

Pervasive Sensing to Model Political Opinions in Face-to-Face Networks

Anmol Madan¹, Katayoun Farrahi^{2,3}, Daniel Gatica-Perez^{2,3}, and Alex (Sandy) Pentland¹

¹ MIT Media Laboratory, Massachusetts Institute of Technology, Cambridge, USA

² IDIAP Research Institute, Martigny, Switzerland

³ Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Abstract. Exposure and adoption of opinions in social networks are important questions in education, business, and government. We describe a novel application of pervasive computing based on using mobile phone sensors to measure and model the face-to-face interactions and subsequent opinion changes amongst undergraduates, during the 2008 US presidential election campaign. We find that self-reported political discussants have characteristic interaction patterns and can be predicted from sensor data. Mobile features can be used to estimate unique individual exposure to different opinions, and help discover surprising patterns of dynamic homophily related to external political events, such as election debates and election day. To our knowledge, this is the first time such dynamic homophily effects have been measured. Automatically estimated exposure explains individual opinions on election day. Finally, we report statistically significant differences in the daily activities of individuals that change political opinions versus those that do not, by modeling and discovering dominant activities using topic models. We find people who decrease their interest in politics are routinely exposed (face-to-face) to friends with little or no interest in politics.

1 Introduction

A central question for social science, as well as for the practical arts of education, sales, and politics, is the mechanism whereby ideas, opinions, innovations and recommendations spread through society. Diffusion is the phenomena of propagation of ideas or opinions within a social network. On the internet, the proliferation of social web applications has generated copious amounts of data about how people behave and interact with each other in online communities, and these data, are being extensively used to understand online diffusion phenomena. However, many important attributes of our lives are expressed primarily in real-world, face-to-face interactions. To model the adoption of these behaviors, we need fine-grained data about face-to-face interactions between people, i.e. who talks to whom, when, where, and how often, as well as data about exogenous variables that may affect the adoption process. Such social sensing of face-to-face interactions that explain social diffusion phenomena is a promising new area for pervasive computing.

Traditionally, social scientists have relied on self-report data to study social networks, but such approaches are not scalable. It is impossible to use these methods with fine resolution, over long timescales (e.g. months or years), or for a large number of people, (e.g. hundreds or thousands). Further, while people may be reasonably accurate in their reports of long term social interaction patterns, it is clear that memory regarding particular relational episodes is quite poor. In a survey of informant accuracy literature, Bernard et.al. have shown that recall of social interactions in surveys used by social scientists is typically 30-50 % inaccurate [5, 8].

A key question is how mobile sensing techniques and machine perception methods can help better model these social diffusion phenomena. This paper describes the use of mobile phone sensors at an undergraduate community to measure and model physical proximity (via bluetooth sensors), phone communication, movement patterns (via 802.11 WLAN access-points) and self-reported political opinions. Our approach provides insight about the adoption of political opinions in this community.

The contributions of this paper are as follows.

1. We devise a mobile sensing platform and observational methodology to capture social interactions and dependent variables during the last three months of the 2008 US Presidential campaigns of John McCain and President Barack Obama, amongst the residents of an undergraduate residence hall at a North American university. This dataset, first of its kind to our knowledge, consists of 132,000 hours of social interactions data and the dependent political opinions measured using monthly surveys.
2. We estimate exposure to diverse political opinions for individual residents, and propose a measure of dynamic homophily that reveals patterns at the community scale, related to external political events.
3. Pervasive-sensing based social exposure features explain individual political opinions on election day, better than self-reported social ties. We also show that ‘political discussant’ ties have characteristic interaction patterns, which can be used to recover such ties in the network.
4. Using an LDA-based topic modeling approach, we study the behavior differences between individuals who change opinions, and those who held their political opinions. We show statistically significant differences in the activities of people who changed their preferred party versus those that did not. People that changed preferred party often discuss face-to-face with their democrat political discussants, and their daily routines included heavy phone and SMS activity. We also find people that decrease their interest in politics often interact with people that have little or no interest in politics

2 Related Work

Sensing Human Behavior using Mobile Devices

There has been extensive work to model various aspects of human behavior, using smartphones [16, 36, 1, 22, 15, 4, 17], wearable sensor badges [44, 30, 41],

video [46], and web-based social media data [7, 48, 52, 33, 34, 3, 39]. Choudhury et. al., used electronic sensor badges to detect social network structure and model turn-taking behavior in face-to-face networks [10, 12]. Eagle et. al. used mobile phones as sensors to characterize social ties for a student community [16]. At larger scales, mobile location and calling data have been used to characterize temporal and spatial regularity in human mobility patterns [22], and individual communication diversity has been used to explain the economic development of cities [15]. Other examples of the use of mobile phones to map human interaction networks include the CENS participatory sensing project at UCLA [1], and the mHealth and Darwin projects at Dartmouth [4, 17]. Electronic sensor badges instrumented with infrared (IR) sensors to capture the direction of face-to-face proximity, have been used by Olguin, Waber and Kim [42, 41] to enhance organizational performance and productivity, for financial institutions and consultants. Vocal analysis has been used to capture nonlinguistic communication and social signaling in different contexts [11, 35, 44, 42].

On the modeling front, Exponential Random Graph Models (ERGMs) and its extensions [45, 50, 24, 26] have been used to understand interaction networks. Topic models have been explored for activity modeling applications [18, 19, 27]. New topic models have also been proposed in the context of blog post response prediction [51] and blog influence [40]. Other types of topic models, like Author-Topic [47, 37] and Group-Topic [49] models have been used for social network analysis when content or contextual data is present. In this paper, we use LDA topic model [6] to better understand the behavior differences between people who changed their opinions.

