
The Bayesian Echo Chamber: Modeling Social Influence via Linguistic Accommodation

Fangjian Guo
Duke University
Durham, NC, USA
guo@cs.duke.edu

Charles Blundell
Gatsby Unit, UCL
London, UK
c.blundell@gatsby.ucl.ac.uk

Hanna Wallach
Microsoft Research
New York, NY, USA
wallach@microsoft.com

Katherine Heller
Duke University
Durham, NC, USA
kheller@stat.duke.edu

Abstract

We present the Bayesian Echo Chamber, a new Bayesian generative model for social interaction data. By modeling the evolution of people’s language usage over time, this model discovers latent influence relationships between them. Unlike previous work on inferring influence, which has primarily focused on simple temporal dynamics evidenced via turn-taking behavior, our model captures more nuanced influence relationships, evidenced via linguistic accommodation patterns in interaction content. The model, which is based on a discrete analog of the multivariate Hawkes process, permits a fully Bayesian inference algorithm. We validate our model’s ability to discover latent influence patterns using transcripts of arguments heard by the US Supreme Court and the movie “12 Angry Men.” We showcase our model’s capabilities by using it to infer latent influence patterns from Federal Open Market Committee meeting transcripts, demonstrating state-of-the-art performance at uncovering social dynamics in group discussions.

1 INTRODUCTION

As increasing quantities of social interaction data become available, often through online sources, researchers strive to find new ways of using these data to learn about human behavior. Most social processes, in which people or groups of people interact with one another in order to achieve specific (and sometimes con-

tradictory) goals, are extremely complex. In order to construct realistic models of these social processes, it is therefore necessary to take into account their structure (e.g., who spoke with whom), content (e.g., what was said), and temporal dynamics (e.g., when they spoke).

When studying social processes, one of the most pervasive questions is “who influences whom?” This question is of interest not only to sociologists and psychologists, but also to political scientists, organizational scientists, and marketing researchers. Since influence relationships are seldom made explicit, they must be inferred from other information. Influence has traditionally been studied by analyzing declared structural links in observed networks, such as Facebook “friendships” [Backstrom et al., 2006], paper citations [de Solla Price, 1965], and bill co-sponsorships [Fowler, 2006]. For many domains, however, explicitly stated links do not exist, are unreliable, or fail to reflect pertinent behavior. In these domains, researchers have used observed interaction dynamics as a proxy by which to infer influence and other social relationships. Much of this work has concentrated (either implicitly or explicitly) on turn-taking behavior—i.e., “who acts next.”

In this paper, we take a different approach: we move beyond turn-taking behavior, and present a new model, the Bayesian Echo Chamber, that uses observed interaction content, in the context of temporal dynamics, to capture influence. Our model draws upon a substantial body of work within sociolinguistics indicating that when two people interact, either orally or in writing, the use of a word by one person can increase the other person’s probability of subsequently using that word. Furthermore, the extent of this increase depends on power differences and influence relationships: the language used by a less powerful person will drift further so as to more closely resemble or “accommodate” the language used by more powerful people. This phenomenon is known as linguistic accommodation [West and Turner, 2010]. We demonstrate that linguistic accommodation can reveal

Appearing in Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS) 2015, San Diego, CA, USA. JMLR: W&CP volume 38. Copyright 2015 by the authors.

more nuanced influence patterns than those revealed by simple reciprocal behaviors such as turn-taking.

The Bayesian Echo Chamber is a new, mutually exciting, dynamic language model that combines ideas from Hawkes processes [Hawkes, 1971] with ideas from Bayesian language modeling. We draw inspiration from Blundell et al.’s model of turn-taking behavior [2012] (described in section 2) to define a new model of the mutual excitation of words in social interactions. This approach, which leverages a discrete analog of a multivariate Hawkes process, enables the Bayesian Echo Chamber to capture linguistic accommodation patterns via latent influence variables. These variables define a weighted influence network that reveals fine-grained information about who influences whom.

