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 popular news websites covering campaign rallies held during the 2016 presidential election campaign. I apply computer vision techniques to photos of supporters, automatically identifying characteristics such as age, gender and emotions of the attendants. Contrary to expectations, news outlets do not seem to have featured more positive pictures of their co-partisans. Even more surprising, and defying the popular notion of a left-leaning media landscape, the quantity and quality of images across both liberal and conservative news outlets actually appears to have favored Donald Trump.

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 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

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

studies on media bias have therefore focused on television or newspapers, but recently, research on news websites has featured much more prominently in scientific publications. Given that online media sources are already preferred by more than a third of Americans, and more importantly, 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. This is where my research design comes in. The images that accompany many news articles on the internet are a treasure trove waiting to be exploited. The phrase "a picture says more than a thousand words" may be trite, but it is also true: Pictures accompanying a news article prime the reader, setting the tone for the entire story. In contrast to traditional newspapers, this is all the more true for online news, where pictures are more numerous, articles are shorter, and most readers only look at the image and the first few lines of text. 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 a specific type of picture that has been ubiquitous in reports on the 2016 presidential election: crowds at campaign rallies. My research question, then, is whether news outlets portray supporters of their preferred candidate in a more positive manner.

These events grant Americans the rare opportunity to come face to face with those who represent them, for the de jure sovereigns to meet the de facto ones. Thus it is no wonder that according to Gulati (2004), politicians like to use pictures of themselves with their constituents as a way to portray themselves as "men of the people". 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.

The Achilles's heel of most of the media bias literature is that it largely relies on human-coded news articles and television shows, and thus ultimately rests on the objectivity of researchers. Machine learning-based approaches have attempted to address this shortcoming: Natural language processing-based techniques have enabled researchers to automatize the coding of news items, and thus yielded a step-up in objectivity. However, the problem of human coder bias ultimately still exists: the training data on which these algorithms are based is still hand-coded. So while the fact that the corpus itself has not been trained by humans does lend these studies greater credibility, the algorithms themselves can still be, and in a number of well-documented cases (Caliskan-islam, Bryson and Narayanan, 2016; Caliskan-Islam, Bryson and Narayanan, 2017), are biased.

My approach is to rely on the distribution of images among news outlets, as well as the use of computer vision techniques to identify faces, age, gender and emotion. Demographic information such as age and gender is automatically encoded by humans and thus forms a critical component for affective processing (Kurzban, Tooby and Cosmides, 2001). Similarly, emotional content (and more specifically, images (Burton et al., 2005)) aids both cognitive processing and memory retention, thus influencing how a political stimulus is perceived (Fazio, 2001; Spezio and Adolphs, 2006). If the news media really is biased along ideological lines, left-leaning outlets should use more photos of smiling supporters at Clinton rallies and angry followers at Trump events, with the reverse being true for the conservative news media.

I proceed as follows: First, I introduce the relevant literature, focusing on the different definitions, forms, sources and objects of media bias. The methods section contains a detailed description of the process by which I acquired the images through web scraping, as well as the stages of pre-processing. Since computer vision has seen very little use in political science, I also provide an introduction to convolutional neural networks. In direct contrast to Donald Trump's lamentations about the "liberal" news media, I show that the photographic coverage of the 2016 presidential election campaign actually favored the Republican candidate in several ways. Ideological media bias along partisan lines with regard to emotions is not observed however, leading me to conclude that news outlets likely use images purely to draw in readers, and rely on text to convey their message.

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 at campaign rallies. 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).

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). 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 supporters as enthusiastic and Trump supporters as angry, with the reverse being true for conservative websites.

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).

My own study builds on the literature in many ways: I analyze the emotional content of faces (Friedman, DiMatteo and Mertz, 1980; Banning and Coleman, 2009; Hehman et al., 2012) at the rallies (Barrett and Peake, 2007) of presidential candidates (Friedman, DiMatteo and Mertz, 1980; Mullen et al., 1986; Moriarty and Garramone, 1986; Moriarty and Popovich, 1991; Banning and Coleman, 2009), with the research subjects reacting to the candidate's actions (Friedman, DiMatteo and Mertz, 1980). For this purpose, I rely online news sources of both traditional newspapers and digital-only publications (Hehman et al., 2012).

This article does however present two innovations: One, to my knowledge, no previous study has utilized the reactions of partisan supporters, instead relying on either professional newscasters or politicians. Recent presidential campaigns have put increasing focus on interaction between the candidate and his followers. Furthermore, the 2016 campaign in particular has frequently featured supporters (particularly those of Donald Trump, but also Bernie Sanders) as the object of coverage, in addition to the candidate himself. Ergo, a candidate's disciples are worthy of studying.

Furthermore, I introduce automatic coding of facial expressions as a way of content analysis that does not suffer from human bias, which neither text nor image studies have been able to do so far. True, I still rely on human-coded training data, but emotions, unlike newspaper articles or faces of presidential candidates are not inherently political.

Methods

Previous approaches to measuring media bias in images have mainly focused on how (presidential) candidates are portrayed in these pictures. To do so, they rely on hand-coded content analyses. I cast my net a little wider by analyzing image distribution, metadata, and emotions of rally participants detected through computer vision. This allows me to evaluate the prospects of both selection and presentation bias. Theoretically, there is no reason to expect that media outlets would limit themselves to one or the other. If a specific outlet really does have a set ideology, and as a result, engages in partisan cheerleading, it only makes sense for it to pursue every avenue. On the one hand, this means that individual writers should be expected to choose pictures at least partially based on the affect they convey (Burton et al., 2005), with more positive emotions for the preferred campaign. But on aggregate, a news outlet would likely also feature a higher number of stories on the candidate it supports, and also be prone to furnish these stories with more and better pictures. Hence, both presentation and selection bias are plausible and need to be examined.

Scraping

The first step towards building a dataset 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 selected their sources based on two main criteria: One, maintaining a healthy 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.[2] 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[3], all news websites can be searched inside Google. The advantage

[2]As noted above, 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, so I use it as a reference category, with the New York Times, MSNBC, CNN and the Huffington Post to its left, and the Chicago Tribune, Fox News, Breitbart and the Wall Street Journal to its right.

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

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[4] to the layout of each website individually. The scraping was carried out between February 16-19, 2017.

The specific search term used is "Trump rally crowd"/"Clinton rally crowd"[5]. 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 I start out with 500 images from each campaign for each website, for a total of 9000 images. At this point, no restrictions are put on the time from which the images may originate, so that the results occasionally contain photos from before and after the campaign. 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. 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. The only other possibility for automatizing 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[6]) associate with images on HTML pages. Unfortunately, not all of the media sources in this study are particularly diligent about following this guideline, so this would drastically reduce and potentially bias the sample. Consequently I filter out undesirable images by hand.

Another problem with images scraped from the web is that not all of them are actually from the correct time frame. As Clinton already ran a primary campaign in 2007 and 2008, a small proportion of photos dates back to this time. Furthermore, some

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

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

[6]`https://www.w3schools.com/tags/att_img_alt.asp`

pictures were taken after the 2016 election, mostly from Trump's victory rallies and the Women's March on Washington. These images are filtered out based on the date on which the accompanying article was published. For most of the websites, this information is contained directly in the URL and can be extracted with regular expressions. For the rest, the dates were scraped from the websites from which the images originated. Unfortunately, the sample was further reduced by the fact that in some cases, the dates simply could not be recovered. The period retained in the sample ranges from June 13, 2015 (the day Hillary Clinton entered the race) to election day, November 8, 2016, for a total of 515 days (however, in some of my analyses, this is further reduced to the length of the general election campaign, starting on June 8, 2016).

'Three's a crowd', so as a rule for filtering out images with too few faces, I omit all photos 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.

At the end of this process, the original 9000 images have been winnowed down to 1,158, which, in the next step, produce 12,825 faces.

Computer vision

Next, four computer vision techniques - face, age, gender and emotion detection - are applied to the remaining images. Before proceeding with these steps, I resize images (using the Python package PIL) with a width of less than 2000 pixels to this size, increasing their height proportionally. This does not improve the quality, but it still raises the chances of small faces being detected, because the facial detection algorithm has a minimum requirement of 36x36 pixels for a face - even though it is perfectly capable of detecting smaller faces once scaled up. The facial detection algorithm then attempts to find every human face contained in a picture. The result is a square bounding box, uniquely identified through its width and height, as well as the space to its left and top (all measured in pixels). For further processing, I crop each face out of its photo, creating a unique picture for each, using the R package magick. Once the faces are identified and separated, three additional algorithms estimate the gender, age and emotions of each. There is some margin of error involved in this process, but validation on widely-used testing data (the IMDB-WIKI dataset for age and gender, and the JAFFE dataset for emotions) returns good results [7]. The fairly large sample of 12,825 faces should also help to alleviate bias.

Face, age, gender and emotion detection are separate techniques, but they all share the same underlying process - convolutional neural networks (CNNs). I begin by explaining

[7] The percent correctly predicted for gender is 0.79. The mean absolute error for age is 8.24 years. Table 2 in the appendix contains a breakdown of the accuracy of the emotion estimation process.

the architecture of CNNs, and then move on to how they are trained. The most important component of a CNN 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.

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. 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.

I implement face, age, gender and emotion detection through Microsoft's Cognitive

Services API[8] (using R and the httr package) - a set of 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[9]) does have some downsides - mainly the fact that replicability may decrease over time as Microsoft improves its product. However, one key feature of neural networks virtually ensures that Microsoft's Cognitive Services is more reliable than anything else available to me: Neural networks scale incredibly well with large sets of training data. Unfortunately, computer vision training data in the public domain cannot rival that of large technology companies, as labeling this data is very labor intensive, and thus quite expensive. For these firms on the other hand, this kind of data is simply generated as a byproduct of their other activities. As a result, the APIs provided by Microsoft yield the lowest error rate, and thus the best possible results I can attain.[10]

For gender classification, the output of the API is a simple binary value. For age, an estimation is given. The results from emotion 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 to create a categorical variable, corresponding to whichever emotion has the highest value for a face.

One question a reader might ask is whether computer vision is the only option for evaluating these images. It is true that theoretically, the pictures could be hand-coded, for example through crowdsourcing. Since the basic emotions used in my research design are universal among all humans (Ekman and Friesen, 1971), the results would likely be fairly accurate (although variance might be higher). However, in order to remove coder ideology from the equation, the faces would have to be cropped out of the images first so that the partisan nature of the photo would be unidentifiable (due to "Make America Great Again!" hats, etc.). This already requires computer vision, drastically reducing the extra effort required to implement the other classification techniques. Furthermore, even then, faces of the same image might share some characteristics such as lighting, that might enable coders to piece them together (consciously or subconsciously), possibly creating a dependence effect. Most importantly however, the use of automated computer vision techniques makes my results considerably more replicable than they would be if I were to rely on human coders.

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

[9] The framework appears to be the company's own open source "Microsoft Cognitive Toolkit": `https://github.com/microsoft/cntk`

[10] Table 2 in the appendix contains a comparison of error rates when the neural network-based API and the theory-based FACS model are applied to the widely used JAFFE training dataset. The neural network performs far more reliably.

