Trait Associations for Hillary Clinton and Donald Trump in News Media: A Computational Analysis

November 15, 2017

Abstract

We study media representations of Hillary Clinton and Donald Trump in the 2016 United States presidential election. In particular, we train models of semantic memory on a large number of news media outlets that published online articles during the course of the election. Based on the structure of word co-occurrence in these media outlets, our models learn semantic representations for the two presidential candidates, as well as for widely studied personality traits. We find that models trained on media outlets most read by Clinton voters and media outlets most read by Trump voters differ in the strength of association between the two candidates' names and trait words pertaining to morality. We observe some differences for trait words pertaining to warmth, but none for trait words pertaining to competence.

Keywords: Social Perception; Semantic Representation; Political Psychology; Computational Models, Media

Introduction

The relationship between trait associations and the perception of individuals and groups has a long history of research in psychological science. Within social psychology, dimensional models of person perception have been dominant, positing that there are two (e.g., Cuddy, Fiske, & Glick, 2008; Fiske, Cuddy, & Glick, 2007; Fiske, Cuddy, Glick, & Xu, 2002) or three (e.g., Brambilla, Rusconi, Sacchi, & Cherubini, 2011; Brambilla, Sacchi, Rusconi, Cherubini, & Yzerbyt, 2012; Goodwin, Piazza, & Rozin, 2014; Landy, Piazza, & Goodwin, 2016; Leach, Ellemers, & Barreto, 2007) key trait dimensions upon which individuals and social groups are evaluated. In the context of political preferences, existing research suggests that the most important trait dimension is competence (Ballew & Todorov, 2007; Cislak & Wojciszke, 2006; Funk, 1997; Todorov, Mandisodza, Goren, & Hall, 2005), although other research also highlights the importance of the morality dimension (Goodwin et al., 2014).

There are two difficulties facing researchers who attempt to estimate the dimensional structure of representations for real political candidates. Firstly, associations between individuals and traits or dimensions may vary as a function of political context (e.g., Goodwin et al., 2014; Judd, James-Hawkins, Yzerbyt & Kashima, 2005). Additionally, determining trait associations with political candidates in the past (i.e., in the context of political events that have already occurred) is nearly impossible using survey-based techniques, making retrospective analyses of person perception in the political domain very difficult.

One alternative approach to standard survey-based techniques, and one that we pursue in the present paper, is to study representations in media coverage pertaining to a particular election. The election that we consider is the 2016 US Presidential election, and we attempt to uncover person representations and corresponding trait associations with Donald Trump and

Hillary Clinton present in a large number of media outlets that published articles during this election. In addition to helping us characterize media representations for Trump and Clinton, this approach allows us to test for differences between media outlets favored by Trump voters and media outlets favored by Clinton voters. In our approach we utilize distributional models of semantic memory and our contribution is therefore to also showcase how this methodology can be applied to studying person perception in the context of geopolitical events. Of course, our analysis also sheds light on some of the psychological factors underlying one of the most turbulent and unusual events in recent US history.

Trait Associations for Political Candidates

In general, much of the existing work on person perception in the political domain has examined how voter preferences depend on the target's trait dimensions (e.g., how competent is a specific candidate perceived to be). In this type of research, participants are often presented, either explicitly or implicitly, with trait information about different individuals. Less is known, however, about the origin of trait-candidate associations, and how these associations are influenced by the immediate political context. Existing evidence strongly suggests that news media play an important role shaping people's voting attitudes and preferences (DellaVigna & Kaplan, 2007; Gerber, Karlan, & Bergen, 2009; also Margetts, 2017). Media coverage is often used as a source of information about social norms, which can shift views of individuals, and lead to the polarization and homophily of attitudes (Bennett, 2012; Slater, 2007). In a similar vein, media coverage is often the only source of information about candidates' personalities and values. Given that political preferences can be shaped by people's perception of the match between their own traits and those of political candidates, it is not surprising that media portrayal of candidates can be very influential (Caprara & Zimbardo, 2004). It is also well known that people exhibit strong confirmation bias in consuming political information, and that selective exposure can drive the emergence of polarization and compartmentalization in political thought. Thus causality is likely to go both ways, with the preferences for selective consumption on one side, and growing personalization and selective marketing by media to meet demands of specific groups of customers on the other (Bakshy, Messing, & Adamic, 2015; Van Aelst, Sheafer, & Stanyer, 2012). More generally, the content of news media both reflects and influences public sentiment and, for that reason, can yield useful insights about the psychological correlates of voters' preferences (Andranik, Sprenger, Sandner, & Welpe, 2010; Hopmann, Vliegenthart, De Vreese, & Albaek, 2010; De Vreese & Semetko, 2004).

