# Measuring Latent Political Ideal Points of Twitter Users from User Description Text Data

Patrick Wu, In Son Zeng, Yanwen Zhai

December 14, 2018

## Introduction

A classic problem in political science is estimating individuals' ideal points. A political ideal point is a value on a (usually) single dimension interpreted as "ideology" that summarizes the individual's political preferences (Clinton et al. [2004], Poole [2005]). If we assume that policies or political actions also have ideal points along this same dimension, the idea is that his or her utility for a policy or political action declines with the Euclidean distance from his or her political ideal point. Most work has focused on estimating ideal points for political actors (see, e.g., Poole [2005], Krehbiel [1998], Martin and Quinn [2002]). The goal of these methods is to go beyond binary classification of individuals as Republicans or Democrats.

Estimating political ideal points for non-political elites is more difficult. The approaches developed in the last few years have used political donations data (Bonica [2013]) or used the politicians each Twitter user follows (Barbera [2015]). This project, instead, seeks to estimate political ideal points of Twitter users using a text-based approach.

## Model/Methods

The model we employ largely resembles the models found in Barbera [2015], Hoff et al. [2002], and Clinton et al. [2004]. Its closest analogue is the Bayesian item-response theory model (Leventhal and Stone [2018]). In particular, our approach focuses on the political words that a user chooses to use in his or her self-provided biography on Twitter. This set of political words is chosen by the analyst, and should be appropriate for the dataset's political time period.

*Model* Suppose that each Twitter user is presented with a choice of mentioning or not mentioning a political keyword. Let $y_{ij} = 1$ if user $i$ mentions word $j$ and let $y_{ij} = 0$ otherwise. We can consider this a function of the squared Euclidean distance in the latent political dimension between user $i$ and word $j$: $-\gamma(\theta_i - \phi_j)^2$, where $\theta_i \in \mathbb{R}$ is the latent political ideal point of Twitter user $i$ along this latent political dimension, $\phi_j \in \mathbb{R}$ is the political ideal point of word $j$ along this same latent political dimension, and $\gamma$ is the discrimination parameter that measures how important this relationship is to estimating the ideal point (Gelman and Hill [2007]). We also let $\beta_i$ be a measure of how political an individual is on Twitter, which takes into account user $i$'s propensity to use any political words at all.

We define the probability that user $i$ uses word $j$ as a logit model:

$$p(y_{ij} = 1|\beta_i, \theta_i, \phi_j, \gamma) = \text{logit}^{-1}\left(\beta_i - \gamma(\theta_i - \phi_j)^2\right) \tag{1}$$

Then, assuming conditional independence between users, the likelihood of our model is

$$p(\mathbf{y}|\theta, \phi, \beta, \gamma) = \prod_{i=1}^{n}\prod_{j=1}^{m}(\text{logit}^{-1}(\pi_{ij}))^{y_{ij}}(1 - \text{logit}^{-1}(\pi_{ij}))^{1-y_{ij}} \tag{2}$$

where $\pi_{ij} = \beta_i - \gamma(\theta_i - \phi_j)^2$. We assume the following priors: $\beta_j \sim N(\mu_\beta, \sigma_\beta^2)$, $\theta_i \sim N(\mu_\theta, \sigma_\theta^2)$, and $\phi_j \sim N(\mu_\phi, \sigma_\phi^2)$. The full joint posterior distribution is

$$p(\theta, \phi, \beta, \gamma|y) \propto \prod_{i=1}^{n}\prod_{j=1}^{m}(\text{logit}^{-1}(\pi_{ij}))^{y_{ij}}(1 - \text{logit}^{-1}(\pi_{ij}))^{1-y_{ij}}\prod_{i=1}^{n}N(\beta_i|\mu_\beta, \sigma_\beta^2)\prod_{i=1}^{n}N(\theta_i|\mu_\theta, \sigma_\theta^2)\prod_{j=1}^{m}N(\phi_j|\mu_\phi, \sigma_\phi^2)$$

*Identification* Notice that model (1) is not identified. We can add any constant $c$ to both $\theta_i$ and $\phi_j$ without changing the probability of $y_{ij} = 1$. We could also multiply $\theta_i$ and $\phi_j$ by any non-zero constant $c$ and divide $\gamma$ by $c^2$ without changing the probability of $y_{ij} = 1$. This is often called additive and multiplicative aliasing (Gelman and Hill [2007]). Moreover, notice that this scale is also reflective invariant: changing the signs of both $\phi_j$ and $\theta_j$ will not change the squared Euclidean distance. We employ the solution Gelman and Hill offer: constrain $\beta_i$ to an informative $N(0, 1)$ prior and constrain either $\phi_j$ or $\theta_i$ to an informative $N(0, 1)$ prior. For computational reasons, we assign $\theta_i \sim N(0, 1)$. To solve the reflection invariance issue, we use our prior knowledge and choose the initial value for each word that matches its expected party association: -1 for Democratic words and 1 for Republican words. Jackman [2001] shows that this is enough to ensure global identification in most cases. Every other hyperparameter in the priors is assigned a flat prior, including $\gamma$. We can impose these informative priors because the ideal points are unitless; only the relative distance between the ideal points matter.