Results

Website	Images	Clinton	Trump	ClintonPCT	TrumpPCT
breitbart	166	28	138	16.87	83.13
chicagotribune	121	35	86	28.93	71.07
cnn	188	96	92	51.06	48.94
foxnews	43	13	30	30.23	69.77
huffingtonpost	108	29	79	26.85	73.15
msnbc	147	70	77	47.62	52.38
nytimes	189	75	114	39.68	60.32
usatoday	134	41	93	30.6	69.4
wsj	62	27	35	43.55	56.45
All	1158	414	744		

Table 1: *Number of images, by news outlet and candidate*

Table 1 provides an overview of the number of images that portray crowds at campaign rallies. The total number of (usable) images varies quite heavily depending on the source, with the New York Times (189 pictures) and Fox News (43 pictures) at opposite ends of the spectrum. The low number of images from Fox News, and to a lesser extent, the Wall Street Journal is striking. While the same number of images was scraped for these news outlets, a much smaller proportion of them actually pertained to rallies held by Clinton and Trump. Instead, these searches returned, in addition 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).

Differences in the number of images between the more left-leaning outlets appear to be driven by seniority and prominence and thus budget. By far the oldest in this group, the New York Times leads CNN, which is followed by MSNBC and the fairly new Huffington Post. Contrary to what might be assumed, an outlet's medium (traditional newspaper, TV network, online-only) does not appear to play a role.

Clear differences also exist with regard to coverage of the two campaigns. With the exception of CNN, all websites host more images of supporters at Trump rallies (table 1, 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, 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. Even so, table 1 provides strong evidence for the existence of selection bias a) along partisan lines and b) in favor of the Republican side.

Among liberal outlets, the latter was likely driven by the greater selling power of stories on Trump. The candidate himself reiterated throughout the campaign that even though the left-leaning media was abhorred by his behavior, it also benefited greatly from the increase in readership the Trump campaign brought with it. Consequently, running more stories on the controversial candidate, and stocking them with more photos of his rallies, made business sense for the liberal outlets just as much as for the conservative ones.

Figure 1: *Average number of pixels per picture, by news outlet and candidate*

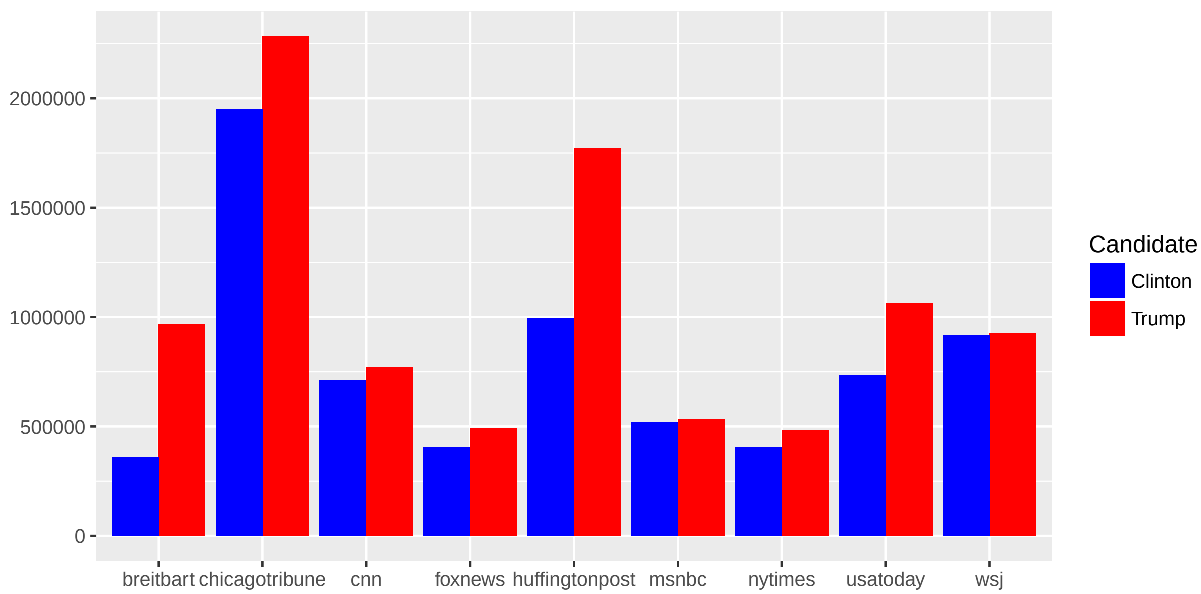

An alternative to analyzing the quantity of pictures is to look at their quality. Figure 1 shows the average number of pixels,[11,12] broken down by news outlet and campaign. Once again, Trump appears to be favored, and depending on the website, quite heavily so. Breitbart, the Huffington Post and USA Today (all of which hosted more than 70% Trump rally images) show a particularly large gap in terms of image quality. Strikingly, not a single website appears to be favoring Clinton.

The reason is likely the same as for the larger share of Trump pictures: 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

[11]Figure 9 in the appendix, plotting image size in kilobytes instead of pixels, shows very similar results.
[12]This analysis was conducted on the original images, before they were resized as described above.

appeal of covering the Trump campaign makes this an easy choice.

Figure 2: *Bipartite graph of websites and the pictures they use*

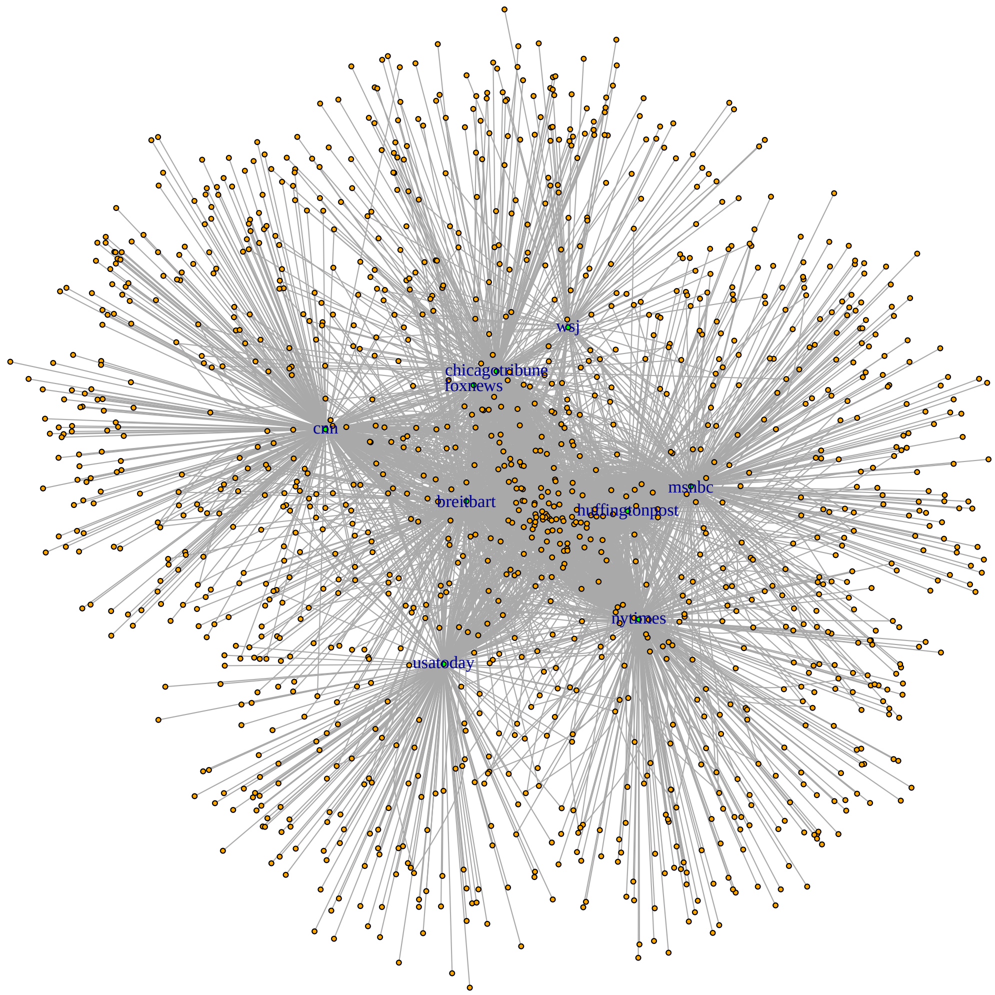

Another way to test whether news outlets can be grouped - ideologically, or in another way - is to measure whether they are using the same pictures. While most of the larger websites covered here do have their own in-house photographers, they still rely heavily on stock photography agencies such as Getty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

As a way of displaying this overlap, the bipartite graph in figure 2 shows connections between two types of vertices - websites and the images they use[13]. Thus, a photo that is 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 2 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 1), 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 10 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 2 are corroborated here.

Having covered results based on image distribution and metadata, I now move on to the content analysis. The simplest quantity to analyze here is the number of faces identified per image. Figure 3 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 above, pictures of Republican rallies are larger, and thus allow more

[13]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, but there should be no systematic bias.

Figure 3: *Average number of faces per picture, by news outlet and candidate*

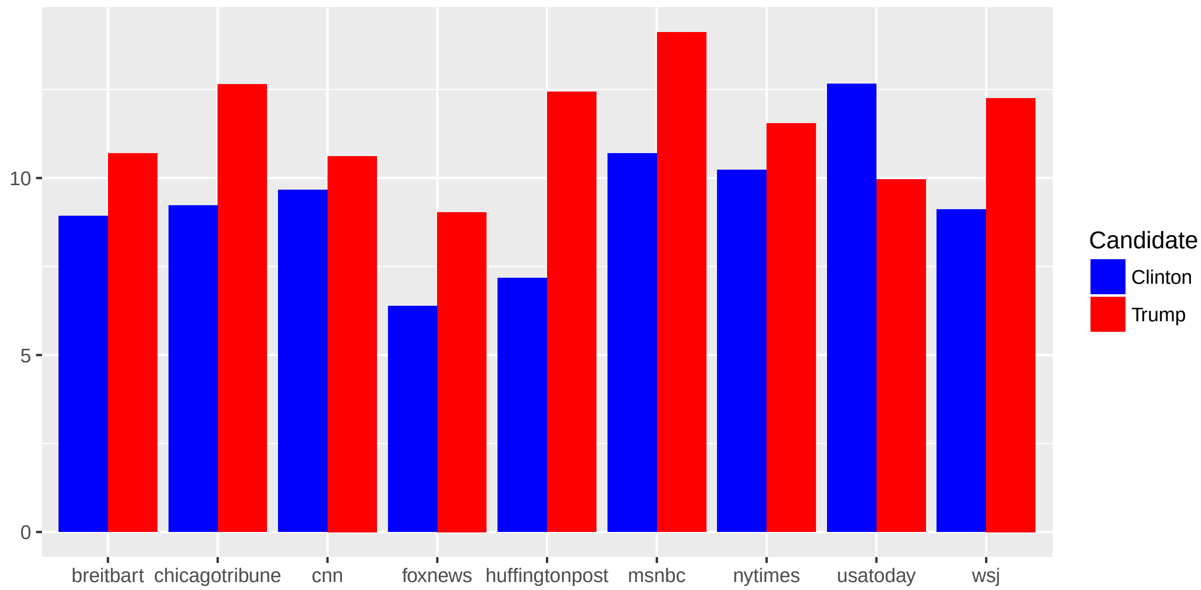

space for faces that are of high enough quality to be detected. However, if this was the only reason, then the large gaps for Breitbart, Huffington Post and USA Today would be carrying over, but this does not appear to be the case. Instead, perception may be playing a role: Trump took particular pride 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.

As a result of both the higher proportion of images for Trump rallies and the higher number of faces detected, my sample for the further analyses below is larger for Republicans.