In order to examine representations for Hillary Clinton and Donald Trump in news media we adopt the insights of distributional models of semantic memory. These models propose that people's semantic representations for objects and concepts can be approximated by examining the statistical structure of the natural language environments that they are exposed to (Griffiths, Steyvers, & Tenenbaum, 2007; Jones & Mewhort, 2007; Landauer & Dumais, 1997; Mikolov et al., 2013; Pennington, Socher, & Manning, 2014; see Firth, 1957, and Harris, 1954, for early theories). This insight has been successfully used to study a range of psycholinguistic and cognitive phenomena. For example, distributional models trained on educational texts have been shown to describe the learning of word knowledge in children (Landaur & Dumais, 1997; Landauer, Foltz, & Laham, 1998). Likewise, models trained on everyday language datasets have been shown to predict judgments of word similarity, semantic priming effects, word categorization effects, behavior in free association tasks, and behavior in recall tasks, for adults (see Bullinaria & Levy, 2007, or Jones, Willits & Dennis, 2015, for a review). The power of this approach extends beyond low-level cognition: Bhatia (2017) has applied distributional models of semantic memory, trained on news media and online encyclopedias, to predict high-level associative judgments, including forecasting and factual judgment. Bhatia (in press) has similarly applied this approach to study prejudice and stereotyping, as learned from news media (see also Lenton, Sedikides & Bruder, 2009 for a related insight). More relevant to the topic of this paper, Holtzman et al. (2011) and Dehghani, Sagae, Sachdeva and Gratch (2014) have applied distributional models to study political bias in traditional news media and social media respectively, and Garten et al. (2016) have recently used this approach to study morality-based representations in social networks. The ability of distributional models to approximate actual human semantic representations suggests that training these models on news media published during the 2016 U.S. presidential election could reveal key similarities and differences in representations and trait associations for the two candidates across different types of media outlets. This could, in turn, shed light on how Trump and Clinton voters represented these two candidates over the course of the election.

Pilot Study of News Preferences

A critical component of our analysis involves studying differences in trait associations for Trump and Clinton in media sources that were read by Trump or Clinton voters. For this purpose, we first ran a pilot study of media consumption preferences.

Methods

In the first session of this study, we asked 200 U.S. citizens, recruited through the online experiments website, Prolific Academic, about the media outlets they visited regularly and trusted for information about politics and current affairs. Participants in this study were shown the names of the 250 media outlets we had in our news corpus (see corpus description in later sections of this paper) and were asked to select as many outlets as they wished. After this task, participants were asked to indicate which of the political candidates they voted for in the election. Due to the imbalance in Clinton vs. Trump voters in this session, we ran a second study session, in which we recruited an additional 100 self-reported Republicans and 100 self-reported Democrats from Prolific Academic (who hadn't taken part in our first session) and asked them about their media preferences in the same way as described above. Thus our eventual study had a total (both sessions) of N = 400 participants (mean age = 33.73, SD age = 11.87, 38% female).

Results

We pooled the data from these two sessions and used it to calculate the degree to which Clinton and Trump voters relied on each media outlet for political information. For each media outlet i we calculated $r^C{}_i$ and $r^T{}_i$, which is the total number of Clinton voters and Trump voters in our study who stated that they relied on that media outlet. We then used $r^C{}_i$ and $r^T{}_i$ to calculate the extent to which Clinton and Trump voters relied on that outlet relative to each of the other media outlets, to obtain $R^C{}_i = r^C{}_i / \sum_{j=1}^{250} r^C{}_j$ and $R^T{}_i = r^T{}_i / \sum_{j=1}^{250} r^T{}_i$. Finally, we calculated a measure of relative voter reliance, $RVR_i = R^C{}_i - R^T{}_i$. Strongly positive values of RVR_i indicate that the outlet i was relied on by Clinton voters and not Trump voters, whereas strongly negative values of RVR_i indicate that the outlet was relied on by Trump voters and not Clinton voters.

Out of the 200 participants in the first session, 41 voted for Trump, 94 voted for Clinton, and the remaining did not vote, voted for a third party candidate, or preferred not to discuss their voting behavior. The data from the second session was more balanced, and contained 91 Clinton voters and 70 Trump voters. Our two samples gave us a considerable spread in terms of preferences for specific media outlets. The ten media outlets with the highest and lowest readership among Clinton and Trump voters (highest and lowest RVR scores, respectively) are listed in Table 1 below.

Study of Trait Associations in Media Outlets

With the relative voter reliance ratings for each media outlet, we now proceed to our main study, in which we apply models of semantic memory to uncover trait associations for Trump and Clinton in each outlet in our sample. For our analysis, we obtained online news articles published by 250 U.S.-based media outlets, over the course of the 2016 election. We trained semantic memory models individually on each of these news outlets. This yielded 250 different models, each with unique representations and associations for a very large set of words and phrases. These words and phrases included *Hillary Clinton*, *Donald Trump*, and a large number of trait words, and thus for each model (i.e., each media outlet) we were able to calculate the strength of association between the two candidates' names and the various trait words and dimensions. We then examined both the absolute associations between the two candidates and trait dimensions, as well as the differences in these associations for media outlets predominantly relied on by Clinton voters and media outlets relied on by Trump voters, based on the ratings from our pilot study.