We provide details of the Bayesian Echo Chamber in section 3, including an MCMC algorithm for inferring the latent influence variables (and other parameters) from real-world data. To validate this algorithm, we provide parameter recovery results obtained using synthetic data. In section 5 we compare to several baseline models on data sets including arguments heard by the US Supreme Court [MacWhinney, 2007] and the transcript of the 1957 movie “12 Angry Men.” We compare influence networks inferred using our model to those inferred using Blundell et al.’s model. We show that by focusing on linguistic accommodation patterns, our model infers different—more substantively meaningful—influence networks than those inferred from turn-taking behavior. We also combine our model with Blundell et al.’s so as to jointly model turn-taking and linguistic accommodation. We investigate the possibility of tying the latent influence parameters to see if a single global notion of influence can be discovered. Finally, we showcase our model’s potential as an exploratory analysis tool for social scientists using recently released transcripts of Federal Reserve’s Federal Open Market Committee meetings.

2 INFLUENCE VIA TURN-TAKING

In this section, we give a brief description of a variant of Blundell et al.’s model for inferring influence from turn-taking behavior. Unlike Blundell et al.’s original paper, which modeled pairwise actions, we concentrate on a broadcast or group discussion setting appropriate for the data that we wish to model. (We also do not cluster participants by their interaction patterns.) This setting, in which every utterance is heard by and thus potentially influences every participant, occurs in many scenarios of interest to social scientists. Furthermore, the comparatively information-impooverished nature of this setting makes it one in which ability to infer influence relationships is deemed extremely valuable.

Blundell et al.’s model specifies a probabilistic generative process for the time stamps $\mathcal{T} = \{\mathcal{T}^{(p)}\}_{p=1}^P$ associated with a set of actions made by P people. In a group discussion setting, these actions correspond to utterances, and the model captures who will speak next and when that next utterance will occur. Letting $N^{(p)}(T)$ denote the total number of utterances made by person p over the entire observation interval $[0, T)$, each utterance made by p is associated with a time stamp indicating its start time, i.e., $\mathcal{T}^{(p)} = \{t_n^{(p)}\}_{n=1}^{N^{(p)}(T)}$. We assume that the duration of each utterance $\Delta t_n^{(p)}$ is observed and that its end time $t_n^{\prime(p)}$ can be calculated from its start time and duration: $t_n^{\prime(p)} = t_n^{(p)} + \Delta t_n^{(p)}$.

Hawkes processes [Hawkes, 1971]—a class of self- and mutually exciting doubly stochastic point processes—form the mathematical foundation of Blundell et al.’s model. A Hawkes process is a particular form of inhomogeneous Poisson process with a conditional stochastic rate function $\lambda(t)$ that depends on the time stamps of all events prior to time t . Blundell et al. model turn-taking interactions using coupled Hawkes processes. For a group discussion setting, we instead define a multivariate Hawkes process, in which each person p is associated with his or her own Hawkes process defined on $(0, \infty)$. Letting $N^{(p)}(\cdot)$ denote the counting measure of person p ’s Hawkes process, which takes as its argument an interval $[a, b)$ and returns the number of utterances made by p during that interval, the stochastic rate function for p ’s Hawkes process is

$$\lambda^{(p)}(t) = \lambda_0^{(p)} + \sum_{q \neq p} \int_0^{t^-} g^{(qp)}(t, u) dN^{(q)}(u) \quad (1)$$

$$= \lambda_0^{(p)} + \sum_{q \neq p} \sum_{n: t_n^{\prime(q)} < t} g^{(qp)}(t, t_n^{\prime(q)}), \quad (2)$$

where $\lambda_0^{(p)}$ is person p ’s base rate of utterances and $g(t, u)$ is a non-negative stationary kernel function that specifies the extent to which an event at time $u < t$ increases the instantaneous rate at time t , as well as the way in which this increase decays over time. Person p ’s rate function is coupled with the Hawkes processes of the other $P - 1$ people via their respective counting measures $\{N^{(q)}(\cdot)\}_{q \neq p}$ and the kernel function $g(t, u)$. Note that time stamp $t_n^{\prime(q)}$ is the *end* time of the n^{th} utterance made by person q . Consequently, an utterance made by person q only causes an increase in p ’s instantaneous rate after q ’s utterance is complete.