Figure 4 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,

Figure 4: *Gender of rally attendants, by news outlet and candiate*

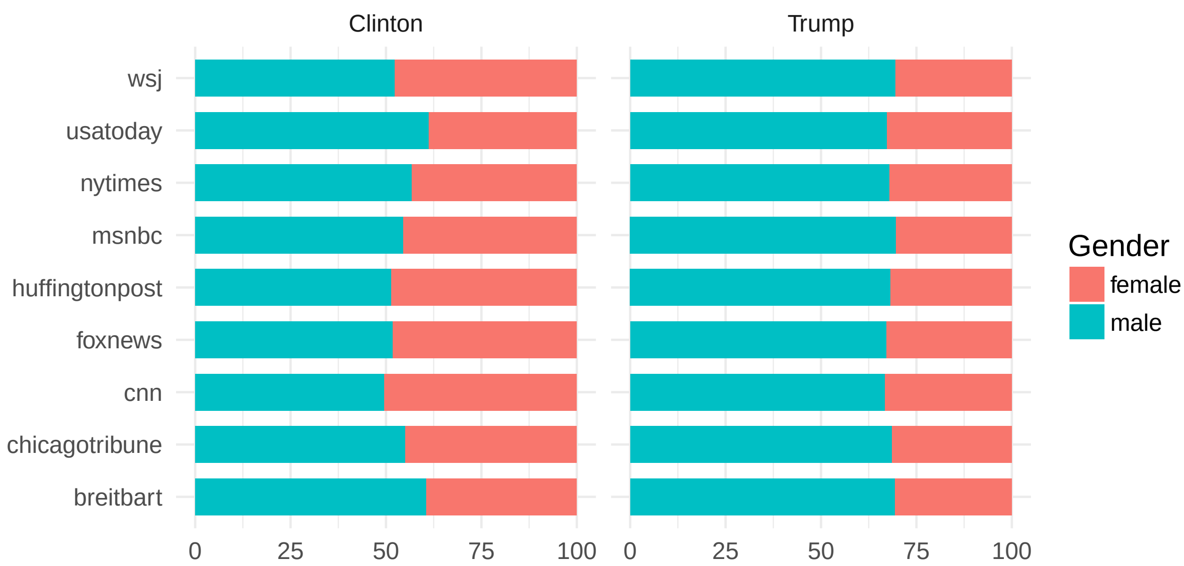

but evidently the conservative media does not emphasize this fact. If such outlets engage in partisan cheerleading, this makes sense: Portraying Trump as a misogynist by overrepresenting males in coverage of his campaign will only hurt him. Liberal outlets on the other hand don't overemphasize males at Trump rallies due to their own norms of gender equality. Furthermore, gender simply isn't an important enough cleavage in American politics - despite Trump's lack of appeal among women, a majority of white women still ended up voting for him (Malone, 2016).

Similarly, media outlets do not appear to report on age in a partisan manner. Figure 5 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 small 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.

In order to assess whether the news media portrays emotions in photos of 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

Figure 5: *Age of rally attendants, by news outlet and candidate*

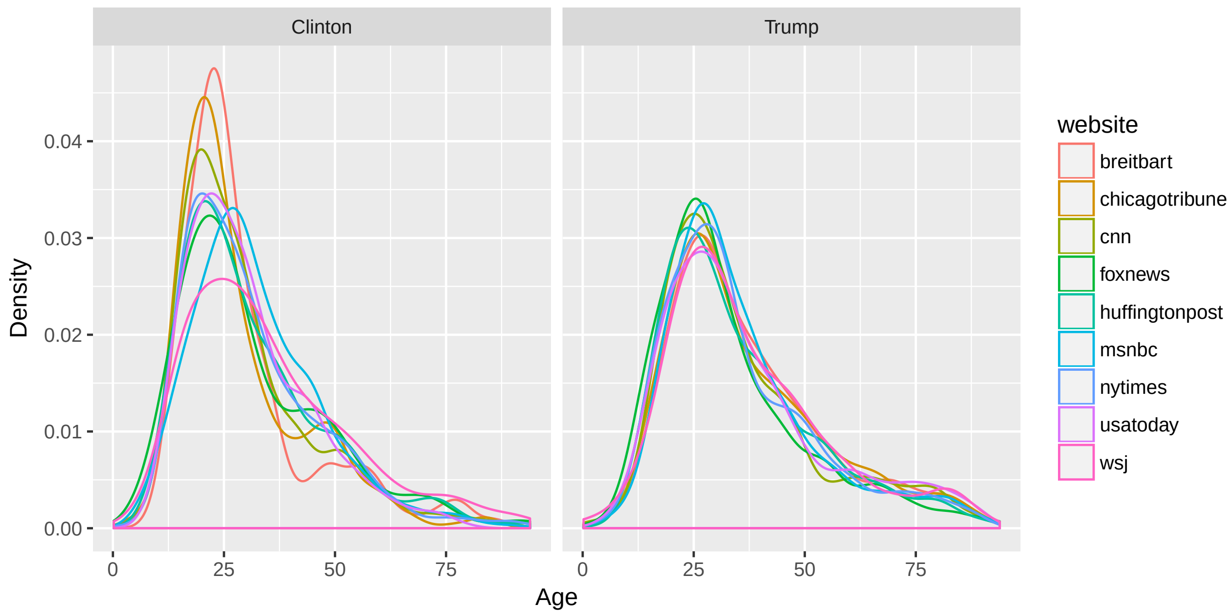

Figure 6: *Predicted probabilitiy of exhibiting happiness, by news outlet and candidate*

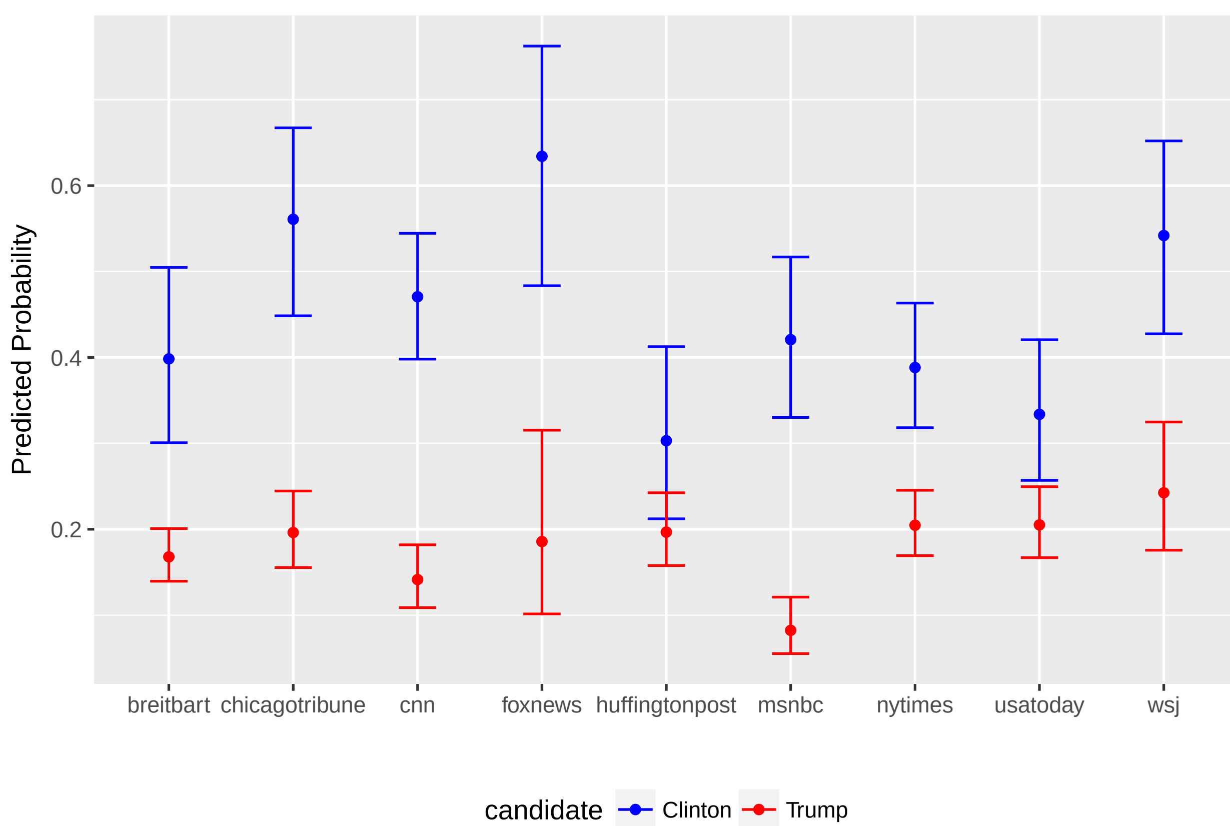

Figure 7: *Predicted probabilitiy of exhibiting anger, by news outlet and candidate*

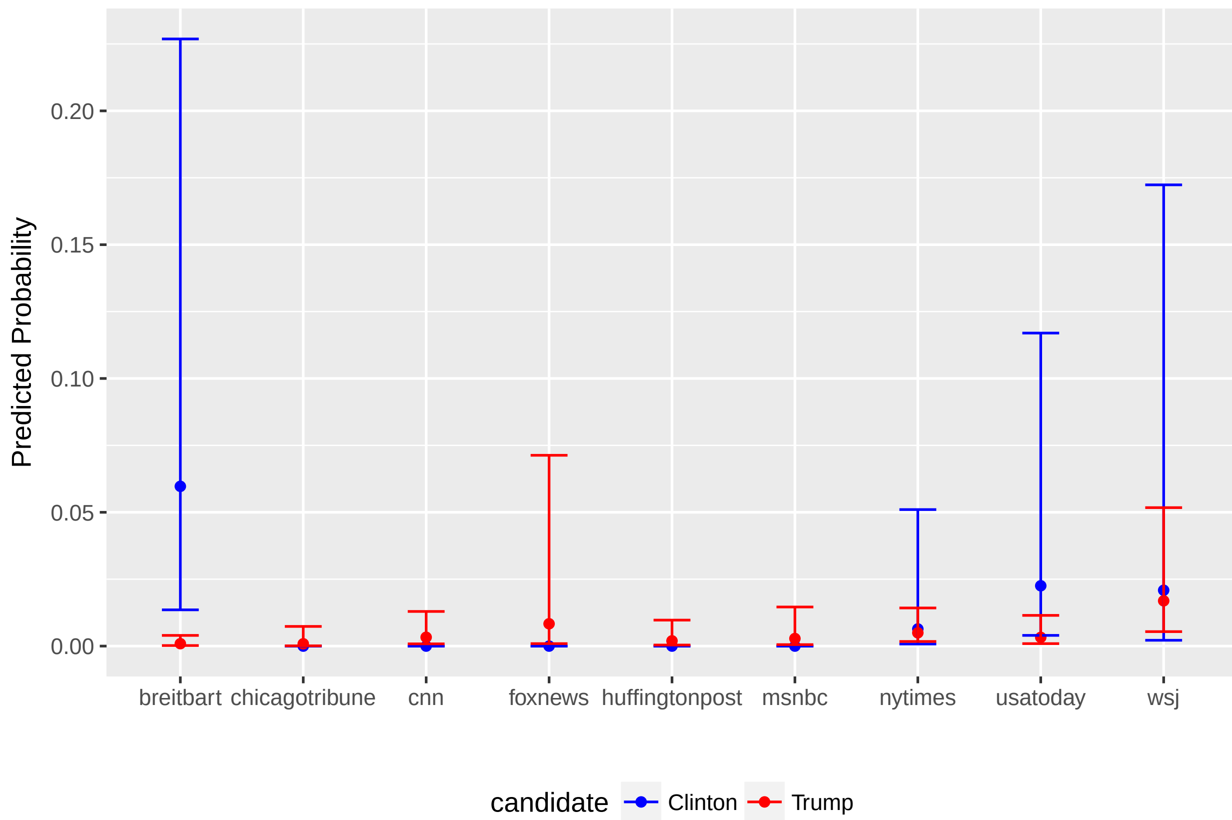

layer of the neural network is a set of probabilities denominating the likelihood of each emotion being displayed.[14] This result is turned into the categorical dependent variable of the model, where the expressed emotion is the one with the highest probability.[15] The independent variables include the estimated age and gender of the rally participant, as well as the number of days until the election (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, I use the distance to the other candidate in the horse race poll average of FiveThirtyEight on that day[16,17]. The variables I am substantively interested in are (1)

[14]Density plots for each of these values, broken down by emotion and news outlet, can be found in figure 11 in the appendix.