Methods

Latent semantic analysis. The semantic model used in this paper is latent semantic analysis (LSA), one of the earliest, simplest, and most prominent computational theories of semantic representation (Landauer & Dumais, 1997; Landauer et al., 1998). The core insight of LSA is that words can be represented using multidimensional vectors, which are obtained by performing a dimensionality reduction on word-distribution data. The relationship or association between words is subsequently captured by the proximity of the vectors for their corresponding words.

More formally, for a natural language environment with N different words occurring in K contexts, we can represent the structure of word distribution using an NxK matrix S. S captures word-context co-occurrence, so that the number of times word n occurs in context k is indicated by the value of the cell in row n and column k of S. LSA recovers vector representations for each of the words by decomposing S into some $M << K$ latent dimensions. This decomposition is achieved through singular value decomposition, and the resulting matrix can be written as $S^*= U \cdot V \cdot W$. Here V is an MxM matrix with the M largest singular values from the decomposition, U is the corresponding NxM matrix of words, and contains a representation of each of the N words as vectors on the M latent dimensions.

The proximity between these vectors provides a quantitative account of word relationship and association. The metric typically used to compute vector proximity, and thus word association, is cosine similarity, so that the proximity between any two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ is given by $\mathrm{sim}(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x} \cdot \boldsymbol{y} / (||\boldsymbol{x}|| \cdot ||\boldsymbol{y}||)$. This metric varies between -1 and +1, with values close to +1 capturing words which are very positively associated, and values close to -1 capturing words which are very negatively associated.

As with other similar approaches like factor analysis and principle components analysis, the recovery of the M latent dimensions from SVD is useful for describing word representations in a manner that accounts for the structure of the variability in word occurrence. For this reason, the ability of LSA to predict human behavior depends critically on the right value of M. If M is too large or if $M = K$ (which is the special case with no dimensionality reduction) the model is able to process only first-degree co-occurrence, and cannot infer relationships between two words that never directly co-occur but occur frequently with a third word. Likewise, if M is too small the model makes spurious inferences regarding the relationships between words. $M = 300$

has been shown to provide the best account of behavioral data in psycholinguistic and cognitive tasks (Landauer & Dumais, 1997; Landauer et al., 1998).

Note that since the work of Landauer and Dumais, there have been a number of technical advancements in uncovering semantic representations from text (e.g., Mikolov et al., 2013; Pennington et al., 2014). These techniques differ in terms of how they specify word context, and how they use contextual similarities in word use in the absence of direct word co-occurrence. We do not expect the key qualitative results of our analysis to vary with the use of different techniques (though more recent approaches may lead to lower variability and better quality word representations).

News corpora. The goal of this paper is to use LSA to understand the associations and representations embedded in news media over the course of the 2016 U.S. presidential election. For this purpose, our analysis used the NOW corpus (http://corpus.byu.edu/now/), a very large dataset of news articles published online and linked to by Google News. Currently, this corpus contains over three billion words of data from 2010 to the present time. However, in order to study news media during the 2016 election, we restricted our analysis to articles published between June 16, 2015 (the date that Donald Trump declared his candidacy) and November 7, 2016 (the day before the election). Additionally, we considered only U.S. media outlets. Finally, for tractability, we examined only the 250 media outlets with the most number of published articles in the NOW corpus during this time period. These outlets include traditionally print or television news media, purely online news media, as well as sports and entertainment media. Our final dataset consisted of 322,699 online articles with a total of 82,784,212 words.

To facilitate our analysis, we lowercased all words in the corpus of news articles, removed all punctuation, and replaced each mention of *donald trump* and *hillary clinton* with *donald_trump* and *hillary_clinton* respectively. We then identified unique words based on white space (using *donald_trump* and *hillary_clinton* ensured that we retained representations for the full names of the candidates rather than splitting them up into first and last names).

Candidate-trait associations. We trained LSA models on each of the 250 media sources separately, yielding 250 models. For each model, we considered each article to be a unique context. Thus for media outlet i we obtained the matrix S_i capturing the occurrence of words across the articles in this outlet. We also applied a term frequency-inverse document frequency (tf-idf) transformation to each matrix S_i (which is recommended in order to control for word frequency effects) prior to performing singular value decomposition. Note that words were not stemmed, however we excluded words present in a common set of stop words. Our decomposition used 300 latent attribute dimensions, as recommended by prior work on LSA (Landauer & Dumais, 1997; Landauer et al., 1998).

Our analysis used a list of 170 personality traits compiled Goodwin et al. (2014). Each of these traits were rated in terms of their relevance to nine broader trait dimensions: ability, agency, character, communion, goodness, grit, morality, strength, and warmth (see Goodwin et al. for details). Goodwin et al.'s data also contains ratings on a general personality dimension (which does not reflect any particular trait, i.e., the question was simply how useful the trait would be in providing information about a person's personality) and ratings of overall valence for each trait word (how positive or negative each word is), but we did not include these ratings because of their non-specificity.