Adoption of Political Opinions

In political science and sociology, an important area of study is how opinions about political candidates and parties, and voting behavior, spread through different interaction networks. Political scientists have proposed two competing models of social influence and contagion [9]. The social cohesion model suggests that influence is proportional to tie strength, while the structural equivalence model [21] proposes that influences exist across individuals with similar roles and positions in networks. Huckfeldt and Sprague [25] studied the interdependence of an individual's political opinions, their political discussant network and context and demographics during the 1984 presidential elections. They found a social dissonance effect in the propagation of political opinions, and also report an 'inverse U' relationship with tie-strength, i.e. discussant effects are stronger for less intimate relationships like acquaintances and frequent contacts than they are for close friends.

In the online context, Adamic and Glance [2] studied political blogs during the 2004 presidential elections, and found that content, discussions and news items on liberal and conservative blogs, connected primarily to separate clusters, with very few cross-links between the two major clusters. Leskovec et. al. [33] tracked the propagation of short, distinctive political phrases during the 2008

elections, and model the news cycle across both mainstream news sources and political blogs.

3 Methodology

In the past, researchers have used Call Data Records (CDRs) provided by mobile operators to better understand human behavior [22, 15]. Our approach, however, is to use pervasive sensing methods for capturing social interactions, and this has several advantages. Firstly, it allows us to sample different sensors and dependent training labels, and not just calling data alone. Secondly, from a privacy perspective, this requires the user’s explicit participation in data collection. Additionally, in the future, it could be used to provide the user immediate feedback on the mobile device itself.

3.1 Privacy Considerations

An important concern with long-term user data collection is securing personal privacy for the participants. This study was approved by the Institutional Review Board (IRB). As financial compensation for completing monthly surveys and using data-collection devices as their primary phones, participants were allowed to keep the devices at the end of the study. The sensing scripts used in the platform capture only hashed identifiers, and collected data is secured and anonymized before being used for aggregate analysis.

3.2 Mobile Sensing Platform

Given the above goals, the mobile phone based platform for data-collection was designed with the following long-term continuous sensing capabilities, using Windows Mobile 6.x devices. Daily captured mobile sensing data was stored on-device on read/write SD Card memory. On the server side, these logs files were merged, parsed and synced by an extensive Python post-processing infrastructure, and stored in MySQL for analysis. This sensing software platform for Windows Mobile 6.x has been released under the LGPLv3 open source license for public use [28].

Proximity Detection (Bluetooth) The software scanned for Bluetooth wireless devices in proximity every 6 minutes (a compromise between sensing short-term social interactions and battery life, [16]). The Windows Mobile phones used in our experiment were equipped with class 2 Bluetooth radio transceivers, with practical indoor sensing range of approximately 10 feet. Scan results for two devices in proximity have a high likelihood of being asymmetric, which is accounted for in our analysis. Due to API limitations of Windows Mobile 6.x, signal strength was not available during scans.

Approximate Location (802.11 WLAN) The software scanned for wireless WLAN 802.11 Access Point identifiers (hereafter referred to as WLAN APs) every 6 minutes. WLAN APs have an indoor range of approximately 125 feet

and the university campus had almost complete wireless coverage. Across various locations within the undergraduate residence, over 55 different WLAN APs with varying signal strengths can be detected.

Communication (Call and SMS Records) The software logged Call and SMS details on the device every 20 minutes, including information about missed calls and calls not completed.

Battery Impact The battery life impact of periodic scanning has been previously discussed [16]. In this study, periodic scanning of Bluetooth and WLAN APs reduced operational battery life by 10-15%, with average usable life between 14-24 hours (varying with handset models and individual usage). Windows Mobile 6.x devices have relatively poorer battery performance than other smartphones, and WLAN usage (web browsing by user) had a bigger impact on battery life than periodic scanning.

3.3 Dataset Characteristics

The mobile phone interaction dataset was collected from 67 participants and consisted of approximately 450,000 bluetooth proximity scans, 1.2 million WLAN access-point scans, 16,900 phone call records and 17,800 SMS text message events. The average duration of phone calls is approx 138 seconds, and 58 percent of phone calls were during weekdays.

3.4 Political Opinions (Dependent Variables)

The dependent political opinions were captured using three monthly web-based surveys, once each in September, October, and November 2008 (immediately following the presidential election). The monthly survey instrument was based on established political science literature, and consisted of questions shown in Table 1. The questions were identical to the survey instrument used by Lazer and Rubineau [32], who measured the monthly political opinions of students across different universities (during the same 2008 election period) and studied the co-evolution of political opinions and self-report friendship networks.

Political scientists have established that shifts in political opinions are gradual [25]. This is observed in our dataset, as approximately 30% of the 67 participants changed their opinions for each of the dependent questions during the three month observation period. Opinion changes were along 1-point or 2-points on the respective 4/7-point Likert scales. Similar variations in our dependent variables were also reported in the analysis of Lazer and Rubineau [32].

For each monthly survey, participants also identified other residents that were political discussants, close friends or social acquaintances, identical to those used here [32]. Baseline information including race, ethnicity, political opinions of the person's parents and religious affiliations was also collected from some of the participants before the start of the experiment, but is not used in our analysis.

Table 1. Political Survey Instrument used to capture different political opinions. All responses were constructed as Likert scales.