[15]Figure 12 in the appendix shows the distribution of this categorical emotion variable.

[16]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.

[17]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. Table 4 and figures 13-15 in the appendix show the results without the polling data variable, for the full sample. The happiness of Clinton supporters is not quite as high here, but with regard to media

Figure 8: *Predicted probabilitiy of exhibiting no emotion, by news outlet and candidate*

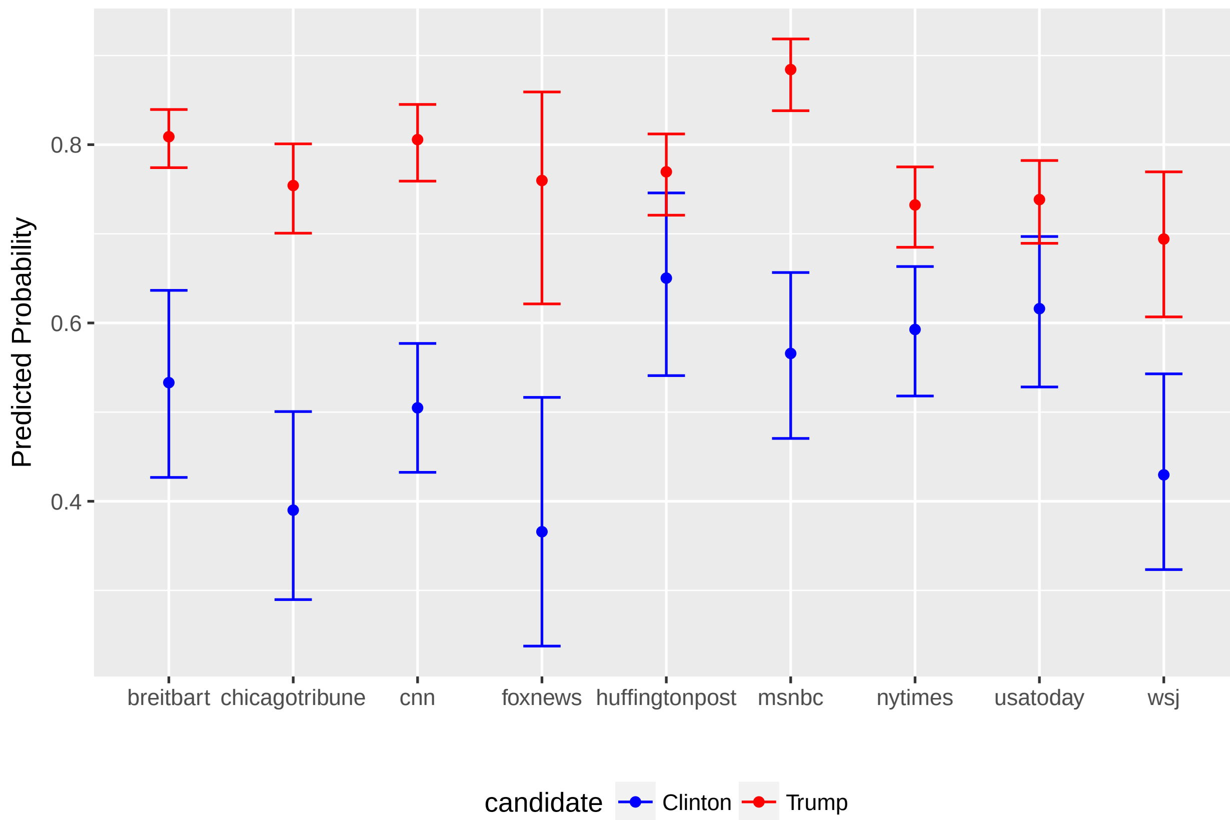

the candidate whose rally a particular supporter is attending and (2) the news outlet the corresponding photo is appearing on. Specifically for assessing media bias, an interaction term between the two measures the effect of partisanship on emotion, contingent on the media source.[18,19]

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, this makes sense. However, there is no clear evidence of media bias: For both Clinton and

bias, nothing changes.

[18]Table 3 in the appendix shows the full regression table.

[19]One possible caveat to this model is that it assesses media bias with regard to emotions at the face level, while journalists make their selections at the image level. Due to the large number of pictures, multiplied by the high number of categories in the dependent variable, accounting for this type of clustering statistically is difficult. Barrett and Barrington (2005) and Hehman et al. (2012) avoid this problem because they conduct a difference-of-means test at the image level, instead of a regression. I discuss the application of this procedure to my data in the appendix and show that it suffers from problems of its own.

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.

The results for anger (figure 7) are less useful, because anger is so rare in general (see appendix, figure 12). There are a few outliers, particularly Breitbart and the Wall Street Journal, but the confidence intervals associated with these findings are huge. Given that the sample of Clinton supporters for these two outlets only consists of 28 and 27 images, respectively, it seems probable that influential outliers drive these results. Overall, anger simply does not occur enough for any meaningful interpretation.

In contrast to anger, the lack of emotions is the most common category of the dependent variable. Here (figure 8), the results are diametrically opposite to happiness. Republicans assume a neutral facial expression much more frequently than Democrats, but there again does not appear to be much of a media effect. On the one hand, this finding simply complements the previous one - if Trump supporters are not as happy, by definition, they have to be something else. Since the other emotions barely occur, the neutral category is all that is left. Furthermore, there is a possibility that it has a substantive significance beyond what the term "neutral" conveys: An important part of Donald Trump's message was to conjure up the image of a country that is in decline. This is consistent with the literature which shows that Republicans are more likely to appeal to anxiety (MacKuen et al., 2007; Ridout and Searles, 2011). It is possible that the "fearmongering" of Donald Trump leads to stone-faced supporters, which the computer vision model interprets as a lack of emotion.

Overall, partisan supporters do show different emotions, but the media does not appear to portray these emotions in a biased manner.[20]

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

[20] As a more intuitive way of presenting the results, see figure 16 in the appendix. Here, the grayscale values (0 to 255) for each all images from a media outlet are averaged. This generates pictures showing the "average" Clinton/Trump supporter per news website. While there are clearly visible differences between the two campaigns, differences between media outlets appear only in cases in which the sample size is fairly low, overemphasizing the impact of specific images. Overall, these "average faces" underscore the findings of the automatic affect recognition method.

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 with only three people on it. 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, their editor, or someone else entirely - a factor which likely also varies between the media sources covered here.

Another important consideration is that the ideological incentives of media outlets may not be 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. Consequently, it is also possible that news outlets do use images to underscore their ideological message, but that message may differ from article to article. This hypothesis could be tested in future research by applying text analysis methods to the articles, and computer vision techniques to the images that accompany them.

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, doesn't make it so. The job of a researcher is not to prove our preconceived notions, but to follow where the evidence leads us.

Furthermore, while I do not observe *partisan* media bias, another form is present nonetheless: Contrary to Donald Trump's lamentations, the media coverage, if anything, favored him. The larger volume of pictures on Trump, the higher quality of these images, and even the greater number of people in them all tell the same story. The Republican candidate played (whether justified or not) the challenger, casting Clinton (as well as his other opponents) in the role of the establishment. The media appears to have followed that narrative, giving greater attention to the more "newsworthy" of the two candidates. The ubiquitousness of this phenomenon throughout even left-leaning outlets underscores the notion that business comes first. Given the lack of evidence for partisan media bias, combined with the positive results from some text-based content analyses, one possible conclusion is that the media relies on images primarily to draw readers in. The partisan message is then left to the writing itself. From this point of view, it makes sense to rely on attention-grabbing images, rather than those that support the outlet's partisan stance.

In addition to these substantive considerations, the methodological aspects of this paper also merit further discussion. 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. For face and emotion detection, this is not as much of an issue, because contours of the face as well as other easily recognizable features are sufficient. The lack of quality did however prevent me from implementing race detection as well as face recognition (i.e. recognizing the candidates as well as other political elites who may be present at the rallies). Controlling for these factors, while not essential, might have provided me with additional empirical leverage.

Computer vision has not been widely used in political science, but there are a number of possible applications. The one major way in which social scientists have relied on this method so far is for the automatic recognition of text, utilizing optical character recognition (OCR). However, many conventional OCR solutions are slow, expensive, and notoriously unreliable, so that outsourcing the work to human coders is often much simpler. Computer science researchers are still making a lot of progress in this area, particularly with the advent of deep neural networks - promising a much better implementation in the future. Beyond this more conventional application, a number of other possibilities exist. Satellite images have been used 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. Object recognition, which has developed tremendously in recent years, has also seen its first use in political science (Anastasopoulos et al., 2016). And of course, affect recognition has a wide range of applications in both experimental and observational studies on emotions, a field of study that has generally become more important in political science (Marcus, 2000; Brader and Marcus, 2013). But the media is an area of research for which deep learning may hold the greatest 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 (convolutional neural networks) or text (recurrent neural networks), 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.

Barrett, Andrew W. and Lowell W. Barrington. 2005. "Bias in newspaper photograph selection." *Political Research Quarterly* 58(4):609–618.

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

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.

Brader, Ted and George E. Marcus. 2013. "Emotion and Political Psychology." *Oxford Handbook of Political Psychology* (February):165–204.

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.

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

Ekman, Paul and Wallace V Friesen. 1971. "Constants across cultures in the face and emotion." *Journal of Personality and Social Psychology* 17(2):124–129.

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.

Gulati, Girish J. 2004. "Members of Congress and Presentation of Self on the World Wide Web." *The International Journal of Press/Politics* 9(1):22–40.

Hatemi, Peter K. and Rose McDermott. 2011. *Man Is by Nature a Political Animal*.

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.

Kurzban, R, J Tooby and L Cosmides. 2001. "Can race be erased? Coalitional computation and social categorization." *Proceedings of the National Academy of Sciences of the United States of America* 98(26):15387–15392.

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.

Lin, Y-R, J P Bagrow and D Lazer. 2011. "More Voices than Ever? Quantifying Bias in Social and Mainstream Media." *International AAAI Conference on Weblogs and Social Media* (September 2015).

MacKuen, M., George E. Marcus, W Russell Neuman and Luke Keele. 2007. The Third Way: The Theory of Affective Intelligence and American Democracy. In *The Affect Effect: Dynamics of Emotion in Political Thinking and Behavior*. The University of Chicago Press pp. 125–151.

Malone, Clare. 2016. "Clinton Couldn't Win Over White Women.".
URL: *https://fivethirtyeight.com/features/clinton-couldnt-win-over-white-women/*

Marcus, G E. 2000. “Emotions in Politics.” *Annual Review of Political Science* 3:221–50.