In order to calculate the association between the presidential candidates and these traits we used the vectors corresponding to the individual trait words (e.g. *honest*, *warm*, *intelligent*) and calculated their cosine similarity with the vectors corresponding to *donald_trump* and *hillary_clinton* (to avoid confounds we did not separately consider vectors for *trump* or *clinton* as these could correspond to references to other members of the Trump and Clinton families). Using these methods, we had, for each outlet and for each trait word, a measure of the association of that trait word with Hillary Clinton and with Donald Trump on that outlet. For a trait j and media outlet i, we write this association as A^C_{ij} for Hillary Clinton and A^T_{ij} for Donald Trump. We used $RA_{ij} = A^C_{ij} - A^T_{ij}$ to specify the relative association of the trait with Clinton vs. Trump for the media outlet. Positive scores of RA_{ij} correspond to a stronger association between a given trait j and Clinton in a given news outlet i. Negative scores of RA_{ij} correspond to a stronger association between a given trait j and Trump in a given news outlet i.

Results

Nine Dimensions

There are a few outlets that were not listed as being relied on by any Trump or Clinton voters, and so we excluded these outlets from our analysis. Additionally, some outlets did not have articles mentioning either Hillary Clinton or Donald Trump, and thus we removed these outlets from our analysis as well. This left 197 media outlets.

Our main objective was to examine differences in the relative strength of association between the candidates and different trait dimensions, and to test whether these differences depended on whether the media outlets were read by Clinton voters or Trump voters. This involved running a regression in which the dependent variable was the relative association between the candidates and a trait j as assessed on the model trained on media outlet i (RA_{ij}). The corresponding independent variables were the rating of trait j on Goodwin et al.'s (2014) trait dimensions (ability, agency, character, communion, goodness, grit, morality, strength, or warmth), the relative reliance of Clinton vs. Trump voters on media outlet i (RVR_i), and the interaction between RVR_i and the trait dimension rating. The main effect of the trait dimension rating indicates whether the trait dimension is more associated with Clinton or Trump, independently of the media outlet in consideration. The interaction term captures how media sources relied on by Clinton voters or Trump voters associate different dimensions with Clinton or Trump. A strong positive interaction term would reveal that media outlets favored by Clinton voters represented Clinton as being associated with the trait dimension more than Trump, and vice versa for media outlets favored by Trump voters. A negative interaction term would indicate the opposite, namely that media outlets favored by Clinton voters represented Trump as being associated with the trait dimension more than Clinton, and vice versa for media outlets favored by Trump voters.

We ran nine regressions (one for each of Goodwin et al.'s nine trait dimensions) with random effects for the media outlets. Firstly, none of the main effects of any dimension were significant, indicating that there were no absolute differences in associations for Clinton and Trump on the nine dimensions, independent of the media outlet. These regressions did however reveal positive and highly significant interaction effects for the character, communion, goodness, and morality dimensions ($p < 0.001$ for these four dimensions; see Table 2). They also revealed a weaker positive interaction effect for the warmth dimension ($p < 0.05$). There were no significant interaction effects for the ability, agency, grit, or strength dimensions. These results indicate that media outlets read by Clinton (Trump) voters were more likely to associate Clinton (Trump) with the character, communion, goodness, and morality dimensions. There is some evidence that these differences persist for the warmth dimension, but no evidence that they do so for the ability, agency, grit, and strength dimensions. Note that the interaction effects for the character, communion, goodness, and morality dimensions survive a Holm–Bonferroni correction for multiple comparisons, but the warmth interaction effect does not.

Two and Three Dimension Models

Research on social perception has suggested that there are two or three core dimensions that best capture person perception. Two dimensional theories posit that these dimensions are warmth (which includes morality) and competence, whereas three dimensional theories posit that these dimensions are warmth, morality, and competence. In order to examine trait associations in terms of these core dimensions, we generated a composite variable for morality/warmth by averaging ratings for the character, communion, goodness, morality, and warmth dimensions, a composite variable for morality by averaging ratings for the character, goodness, and morality dimensions, and a composite variable for competence by averaging the ratings for ability and agency dimensions, for each trait. We then ran individual regressions testing for the effect of our composite dimensions on trait associations. Again, the dependent variable in these regressions was the relative association between a candidate and a trait on a media outlet (RA_{ij}), and the independent variables were the relative voter reliance on the media outlet (RVR_i), the rating of the trait on the composite dimension in consideration, and the interaction between relative voter reliance and the dimension composite (we also included random effects for the media source). These three individual regressions revealed that both the morality/warmth composite and the morality composite had strong significant interaction effects with the relative voter reliance on the media outlet ($p < 0.001$; see Table 2 for details), that survive the Holm–Bonferroni correction for multiple comparisons. In contrast, the competence composite did not display this interaction effect.