Survey Question	Possible Responses
Are you liberal or conservative?	7-point Likert scale Extremely conservative to extremely liberal
How interested are you in politics	4-point Likert scale Not interested to very interested
What is your political party preference?	7-point Likert scale Strong Democrat to strong Republican
Which candidate are you likely to vote for? (Sept and Oct)	Choice between leading Republican and Democrat nominees
Which candidate did you vote for? (Nov)	Choice between B. Obama and J. McCain
Are you going to vote in the upcoming election? (Sept and Oct)	4-point Likert scale
Did you vote in the election? (Nov)	Yes or No

4 Analysis

4.1 Individual Exposure to Diverse Opinions

What is an individual’s social exposure to diverse ideas and opinions? Threshold and cascade models of diffusion [23, 29] assume that all individuals in a population have a uniform exposure, or that the underlying distribution of exposure to different opinions is known. While exposure to different opinions is dynamic and characteristic for every individual, it has previously not been incorporated into empirical social diffusion models. Dynamic exposure to different opinions can be estimated for each participant, on a daily or hourly basis. Contact between two individuals can be given as a function of physical proximity counts (blue-tooth), phone call and SMS counts, total duration of proximity, total duration of phone conversation, or other measures of tie-strength. These features represent the time spent with others having different opinions in classes, at home, and in phone communication.

Normalized exposure, N_i represents the average of all opinions a person is exposed to on a daily basis, weighted by the amount of exposure to different individuals and their self-reported opinions, where O_j represents the opinion response for person j for a particular question in Table 1, $contact_{ij}$ is the blue-tooth proximity counts between i and j (tie-strength), and $Nbr(i)$ is the set of neighbors for i in the interaction network.

$$N_i = \sum_{j \in Nbr(i)} contact_{ij} \cdot O_j / \sum_j contact_{ij} \quad (1)$$

Cumulative exposure, C_i to a particular political opinion O , represents the magnitude of a particular opinion that a person is exposed to on a daily basis, and is a function of the amount of contact with different individuals and their self-reported opinion. $contact_{ij}$ can be estimated from other mobile interaction features, like counts for calling, SMS, and 802.11 WLAN co-location. In Section 4.4, N_i is used for future opinion prediction and in Section 5, C_{iO} from both bluetooth and call features are used for change of opinion modeling.

$$C_{iO} = \delta_j \cdot \sum_{j \in Nbr(i)} contact_{ij} \quad (2)$$

$\delta_j = 1$ only if person j holds opinion O , and 0 otherwise. Figure 1 shows cumulative and normalized exposure for one participant during the election campaign period. This individual did not have much exposure to republicans, though was often in proximity to democrats, some days much more than others.

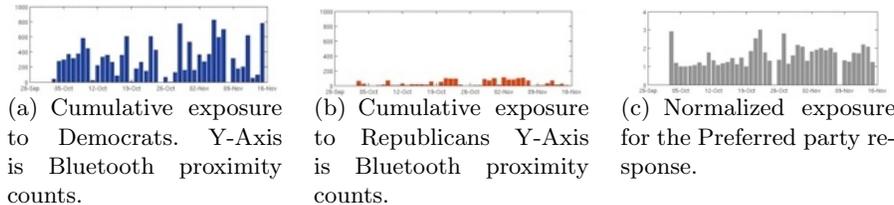


Fig. 1. Characteristic daily normalized and cumulative exposure for one resident during the election period (Oct-Nov 2008). Contact is Bluetooth physical proximity. X-Axis is days for all graphs. This individual had much more exposure to democratic opinions than republican during this period. Incidentally, this person did not show an opinion shift for the preferred-party response during the study (not shown).

4.2 Pervasive Reflection of Dynamic Homophily

Homophily, or the idea of “birds of a feather flock together”, [31] is a fundamental and pervasive phenomenon in social networks, and refers to the tendency of individuals to form relationships with others that have similar attributes, behaviors or opinions. McPherson and Smith [38] provide an in depth review of homophily literature. The emergence of homophily during network formation has been explained using agent based models, and in economics [14] by incorporating chance, choice, and tie formation costs. In this section we define a measure of dynamic homophily based on mobile phone interaction features.

In sociological literature, homophily is estimated using the homophily index H_i , and Coleman’s inbreeding homophily index, IH_i , which are a function of the relative fraction of social ties expressed between people who hold similar opinions, and those that hold different opinions. These homophily indices are explained in more detail in [14, 38]. These sociological measures of homophily are useful for static networks, but do not capture the dynamics of the underlying phenomena. To overcome these limitations, we propose a measure of dynamic homophily based on social exposure features for the daily timescale, given as,

$$\Delta_i(t) = \left| O_i - \sum_{j \in Nbr(i)} contact_{ij} \cdot O_j / \sum_j contact_{ij} \right| \quad (3)$$

$$H(t) = \sum_i \Delta_i(t)/n \quad (4)$$

where $\Delta_i(t)$ is the difference between the individual’s opinions and the opinions he/she is exposed to, $H(t)$ is a daily measure of dynamic homophily for the entire community, and O_i are an individuals political opinion responses, on the full-range of the 4 or 7-point scale, i.e., for the political interest response, O_i ranges from 1 (“Very Interested”) to 4 (“Not at all interested”) and for the preferred party response, O_i ranges from 1 (“Strong Democrat”) to 7 (“Strong Republican”). Unlike the static homophily measures above, O_i is not based on the redistributed classes.

Daily variations in $H(t)$ are due to changes in mobile phone interaction features, that capture how participants interact with others. A negative slope in $H(t)$ implies that residents have more social exposure to individuals sharing similar opinions, in comparison to the previous day or week. Similarly, an upward slope implies that residents have decreasing social exposure with others having similar opinions.

This daily measure captures dynamic homophily variations during the election period, not captured using existing static measures of homophily. For a few days around the election day and final debates, participants show a higher tendency overall to interact with like-minded individuals. Statistical validation of these variations using repeated-measures ANOVA for different political opinions for three relevant conditions (periods) are given in Table 2 (plots in Figure 2).