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.

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.

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.

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.

Appendix

Emotion	FACS	Microsoft
Happiness	0.57	0.79
Sadness	-0.05	0.63
Surprise	0.08	0.76
Anger	0.08	0.33
Disgust		0.30
Fear	0.16	0.49

Table 2: *Validation of emotion detection methods against JAFFE dataset (values are Pearson's R)*

Barrett and Barrington (2005) and Hehman et al. (2012) rely on difference-of-means tests to assess the existence of media bias in the selection of images with regard to emotions. Specifically, they compare the mean level of positivity (i.e. happiness, warmth, competence, etc.) of photos depicting an outlet's co-partisan, to those showing their opponent. In an attempt to provide an alternative analysis at the image rather than the face level, I do the same. For this test, the data is aggregated as following: The level of

Table 3: *Multinomial regression, general election campaign*

	Dependent variable:					
	anger	contempt	disgust	happiness	neutral	surprise
	(1)	(2)	(3)	(4)	(5)	(6)
Trump rally	-6.591^{***}	6.799^{***}	8.083^{***}	-3.273^{***}	-1.992^{**}	-1.602^{*}
	(1.231)	(0.478)	(0.682)	(0.809)	(0.799)	(0.939)
Age	-0.006	-0.025	-0.028	-0.032^{***}	-0.048^{***}	-0.045^{***}
	(0.011)	(0.033)	(0.053)	(0.008)	(0.008)	(0.009)
Male	0.375	-0.399	3.553^{***}	-0.360	0.507	0.034
	(0.506)	(1.512)	(1.127)	(0.341)	(0.339)	(0.377)
Chicago Tribune	-17.592^{***}	-6.193^{***}	-4.004^{***}	-4.725^{***}	-5.379^{***}	-3.595^{***}
	(0.582)	(0.001)	(0.009)	(0.354)	(0.356)	(0.620)
CNN	-17.816^{***}	-5.564^{***}	-3.700^{***}	-3.969^{***}	-4.191^{***}	-3.303^{***}
	(0.422)	(0.004)	(0.179)	(0.711)	(0.708)	(0.724)
Fox News	-4.368^{***}	5.450^{***}	8.273^{***}	12.541^{***}	11.699^{***}	-7.378^{***}
	(0.468)	(0.00005)	(0.0005)	(0.248)	(0.248)	(0.325)
Huffington Post	-22.193^{***}	-9.495^{***}	-3.650^{***}	-5.864^{***}	-5.392^{***}	-4.421^{***}
	(0.487)	(0.0001)	(0.009)	(0.725)	(0.710)	(0.787)
MSNBC	-11.321^{***}	6.278^{***}	5.480^{***}	9.194^{***}	9.199^{***}	9.561^{***}
	(0.607)	(0.0003)	(0.001)	(0.515)	(0.511)	(0.667)
New York Times	-5.905^{***}	-2.836^{***}	-3.004^{***}	-3.700^{***}	-3.568^{***}	-3.556^{***}
	(0.792)	(0.713)	(0.794)	(0.460)	(0.452)	(0.550)
USA Today	-5.885^{***}	-6.050^{***}	-4.749^{***}	-5.089^{***}	-4.767^{***}	-4.178^{***}
	(0.859)	(0.002)	(0.044)	(0.586)	(0.577)	(0.738)
Wall Street Journal	4.068^{***}	-2.405^{***}	0.592^{***}	5.427^{***}	4.903^{***}	4.963^{***}
	(0.930)	(0.00000)	(0.00004)	(0.608)	(0.605)	(0.832)
Days until election	-0.003	0.030	-0.033	0.002	0.001	0.005
	(0.005)	(0.019)	(0.033)	(0.004)	(0.004)	(0.004)
Poll difference to other candidate	-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 rally × Chicago Tribune	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 rally × CNN	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 rally × Fox News	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 rally × Huffington Post	20.828^{***}	-4.510^{***}	2.478^{***}	3.884^{***}	3.204^{***}	2.532^{***}
	(0.487)	(0.00000)	(0.018)	(0.889)	(0.873)	(0.918)
Trump rally × MSNBC	11.313^{***}	-15.469^{***}	-7.219^{***}	-11.051^{***}	-10.254^{***}	-10.349^{***}
	(0.607)	(0.0002)	(0.002)	(0.568)	(0.546)	(0.694)
Trump rally × New York Times	5.823^{***}	1.186^{*}	7.371^{***}	2.121^{***}	1.692^{***}	2.630^{***}
	(0.860)	(0.712)	(0.730)	(0.537)	(0.525)	(0.557)
Trump rally × USA Today	4.686^{***}	-6.087^{***}	0.466^{***}	2.814^{***}	2.202^{***}	2.510^{***}
	(1.040)	(0.00003)	(0.001)	(0.699)	(0.687)	(0.815)
Trump rally × Wall Street Journal	-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)
N	5,287	5,287	5,287	5,287	5,287	5,287

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

Table 4: *Multinomial regression, primaries and general election campaign*

	Dependent variable:						
	anger	contempt	disgust	fear	happiness	sadness	surprise
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Trump rally	-1.609^{***}	5.057^{***}	-1.110^{***}	0.427^{***}	-0.334^{**}	7.474^{***}	0.726^{***}
	(0.265)	(0.227)	(0.001)	(0.009)	(0.143)	(0.249)	(0.150)
Age	0.029^{***}	0.004	0.010	0.047^{*}	0.013^{***}	0.049^{***}	0.007^{*}
	(0.005)	(0.019)	(0.023)	(0.024)	(0.001)	(0.005)	(0.003)
Male	0.283	-0.627^{***}	-1.208^{***}	-4.315^{***}	-0.958^{***}	-0.717^{***}	-0.692^{***}
	(0.265)	(0.088)	(0.0003)	(0.0003)	(0.042)	(0.223)	(0.115)
Chicago Tribune	-14.392^{***}	-1.047^{***}	2.209^{***}	-0.972^{***}	0.473^{***}	8.840^{***}	1.565^{***}
	(0.193)	(0.311)	(0.00001)	(0.00000)	(0.175)	(0.162)	(0.311)
CNN	-18.783^{***}	-4.528^{***}	0.306^{***}	0.203^{***}	0.360^{**}	8.015^{***}	1.223^{***}
	(0.183)	(0.00000)	(0.00000)	(0.0001)	(0.146)	(0.153)	(0.232)
Fox News	0.263	-3.410^{***}	0.145^{***}	0.047^{***}	0.724^{***}	-5.955^{***}	1.279^{***}
	(0.317)	(0.00001)	(0.00000)	(0.00001)	(0.261)	(0.211)	(0.179)
Huffington Post	-14.042^{***}	-4.550^{***}	0.580^{***}	6.254^{***}	0.095	9.648^{***}	1.058^{***}
	(0.171)	(0.00000)	(0.00000)	(0.011)	(0.201)	(0.183)	(0.411)
MSNBC	-2.110^{***}	-2.346^{***}	0.679^{***}	-1.020^{***}	0.202	-15.492^{***}	1.251^{***}
	(0.188)	(0.0002)	(0.00000)	(0.008)	(0.151)	(0.209)	(0.247)
New York Times	-1.520^{***}	4.984^{***}	0.887^{***}	-1.145^{***}	-0.006	8.031^{***}	0.286
	(0.136)	(0.364)	(0.001)	(0.00002)	(0.152)	(0.165)	(0.329)
USA Today	-0.386^{**}	-4.721^{***}	1.094^{***}	0.020^{***}	0.165	8.199^{***}	1.238^{***}
	(0.173)	(0.00000)	(0.00000)	(0.0001)	(0.161)	(0.162)	(0.279)
Wall Street Journal	-0.950^{***}	-3.212^{***}	0.427^{***}	-0.800^{***}	0.732^{***}	8.833^{***}	1.272^{***}
	(0.181)	(0.00001)	(0.00000)	(0.00001)	(0.183)	(0.192)	(0.395)
Days until election	-0.001	0.004	-0.004	0.006	0.0002	-0.001	-0.001^{**}
	(0.001)	(0.003)	(0.005)	(0.005)	(0.0002)	(0.001)	(0.0005)
Trump rally × Chicago Tribune	14.979^{***}	0.663^{**}	-0.980^{***}	-1.436^{***}	-0.153	-7.682^{***}	-1.032^{***}
	(0.193)	(0.310)	(0.00000)	(0.00000)	(0.199)	(0.145)	(0.352)
Trump rally × CNN	19.569^{***}	-5.273^{***}	0.528^{***}	0.068^{***}	-0.398^{**}	-6.520^{***}	-0.709^{**}
	(0.183)	(0.00000)	(0.00000)	(0.00004)	(0.179)	(0.135)	(0.291)
Trump rally × Fox News	0.866^{***}	-4.658^{***}	-0.342^{***}	-0.535^{***}	-0.277	8.119^{***}	-0.122
	(0.255)	(0.00001)	(0.00000)	(0.00001)	(0.303)	(0.211)	(0.166)
Trump rally × Huffington Post	15.002^{***}	-6.830^{***}	-0.724^{***}	-7.372^{***}	0.009	-8.598^{***}	-0.859^{*}
	(0.171)	(0.00000)	(0.00000)	(0.00001)	(0.225)	(0.165)	(0.454)
Trump rally × MSNBC	2.747^{***}	-3.576^{***}	1.010^{***}	4.787^{***}	-0.449^{**}	15.865^{***}	-1.241^{***}
	(0.174)	(0.0002)	(0.00000)	(0.008)	(0.183)	(0.209)	(0.321)
Trump rally × New York Times	3.100^{***}	-5.639^{***}	7.687^{***}	0.897^{***}	0.255	-6.997^{***}	0.180
	(0.124)	(0.170)	(0.001)	(0.00003)	(0.177)	(0.146)	(0.364)
Trump rally × USA Today	1.920^{***}	-6.642^{***}	0.078^{***}	0.777^{***}	0.179	-6.897^{***}	-0.539^{*}
	(0.171)	(0.00000)	(0.00000)	(0.0001)	(0.189)	(0.145)	(0.326)
Trump rally × Wall Street Journal	2.865^{***}	-4.810^{***}	-0.256^{***}	-0.559^{***}	-0.238	-6.981^{***}	-0.117
	(0.168)	(0.00001)	(0.00000)	(0.00001)	(0.223)	(0.166)	(0.447)
Constant	-5.041^{***}	-11.835^{***}	-12.647^{***}	-13.152^{***}	-0.822^{***}	-14.496^{***}	-4.024^{***}
	(0.361)	(0.033)	(0.001)	(0.002)	(0.138)	(0.324)	(0.194)
N	12,825	12,825	12,825	12,825	12,825	12,825	12,825