Blundell et al. use a standard exponential kernel function of the form $g^{(qp)}(t, u) = \nu^{(qp)} \exp\left(-\frac{(t-u)}{\tau_T^{(p)}}\right)$, although the shape of the non-negative kernel function $g^{(qp)}(t, u)$ could instead be learned in a non-parametric fashion [Zhou et al., 2013] if desired. Non-negative parameter $\nu^{(qp)}$

controls the degree of instantaneous excitation from person q to person p , while $\tau_T^{(p)}$ is a time decay parameter specific to person p that characterizes how fast excitation decays. Since the goal is to model influence between people, self-excitation is prohibited by enforcing $\nu^{(pp)} = 0$. Since a larger value of $\nu^{(qp)}$ will result in a higher instantaneous rate of utterances for person p , the non-negative parameters $\{\{\nu^{(qp)}\}_{q \neq p}\}_{p=1}^P$ define a weighted influence network that reflects conversational turn-taking behavior—i.e., who is likely to speak next and when that next utterance will occur. Details of an inference algorithm and appropriate priors for this turn-taking-based model can be found in the supplementary material.

3 INFLUENCE VIA LINGUISTIC ACCOMMODATION

In this section, we present our new dynamic Bayesian language model, the Bayesian Echo Chamber. This model specifies a probabilistic generative process for the words that occur in a set of utterances $\{\mathcal{W}^{(p)}\}_{p=1}^P$ made by P people, conditioned on the utterance start times and durations. Letting $N^{(p)}(T)$ denote the total number of utterances made by person p over the interval $[0, T)$, each utterance made by p consists of $L_n^{(p)}$ word tokens, i.e., $\mathcal{W}^{(p)} = \{\{w_{l,n}^{(p)}\}_{l=1}^{L_n^{(p)}}\}_{n=1}^{N^{(p)}(T)}$. Each token is an instance of one of V unique word types.

The generative process for each token draws upon ideas from both dynamic Bayesian language modeling and multivariate Hawkes processes. The l^{th} token in the n^{th} utterance made by person p is drawn from categorical distribution specific to that utterance: $w_{l,n}^{(p)} \sim \text{Categorical}(\phi_n^{(p)})$, where $\phi_n^{(p)}$ is a V -dimensional discrete probability vector. Each such probability vector is in turn drawn from a Dirichlet distribution with a person-specific concentration (or precision) parameter and an utterance-specific base measure: $\phi_n^{(p)} \sim \text{Dirichlet}(\alpha^{(p)}, \mathbf{B}_n^{(p)})$. Concentration parameter $\alpha^{(p)}$ is a positive scalar that determines the variance of the distribution, while base measure $\mathbf{B}_n^{(p)}$ is a V -dimensional discrete probability vector that specifies the mean of the distribution and satisfies

$$B_{v,n}^{(p)} \propto \beta_v^{(p)} + \sum_{q \neq p} \rho^{(qp)} \psi_{v,n}^{(qp)} \quad \text{and} \quad \sum_{v=1}^V B_{v,n}^{(p)} = 1. \quad (3)$$