Table 2. Statistically significant variations in Dynamic Homophily around the final election debate period (15th Oct 2008) and election day (4th Nov 2008) period. Dynamic homophily is calculated using bluetooth proximity (phone calling and SMS are not significant for any self-reported political opinions). For each period, the average dynamic homophily for the 5-day period per participant was estimated. This analysis was first done for all participants, and then repeated for freshmen-only, who had only been in the community for a month before start of the study, and where stronger effects are observed. The three experimental conditions (periods) chosen for validating the main effect were (a) Baseline Period (1st condition), i.e., 4th October to 10th October 2008 (b) Final election debate Period (2nd condition), 12th October to 18th October 2008 and (c) Election period (3rd condition): 1st November to 7th November 2008.

Opinions Evaluated for main effects over three periods (conditions)	Result Summary (plots in Figure 2)
Political Interest for all participants	Significant effect, higher tendency to interact with like-minded individuals during debate and final election period as compared to baseline period, $F - value = 8.49, p < 0.0004$
Political Interest for freshmen only	Significant effect, higher tendency to interact with like-minded individuals during debate and final election period as compared to baseline period, $F - value = 3.43, p = 0.04$
Party preference for all participants	Not a significant effect, $F - value = 0.87, p < 0.42$
Liberal-conservative tendency for all participants	Significant effect, higher tendency to interact with like-minded individuals during debate and final election period as compared to baseline period, $F - value = 6.26, p < 0.003$

Figure 2 (a) shows $H(t)$ for political interest for all participants, where daily network structure is estimated on the basis of Bluetooth proximity counts. The first dip in this graph corresponds to the period of the final election debate during the campaign, 14th Oct 2008. The difference between the three conditions is statistically significant (F -value = 8.49, $p < 0.0004$). Figure 2(b) and (c) show similar dips for the preferred-party and liberal-conservative responses. Figure 2(d) shows $H(t)$ for political interest only for freshmen, based on daily bluetooth proximity networks. The dynamic homophily effects for freshmen, who only had a month to form ties in this community at this point, are visually pronounced, and a second dip is seen around 4th November 2008 (Election day, F -value = 3.43, $p = 0.04$). We find that these behavior changes related to external events are seen in bluetooth proximity data, but not in calling and SMS interactions. This suggests that exposure to different opinions based on physical proximity plays a more important role than exposure to opinions via phone communication. Similar results are also observed for the preferred party responses and liberal-conservative responses with respect to phone calling patterns.

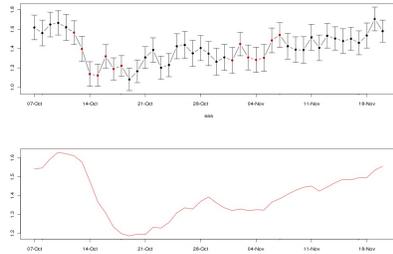
4.3 Inferring Political Discussants

What are the behavioral patterns of political discussants? In monthly self-reported survey responses, only 39.6% of political discussants are also close friends. Similarly, it is found that having similar political opinions does not increase the likelihood that two individuals will be political discussants in this dataset.

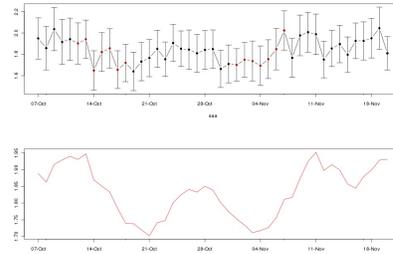
While these political discussants do not fit the mould of ‘close friends’ or individuals with similar political opinions. we find that it is possible to identify political discussants from their interaction patterns. Classification results based on mobile phone interaction features – total communication; weekend/late-night communication; total proximity; and late-night/weekend proximity, that characterize a political discussant are shown in Table 3. Two different approaches are used for comparison, an AdaboostM1 based classifier [20] and a Bayesian network classifier [13], where each input sample represents a possible tie, and both show similar results. Cost-sensitive approaches are used in both cases, to account for unbalanced classes. Political discussants are treated as unidirectional ties. Precision and recall of the discussant class are similar if self-reported the training labels are converted to bi-directional ties.

4.4 Exposure and Future Opinions

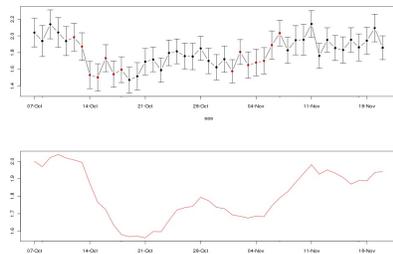
Exposure based features described in the previous section can be used as a feature to train a linear predictor of future opinions. The coefficients used in a linear model of opinion change include normalized exposure during the period, the persons opinion at the start of the study (September 2008), and a constant term that represents a linearly increasing amount of media influence as we get closer to the election date (Nov. 2008). For the various political opinion questions, regression values are in the $R^2 = 0.8, p < 0.01$ region. Using exposure



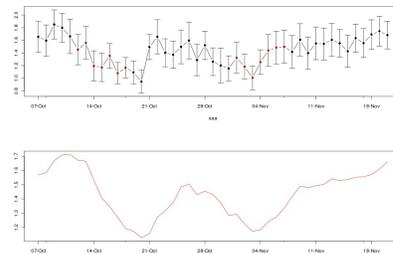
(a) Dynamic homophily of **political interest responses** (using bluetooth proximity) for all participants. Notice the decline, i.e. tendency to interact with others having similar opinions, lasting for a few days, around Oct 15th 2008, which was the last presidential debate.



(b) Dynamic homophily of **preferred party responses** (using bluetooth proximity) for all participants.



(c) Dynamic homophily of **liberal-conservative responses** (using bluetooth proximity) for all participants.



(d) Dynamic homophily of **political interest responses** (using bluetooth proximity) only for **Freshmen**. There are two periods of decline, each lasting for a few days. The first is around Oct 15th (last presidential debate) and the second is around 4th Nov, Election Day.