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

Figure 9: *Average filesize per picture, by news outlet and candidate*

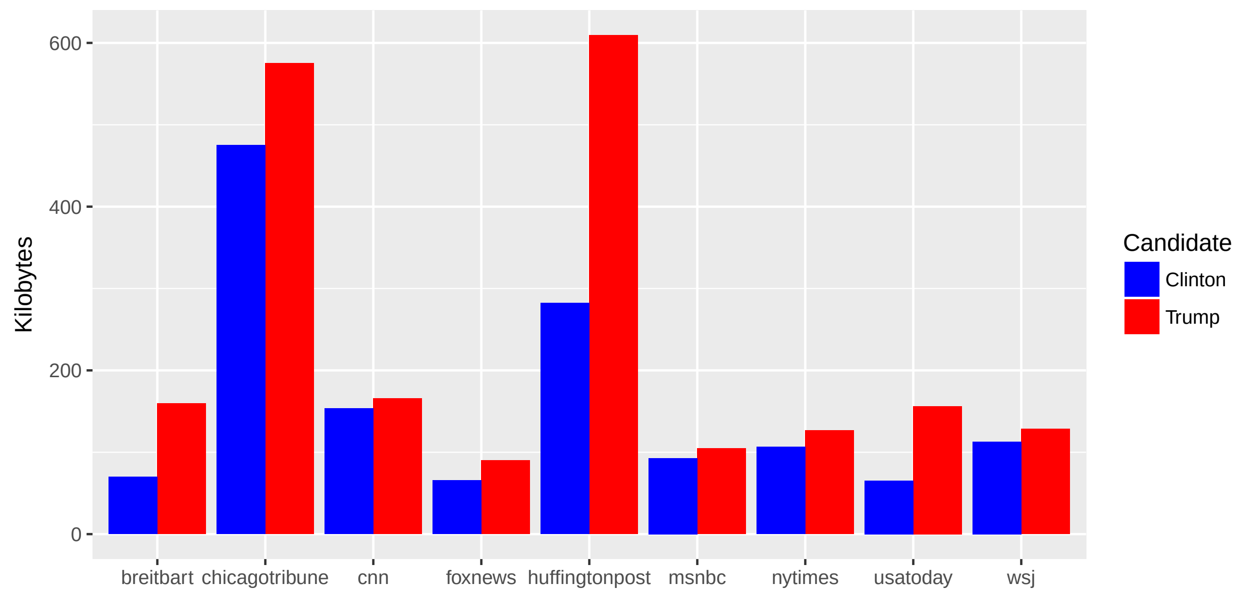

happiness in an image is equivalent to the proportion of faces in it whose detected emotion is happiness. Then, I conduct a two sample t-Test, where sample 1 is the mean happiness of images that correspond to media outlets' ideological leaning (for this purpose, the New York Times, MSNBC, CNN and the Huffington Post are considered liberal, and the Chicago Tribune, Fox, Breitbart and the Wall Street Journal conservative. Since USA Today is used as the reference category in the analyses above and does not have a clear partisan leaning, it is omitted here), and sample 2 contains the mean happiness of images that portray an outlet's ideological opponent.

For images with a shared political orientation, the mean proportion of happiness is $M = 0.33$, whereas for the other sample, $M = 0.29$. With a p-value of $p = 0.001$, this difference is statistically significant at the 1%-level, exactly the result one would expect under the hypothesis of partisan media bias. Similarly, when the sample is restricted to only liberal websites, a greater degree of happiness is observed among Clinton supporters ($M = 0.37$), compared to Trump supporters ($M = 0.24$, $p = 0.001$). However, the result for conservative news sources does not fit the bill: Here, Clinton supporters ($M = 0.45$) still appear to be happier than Trump supporters ($M = 0.3$, $p = 0.001$), even though theory dictates that conservative outlets should favor their co-partisans instead.

So how is it possible for partisan media bias among liberal outlets and reverse partisan media bias among conservative websites to add up to the expected result of partisan media bias in the full sample? Why don't these two opposite results simply cancel out? The reason lies in the fact that there are considerably more pictures in the sample of

Figure 10: *Number of shared images between news outlets. Larger/darker circles indicate more overlap.*

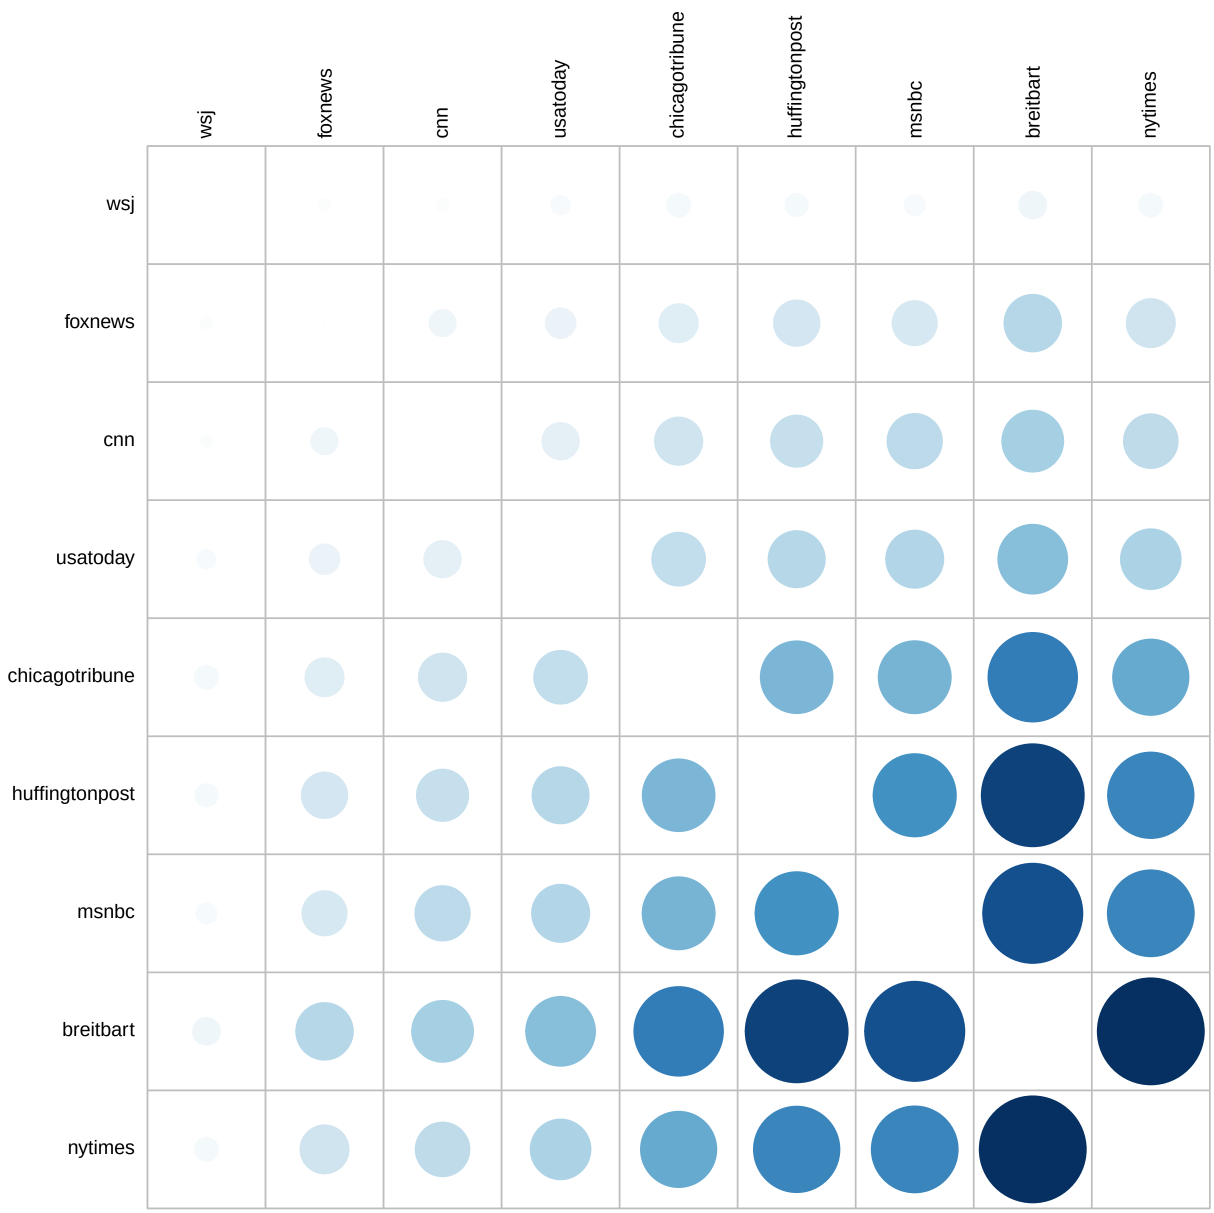

liberal websites ($N = 632$), compared to the conservative ones ($N = 392$), so the effect of the former drowns out that of the latter.

This means that in the two sub-samples, I am not observing media bias - I am observing the fact that Clinton supporters actually are happier. But when aggregated, it looks like there really is a media bias effect, even though none exists. It is quite possible that the positive findings in Barrett and Barrington (2005) and Hehman et al. (2012) are merely a statistical artifact, stemming from the same problem. Nevertheless, the result from the t-tests mirror and thus reinforce those of the multinomial model once they are disaggregated. They also show that it doesn't matter whether the analysis is conducted with faces or images as the unit of analysis - the results are the same.

Figure 11: *Probabilities for emotions outputted by the CNN's fully connected layer*