We also ran variants of the above regressions directly comparing the composite dimensions against each other. Our first combined regression compared the morality/warmth composite against the competence composite. Here, the dependent variable was RA_{ij}, and the independent variables were RVR_i, the rating of the trait on the warmth/morality composite, the interaction between RVR_i and the morality/warmth composite, the rating of the trait on the competence composite, and the interaction between RVR_i and the competence composite (as before, we also included random effects for the media outlet). This regression revealed a strong interaction effect for the morality/warmth composite ($p < 0.001$) but not the competence composite. We also repeated the above analysis but with a separate morality composite and warmth composite (instead of a combined morality/warmth composite). We therefore ran the regression with the morality, competence and warmth composites, and their respective interactions with the RVR_i. This analysis revealed a significant positive interaction effect for the morality composite ($p < 0.05$) but not the warmth dimension or the competence composite. Overall, these results indicate that media outlets read by Clinton (Trump) voters were more likely to associate Clinton (Trump) with traits rated highly on the morality composite. These is no evidence for differences on the warmth or competence composites. The results of these two regressions are also summarized in Table 2. Again, there were no significant main effects of the dimensional composites on relative associations.

Valence Analysis

Goodwin et al.'s (2014) dataset contains negative trait words that occasionally map on quite strongly to dimensions like morality and warmth (e.g., *dishonest* has a high rating on the morality dimension as it is strongly reflective of the morality of an individual). It is necessary to rerun the above analysis excluding the negatively valenced words, to ensure that the effects documented above did not emerge due to a perverse association between Clinton and immorality-related words in pro-Clinton outlets, and Trump and immorality-related words in pro-Trump outlets.

We did this by performing a median-split on the valence ratings in Goodwin et al.'s data and excluding trait words below the median valence. This regression did not change findings for the nine dimensions, except that the interaction effect of character decreased in significance to $p = 0.02$, failing the Holm-Bonferroni correction for multiple comparisons. The interaction effect of warmth also increased in significance to $p = 0.01$, however it still failed the Holm–Bonferroni correction. We still found a significant interaction effect for the morality/warmth composite ($p < 0.01$), in a regression with the morality/warmth composite and the competence composite. The significant interaction effect of a separate morality composite in a variant of this regression including warmth and the competence composite also persisted ($p < 0.05$). A related set of regressions for only negatively valenced words (words below the median valence on Goodwin et al.'s trait ratings) did not yield any significant differences for any dimensions of interest ($p > 0.05$). This result could be due to differences in frequency of word use for positively and negatively valenced traits, in news media.

Dimension Independence Analysis

Results reported so far suggest that semantic representations of Clinton and Trump present in media read by Clinton and Trump voters varied primarily in terms of morality associations. There seems to be some evidence that warmth associations could have varied as well. However, the warmth dimension is positively correlated with the morality dimension (as well as the morality composite), so this relationship could be incidental. Our analysis testing for the independent effects of both the morality composite and warmth simultaneously found no effect of warmth (but a significant positive effect of morality).

Another way to attempt this analysis is to use only a subset of the trait words that Goodwin et al. compiled that are either high in morality and neutral/low in warmth or neutral/low in morality and high in warmth. There are eight such high morality/low warmth words and eight low morality/high warmth words in Goodwin et al.'s Study 3, and we attempted to replicate our above analysis using these words. For comparability we also performed this analysis with the eight high competence words in Goodwin et al. The words used in this analysis are summarized in Table 3.

First we ran three separate regressions, examining the individual effects of these sets of words on trait associations. In the first regression, our dependent variable was relative association RA_{ij}, and the independent variables were the relative voter reliance RVR_i, whether or not the word was a high morality/low warmth word (1 if high morality/low warmth, and 0 otherwise), and the interaction between RVR_i and the high morality/low warmth binary variable. The other two regressions replaced the high morality/low warmth binary variable and its interaction with the low morality/high warmth and the high competence binary variables. As above, the regressions included random effects for the media source.

Consistent with the other results, we found a significant interaction effect ($p < 0.05$) for the high morality/low warmth regression, but no interaction effects for the low morality/high warmth or high competence regressions. We also performed a combined regression, in which each of the three binary variables and their interaction with RVR_i were included together as independent variables. This regression also found a significant interaction effect ($p < 0.05$) for the high morality/low warmth variable, but not for the other two variables (there were also no significant main effects for the morality, warmth, and competence variables in any of the regressions). The results of these regressions are provided in Table 4. Taken together, this additional analysis supports findings reported in Table 2, showing that only trait associations for the morality dimension for Trump and Clinton are reliably different between media outlets read by Trump and Clinton voters.