<u>*Why a Bayesian Approach?*</u> The justification for a Bayesian approach is twofold. First, we want to incorporate previous knowledge (such as the findings in Bonica [2013] and Barbera [2015]) of the distribution of ordinary citizens in our priors. Second, the number of parameters to estimate is very large: there is one $\beta_i$ for each user, one $\theta_i$ for each user, one $\phi_j$ for each word, and a $\gamma$ parameter. A Bayesian approach turns what is typically a very difficult problem in classical estimation to a routine application of MCMC.

## Analysis

<u>*Data*</u> Our Twitter user biographies come from some of the dissertation work of Patrick Wu. It was collected a month before the 2016 general U.S. election on November 8, 2016. For this period, we chose seven Democratic keywords and seven Republican keywords. The Democratic keywords are

Clinton, Hillary, Democrat, HillaryClinton, StrongerTogether, NeverTrump, and ImWithHer. The Republican keywords are Trump, Donald, Republican, RealDonaldTrump, MAGA, NeverHillary, and AlwaysTrump. These are all names, campaign slogans, Twitter handles, and hashtags commonly appeared on Twitter during the campaign period.

To process the data, we stemmed all words in the user biographies and performed exact matching based on the stemmed words for the 14 political keywords. We only looked at users who used at least one of the political keywords. Our final dataset contains $9,190$ user biographies.

*Analysis of Posterior Distributions* Table 1 illustrates the posterior mean, standard deviation, 95% credible interval, and median for keywords $\phi_j$, individual effects $\beta_i$ and $\theta_i$, and discrimination parameter $\gamma$.

[Table 1 goes here]

We also plot the posterior means of the political keywords.

[Figure 1 goes here]

As Table 1 and Figure 1 show, the 14 keywords fall on the expected sides of the dimensions, with AlwaysTrump being the strongest Republican keyword, while StrongerTogether was the strongest Democratic keyword. Figure 2 further confirms the validity of the ideological placement of the keywords.

[Figure 2 goes here]

Democratic words have strong positive correlations with other Democratic words, and strong negative correlations with Republican words. The opposite is found with Republican words.

The summary statistics of the posterior means of $\beta_i$ and $\theta_i$ are also included in Table 1. Since there are over 9000 $\beta_i$ and $\theta_i$ parameters to individually analyze, we plot the distribution of the posterior means of $\beta_i$ and $\theta_i$ across all users for analysis.

[Figure 3 goes here]

The distribution of the posterior means of $\beta$ shows that most users are not very political, but there is a group of users who are quite political. What is more interesting is the distribution of the posterior means of $\theta_i$ across users: it seems like the users who lean Democrat are more unified, while there are two subgroups, one more extreme than the other, of Republicans. This was apparent in the 2016 election: Democrats were more unified, while Republicans tended to be more divided on Donald Trump.

*Model Diagnostics* The traceplots for all $\phi_j$ parameters and the $\gamma$ parameter present evidence that the chain converges. Figure 4 shows the traceplot of $\phi_{14}$, which is very characteristic of all traceplots.

[Figure 4 goes here]

Using the Geweke diagnostic test, we see that the Geweke statistics for the 14 keywords $\phi_j$ range from -1.661 (for $\phi_{13}$) to 1.160 (for $\phi_2$), the absolute values of which are all less than 2, indicating convergence. Among all 18,396 parameters estimated, less than 5% had an absolute Geweke statistic greater than 2.

We confirmed this with the Gelman-Rubin statistic, using 5 chains. All Gelman-Rubin statistics were between 0.999 and 1.022 for all 18,396 parameters, indicating good convergence.

Lastly, to check our prior specification, we looked at the $p_D$ (effective number of parameters), which is 15,830.61. This quantity is less than the 18,396 parameters estimated, meaning that our prior specification is reasonable.