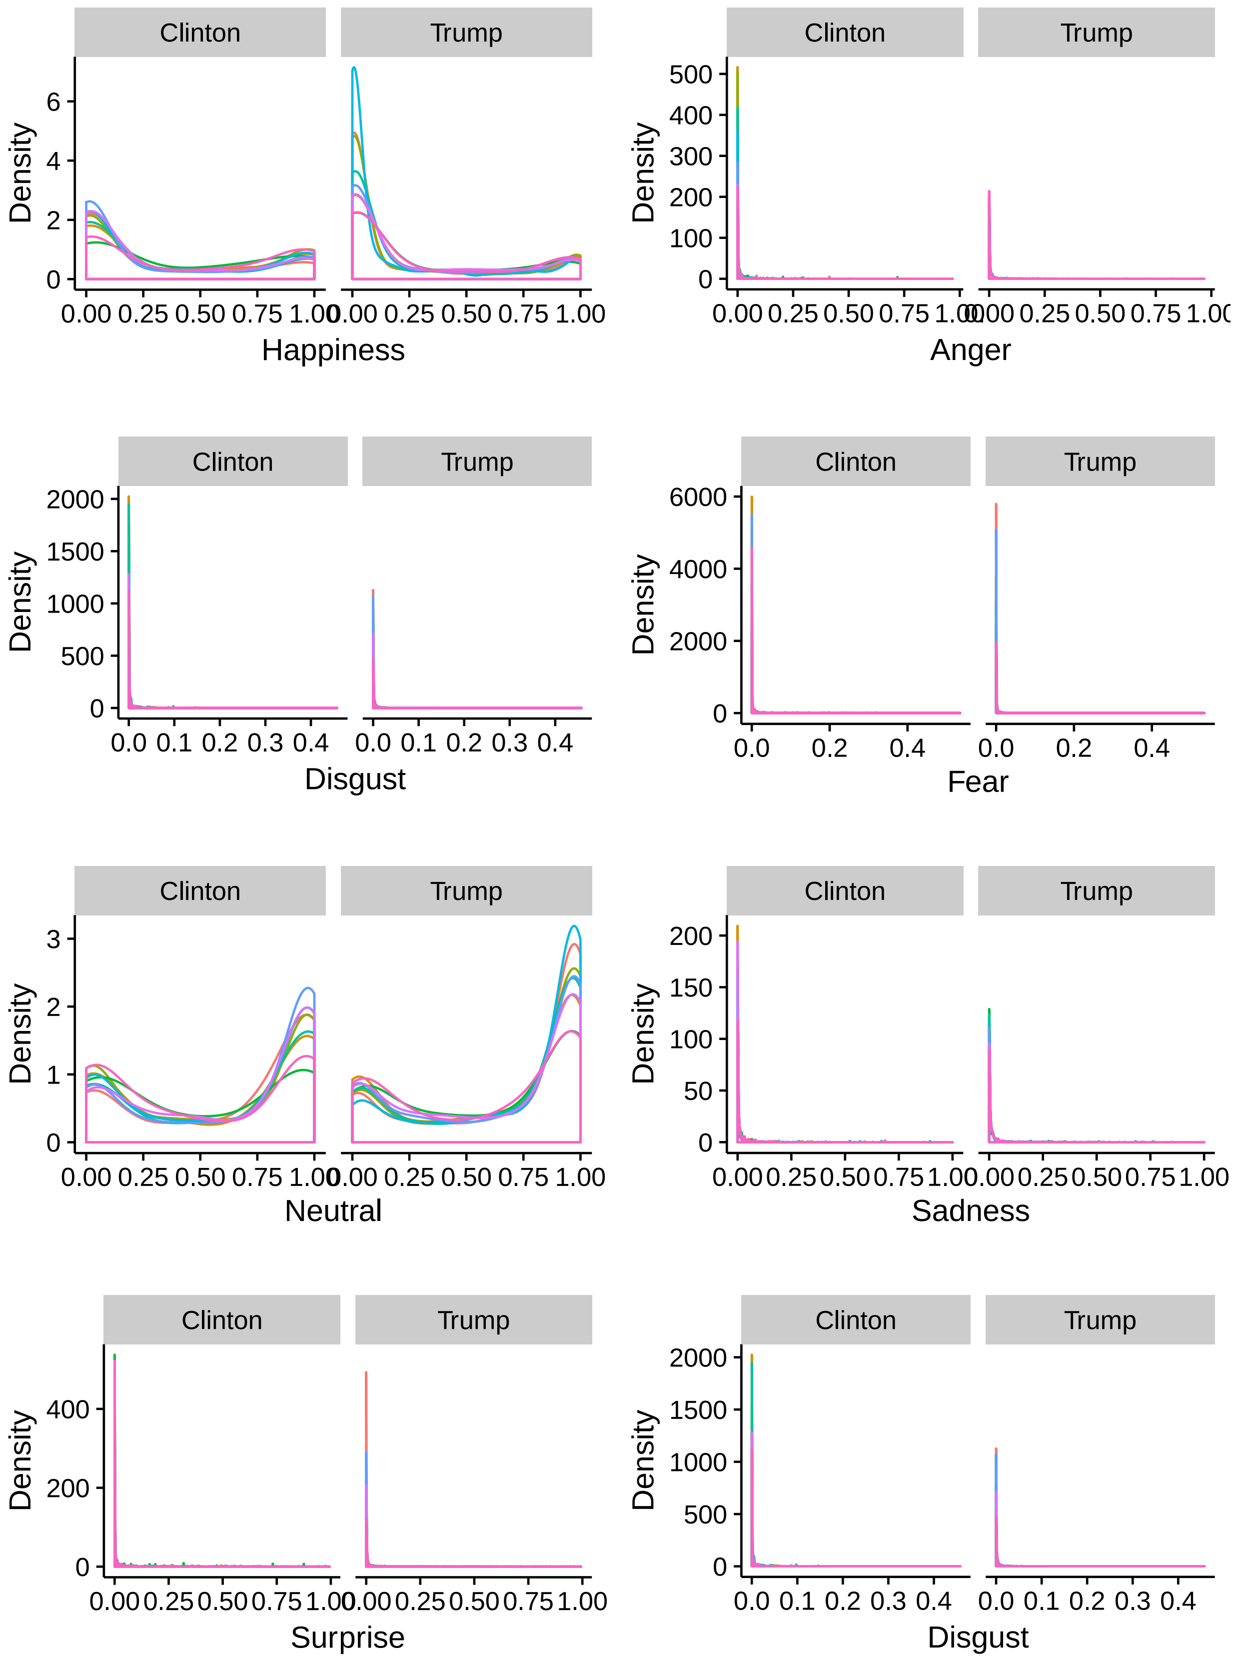

Figure 12: *Distribution of the categorical emotion variable*

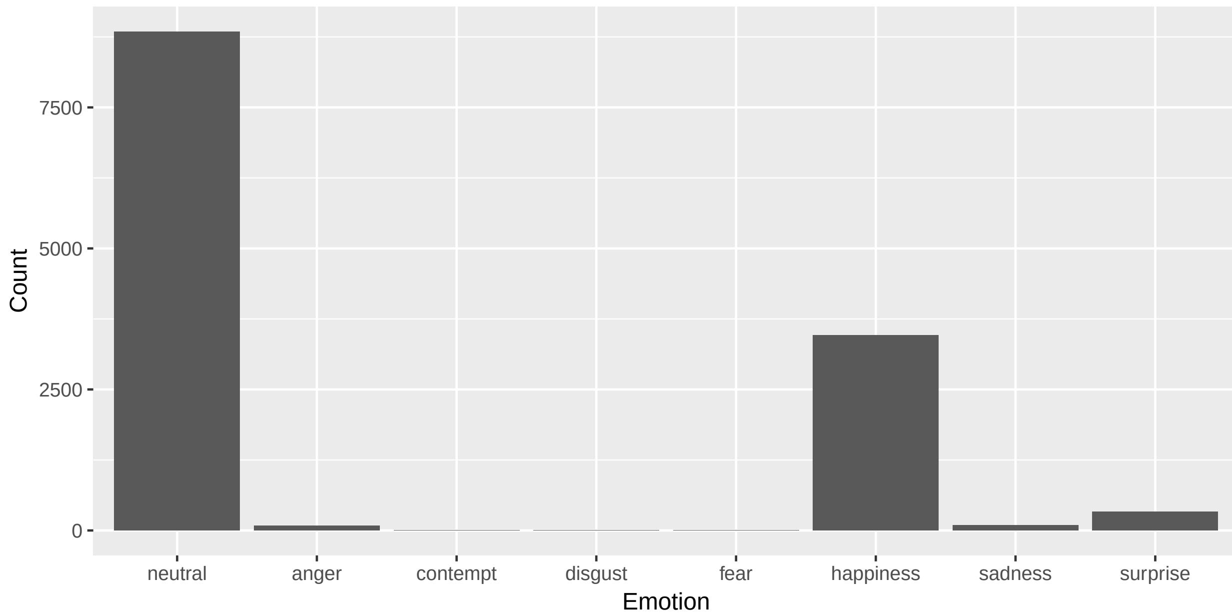

Figure 13: *Predicted probabilitiy of exhibiting happiness, by news outlet and candidate, primaries and general election campaign*

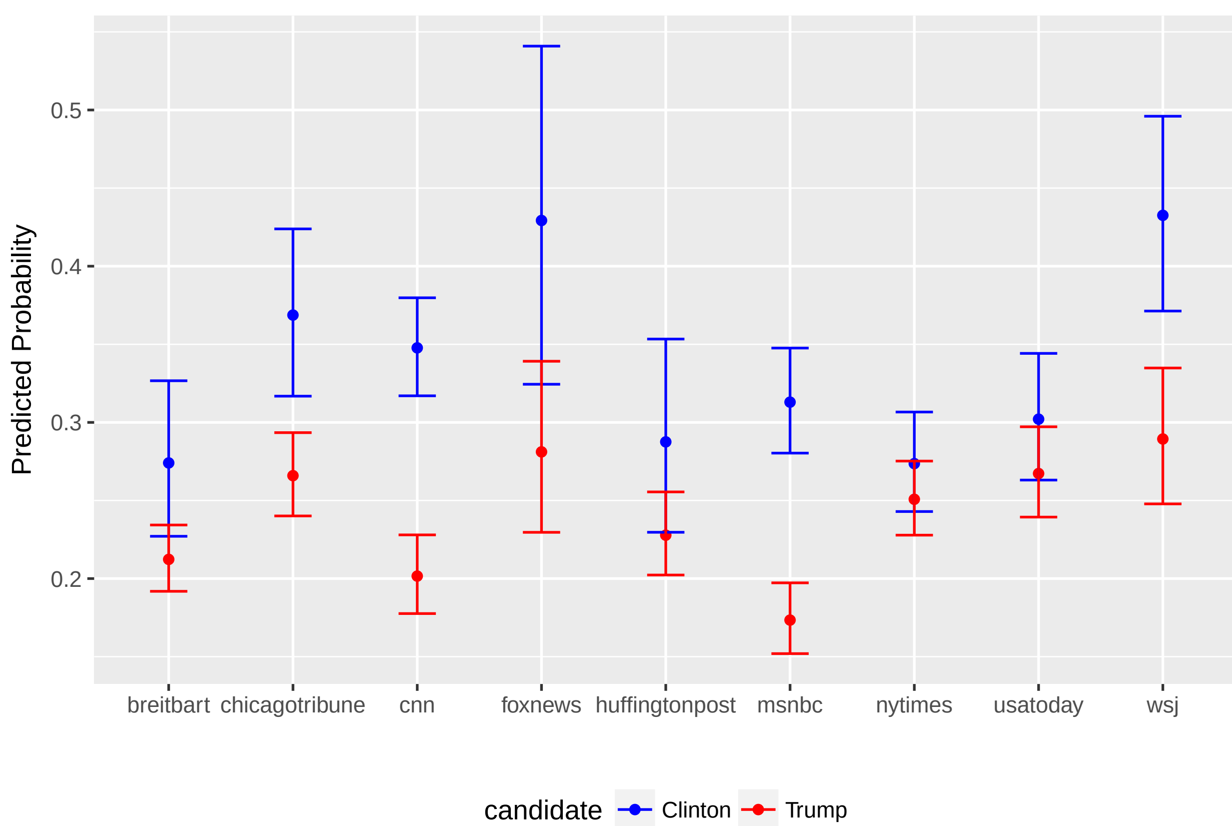

Figure 14: *Predicted probabilitiy of exhibiting anger, by news outlet and candidate, primaries and general election campaign*