V -dimensional vector $\beta^{(p)} \in \mathbb{R}_+^V$ characterizes person p 's inherent language usage. Non-negative parameter $\rho^{(qp)}$ controls the degree of linguistic excitation from person q to person p . Self-excitation is prohibited by enforcing $\rho^{(pp)} = 0$. Finally, $\psi_n^{(qp)} \in \mathbb{R}_+^V$ is a V -dimensional vector of decayed excitation pseudo-counts, constructed from all utterances made by per-

son q prior to person p 's n^{th} utterance, satisfying

$$\psi_{v,n}^{(qp)} = \sum_{m: t'_m^{(q)} < t_n^{(p)}} \left(\sum_{l=1}^{L_m^{(q)}} \mathbf{1}(w_{l,m}^{(q)} = v) \right) \times \exp\left(-\frac{t_n^{(p)} - t'_m^{(q)}}{\tau_L^{(p)}}\right), \quad (4)$$

where $\mathbf{1}(\cdot)$ is the indicator function. The inner sum is therefore equal to the number of tokens of type v in person q 's m^{th} utterance. Note that $t'_m^{(q)}$ is the end time of that utterance. Consequently, an utterance made by q only affects $\psi^{(qp)}$, and hence base measure $\mathbf{B}_n^{(p)}$, after q 's utterance is complete. Finally, $\tau_L^{(p)}$ is a time decay parameter specific to person p that characterizes how fast excitation decays. Since a larger value of $\rho^{(qp)}$ will increase the probability of person p using word types previously used by person q , the parameters $\{\{\rho^{(qp)}\}_{q \neq p}\}_{p=1}^P$ define a weighted influence network that reflects linguistic accommodation. A graphical model depicting the dependencies between utterances is in the supplementary material.

3.1 Inference

For real-world group discussions, the utterance contents $\mathcal{W} = \{\mathcal{W}^{(p)}\}_{p=1}^P$, start times \mathcal{T} , and durations \mathcal{D} are observed, while parameters $\Theta = \{\{\phi_n^{(p)}\}_{n=1}^{N^{(p)}(T)}, \alpha^{(p)}, \beta^{(p)}, \{\rho^{(qp)}\}_{q \neq p}, \tau_L^{(p)}\}_{p=1}^P$ are unobserved; however, information about the values of these unobserved parameters can be quantified via their posterior distribution given \mathcal{W} , \mathcal{T} , and \mathcal{D} , i.e., $P(\Theta | \mathcal{W}, \mathcal{T}, \mathcal{D}) \propto P(\mathcal{W} | \Theta, \mathcal{T}, \mathcal{D}) P(\Theta)$. The likelihood term can be factorized into the following product due to our model's independence assumptions:

$$P(\mathcal{W} | \Theta, \mathcal{T}, \mathcal{D}) = \prod_{p=1}^P \prod_{n=1}^{N^{(p)}(T)} P(\mathbf{w}_n^{(p)} | \{\{w_m^{(q)}\}_{m: t'_m^{(q)} < t_n^{(p)}\}_{q \neq p}, \Theta).$$

Using Dirichlet–multinomial conjugacy, the probability vectors $\{\{\phi_n^{(p)}\}_{n=1}^{N^{(p)}(T)}\}_{p=1}^P$ can be integrated out:

$$P(\mathcal{W} | \Theta, \mathcal{T}, \mathcal{D}) = \prod_{p=1}^P \prod_{n=1}^{N^{(p)}(T)} \prod_{l=1}^{L_n^{(p)}} \frac{\sum_{l'=1}^{l-1} \mathbf{1}(w_{l',n}^{(p)} = w_{l,n}^{(p)}) + \alpha^{(p)} B_{w_{l,n}^{(p)},n}^{(p)}}{l-1 + \alpha^{(p)}}.$$

To complete the specification of $P(\Theta)$, we place gamma priors over the remaining parameters $\{\alpha^{(p)}, \beta^{(p)}, \{\rho^{(qp)}\}_{q \neq p}, \tau_L^{(p)}\}_{p=1}^P$. The resultant posterior distribution $P(\Theta | \mathcal{W}, \mathcal{T}, \mathcal{D})$ is analytically intractable; however, posterior samples of

$\{\alpha^{(p)}, \beta^{(p)}, \{\rho^{(qp)}\}_{q \neq p}, \tau_L^{(p)}\}_{p=1}^P$ can be obtained using a collapsed slice-within-Gibbs sampling algorithm [Neal, 2003]. Additional details, including pseudocode, are given in the supplementary material.