Fig. 2. Top: actual values of $H(t)$ with standard error bars. Bottom: Moving average.

based features explains an additional 15% - 30% variance across different political opinion questions. The effects for freshmen are approximately twice as strong as compared to the entire population, similar to the variations in dynamic homophily related to external events. In the context of social science literature, this is a relevant effect.

Table 3. Identifying Political discussants based on exposure features. Classification results using Meta-cost AdaboostM1 (individual classifiers are decision stumps), 5-fold cross validation

Class	Precision	Recall	F-Measure
Non-discussants	0.87	0.62	0.72
Political discussants	0.35	0.67	0.46

Table 4. Identifying Political discussants based on exposure features. Classification results using cost-sensitive Bayesian Network classifier, 5-fold cross validation and K2 hill-climbing structure learning

Class	Precision	Recall	F-Measure
Non-discussants	0.87	0.61	0.72
Political discussants	0.35	0.70	0.46

Table 5. Least squares regression results for the opinion change model. The dependent variable in all cases is the self-reported political opinion in November. The independent regression variables are averaged opinion of self-reported close friends relationships and political discussants (I), normalized bluetooth exposure (II), and normalized exposure combined with past opinion (III). As seen, automatically captured mobile phone features substantially outperform self-reported close friends or political discussants.

	I	II	III
Political Opinion Type	Self-reported Disc. / Close Friends	Normalized Exp. Only	Normalized Exp. & Sept Opinion
Preferred Party	n.s. / n.s.	0.21**	0.78***
Liberal or Conservative	n.s. / n.s.	0.16*	0.81***
Interest in Politics	n.s. / 0.07*	0.24**	0.74***
Preferred Party (freshmen only)	n.s. / n.s.	0.46*	0.83*
Interest in Politics (freshmen only)	n.s. / n.s.	0.21**	0.78***

All values are R^2 n.s.: not significant
 * : $p < 0.05$ ** : $p < 0.01$ *** : $p < 0.001$

5 Modeling Opinion Change with Topic Models

An important question in sociology is ‘what influences opinion change’? Is there an underlying mechanism resulting in the opinion change for some people? Can we measure this mechanism, and if so, can we predict future opinion changes from observed behavior? In this section, we propose a method for activity modeling based on the Latent Dirichlet Allocation (LDA) [6] topic model, to contrast the activities of participants that change opinions, with those that do not. We discover in an unsupervised manner, the dominating routines of people in the dataset, where routines are the most frequently co-occurring political opinion exposure patterns also referred to as topics.

5.1 Latent Dirichlet Allocation

Topic models can be used to discover a set of underlying (or latent) topics from which a corpus of M documents d is composed via unsupervised learning. They are generative models initially developed for textual content analysis. In LDA [6], a word w is generated from a convex combination of K topics z . The number of latent topics, K , must be chosen by the user. The probability of word w_t from a vocabulary of V words in document d_i is $p(w_t|d_i) = \sum_k p(w_t|z_k)p(z_k|d_i)$, $\sum_k p(z_k|d_i) = 1$. LDA assumes a Dirichlet prior probability distribution on $\Theta = \{\{p(z_k|d_i)\}_{k=1}^K\}_{i=1}^M$ and $\Phi = \{\{p(w_t|z_k)\}_{t=1}^V\}_{k=1}^K$ to provide a complete generative model for documents. Words are considered to be exchangeable, meaning they are independent given topics. The objective of LDA inference is to obtain (1) the distribution of words given topics Φ and (2) the distribution of topics over documents Θ .

When considering behavioral data, what we refer to as ‘multimodal exposure (MME) features’ can be seen as analogous to text words and a user is analogous to a document. Further, latent topics are analogous to human routines, where Φ gives an indication of how probable topics are for users, and Θ results in a distribution of exposure features given topics.

5.2 Multimodal Exposure (MME) Features and Topics

We formulate a multimodal vector of exposure features (MME features) encompassing four components: (1) time (2) political opinion (3) type + amount of interaction and (4) relationship. Overall, a MME feature captures the exposure to a particular political opinion, including details such as time and relationship. Given a survey question from Table 1, a MME feature has the following structure $(t, p_o, b, c, f, s, p_d)$. Component (1) is the time where $t \in \{10 \text{ pm} - 2 \text{ am (late night = LN), } 2 - 8 \text{ am (early morning = EM), } 8 \text{ am} - 5 \text{ pm (day = D), } 5 - 10 \text{ pm (evening = E)}\}$. These 4 time intervals in the day are specific to the overall daily activities of the users in the dataset. Component (2) is the political opinion $p_o \in o$ and o is the set of possible responses from Table 1 for the survey question chosen. Component (3) is the type and amount of interaction where b is a measure of the cumulative exposure (Equation 2) from bluetooth proximity to opinion p_o and c is the cumulative exposure from the mobile phone logs to opinion p_o . Cumulative exposure is quantized into the following bins: $b \in \{0, 1 - 2, 2 - 9, 9+\}$, $c \in \{0, 1 - 2, 3+\}$ to limit the vocabulary size. $b = 0$ implies no proximity interaction in the time interval t with political opinion p_o and $c = 3+$ implies 3 or more calls and/or SMS with political opinion p_o during time interval t . Finally, the relationship metric is defined by $f \in [\text{friend, not friend}]$, $s \in [\text{socialize, do not socialize}]$, and $p_d \in [\text{political discussants, not political discussants}]$. Topics are essentially clusters of dominating ‘opinion exposures’ present over all individuals and days in the real-life data collection, described in terms of MME features.

5.3 Model Selection with Statistics

In order to choose the optimal number of topics, K , for the model, we consider statistical significance measures over the entropy of topic distributions. We chose entropy of topic distributions as it (1) enables the computation of statistical significance over a vector of probability distributions and (2) summarizes the probability distributions of user behaviors.

