the Real World." *The Journal of Politics* 74(2):514–528.

Spezio, Michael L and Ralph Adolphs. 2006. "Emotional Processing and Political Judgment: Toward Integrating Political Psychology and Decision Neuroscience.".

Sutton, P, D Roberts, C Elvidge and K Baugh. 2001. "Census from Heaven: an estimate of the global human population using night-time satellite imagery." *Int. J. Remote Sensing* 22(16):3061–3076.

Sutton, Pc, Cd Elvidge and Tilottama Ghosh. 2007. "Estimation of gross domestic product at sub-national scales using nighttime satellite imagery." *International Journal of Ecological Economics & Statistics* 8(S07):5–21.

Torres, Michelle. 2018. "Give me the full picture: Using computer vision to understand visual frames and political communication.".

Vigil, J. M. 2010. "Political leanings vary with facial expression processing and psychosocial functioning." *Group Processes & Intergroup Relations* 13(5):547–558.

Waldman, Paul and James Devitt. 1998. "Newspaper Photographs and the 1996 Presidential Election: The Question of Bias." *Journal of Mass Communication* 75(2):302–311.

Weber, C. 2013. "Emotions, Campaigns, and Political Participation." *Political Research Quarterly* 66(2):414–428.

Westen, Drew, Pavel S Blagov, Keith Harenski, Clint Kilts and Stephan Hamann. 2006. "Neural bases of motivated reasoning: an FMRI study of emotional constraints on partisan political judgment in the 2004 U.S. Presidential election." *Journal of cognitive neuroscience* 18(11):1947–1958.

Zhang, Shifeng, Xiangyu Zhu, Zhen Lei, Hailin Shi, Xiaobo Wang and Stan Z. Li. 2017. "S3FD: Single Shot Scale-Invariant Face Detector." *Proceedings of the IEEE International Conference on Computer Vision* 2017-Octob:192–201.

Appendix 1 – Tables & Figures

Table 3: *Candidates dataset – multinomial regression. The table shows that Donald Trump is generally less happy than Hillary Clinton, but there is no media bias pertaining to the emotions of the candidates.*

	Dependent variable:					
	Disgust	Fear	Happy	Neutral	Sad	Surprise
	(1)	(2)	(3)	(4)	(5)	(6)
Trump	−2.096**	−1.046**	−0.866**	−0.347	−0.165	−0.484
	(0.830)	(0.481)	(0.406)	(0.456)	(0.473)	(1.280)
Chicago Tribune	0.870	1.583	1.110	1.157	1.191	−9.353***
	(1.200)	(1.070)	(1.048)	(1.099)	(1.122)	(0.0002)
CNN	−0.580	0.428	0.274	0.307	0.599	−0.222
	(0.596)	(0.452)	(0.421)	(0.466)	(0.477)	(1.290)
Fox News	0.095	0.562	0.254	0.741	−0.190	0.001
	(0.573)	(0.481)	(0.452)	(0.486)	(0.553)	(1.299)
Huffington Post	−14.121***	0.132	0.146	−0.224	0.320	1.387
	(0.00003)	(0.852)	(0.784)	(0.915)	(0.893)	(1.436)
MSNBC	−0.118	0.467	0.568	0.600	0.609	0.576
	(0.569)	(0.466)	(0.432)	(0.472)	(0.491)	(1.084)
New York Times	−0.223	−13.524***	0.388	−13.270***	−0.594	2.082
	(1.454)	(0.733)	(1.068)	(0.636)	(1.448)	(1.601)
USA Today	−0.917	−0.204	0.050	−0.001	0.320	1.387
	(1.270)	(0.885)	(0.787)	(0.887)	(0.893)	(1.436)
Wall Street Journal	0.113	0.266	0.660	0.113	0.398	0.876
	(0.533)	(0.456)	(0.418)	(0.472)	(0.484)	(0.998)
Trump x Chicago Tribune	−10.612***	0.294	−0.189	0.675	−0.518	−2.052***
	(0.002)	(1.562)	(1.516)	(1.556)	(1.650)	(0.00001)
Trump x CNN	1.760	−0.250	−0.050	0.313	0.224	0.148
	(1.071)	(0.692)	(0.599)	(0.656)	(0.668)	(1.952)
Trump x Fox News	2.001*	0.314	0.353	0.298	0.858	0.772
	(1.027)	(0.704)	(0.629)	(0.676)	(0.737)	(1.830)
Trump x Huffington Post	−0.610***	−13.466***	−0.027	0.060	0.166	−13.078***
	(0.00001)	(0.0001)	(1.150)	(1.338)	(1.274)	(0.00004)
Trump x MSNBC	1.623	−0.272	−0.100	−0.556	0.203	0.810
	(1.043)	(0.708)	(0.610)	(0.680)	(0.684)	(1.605)
Trump x New York Times	−8.283	13.449***	0.527	13.392***	1.772	0.481
	(179.156)	(0.733)	(1.529)	(0.636)	(1.850)	(2.375)
Trump x USA Today	2.791	1.924	0.763	0.532	1.448	−10.299***
	(2.047)	(1.449)	(1.352)	(1.497)	(1.427)	(0.0005)
Trump x Wall Street Journal	0.961	−0.391	−0.271	−0.015	0.202	−0.614
	(1.011)	(0.677)	(0.574)	(0.647)	(0.655)	(1.619)
Constant	0.224	1.120***	2.252***	0.917***	0.595*	−2.080***
	(0.335)	(0.288)	(0.263)	(0.296)	(0.311)	(0.750)
Akaike Inf. Crit.	8,047.323	8,047.323	8,047.323	8,047.323	8,047.323	8,047.323

Note: *p<0.1; **p<0.05; ***p<0.01

Table 4: *Rallies dataset – multinomial regression. The table shows that supporters of Donald Trump are generally less happy than those of Hillary Clinton, but there is no media bias pertaining to the emotions of the rally attendants.*

	Dependent variable:					
	anger	contempt	disgust	happiness	neutral	surprise
	(1)	(2)	(3)	(4)	(5)	(6)
Trump	-6.591^{***}	6.799^{***}	8.083^{***}	-3.273^{***}	-1.992^{**}	-1.602^{*}
	(1.231)	(0.478)	(0.682)	(0.809)	(0.799)	(0.939)
Chicago Tribune	-0.006	-0.025	-0.028	-0.032^{***}	-0.048^{***}	-0.045^{***}
	(0.011)	(0.033)	(0.053)	(0.008)	(0.008)	(0.009)
CNN	0.375	-0.399	3.553^{***}	-0.360	0.507	0.034
	(0.506)	(1.512)	(1.127)	(0.341)	(0.339)	(0.377)
Fox News	-17.592^{***}	-6.193^{***}	-4.004^{***}	-4.725^{***}	-5.379^{***}	-3.595^{***}
	(0.582)	(0.001)	(0.009)	(0.354)	(0.356)	(0.620)
Huffington Post	-17.816^{***}	-5.564^{***}	-3.700^{***}	-3.969^{***}	-4.191^{***}	-3.303^{***}
	(0.422)	(0.004)	(0.179)	(0.711)	(0.708)	(0.724)
MSNBC	-4.368^{***}	5.450^{***}	8.273^{***}	12.541^{***}	11.699^{***}	-7.378^{***}
	(0.468)	(0.00005)	(0.0005)	(0.248)	(0.248)	(0.325)
New York Times	-22.193^{***}	-9.495^{***}	-3.650^{***}	-5.864^{***}	-5.392^{***}	-4.421^{***}
	(0.487)	(0.0001)	(0.009)	(0.725)	(0.710)	(0.787)
USA Today	-11.321^{***}	6.278^{***}	5.480^{***}	9.194^{***}	9.199^{***}	9.561^{***}
	(0.607)	(0.0003)	(0.001)	(0.515)	(0.511)	(0.667)
Wall Street Journal	-5.905^{***}	-2.836^{***}	-3.004^{***}	-3.700^{***}	-3.568^{***}	-3.556^{***}
	(0.792)	(0.713)	(0.794)	(0.460)	(0.452)	(0.550)
Trump x Chicago Tribune	-5.885^{***}	-6.050^{***}	-4.749^{***}	-5.089^{***}	-4.767^{***}	-4.178^{***}
	(0.859)	(0.002)	(0.044)	(0.586)	(0.577)	(0.738)
Trump x CNN	4.068^{***}	-2.405^{***}	0.592^{***}	5.427^{***}	4.903^{***}	4.963^{***}
	(0.930)	(0.00000)	(0.00004)	(0.608)	(0.605)	(0.832)
Trump x Fox News	-0.003	0.030	-0.033	0.002	0.001	0.005
	(0.005)	(0.019)	(0.033)	(0.004)	(0.004)	(0.004)
Trump x Huffington Post	-0.165	0.402	1.290	-0.035	0.048	0.042
	(0.146)	(0.406)	(0.920)	(0.103)	(0.102)	(0.116)
Trump x MSNBC	15.218^{***}	-6.280^{***}	-2.010^{***}	2.543^{***}	2.972^{***}	1.987^{***}
	(0.582)	(0.00002)	(0.0002)	(0.345)	(0.340)	(0.596)
Trump x New York Times	16.349^{***}	-8.490^{***}	-0.587^{***}	1.040	1.429^{*}	1.191
	(0.422)	(0.00001)	(0.001)	(0.828)	(0.818)	(0.819)
Trump x USA Today	14.503^{***}	-3.839^{***}	-2.528^{***}	-4.520^{***}	-3.841^{***}	16.150^{***}
	(0.468)	(0.00005)	(0.0001)	(0.321)	(0.293)	(0.325)
Trump x Wall Street Journal	20.828^{***}	-4.510^{***}	2.478^{***}	3.884^{***}	3.204^{***}	2.532^{***}
	(0.487)	(0.00000)	(0.018)	(0.889)	(0.873)	(0.918)
Constant	11.313^{***}	-15.469^{***}	-7.219^{***}	-11.051^{***}	-10.254^{***}	-10.349^{***}
	(0.607)	(0.0002)	(0.002)	(0.568)	(0.546)	(0.694)
candidateTrump:websitenytimes	5.823^{***}	1.186^{*}	7.371^{***}	2.121^{***}	1.692^{***}	2.630^{***}
	(0.860)	(0.712)	(0.730)	(0.537)	(0.525)	(0.557)
candidateTrump:websiteusatoday	4.686^{***}	-6.087^{***}	0.466^{***}	2.814^{***}	2.202^{***}	2.510^{***}
	(1.040)	(0.00003)	(0.001)	(0.699)	(0.687)	(0.815)
candidateTrump:websitewsj	-3.031^{***}	-11.784^{***}	-6.525^{***}	-6.944^{***}	-6.941^{***}	-6.105^{***}
	(1.039)	(0.00000)	(0.0001)	(0.675)	(0.663)	(0.882)
Constant	6.693^{***}	-6.105^{***}	-9.853^{***}	9.860^{***}	10.316^{***}	6.180^{***}
	(0.854)	(0.467)	(1.078)	(0.568)	(0.561)	(0.693)
Akaike Inf. Crit.	7,951.389	7,951.389	7,951.389	7,951.389	7,951.389	7,951.389
Note:						*p<0.1; **p<0.05; ***p<0.01