4 RELATED WORK

Several recent probabilistic models use point processes as a foundation for inferring influence and other social relationships from temporal dynamics [Simma and Jordan, 2010; Blundell et al., 2012; Perry and Wolfe, 2013; Iwata et al., 2013; DuBois et al., 2013; Zhou et al., 2013; Linderman and Adams, 2014]. Hawkes processes play a central role in some of these models. Most relevant to this paper is the work of Linderman and Adams [2014], who used Hawkes processes to study gang-related homicide in Chicago. Although temporal dynamics can reveal some social relationships, others may be more readily evidenced by also modeling interaction content. In this vein, Danescu-Niculescu-Mizil et al. [2012] analyzed discussions among Wikipedians and arguments before the US Supreme Court to uncover power differences, while Gerrish and Blei [2010] took a language-based approach to measuring scholarly impact, identifying influential documents by analyzing changes to thematic content over time. These models differ significantly from ours, and have not been used in a comparative analysis of different approaches to characterizing influence. This paper compares approaches, demonstrating that influence networks inferred from linguistic accommodation can be more substantively meaningful than those inferred from turn-taking. We also move beyond the work of Danescu-Niculescu-Mizil et al. and Gerrish and Blei by defining a generative, dynamic Bayesian language model that captures the mutual excitation of words in social interactions.

5 EXPERIMENTS

In this section, we showcase the Bayesian Echo Chamber’s ability to model transcripts of oral arguments heard by the US Supreme Court, the transcript of the 1957 movie “12 Angry Men,” and meeting transcripts from the Federal Reserve’s Federal Open Market Committee. We compare our model with competing approaches using the probability of held-out data, and demonstrate that our model can recover meaningful influence patterns from these data sets. We also compare the linguistic accommodation-based influence networks inferred by our model to the turn-taking-based networks inferred by Blundell et al.’s model. Finally, we combine our model with Blundell et al.’s in order to jointly model turn-taking and linguistic accommodation. We also investigate tying the mod-

els’ latent influence parameters to determine whether a single global notion of influence can be discovered.

The US Supreme Court consists of a chief justice and eight associate justices. Each oral argument heard by the Court therefore involves up to nine justices (some may recuse themselves) plus attorneys representing the petitioner and the respondent. The format of each argument is formulaic: the attorneys for each party have 30 minutes to present their argument, with those representing the petitioner speaking first. Justices routinely interrupt the attorneys’ presentations to make comments or ask questions of the attorneys. Sometimes additional attorneys, known as “amiciae curae,” also present arguments in support of either the petitioner or the respondent. We used the time-stamped transcripts¹ from three controversial Supreme Court cases [MacWhinney, 2007]: *Lawrence and Garner v. Texas*, *District of Columbia v. Heller*, and *Citizens United v. Federal Election Commission* (re-argument).

“12 Angry Men” is a movie about a jury’s deliberations regarding the guilt or acquittal of a defendant. Unlike Supreme Court arguments, the dialog is informal and intended to seem natural. The movie is unique in its limited cast of 12 people and in the fact that it is set almost entirely in one room. These qualities, combined with the fact that the movie explicitly focuses on discussion-based consensus building in a group setting, make its time-stamped transcript an ideal data set for exploring the strengths of our model. We generated an appropriate transcript from the movie subtitles by hand-labeling the person who made each utterance.