In Figure 3, statistical significance test results are displayed for various survey questions (Table 1) (e.g. interest in politics (I)) as a function of the number of topics (x-axis) (a) for the groups 'changed opinion' versus 'did not change' (b) considering all possible opinions and change of opinions as groups. The difference in group entropies is mostly statistically significant for the preferred party opinion when considering the 2 group case in (a), however not for all values of K . In Figure 3(a), the first two points for which statistical significance occurs are at $K = 13$ and $K = 14$ and in the case of Figure 3(b) at $K = 17$. For the opinion interest in politics (I) and the 2 group case in plot (a) at $K = 22$ the p-value reaches its minimum. We consider $K = 14$ for PP (4-point scale) and $K = 22$ for I, points which are statistically significant, in analyzing opinion change in the results.

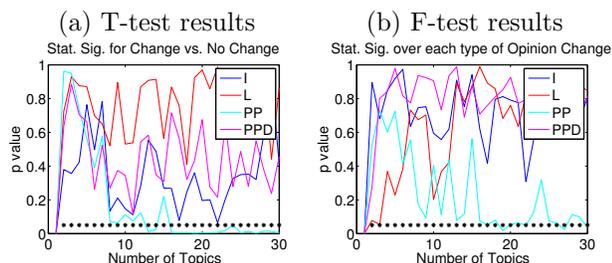


Fig. 3. Significance results (a) for 'changed opinion' versus 'did not change opinion' for interest in politics (I), liberal/conservative (L), preferred party (PP) (4-point scale) and (PPD) (7-point scale) (b) considering all possible opinions and change of opinions as groups.

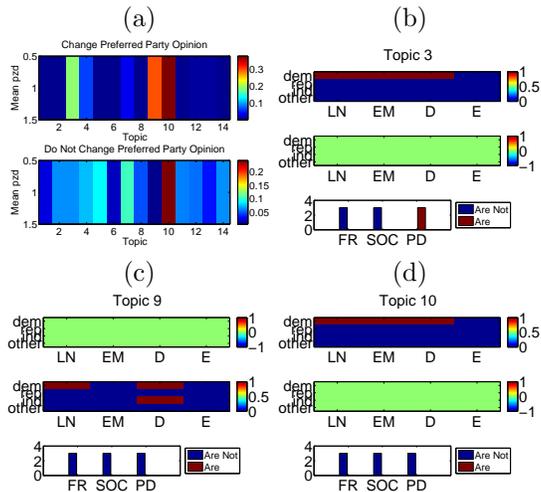
5.4 Results

Interaction Patterns of People who Change Opinion. The goal is to determine the difference in the interaction patterns of these two groups and we do this by comparing the most probable topics, averaged over all the users of each group. In Figure 3 for preferred party (PP) at $K = 14$, we observed the two groups 'people who changed opinion' and 'people who did not change opinion' was statistically significant with $p = 0.026$. Note, 5 users changed PP and 44 users did not. In Figure 4(a), the top plot shows the mean Φ for those that changed opinions and the bottom is for those that did not. The most probable

topics (dominating routine) for users that changed opinion was topic 3, 9, and 10 visualized by (b), (c), and (d), respectively. The most dominant topic for users that did not change was topic 10, which dominated in both groups. For a given topic ((b)-(d)), we display the 3 most probable words' (top) face-to-face interaction features (middle) phone interaction features and (bottom) relationship statistics, abbreviated by FR for friends, SOC for socialize and PD for political discussants.

Looking at Topic 3 (plot (b)), we can see that users that changed opinion predominantly had face-to-face interactions with PD, that were non-friends and not people they socialize with. The preferred party of these political discussants was democrat and this interaction occurred predominantly between 10pm-5pm (LN to D time components). Further, people who changed opinion also had heavy phone call and SMS activity with democrats as well as independents, as seen by Topic 9.

Fig. 4. (a) Mean topic distribution of users that changed opinion (top) and users that did not (bottom). Users that changed PP (4-point scale) had a high probability of topics 3, 9, 10, whereas users that did not change had a high probability of topic 10. By looking at the features of the 3 most probable words for these topics, we can see that users that changed opinion displayed (b) heavy face-to-face interactions with political discussants, and (c) they also had heavy phone call activity with non-friends.



Different Exposure for Increased vs. Decreased Interest in Politics.

We considered the difference in daily routines of users which increased their interest in politics as opposed to those that decreased their interest. Figure 5(a) shows the T-test results for the entropy of topic distributions of both groups, with $p = 0.06$. Figure 5(b) is the mean probability distribution of topics given the users from the two groups with $K = 22$ topics. The mean topic distribution $p(z|d)$ is shown for (top) all users that increased their interest, and (bottom) all users that decreased their interest. Note, 14 users increased their interest and 34 users decreased their interest in politics. Plots (c)-(e) show the most probable words for the dominating topics in both groups. Due to space constraints, we show topics which differed between the groups and disregard topics which were highly probable for both groups. Topic 14 (c) is highly probable for users that increased

their interest. Topic 8 and 18 are highly probable for users that decreased their interest. People who displayed increased interest were communicating most often by phone during the day. The group which decreased their interest had only face-to-face interactions (no phone communication) dominating their daily routines and it included interaction with people with little and no interest as seen by topics 8 and 18. There was heavy face-to-face interactions with friends in the early morning (EM) who had no interest in politics, for the group that decreased their interest.