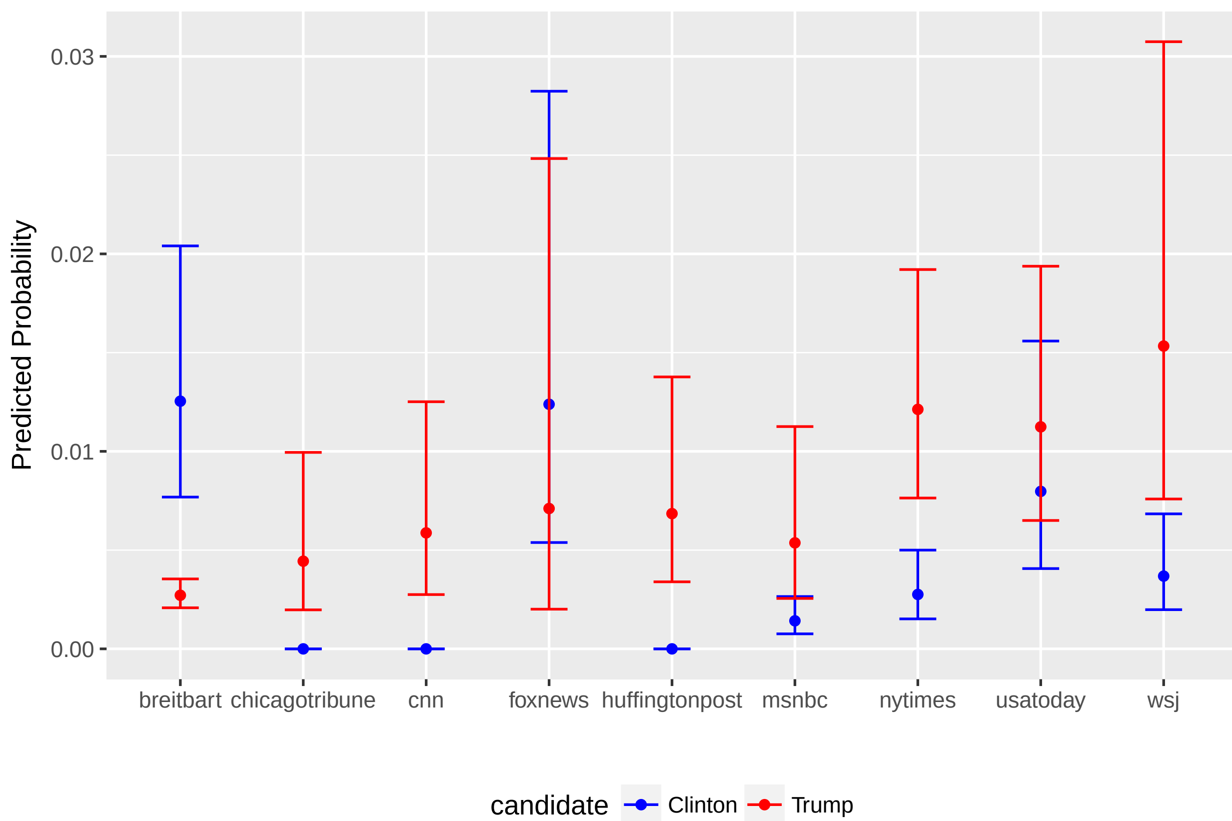

Figure 15: *Predicted probabilitiy of exhibiting no emotion, by news outlet and candidate, primaries and general election campaign*

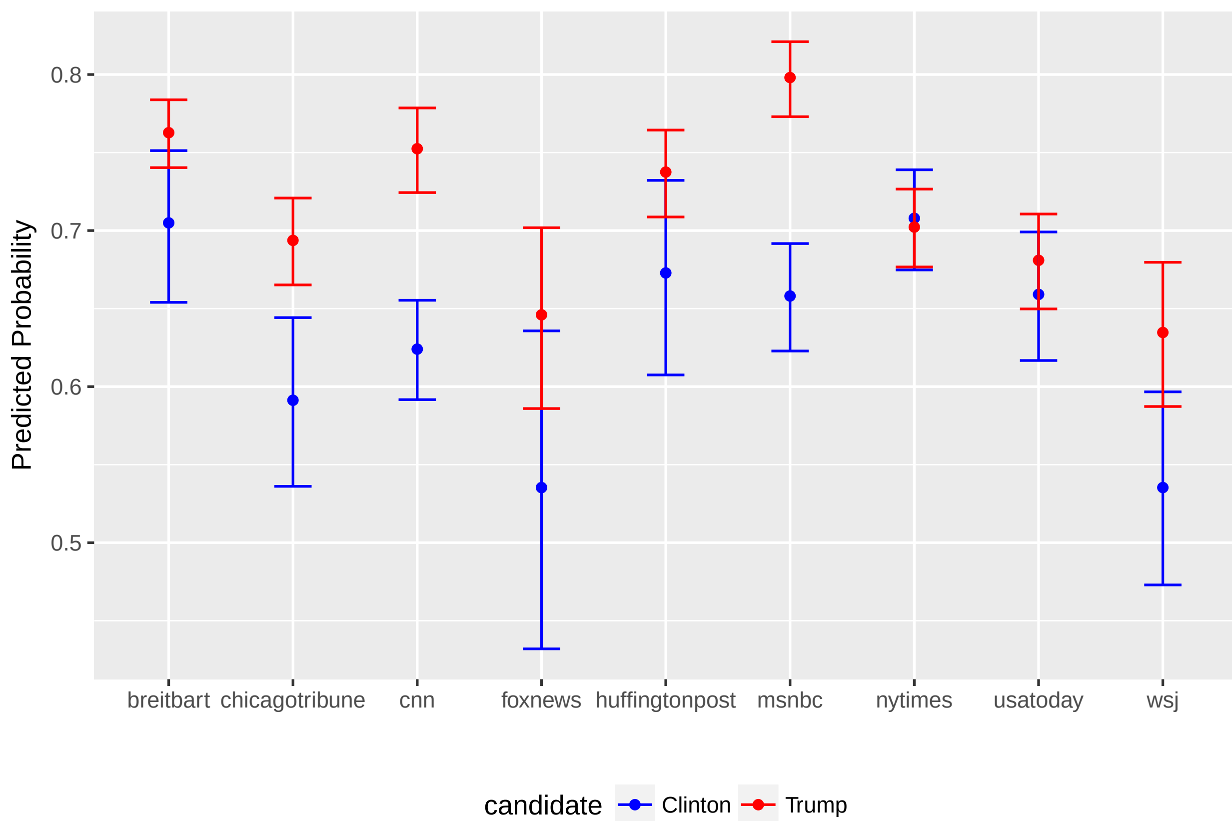

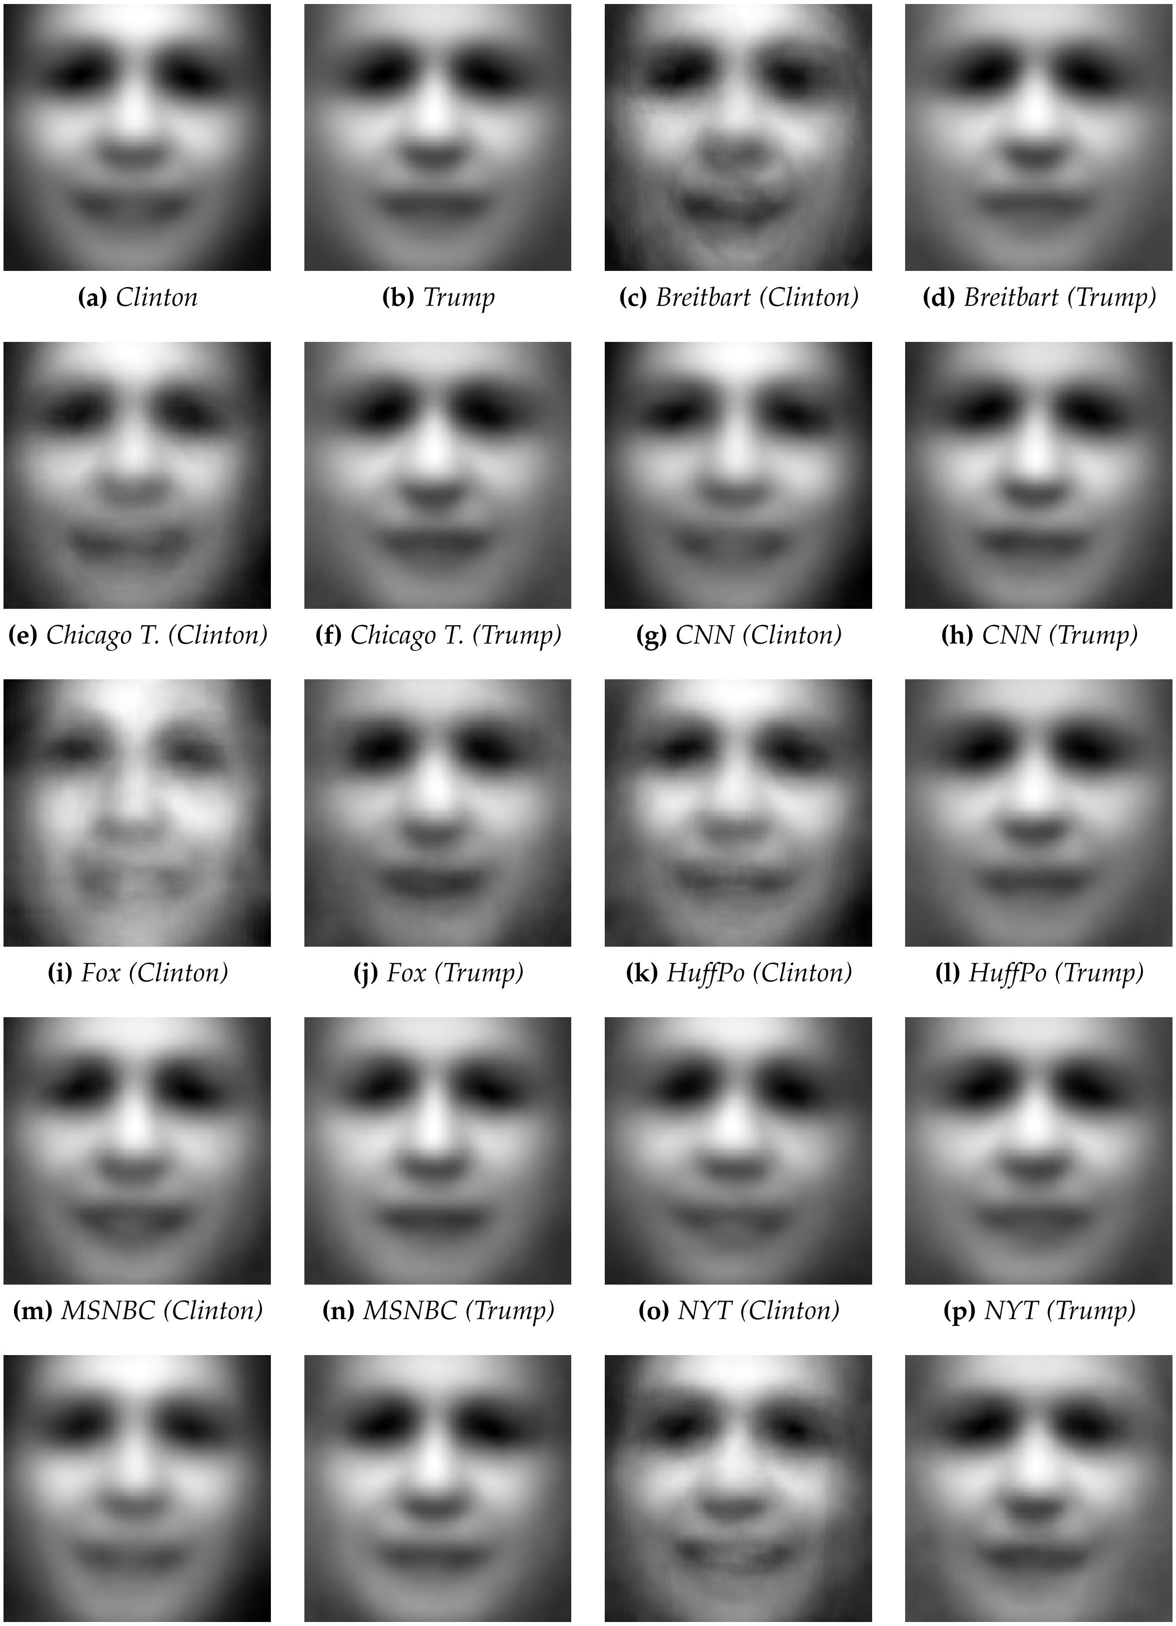

(a) *Clinton* **(b)** *Trump* **(c)** *Breitbart (Clinton)* **(d)** *Breitbart (Trump)*

(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 16: *Average face, by news outlet and candidate*