General Discussion

The 2016 U.S. presidential election was unusual in many respects, including its acrimony, and the tendency for its two major candidates to focus on the moral failings of their opponent (Trump's preferred nickname for Clinton was *Crooked Hillary*, which involves a distinctly moral indictment). Such associations were likely reflected in media representations of the political candidates. We formally studied these representations using computational theories of semantic memory and social psychological approaches to understanding the dimensional structure of social perception. In particular, we trained models of latent semantic analysis (Landauer & Dumais, 1997) individually on a large number of news sources that published online articles over the 2016 election, and tested the associations possessed by these models between Hillary Clinton, Donald Trump, and various words describing personal traits. Comparing media outlets favored by Clinton voters against media outlets favored by Trump voters, we found that differences in representations pertained primarily to morality, with Clinton outlets more strongly associating Clinton with moral traits relative to Trump, and Trump outlets more strongly associating Trump with moral traits relative to Clinton. We also observed some differences for warmth traits, though these differences disappeared when the positive correlation between warmth and morality was controlled for. There were no differences across media outlets for associations with competence traits. Finally, we found no main effects for the trait dimensions which shows that there were no absolute differences in the extent to which morality., warmth, or competence traits were associated with the two candidates, independent of the underlying media outlet.

Recent work has argued that moral traits are powerful determinants of person evaluation, and that morality is distinguishable from other core dimensions such as warmth (Brambilla et al., 2011; Brambilla et al., 2012; Goodwin et al., 2014; Landy et al., 2016; Leach et al., 2007). Our results support this conclusion, and suggest that the differences in moral associations in news media during the 2016 presidential election could have correlated with voter preferences and subsequently the outcomes of this election. Our finding that morality was predominant, distinguishes our work from past research which has tended to stress competence as the primary predictor underlying impressions of political candidates (e.g., Ballew & Todorov, 2007; Cislak & Wojciszke, 2006; Funk, 1997; Todorov et al., 2005). The salience of morality in this particular election could have been in part driven by the fact that Clinton was the first female presidential candidate in the U.S. Previous research shows that women candidates receive more, not less coverage than their male running rival but that this coverage often focuses on stereotypes, including assumptions that female candidates are more compassionate and honest (Fridkin & Kahn, 1992). One possible, albeit speculative, interpretation of our findings is that the focus on morality that dominated the 2016 election was particularly harmful to Clinton due to the existing stereotypes that set a biased expectation about the moral standing of a female candidate for office. Such an explanation is in line with the congruity theory (Eagly & Karau, 2002), by which positive evaluations of a target individual emerge mainly from the consistency between their leadership and social roles. In fact, in testing a prediction of this particular model, Gervais and Hillard (2011) showed that the exceptional level of prejudice expressed towards Clinton was in part due to her behavior, which violates traditional gender roles. We believe that our results complement these findings by showing dimensional differences between associations with the candidates' traits. However, further work is necessary to establish how these differences emerge.

The semantic models we have applied are widely considered to mimic human semantic learning and representation processes (see Jones et al. 2015, for a review). Thus, we would expect individuals selectively exposed to the media outlets considered in this paper to have developed the types of biased moral associations for the two candidates observed in our tests. Although we have not tested this causal link in our work, there is evidence that media bias has this type of causal role (DellaVigna & Kaplan, 2007; Gerber, Karlan, & Bergen, 2009). However, it is also possible that the media biases reflected pre-held voter beliefs, and that the observed differences in moral associations for Clinton and Trump across outlets favored by Clinton vs. Trump voters were the result of the preexisting associations of the readers of these outlets. Regardless of the causal direction, our analysis identifies core differences in representations and associations across media outlets favored by different types of voters, and thus suggests a relationship between moral trait associations in peoples' information environments and voting behavior.

Of course, a more rigorous test of this relationship would involve participant data on trait ratings for Clinton and Trump. Unfortunately, such data is hard to obtain retrospectively: The election outcomes and subsequent political events may have changed peoples' associations with Clinton and Trump, and it is not clear how survey-based methods could control for these changes. We are also not aware of any studies that elicit views about Trump and Clinton using the list of traits that we used in the present study. In fact, the strength of our approach lies in the fact that our models are suited for testing any of a range of dimensional theories, as the analysis is not constrained by the choice of trait words. By obtaining a very large corpus of news media data and leveraging theoretical and technical advances in semantic memory research, the present research is able to uncover the representations present in different real-world information environments. This approach can be extended to examine more than just trait associations with political candidates. We could use a nearly identical analysis on our dataset to uncover differences in trait associations with different social groups (e.g., Muslims and immigrants) across outlets favored by Trump and Clinton voters, during the 2016 election. With an expanded media dataset, we could extend such analyses to prior presidential elections, as well as to social phenomena outside of politics. The increased digitization of information and the corresponding growth in computational resources and technologies for analyzing this information has made it possible to rigorously study the social representations present in peoples' everyday information environments. This has exciting possibilities for our understanding of real-world social cognition and judgment (see Gilovich, 2015; Jones et al., 2015; Margetts, 2017), and we look forward to further contributing to this approach.

Table 1. The ten media outlets with the highest and lowest readership among Clinton and Trump voters. Readership is quantified using Relative Voter Reliance (*RVR*) scores, explained in the main text. Positive (negative) values indicate that the outlet was relied on by Clinton (Trump) voters and not Trump (Clinton) voters.