*Validation* The credible intervals for each individual's ideal points are quite wide; we suspect that the wide credible intervals are due to the relatively small number of observations rather than incorrect or biased estimations. To validate this, we collected each user's number of retweets of Democratic-associated accounts and Republican-associated accounts. If $\theta$ is valid, then an increase in $\theta$ should yield an increase in the retweets of Republican-associated accounts and a decrease in the retweets of Democratic-associated accounts. To study this pattern, we run two (non-Bayesian) zero-inflated negative binomial models.[1]

[Table 2 goes here]

The results show the exact hypothesized pattern. These results are replicated if we use retweets of members of Congress, retweets of the 2016 presidential candidates, and favorites of Democratic and Republican tweets.

## Discussion

Through the Bayesian model, we have obtained and validated the political ideal points of 14 political keywords and 9,190 individuals with reasonable convergence. We have also used the posterior means of $\theta_i$ to predict other political behaviors on Twitter. In the future, we hope to extend estimation of ideal points of people who do not use political keywords. Furthermore, as there are large population of Twitter users, we can reduce the computational costs of estimating ideal points by implementing a Metropolis-Hasting approach that fixes all ideal points of words.

[1]We use zero-inflated negative binomial models to account for many 0's in the number of retweets from users and for overdispersion concerns

# References


Pablo Barbera. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political Analysis*, 23:76–91, 2015.

Adam Bonica. Ideology and interests in the political marketplace. *American Journal of Political Science*, 57, 2013.

Joshua Clinton, Simon Jackman, and Douglas Rivers. The statistical analysis of roll call data. *American Political Science Review*, 2:355–370, 2004.

Andrew Gelman and Jennifer Hill. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, 2007.

Peter D. Hoff, Adrian E. Raftery, and Mark S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97:1090–1098, 2002.

Simon Jackman. Multidimensional analysis of roll call data via bayesian simulation: Identification, estimation, inference, and model checking. *Political Analysis*, 9:227–241, 2001.

Keith Krehbiel. *Pivotal Politics*. The University of Chicago Press Books, 1998.

Brian C. Leventhal and Clement A. Stone. Bayesian analysis of multidimensional item response theory models: A discussion and illustration of three response style models. *Measurement: Interdisciplinary Research and Perspectives*, 16:114–128, 2018.

Andrew D. Martin and Kevin M. Quinn. Dynamic ideal point estimation via markov chain monte carlo for the u.s. supreme court, 1953-1999. *Political Analysis*, 10:134–153, 2002.

Keith Poole. *Spatial Models of Parliamentary Voting*. Cambridge University Press, 2005.

## Figures and Tables

Table 1: Summary Statistics of the Posterior Distributions of $\phi_j$ (Ideal Point of Political Keywords), $\beta_i$ (Politicalness of Users), $\theta_i$ (Ideal Point of Users), and $\gamma$ (Discrimination parameter)

| **Keyword** | mean | 0.01 | **95% cred. interval** | median | 0.99 | SD |
|---:|---:|---:|---:|---:|---:|---:|
| Trump | 2.199 | 2.068 | (2.085,2.318) | 2.198 | 2.340 | 0.060 |
| Clinton | -4.594 | -4.832 | (-4.791,-4.415) | -4.592 | -4.384 | 0.096 |
| Donald | 4.686 | 4.477 | (4.508,4.875) | 4.683 | 4.914 | 0.096 |
| Hillary | -3.757 | -3.947 | (-3.920,-3.600) | -3.756 | -3.578 | 0.082 |
| Republican | 4.017 | 3.836 | (3.859,4.186) | 4.016 | 4.221 | 0.084 |
| Democrat | -3.097 | -3.274 | (-3.246,-2.957) | -3.096 | -2.933 | 0.073 |
| RealDonaldTrump | 5.411 | 5.167 | (5.207, 5.628) | 5.410 | 5.675 | 0.110 |
| HillaryClinton | -4.874 | -5.114 | (-5.074,-4.683) | -4.872 | -4.652 | 0.100 |
| MAGA | 3.831 | 3.651 | (3.679,3.993) | 3.829 | 4.021 | 0.080 |
| StrongerTogether | -5.200 | -5.454 | (-5.414,-4.993) | -5.197 | -4.961 | 0.108 |
| NeverHillary | 4.389 | 4.189 | (4.216, 4.574) | 4.387 | 4.604 | 0.090 |
| NeverTrump | -4.145 | -4.356 | (-4.323,-3.981) | -4.143 | -3.951 | 0.088 |
| AlwaysTrump | 6.476 | 6.141 | (6.189,6.789) | 6.473 | 6.857 | 0.151 |
| ImWithHer | -3.166 | -3.338 | (-3.311,-3.030) | -3.164 | -3.005 | 0.072 |
| **Parameter** | | | | | | |
| $\beta_{avg}$ | 0.01 | -1.80 | (-1.50,1.40) | 0.03 | 1.64 | 0.74 |
| $\theta_{avg}$ | 0.00 | -1.64 | (-1.37,1.36) | 0.00 | 1.63 | 0.70 |
| $\gamma$ | 0.18 | 0.17 | (0.17,0.20) | 0.18 | 0.20 | 0.01 |