Figure 7: *Candidates dataset – predicted probabilitiy of exhibiting anger, by news outlet and candidate. There are not noticeable differences between candidates or media outlets.*

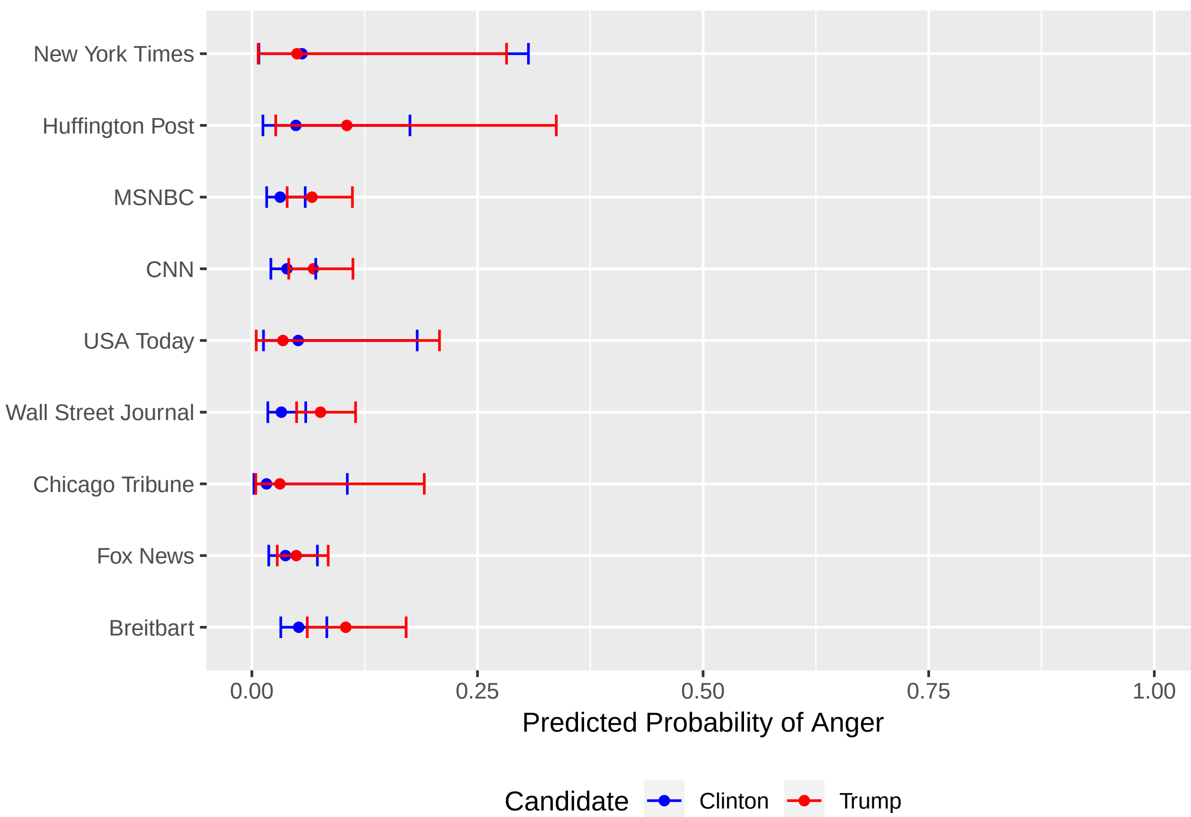

Figure 8: *Candidates dataset – predicted probabilitiy of exhibiting no emotion, by news outlet and candidate. There are not noticeable differences between candidates or media outlets.*

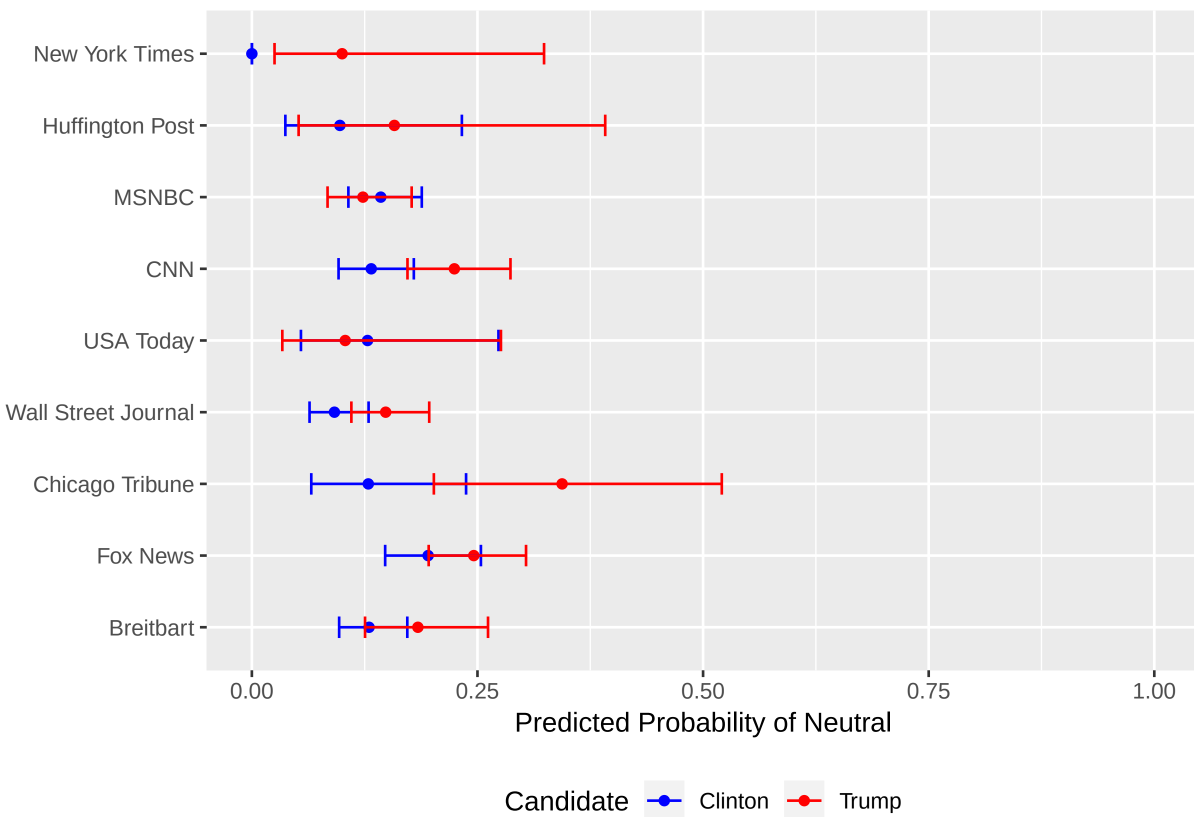

Figure 9: *Rallies dataset – predicted probabilitiy of exhibiting anger, by news outlet and candidate. There are not noticeable differences between candidates or media outlets.*