Media Outlet	*RVR*
CNN	0.027
Washington Post	0.021
NPR	0.021
The New Yorker	0.014
The Atlantic	0.014
…	
CBS Local	-0.009
FOX Sports	-0.009
NFL.com	-0.010
Breitbart News	-0.025
Fox News	-0.047

Table 2: Interaction effects between relative voter reliance on the media outlet and the dimension rating of the trait word for predicting the association between that trait and Clinton vs. Trump, in the media outlet. The regressions reported here use all trait words from Goodwin et al. (2014). Individual regressions analyze each dimension or each dimensional composite separately, whereas the combined regressions use all dimensional composites simultaneously. All regressions include random effects for media source, as well as main effects for dimension rating and relative voter reliance (not reported here).

	β	SE	*z*	*p*	95CI-L	95CI-H
Individual Regressions						
Ability	-0.05	0.07	-0.70	0.49	-0.18	0.09
Agency	0.05	0.07	0.63	0.53	-0.10	0.19
Character	0.21	0.06	3.33	< 0.001	0.09	0.34
Communion	0.23	0.07	3.43	< 0.001	0.10	0.35
Goodness	0.18	0.05	3.60	< 0.001	0.08	0.28
Grit	0.04	0.07	0.56	0.57	-0.10	0.18
Morality	0.18	0.05	3.52	< 0.001	0.08	0.28
Strength	0.06	0.07	0.94	0.35	-0.07	0.20
Warmth	0.13	0.05	2.54	< 0.05	0.03	0.24
Individual Regressions (Composites)						
Morality/Warmth Composite	0.04	0.01	3.60	< 0.001	0.02	0.07
Morality Composite	0.07	0.02	3.60	< 0.001	0.03	0.10
Competence Composite	0.00	0.04	-0.07	0.95	-0.08	0.07
Combined Regression						
Morality/Warmth Composite	0.05	0.01	3.66	< 0.001	0.02	0.07
Competence Composite	0.02	0.04	0.63	0.53	-0.05	0.10
Combined Regression						
Morality Composite	0.06	0.03	2.36	< 0.05	0.01	0.12
Warmth	0.02	0.08	0.27	0.79	-0.14	0.18
Competence Composite	0.02	0.04	0.42	0.68	-0.07	0.10

Table 3: The set of high morality and low/neural warmth, low/neutral morality and high warmth, and high competence words, used to isolate the effects of the morality, warmth, and competence, on trait ratings.

High Morality Low/Neutral Warmth	Low/Neutral Morality High Warmth	High Competence
courageous	warm	athletic
fair	sociable	musical
principled	happy	creative
responsible	agreeable	innovative
just	enthusiastic	intelligent
honest	easy-going	organized
trustworthy	funny	logical
loyal	playful	clever

Table 4: Interaction effects between relative voter reliance on the media outlet and the dimension rating of the trait word for predicting the association between that trait and Clinton vs. Trump, in the media outlet. The regressions reported here use the subset of trait words from Goodwin et al. (2014) that are either high in morality and neutral/low in warmth, low/neutral in morality and high in warmth, or high in competence.

	β	SE	*z*	*p*	95CI-L	95CI-H
Individual Regressions						
Morality	0.84	0.35	2.39	< 0.05	0.15	1.53
Warmth	0.12	0.39	0.30	0.77	-0.66	0.89
Competence	-0.48	0.33	-1.46	0.15	-1.13	0.17
Combined Regression						
Morality	0.82	0.35	2.31	< 0.05	0.13	1.51
Warmth	0.14	0.40	0.34	0.73	-0.64	0.91
Competence	-0.43	0.33	-1.29	0.20	-1.08	0.22

References

Bakshy, E., Messing, S., & Adamic, L. A. (2015). Political science. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–2.

Ballew, C. C., & Todorov, A. (2007). Predicting political elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences*, 104(46), 17948-17953.

Bennett, W. L. (2012). The Personalization of Politics. *The ANNALS of the American Academy of Political and Social Science*, 644(1), 20–39

Bhatia, S. (in press). The semantic representation of prejudice and stereotypes. *Cognition*.

Bhatia, S. (2017). Associative judgment and vector space semantics. *Psychological Review*, 124(1), 1-20.

Brambilla, M., Sacchi, S., Rusconi, P., Cherubini, P., & Yzerbyt, V. Y. (2012). You want to give a good impression? Be honest! Moral traits dominate group impression formation. *British Journal of Social Psychology*, 51(1), 149-166.

Brambilla, M., Rusconi, P., Sacchi, S., & Cherubini, P. (2011). Looking for honesty: The primary role of morality (vs. sociability and competence) in information gathering. *European Journal of Social Psychology*, 41(2), 135-143.

Bullinaria, J. A., & Levy, J. P. (2007). Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior Research Methods*, 39, 510–526.

Cislak, A., & Wojciszke, B. (2006). The role of self-interest and competence in attitudes toward politicians. *Polish Psychological Bulletin*, 4(37), 203-212.

Cuddy, A. J., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in Experimental Social Psychology*, 40, 61-149.