Note that $\beta_{avg}$ and $\theta_{avg}$ are the average of the posterior means of $\beta_i$ and $\theta_i$ across all users. There are 9,190 $\beta_i$ parameters and 9,190 $\theta_i$ parameters, which is too much to list out the summary statistics for each posterior distribution.

Table 2: Validation Using Zero-Inflated Negative Binomial Models: Number of Retweets of Democratic-Associated Accounts and Number of Retweets of Republican-Associated Accounts vs. $\theta$

| **Count** | *Dem. Accts* | *Rep. Accts* |
|---|---|---|
| (Intercept) | 3.12 (0.02) | 4.49 (0.02) |
| $\theta$ | -1.22 (0.03) | 1.57 (0.03) |
| log(theta) | -0.65 (0.02) | -1.00 (0.02) |
| **Zero-Inflated** | | |
| (Intercept) | 4.48 (0.17) | 5.18 (0.20) |
| $\theta$ | 2.25 (0.08) | -2.03 (0.12) |
| log(1 + RT Count) | -0.91 (0.03) | -1.27 (0.04) |
| log likelihood | -27870 | -38090 |

NOTE: log(theta) denotes the overdispersion parameter

Figure 1: Posterior Distribution of $\phi_j$: $\phi_j$ show the ideal point for keywords on the political continuum. The deeper colour represents the greater magnitude of the corresponding keyword.

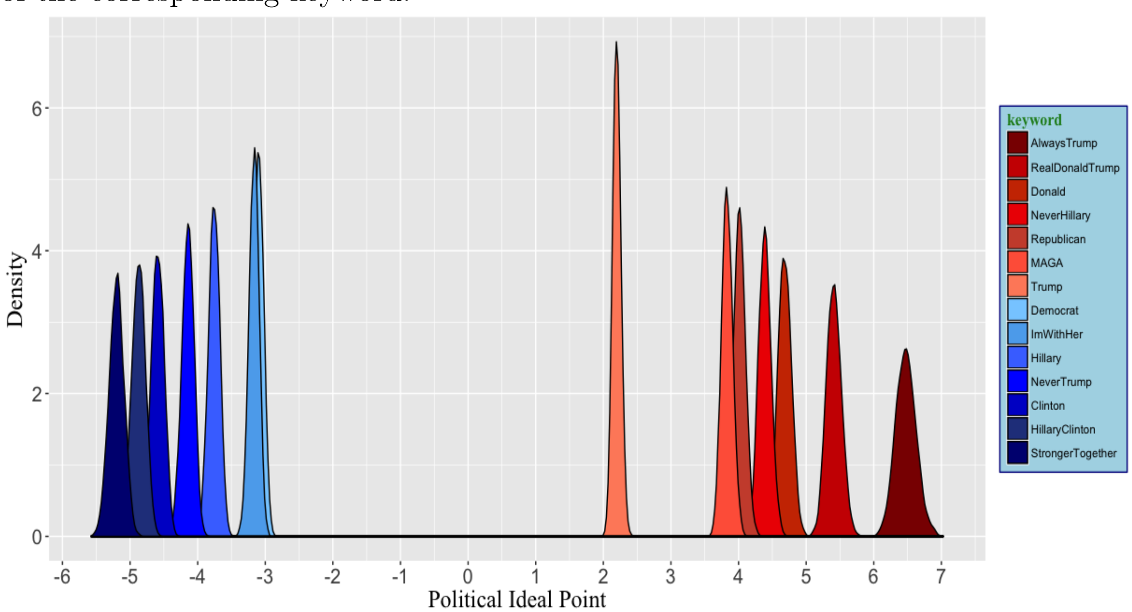


**Left:** Democratic (blue) and **Right:** Republican (red) party affiliated words.

Figure 2: Correlation Plot of $\phi_j$ for Every Keyword: The correlation of $\phi_j$ appears to be positive if the two keywords are affiliated with the same party in prediction. In contrast, the correlation becomes negative if the two keywords belong to the opposing parties. The correlation fluctuates around 0.65 to 0.8 in absolute sense, indicating a fairly strong correlation between keywords.