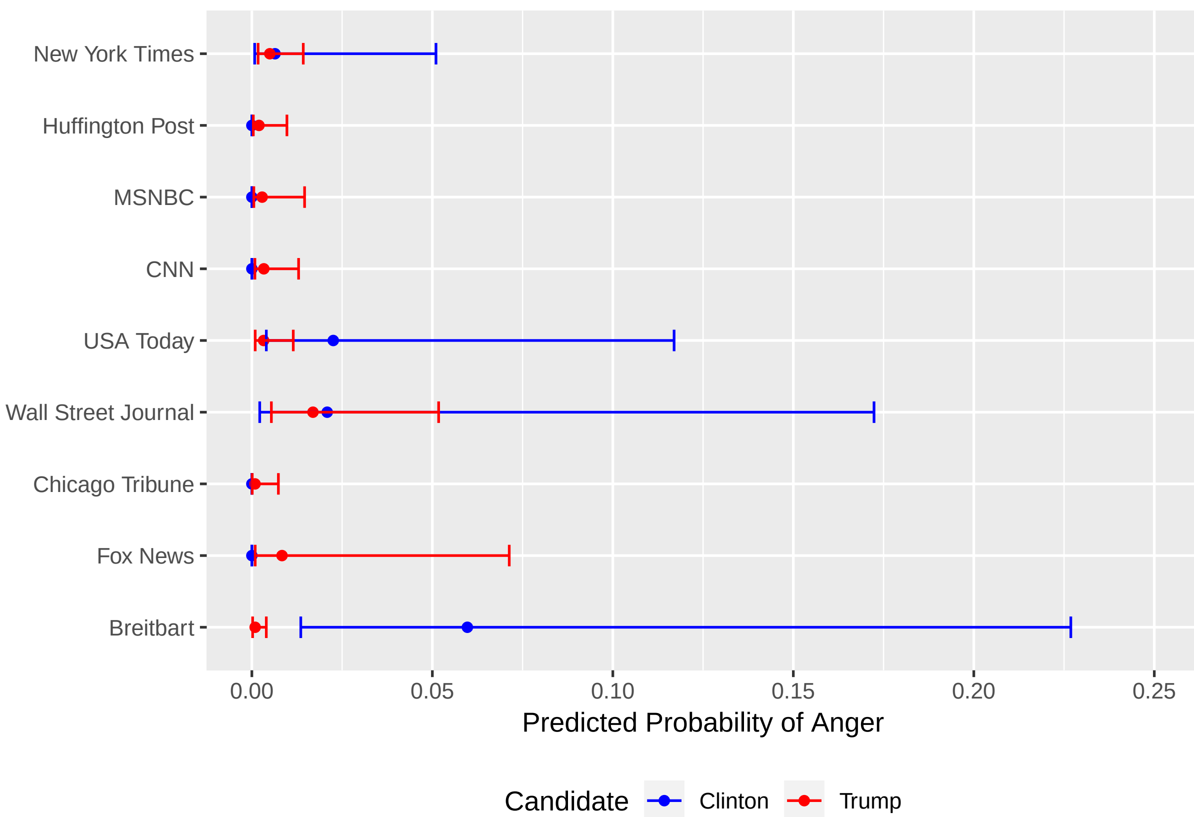

Figure 10: *Rallies dataset – predicted probabilitiy of exhibiting no emotion, by news outlet and candidate. Across all but two outlets (Huffington Post and USA Today), supporters at Trump rallies are significantly more likely to exhibit no emotions. There are no differences between media outlets.*