Dehghani, M., Sagae, K., Sachdeva, S., & Gratch, J. (2014). Analyzing political rhetoric in conservative and liberal weblogs related to the construction of the "Ground Zero Mosque". *Journal of Information Technology & Politics*, 11(1), 1-14.

DellaVigna, S., & Kaplan, E. (2007). The Fox News Effect: Media Bias and Voting. The *Quarterly Journal of Economics*, 1187-1234.

Dehghani, M., Johnson, K., Hoover, J., Sagi, E., Garten, J., Parmar, N. J., ... & Graham, J. (2016). Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3), 366.

Dhillon, P., Foster, D. P., & Ungar, L. H. (2011). Multi-view learning of word embeddings via CCA. In *Advances in Neural Information Processing Systems* (pp. 199-207).

Eagly, A. H., & Karau, S. J. (2002). Role congruity theory of prejudice toward female leaders. *Psychological Review*, 109(3), 573–98.

Firth, J. R. (1957). *Papers in Linguistics*. London, England: Oxford University Press.

Fiske, S. T., Cuddy, A. J., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77-83.

Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, 82(6), 878.

Funk, C. L. (1997). Implications of political expertise in candidate trait evaluations. *Political Research Quarterly*, 50(3), 675-697.

Gerber, A. S., Karlan, D., & Bergan, D. (2009). Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics*, 1(2), 35-52.

Gervais, S. J., & Hillard, A. L. (2011). A Role Congruity Perspective on Prejudice Toward Hillary Clinton and Sarah Palin. *Analyses of Social Issues and Public Policy*, 11(1), 221–240.

Goodwin, G. P., Piazza, J., & Rozin, P. (2014). Moral character predominates in person perception and evaluation. *Journal of Personality and Social Psychology*, 106(1), 148-168.

Griffiths, T. L. (2015). Manifesto for a new (computational) cognitive revolution. *Cognition*, 135, 21-23.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114(2), 211.

Harris, Z. S. (1954). Distributional structure. *Word*, 2, 146–62.

Holtzman, N. S., Schott, J. P., Jones, M. N., Balota, D. A., & Yarkoni, T. (2011). Exploring media bias with semantic analysis tools: Validation of the Contrast Analysis of Semantic Similarity (CASS). *Behavior Research Methods*, 43(1), 193-200.

Hopmann, D. N., Vliegenthart, R., De Vreese, C., & Albæk, E. (2010). Effects of Election News Coverage: How Visibility and Tone Influence Party Choice. *Political Communication*, 27(4), 389–405.

Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114(1), 1-37.

Jones, M. N. (2017). *Big Data in Cognitive Science*. Psychology Press: Taylor & Francis.

Jones, M. N., Willits, J. A., & Dennis, S. (2015). Models of semantic memory. In J. R. Busemeyer & J. T. Townsend (Eds.) *Oxford Handbook of Mathematical and Computational Psychology* (pp 232-254), Oxford University Press.

Judd, C. M., James-Hawkins, L., Yzerbyt, V., & Kashima, Y. (2005). Fundamental dimensions of social judgment: understanding the relations between judgments of competence and warmth. *Journal of Personality and Social Psychology*, 89(6), 899-913.

Kahn, K. F. (1992). Does Being Male Help? An Investigation of the Effects of Candidate Gender and Campaign Coverage on Evaluations of U.S. Senate Candidates. *The Journal of Politics*, 54(2), 497–517.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211-240.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). An introduction to latent semantic analysis. *Discourse Processes*, 25(2-3), 259-284.

Landy, J. F., Piazza, J., & Goodwin, G. P. (2016). When It's Bad to Be Friendly and Smart: The Desirability of Sociability and Competence Depends on Morality. *Personality and Social Psychology Bulletin*.

Leach, C. W., Ellemers, N., & Barreto, M. (2007). Group virtue: the importance of morality (vs. competence and sociability) in the positive evaluation of in-groups. *Journal of Personality and Social Psychology*, 93(2), 234.

Lenton, A. P., Sedikides, C., & Bruder, M. (2009). A latent semantic analysis of gender stereotype-consistency and narrowness in American English. *Sex Roles*, 60(3-4), 269-278.

Margetts, H. (2017). Political behaviour and the acoustics of social media. *Nature Human Behaviour*, 1, 0086.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems* (pp. 3111-3119).

Pennington, J., Socher, R., & Manning, C. D. (2014, October). Glove: Global Vectors for Word Representation. In *Empirical Methods on Natural Language Processing* (pp. 1532-1543).

Slater, M. D. (2007). Reinforcing Spirals: The Mutual Influence of Media Selectivity and Media Effects and Their Impact on Individual Behavior and Social Identity. *Communication Theory*, 17(3), 281–303.

Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728), 1623-1626.

Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2010). Predicting elections with Twitter: What 140 characters reveal about political sentiment. *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media.*

Vreese, C. H. DE., & Sementko, H. A. (2004). News matters: Influences on the vote in the Danish 2000 euro referendum campaign. *European Journal of Political Research*, 43(5), 699–722