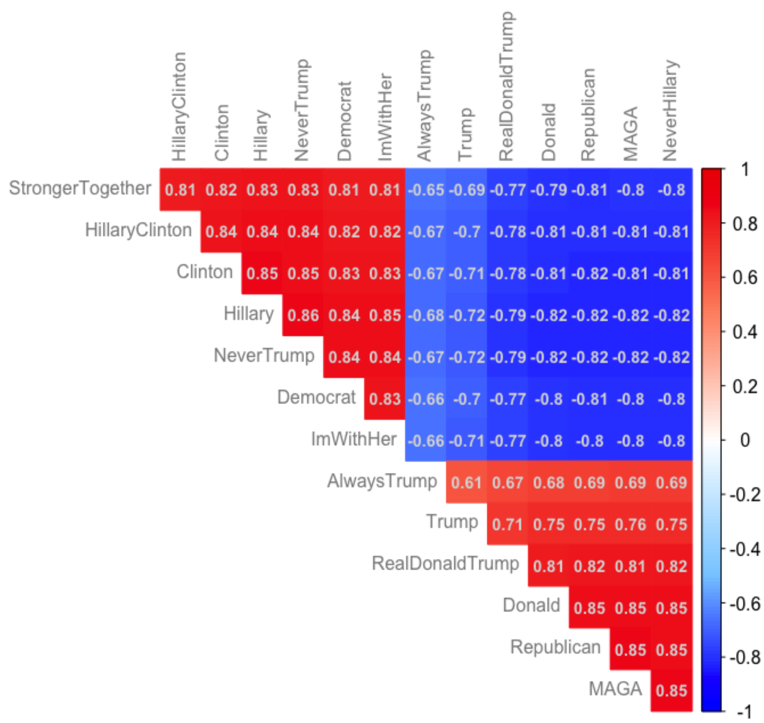

Figure 3: Distribution of Posterior Mean of $\beta_i$ and $\theta_i$ for Individual Effects: The distribution of posterior mean of $\beta$ shows a big cluster at negative side and a small cluster in positive side, indicating that most Twitter users are not politically inclined to use the keywords, while a few of them are quite political. The distribution of posterior mean of $\theta_i$ demonstrates a cluster on negative side and two clusters on positive side. One interpretation for this phenomenon is that for the supporters for Democratic side (left) are quite unified, while the supporters for Republican side (right) are more divided.

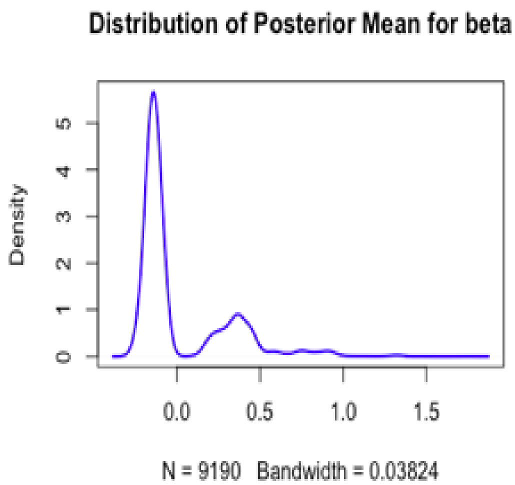


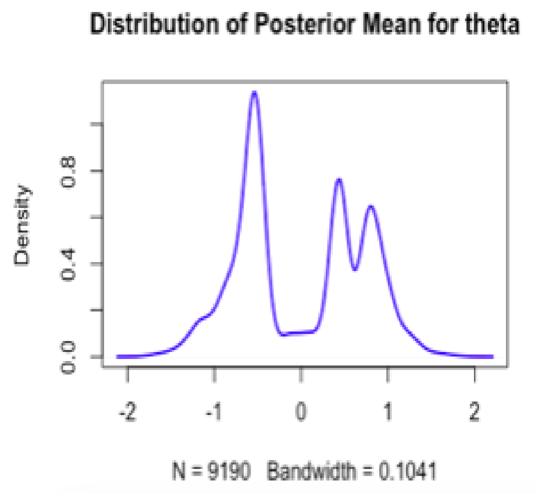


Figure 4: Traceplot for $\phi_{14}$, which captures the ideal point of political keyword ImWithHer. This traceplot presents evidence of convergence.