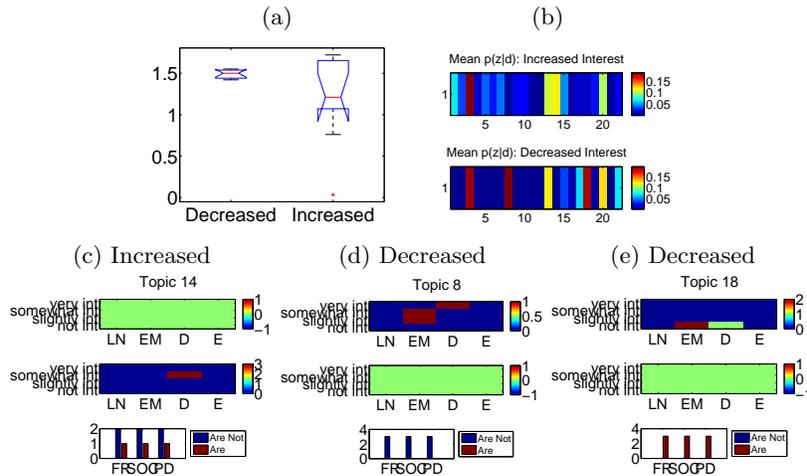


Fig. 5. Routines of people who increased their interest in politics versus those that decreased their interest. (a) T-test results reveal the difference in the entropy of topic distributions for these groups is statistically significant. (b) Mean distribution of topics for users of both groups. (c)-(e) Topics which best characterized users' daily life patterns in both groups. People who increased their interest often communicated by phone (c) and those that decreased interest had many face-to-face interactions with people with little/no interest in politics (d-e).

6 Conclusion

In this paper we describe a novel application of pervasive sensing using mobile phones— modeling the spread of political opinions in real-world face-to-face networks. Using mobile phone sensors, we estimate exposure to different opinions for individuals, find patterns of dynamic homophily at the community scale, recover 'political discussant' ties in the network, and explain individual political opinions on election day. We use an LDA-based model to study specific behaviors of people who changed their political opinions.

There are however, several limitations of our approach. We use bluetooth sensors to identify when people are in physical proximity, but there are many cases where individuals are proximate in the same room or space, but not necessarily interacting, as discussed in [43]. Another limitation is that our dataset consists of interactions captured using mobile phone sensors, and does not account for exposure to political opinions in mass media, e.g. via television or internet blog posts. The ability to estimate future changes in political opinions would certainly improve if such data was available. Mass-media access to political information, however, has been shown to be correlated with the self-reported party preference and political interest responses of the individual. On the technical front, our Windows Mobile platform at the time did not support GPS hardware sensors.

There are several fascinating future extensions of this work. In addition to political opinions, it would be important to understand if pervasive sensing methods can help understand the propagation of other types of opinions and habits in face-to-face networks, e.g., those related to health or purchasing behavior, both in our current dataset and also in other observational data. With the constant improvement in sensing technologies, future projects could use global positioning system (GPS) or infra-red (IR) sensors for better location and proximity sensing. Overall, our quantitative analysis sheds more light on long-standing open questions in political science and other social sciences, about the diffusion mechanism for opinions and behaviors.

7 Acknowledgments

The authors would like to thank Dr. David Lazer, Iolanthe Chronis and various undergraduate contributors for their help with the experimental deployment. Anmol Madan and Alex Pentland were partially sponsored by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053, and by AFOSR under Award Number FA9550-10-1-0122. Views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation. Katayoun Farrahi and Daniel Gatica-Perez were supported by the Swiss National Science Foundation through the MULTI and HAI projects.

References

1. T. Abdelzaher, Y. Anokwa, P. Boda, J. Burke, D. Estrin, L. Guibas, A. Kansal, S. Madden, and J. Reich. Mobiscopes for human spaces. *IEEE Pervasive Computing - Mobile and Ubiquitous Systems*, 6(2), 2007.
2. L. Adamic and N. Glance. The political blogosphere and the 2004 US election: divided they blog. In *Proc. of the 3rd Intl. Workshop on Link discovery*, 2005.
3. N. Agarwal and H. Liu. *Modeling and Data Mining in Blogosphere*, volume 1. Morgan & Claypool Publishers, 2009.

4. S. Avancha, A. Baxi, and D. Kotz. Privacy in mobile technology for personal healthcare. *Submitted to ACM Computing Surveys*, 2009.
5. Bernard H.R., Killworth P., Kronenfeld D., Sailer L. The Problem of Informant Accuracy: The Validity of Retrospective Data. *Annual Reviews in Anthropology*, 1984.
6. D. M. Blei, A. Y. Ng, M. I. Jordan, and J. Lafferty. Latent dirichlet allocation. *JMLR*, 3:2003, 2003.
7. D. Boyd and N. Ellison. Social network sites: Definition, history, and scholarship. *Journal of Computer Mediated Communication*, 13(1):210, 2007.
8. D. Brewer and C. Webster. Forgetting of friends and its effects on measuring friendship networks. *Social Networks*, 21(4):361 – 373, 2000.
9. R. Burt. Social contagion and innovation: Cohesion versus structural equivalence. *American Journal of Sociology*, 1987.
10. T. Choudhury. *Sensing and Modeling Human Networks*. PhD thesis, M.I.T., 2003.
11. T. Choudhury. Characterizing social networks using the sociometer. *Association of Computational Social and Organizational Science*, 2004.
12. T. Choudhury and S. Basu. Modeling Conversational Dynamics as a Mixed Memory Markov Process. *NIPS*, 2004.
13. Christopher Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
14. S. Currarini, M. Jackson, and P. Pin. An Economic Model of Friendship: Homophily, Minorities, and Segregation. *Econometrica*, 77(4):1003–1045, 2009.
15. N. Eagle, M. Macy, and R. Claxton. Network Diversity and Economic Development. *Science*, 328(5981):1029–1031, 2010.
16. N. Eagle, A. Pentland, and D. Lazer. Inferring Social Network Structure Using Mobile Phone Data. *PNAS*, 106(36):15274–15278, 2009.
17. E. M. et al. Darwin Phones: The Evolution of Sensing and Inference on Mobile Phones. *Mobisys*, 2010.
18. K. Farrahi and D. Gatica-Perez. What did you do today? discovering daily routines from large-scale mobile data. *ACM MM*, 2008.
19. K. Farrahi and D. Gatica-Perez. Probabilistic mining of socio-geographic routines from mobile phone data. *IEEE J-STSP*, 4(4):746–755, Aug. 2010.
20. Freund, Y. and Schapire, R.E. A short introduction to boosting. *Japanese Society for Artificial Intelligence*, 5:771–780, 1999.
21. Friedkin N.E. *A Structural Theory of Social Influence*. Cambridge University Press, 1998.
22. M. Gonzalez, C. Hidalgo, and A.-L. Barabasi. Understanding Individual Human Mobility Patterns. *Nature*, 453:779–782, 2008.
23. Granovetter M. Threshold models of collective behavior. *American Journal of Sociology*, 83:1420–1443, 1978.
24. F. Guo, S. Hanneke, W. Fu, and E. P. Xing. Recovering temporally rewiring networks: A model-based approach. In *ICML*, 2007.
25. Huckfeldt, R. and Sprague, J. Discussant Effects on Vote Choice: Intimacy, Structure and Interdependence. *The Journal of Politics*, 53:122–158, 1991.
26. D. R. Hunter and M. S. Handcock. Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics*, 15:565–583, 2006.
27. T. Huynh, M. Fritz, and B. Schiele. Discovery of activity patterns using topic models. In *UbiComp*, pages 10–19, Seoul, Korea, 2008.
28. Intentionally Blank. Intentionally blank. **blank**.
29. Kempe D., Kleinberg J., Tardos E. Maximizing the Spread of Influence in a Social Network. In *KDD*, Washington DC., 2003.