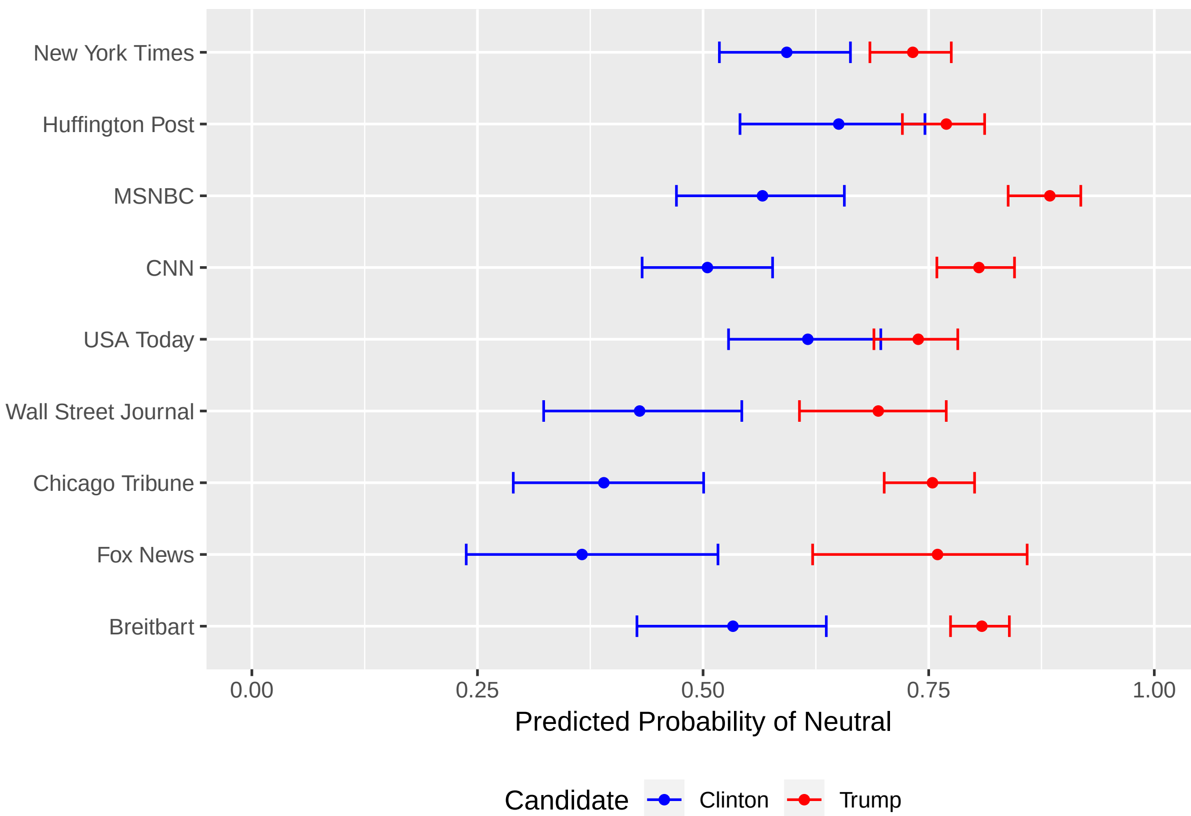

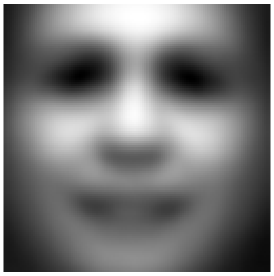

(a) *Clinton*

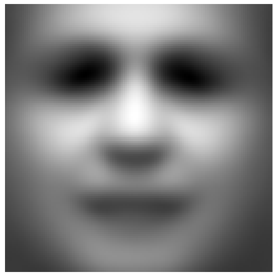

(b) *Trump*

(c) *Breitbart (Clinton)*

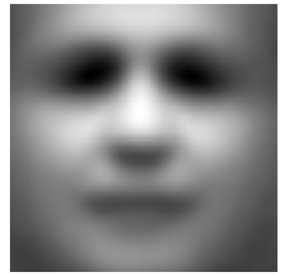

(d) *Breitbart (Trump)*

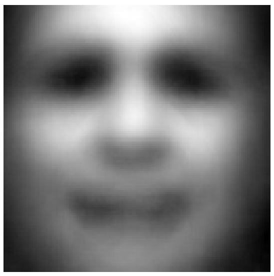

(e) *Chicago T. (Clinton)*

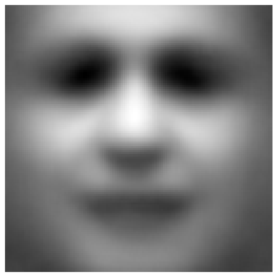

(f) *Chicago T. (Trump)*

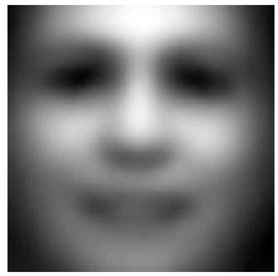

(g) *CNN (Clinton)*

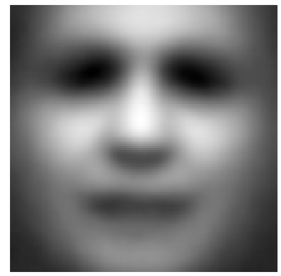

(h) *CNN (Trump)*

(i) *Fox (Clinton)*

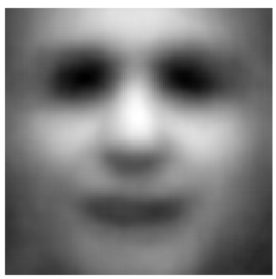

(j) *Fox (Trump)*

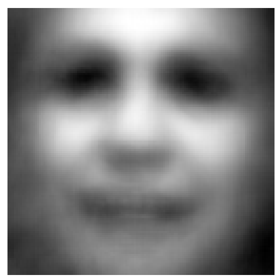

(k) *HuffPo (Clinton)*

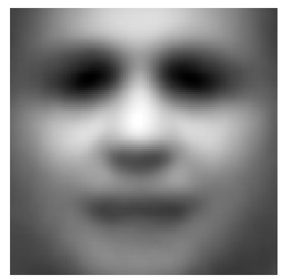

(l) *HuffPo (Trump)*

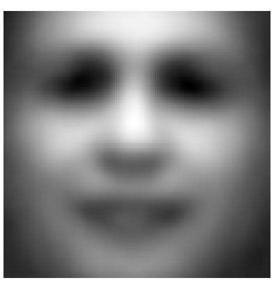

(m) *MSNBC (Clinton)*

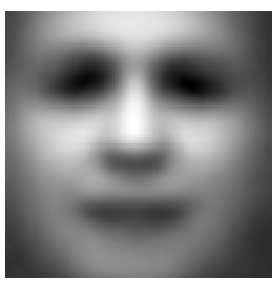

(n) *MSNBC (Trump)*

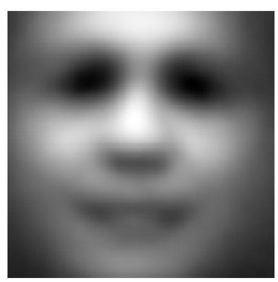

(o) *NYT (Clinton)*

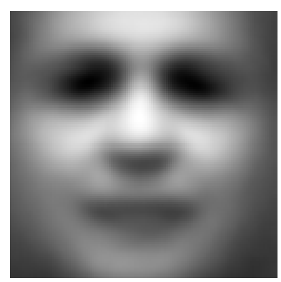

(p) *NYT (Trump)*

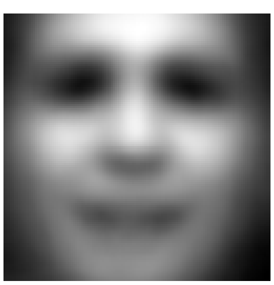

(q) *USA Today (Clinton)*

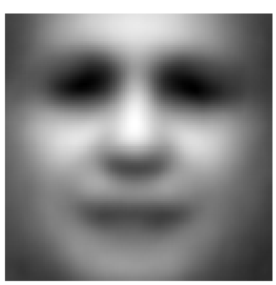

(r) *USA Today (Trump)*

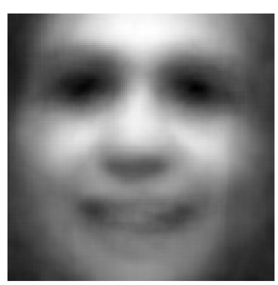

(s) *WSJ (Clinton)*

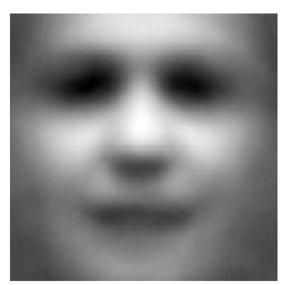

(t) *WSJ (Trump)*

Figure 11: *Average faces of rally supporters, by news outlet and candidate. The images were constructed by coercing all images to the same size and averaging the grayscale value of each pixel.*

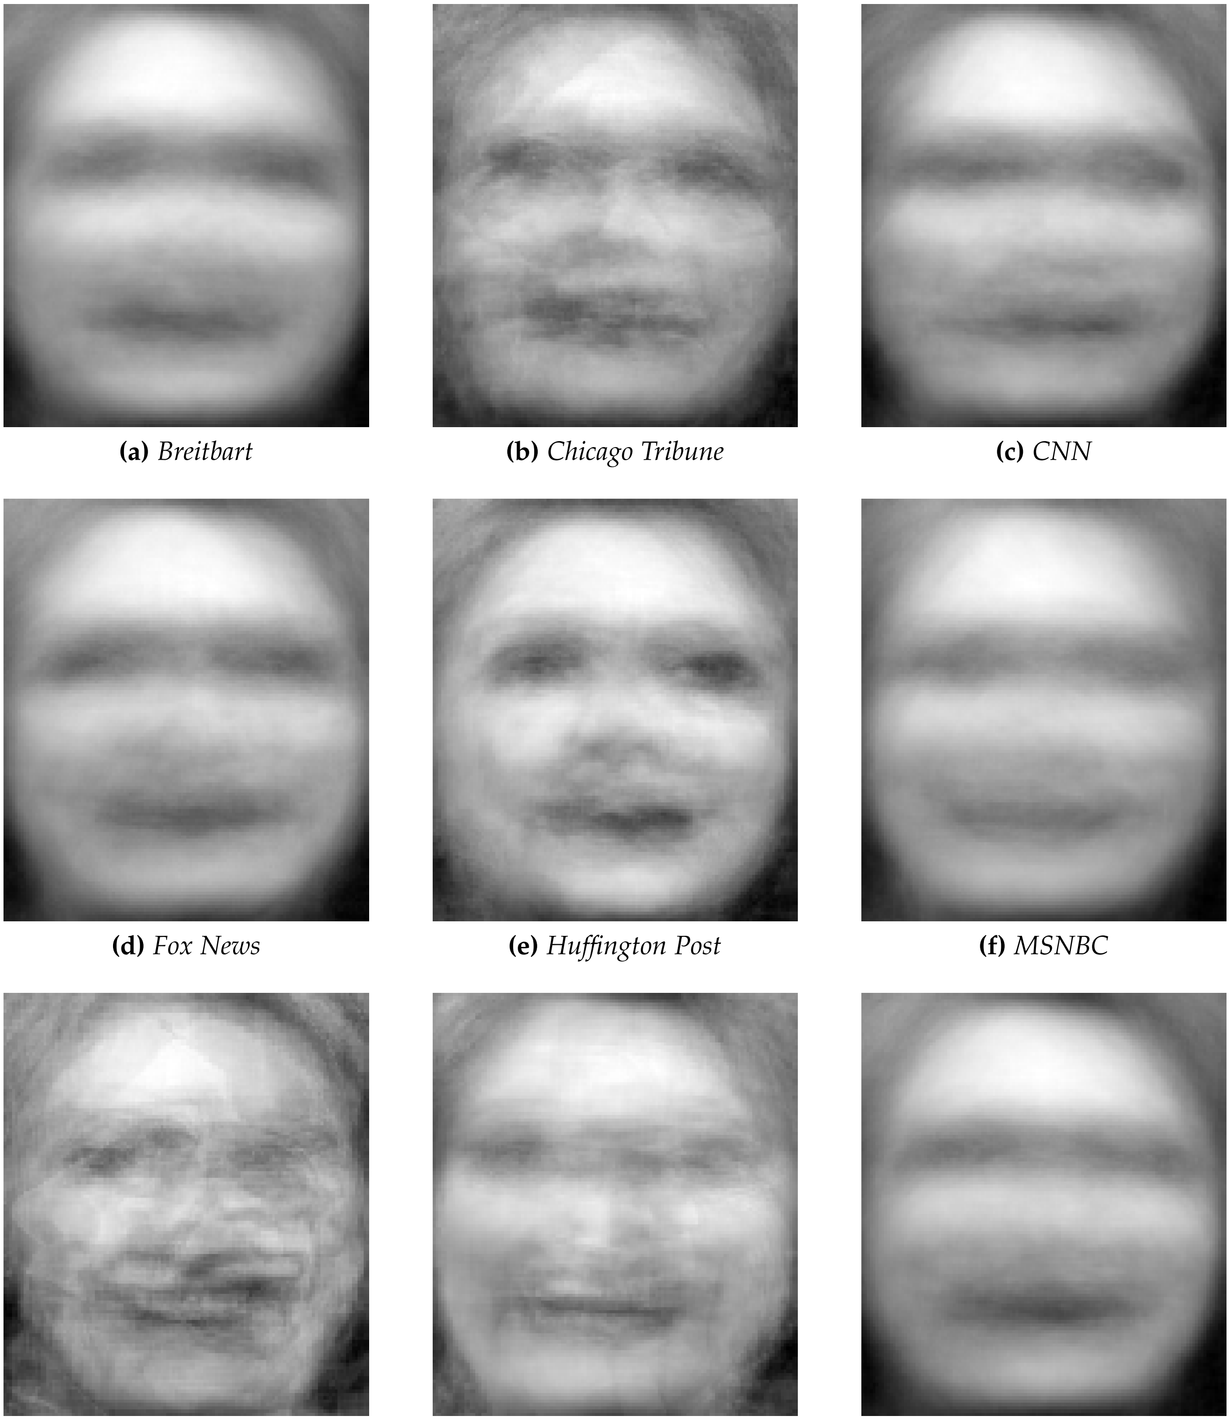

(a) *Breitbart* **(b)** *Chicago Tribune* **(c)** *CNN*

(d) *Fox News* **(e)** *Huffington Post* **(f)** *MSNBC*

(g) *The New York Times* **(h)** *USA Today* **(i)** *WSJ*

Figure 12: *Average faces of Clinton, by news outlet. The images were constructed by coercing all images to the same size and averaging the grayscale value of each pixel.*

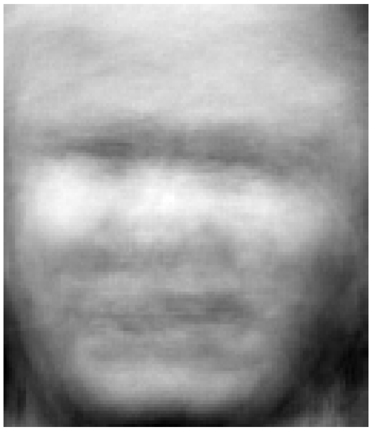

(a) *Breitbart*

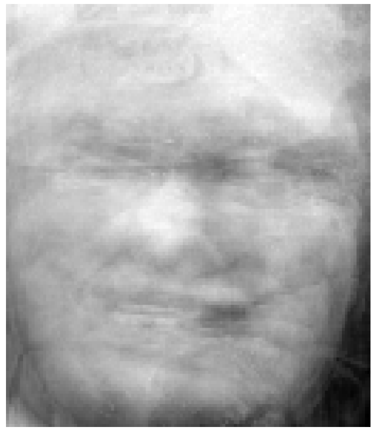

(b) *Chicago Tribune*

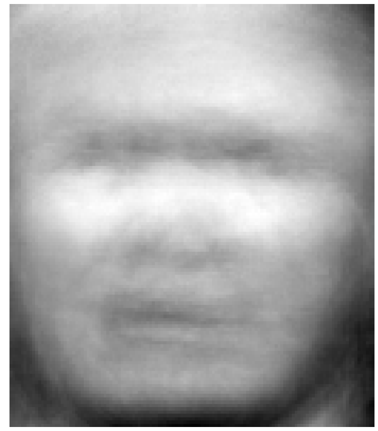

(c) *CNN*

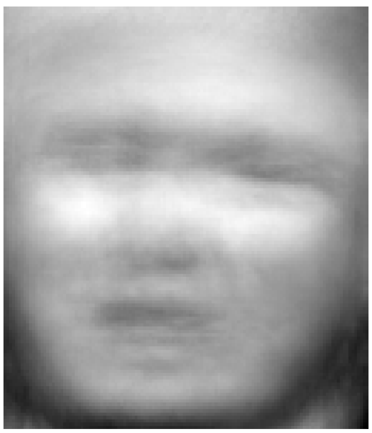

(d) *Fox News*

(e) *Huffington Post*

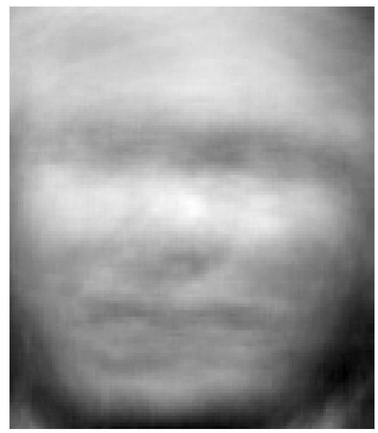

(f) *MSNBC*

(g) *The New York Times*

(h) *USA Today*

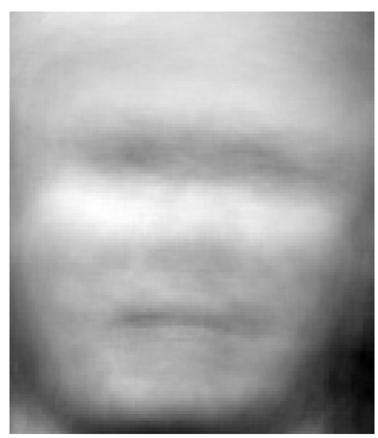

(i) *WSJ*

Figure 13: *Average faces of Trump, by news outlet. The images were constructed by coercing all images to the same size and averaging the grayscale value of each pixel.*

Figure 14: *Filters of the first convolutional layer of the emotion detection classifier. These filters are slid across the images, comparing each part of it to the shape of the filter.*

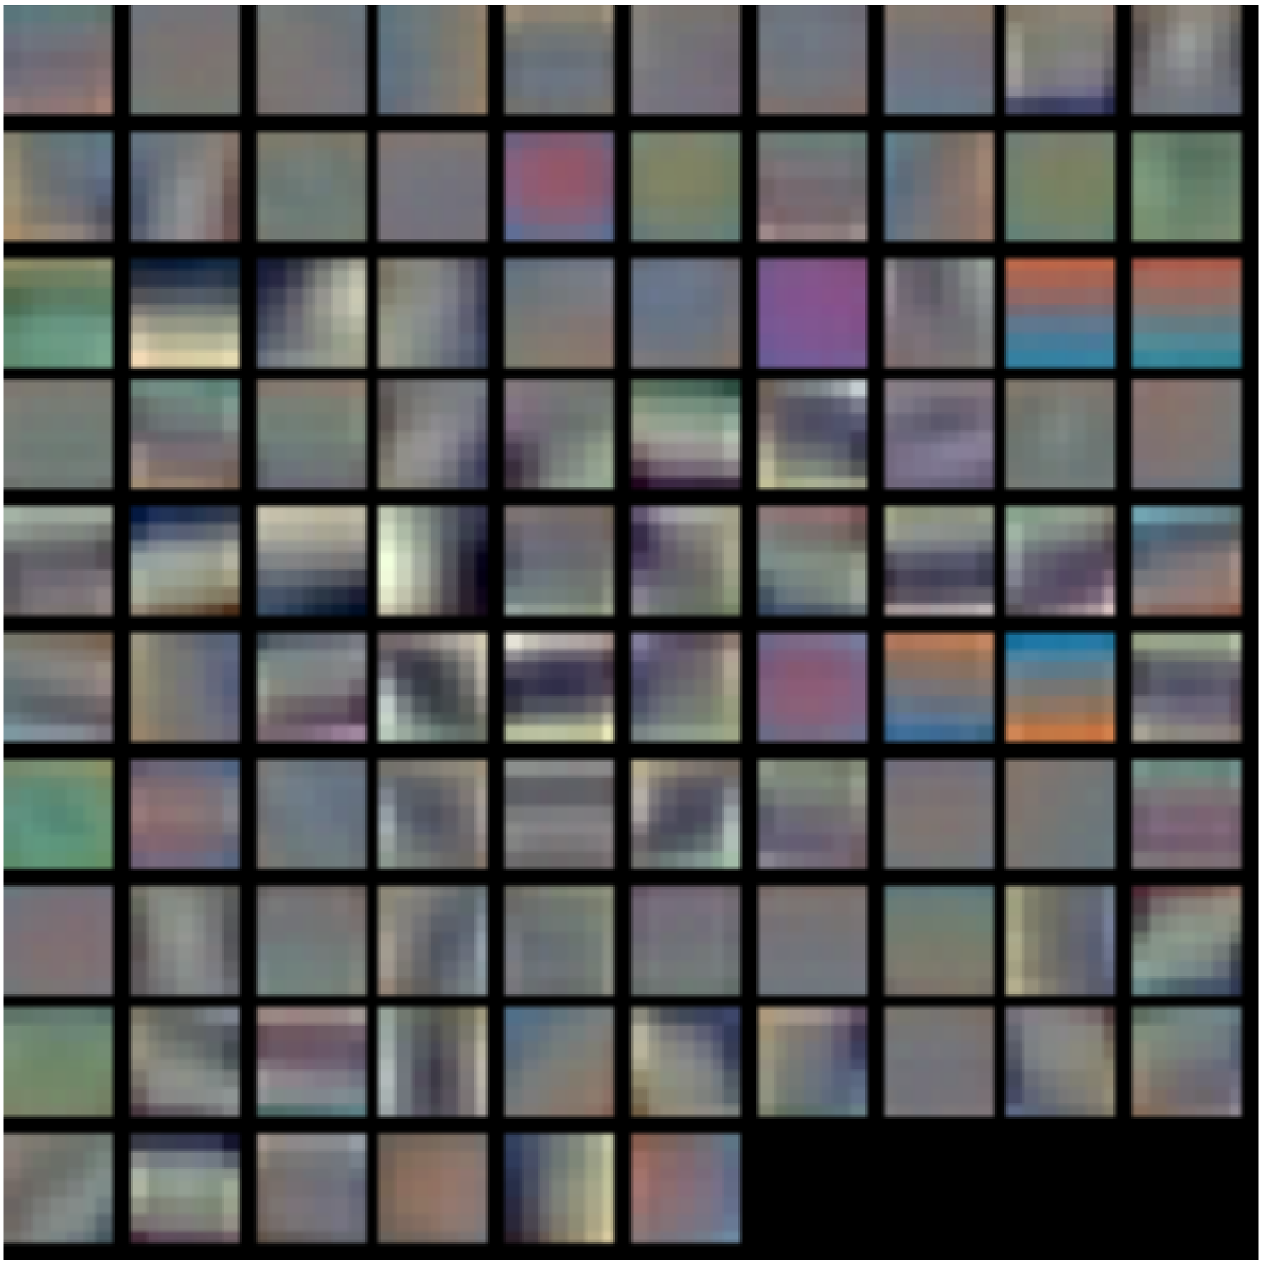

Figure 15: *Rectified responses of the filters (shown in figure 14) of the first convolutional layer of the emotion detection classifier. The graph shows which parts of the face the different components of the neural network pay attention to. Note that for easier interpretation, this is the RGB rather than the LBP version of the network.*