30. M. Laibowitz, J. Gips, R. Aylward, A. Pentland, and J. Paradiso. A sensor network for social dynamics. *IPSN*, page 491, 2006.
31. P. Lazarsfeld and R. K. Merton. Friendship as a Social Process: A Substantive and Methodological Analysis. *Freedom and Control in Modern Society*, 1954.
32. D. Lazer, B. Rubineau, N. Katz, C. Chetkovich, and M. A. Neblo. Networks and political attitudes: Structure, influence, and co-evolution. Working paper series, Sep 2008.
33. J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the Dynamics of the News Cycle. *ACM SIGKDD*, 2009.
34. H. Liu, J. Salerno, and M. Young. *Social Computing, Behavioral Modeling and Prediction*. Springer, 2008.
35. A. Madan and A. Pentland. Vibephones: Socially aware mobile phones. In *ISWC*, 2006.
36. A. Madan and A. Pentland. Modeling Social Diffusion Phenomena Using Reality Mining. In *AAAI Spring Symposium on Human Behavior Modeling*, 2009.
37. A. Mccallum, A. Corrada-emmanuel, and X. Wang. The author-recipient-topic model for topic and role discovery in social networks: Experiments with enron and academic email. Technical report, 2004.
38. M. McPherson, L. Smith-Lovin, , and J. M. Cook. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27:415–444, 2001.
39. S. Moturu. *Quantifying the Trustworthiness of Social Media Content: Content Analysis for the Social Web*. Lambert Academic Publishing, 2010.
40. R. Nallapati and W. Cohen. Link-plsa-lda: A new unsupervised model for topics and influence in blogs. *ICWSM*, 2008.
41. D. Olguin Olguin, P. Gloor, and A. Pentland. Wearable sensors for pervasive healthcare management. *PCT Healthcare*, 2009.
42. D. Olguin Olguin, B. Waber, T. Kim, A. Mohan, K. Ara, and A. Pentland. Sensible organizations: Technology and methodology for automatically measuring organizational behavior. *IEEE Transactions on Systems, Man, and Cybernetics-B*, 2009.
43. S. Patel, J. Kientz, J. Hayes, S. Bhat, and G. Abowd. Farther Than You May Think: An Empirical Investigation of the Proximity of Users to their Mobile Phones. In *Proceedings of Ubicomp*, pages 123–140, 2006.
44. A. Pentland. Socially aware computation and communication. In *ICMI*, 2005.
45. G. Robins, T. Snijders, P. Wang, M. Handcock, and P. Pattison. Recent developments in exponential random graph (p*) models for social networks. *Social Networks*, 29(2):192–215, 2007.
46. D. Roy, R. Patel, P. DeCamp, R. Kubat, M. Fleischman, B. Roy, N. Mavridis, S. Tellex, A. Salata, J. Guinness, M. Levit, and P. Gorniak. The human speechome project. *Lecture Notes in Computer Science*, 4211:192, 2006.
47. M. Steyvers, P. Smyth, M. Rosen-Zvi, and T. Groffiths. Probabilistic author-topic models for information discovery. In *ACM SIGKDD*, pages 306–315, 2004.
48. F. Wang, K. Carley, D. Zeng, and W. Mao. Social computing: From social informatics to social intelligence. In *IEEE Intelligent Systems*, pages 79–83, 2007.
49. X. Wang, N. Mohanty, and A. Mccallum. Group and topic discovery from relations and text. In *ACM SIGKDD (Workshop on Link Discovery)*, 2005.
50. D. Wyatt, T. Choudhury, and J. Bilmes. Discovering Long Range Properties of Social Networks with Multi-Valued Time-Inhomogeneous Models. In *AAAI*, 2010.
51. T. Yano, W. W. Cohen, and N. A. Smith. Predicting response to political blog posts with topic models. In *NAACL09*, pages 477–485, 2009.
52. D. Zeng, F. Wang, and K. Carley. Guest Editors' Introduction: Social Computing. *IEEE Intelligent Systems*, pages 20–22, 2007.