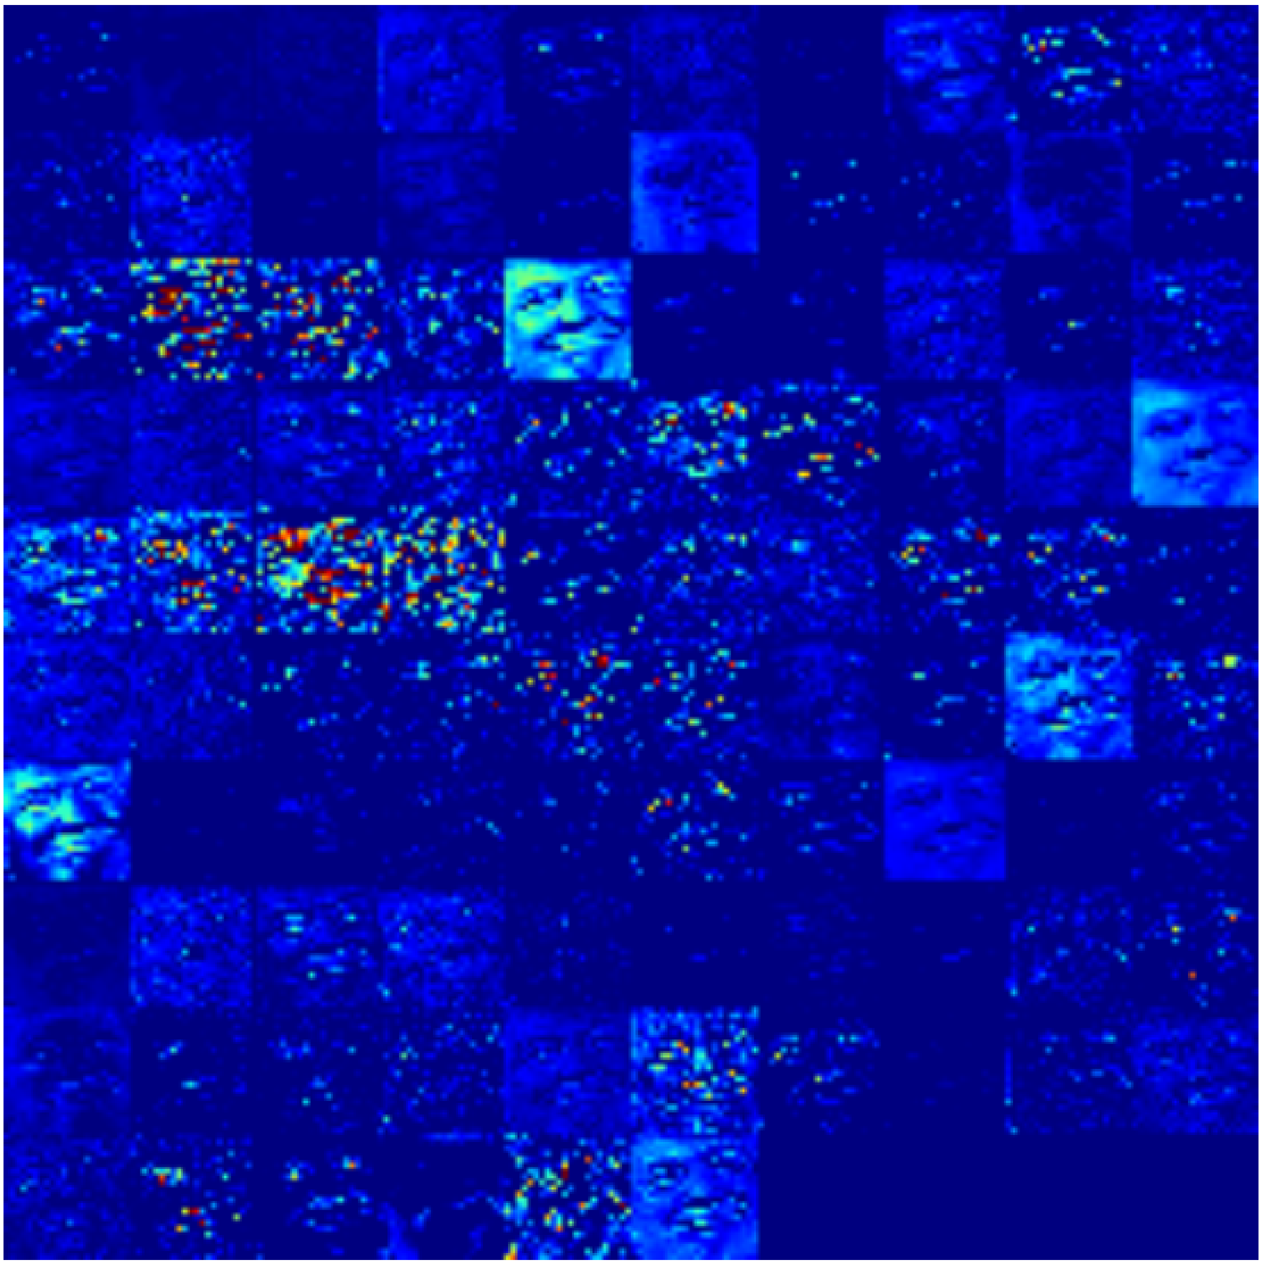

Appendix 2 – Additional Descriptive Statistics

Shared Images

Figure 16: *Number of shared images between news outlets. Larger/darker circles indicate more overlap.*

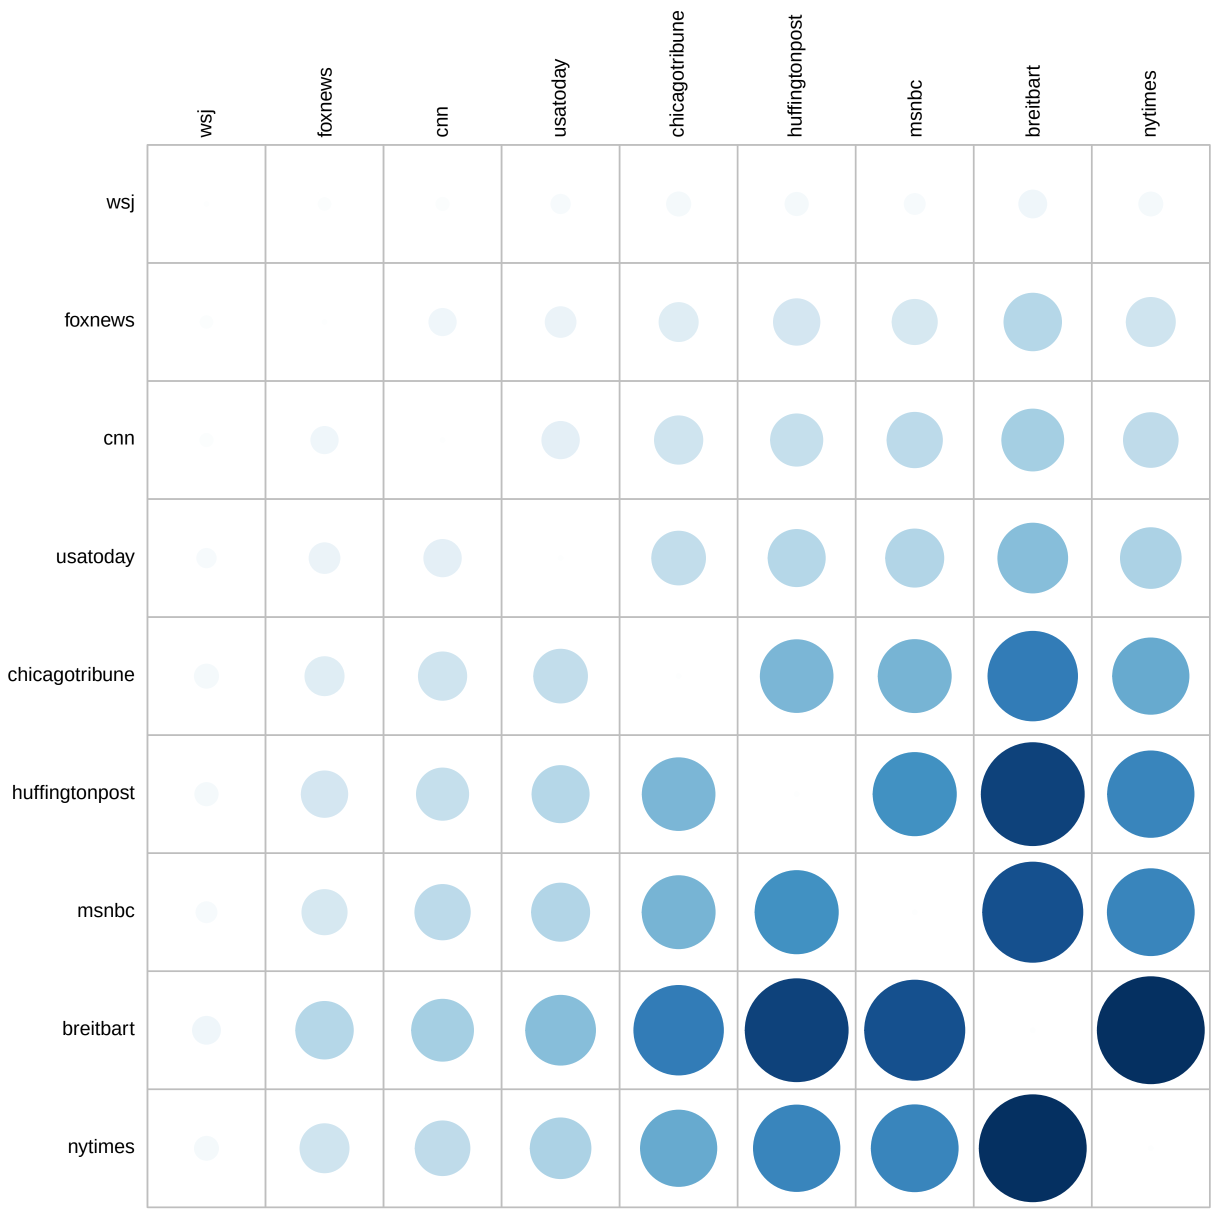

Another way to test whether news outlets can be grouped - ideologically, or in another way - is to measure whether they are using the same pictures. While most of the larger websites covered here do have their own in-house photographers, they still rely heavily on stock photography agencies such as Getty Images or the Associated Press. If two outlets

Figure 17: *Rallies dataset – bipartite graph of websites and the pictures they share.*

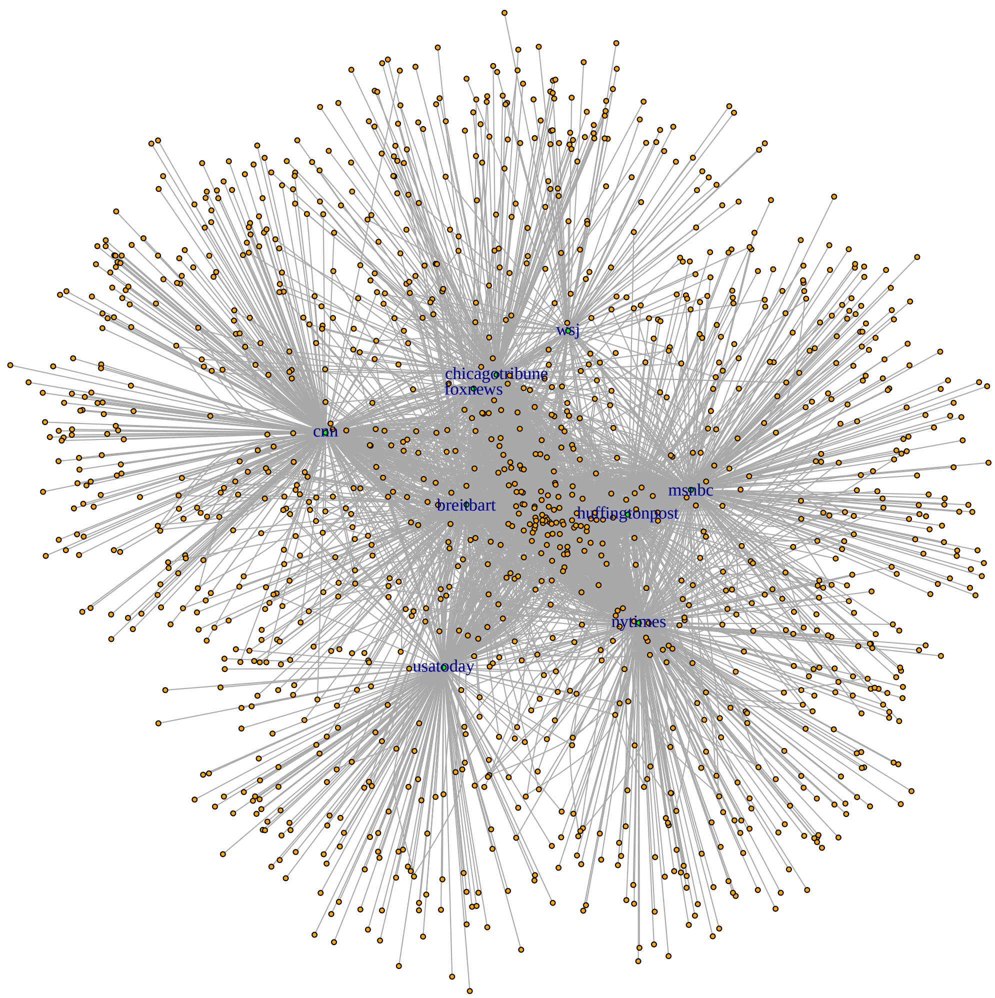

have concordant political views, it should be more likely that they use the same images - either because they want to broadcast the same message and thus coincidentally use the same photos, or because they actually copy from each other. This hypothesis is tested on the rally dataset.

As a way of displaying this overlap, the bipartite graph in figure 17 shows connections between two types of vertices - websites and the images they use[30]. Thus, a photo that is

[30]Whether two images by two different websites are the same is measured through image wavelet hashes. This technique, applied with the Python package "imagehash", is able to classify two images of different resolutions as the same. This is not an exact process, so on occasion, false positives or negatives are possible,

used by two websites has two edges connecting to it. The Fruchterman-Reingold algorithm plots this graph in such a way that the number of overlapping edges is minimized. As a result, websites with few images that are also used by other news outlets are placed at the periphery of the graph, and heavily-connected media sources on the inside.

Figure 17 shows that Breitbart lies at the center of the graph, which means that it shares the most images with other outlets. Interestingly enough, it connects to the same images as the liberal outlets, and is not confined to the "conservative corner" on the top of the graph, where Fox News, the Wall Street Journal, and the Chicago Tribune reside. The Huffington Post also sits in a pretty central position, wedged between MSNBC and the New York Times, suggesting some overlap between these media sources. Given their similar ideological leaning, as well as the fact that the Huffington Post is a news aggregator, and thus should be expected to share more pictures with other outlets, this does make sense. The fact that the network's highest amount of shared images seems to be between Breitbart and the Huffington Post is somewhat surprising though, given that they should be diametrical opposites with regard to ideology. Perhaps the fact that Andrew Breitbart also played a role in the creation of the Huffington Post (Sandoval, 2005) makes them more similar than they outwardly appear to be. While there appears to be a conservative cluster at the top, it may not be as meaningful as it seems at first sight. Fox News has very few pictures in general (see table 2), so its close position with the Chicago Tribune is somewhat deceptive - the Fruchterman-Reingold can easily fit it in this position not because a lot of overlap, but because it has such a low degree (i.e. edges connecting to it). The same might also be true for the Wall Street Journal. CNN and USA Today appear to lie somewhere in the middle with regard to shared connections. They clearly use a large amount of unique pictures, but on the other hand, also connect a lot to the center of the graph. Figure 16 in the appendix displays similar information to the bipartite plot, but in a different way: the higher the number of shared images between two outlets, the larger and darker the circle. The findings of figure 17 are corroborated here.

Number of Faces, Gender and Age

Figures 18 to 20 illustrate a number of findings pertaining to the rallies dataset which are not central to my hypothesis, but could potentially have provided news outlets with a different way to exhibit media bias, be it through the number of supporters shown at rallies, their gender or their age.

Figure 18 shows that photos of Trump rallies consistently feature more people, with USA Today as the only exception. One possible explanation for this is that as shown

but there should be no systematic bias.

Figure 18: *Rallies dataset – Average number of faces per picture, by news outlet and candidate. For all outlets except USA Today, photos of Trump rallies contain a higher average number of faces.*

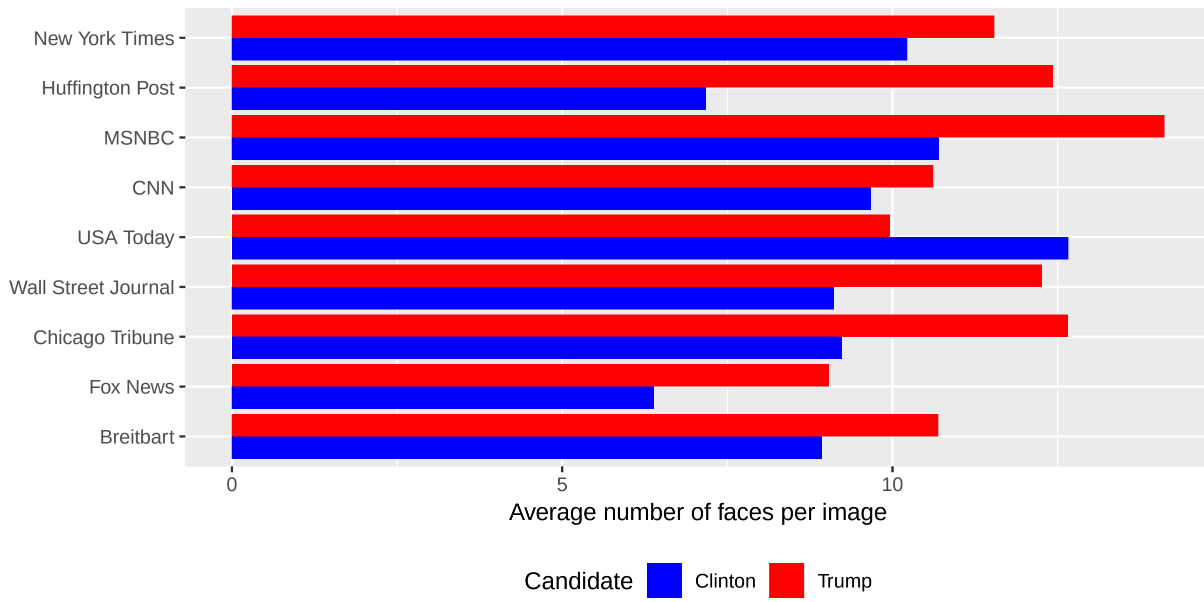

above, pictures of Republican rallies are larger, and thus allow more space for faces that are of high enough quality to be detected. However, if this was the only reason, then the large gaps for Breitbart, Huffington Post and USA Today would be carrying over, but this does not appear to be the case. Instead, perception may be playing a role: Trump took particular pride in the fact that his rallies were attended by huge masses of supporters. The media appears to have illustrated this phenomenon by featuring pictures with larger crowds of supporters. Another possibility is the fact that the Republican candidate sometimes confined reporters to so-called "press pens", to be jeered at by his supporters. It is plausible that one surreptitious reason for this practice was to allow campaign staffers to exert a greater degree of control on the press, placing photographers in locations where pictures would capture larger parts of the crowd.

Figure 19 shows the gender distribution of attendants at partisan rallies. One fact that appears to be immediately obvious is that males are considerably more prominent at Trump rallies, making up almost three fourths of those present. Differences across media outlets are minimal here, however. Women make up a higher share of the crowd at Clinton rallies. Disagreement between websites is slightly more pronounced here, but does not appear to be systematically correlated with partisan leaning (although Breitbart does feature the highest proportion of males).

Republican candidates and Trump in particular may have greater appeal to males, but evidently the conservative media does not emphasize this fact. If such outlets engage in partisan cheerleading, this makes sense: Portraying Trump as a misogynist by

Figure 19: *Rallies dataset – gender of rally attendants, by news outlet and candiate. Trump supporters consist of a higher proportion of males, but there is no media bias in the portrayal of gender.*

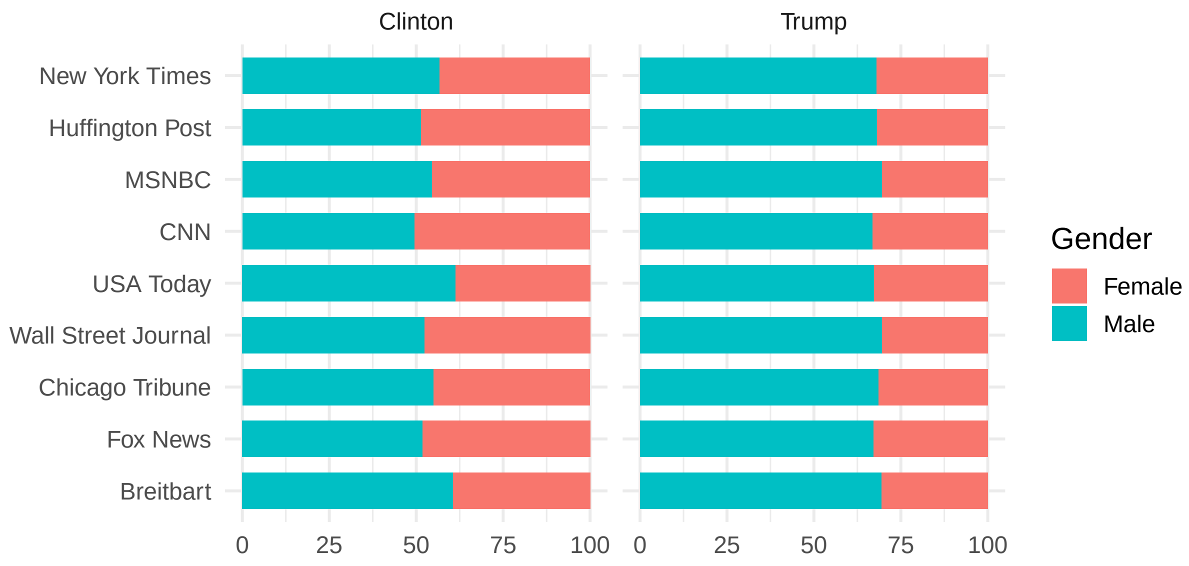

overrepresenting males in coverage of his campaign would only have hurt him. Liberal outlets on the other hand likely do not overemphasize males at Trump rallies due to their own norms of gender equality.

Figure 20: *Rallies dataset – age of rally attendants, by news outlet and candidate. Trump supporters appear older, but there is no media bias in the portrayal of age.*

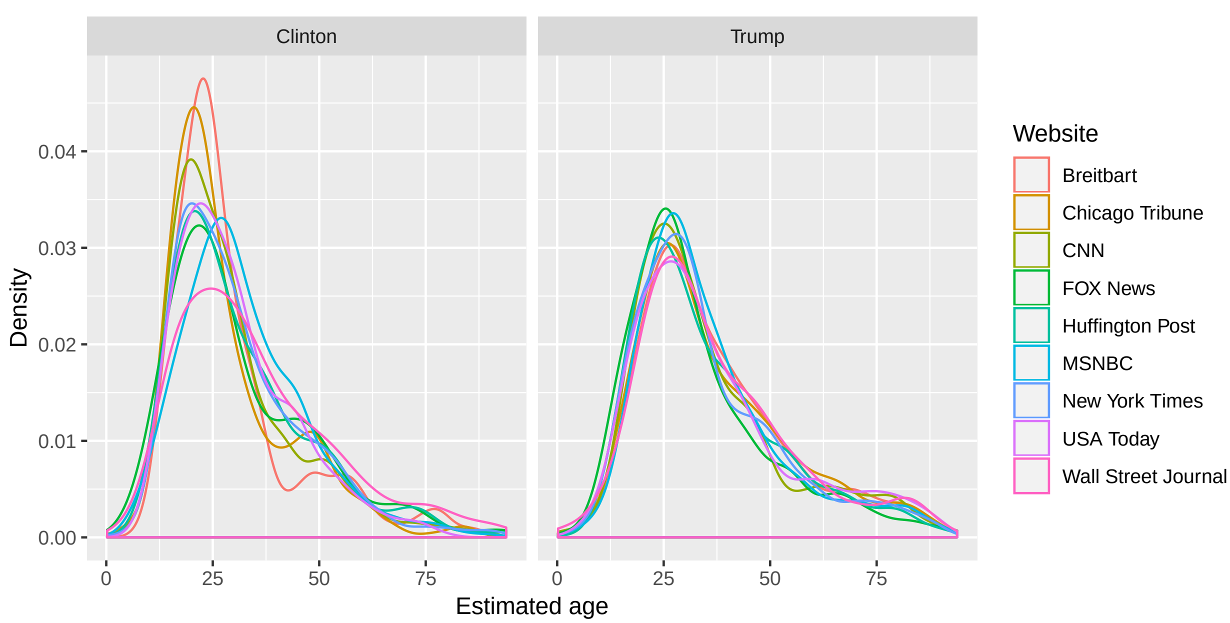

Similarly, media outlets do not appear to report on age in a partisan manner. Figure 20 presents density plots of the estimated age of rally attendants, showing remarkably little differences between websites as well as the two campaigns. Even though Republicans generally enjoy larger support among older people, the distributions between the two candidates are similarly shaped, with medians slightly above age 25 for Republicans and slightly below 25 for Democrats. The only noteworthy exceptions are Breitbart and Chicago Tribune, which show a slightly higher concentration of younger voters in the Clinton campaign than the other media sources. However, the smaller sample size of Clinton rally pictures for these two outlets may be playing a role here. The reasons for the lack of differences between liberal and conservative outlets are likely the same as those for gender.