The Federal Reserve’s Federal Open Market Committee oversees the US’s open market operations and sets the US national monetary policy. The Committee consists of 12 voting members: the seven members of the Federal Reserve Board and five of the 12 Federal Reserve Bank presidents. By law, the FOMC must meet at least four times a year, though it typically meets every five to eight weeks. At each meeting, the Committee votes on the policy (tightening, neutrality, or easing) to be carried out until the next meeting. Meeting transcripts are embargoed for five years; as a result, transcripts from the meetings surrounding the 2007–2008 financial crisis have only recently been released. We used transcripts from 32 meetings ranging from March 27, 2006 to December 15, 2008, inclusive.²

For all data sets, we concatenated consecutive utterances by the same person, discarded contributions from people with fewer than ten (post-concatenation) utterances, and rescaled all time stamps to the interval $(0, T = 100]$. For each data set, we also restricted

¹<http://talkbank.org/data/Meeting/SCOTUS/>

²<http://poliinformatics.org/data/>




Figure 1: Parameter Recovery Results.

the vocabulary to the $V = 600$ most frequent stemmed word types. We did not remove stop words, since they can carry important information about influence relationships [Danescu-Niculescu-Mizil et al., 2012]. The salient characteristics of each data set, after preprocessing, are provided in the supplementary material.

5.1 Parameter Recovery

In this section, we present the results of a parameter recovery experiment, conducted in order to validate our inference algorithm. We used the generative process described in section 3 to generate 300 utterances made by $P = 3$ people. In total, these utterances contain 15,070 tokens spanning $V = 20$ types. We drew the length of each utterance, i.e., $L_n^{(p)}$, from a Poisson distribution with a mean of 50. We generated the start times and durations by assuming a round-robin approach to turn-taking and setting the duration of each utterance to a value proportional to its length in tokens. Figure 1 shows the true and inferred parameter values, depicted using blue triangles and red circles, respectively. The inferred parameter values were obtained by averaging 3,000 samples from the posterior distribution. The error bars indicate one standard deviation. Our proposed inference algorithm does well at accurately recovering the true parameter values.

5.2 Probability of Held-Out Data

The predictive probability of held-out data, sometimes expressed as perplexity, is a standard metric for evaluating statistical language models—the higher the probability, the better the model. We compared predictive probabilities obtained using the Bayesian Echo Chamber and several real-world data sets to those obtained using two comparable language models.

To compute the predictive probability of held-out data, we divided each data set into a training set $\{\mathcal{W}_{\text{train}}, \mathcal{T}_{\text{train}}, \mathcal{D}_{\text{train}}\}$ and a held-out or test set $\{\mathcal{W}_{\text{test}}, \mathcal{T}_{\text{test}}, \mathcal{D}_{\text{test}}\}$. We formed each training set by selecting those utterances that occurred before some time t^* where t^* was chosen to yield either a 90%–10% or 80%–20% training–testing split, i.e., $\mathcal{W}_{\text{train}} = \{\mathbf{w}_n^{(p)} : t_n^{(p)} \leq t^*\}_{p=1}^P$ and $\mathcal{W}_{\text{test}} = \{\mathbf{w}_m^{(p)} : t_m^{(p)} > t^*\}_{p=1}^P$. The predictive probability of held-out data

is then $P(\mathcal{W}_{\text{test}} | \mathcal{T}_{\text{test}}, \mathcal{D}_{\text{test}}, \mathcal{W}_{\text{train}}, \mathcal{T}_{\text{train}}, \mathcal{D}_{\text{train}})$. Although this probability is analytically intractable, its logarithm can be approximated via the lower bound $\frac{1}{S} \sum_{s=1}^S \log P(\mathcal{W}_{\text{test}} | \mathcal{T}_{\text{test}}, \mathcal{D}_{\text{test}}, \Theta^{(s)})$ where $\Theta^{(s)}$ denotes a set of sampled parameter values drawn from the posterior distribution $P(\Theta | \mathcal{W}_{\text{train}}, \mathcal{T}_{\text{train}}, \mathcal{D}_{\text{train}})$.

Approximate log probabilities obtained using the Bayesian Echo Chamber (with $S = 3000$ samples after 1000 burn-in sampling iterations), a unigram language model, and Blei and Lafferty’s dynamic topic model [2006] are provided in table 1. Log probabilities for additional data sets are provided in the supplementary material. The unigram language model is equivalent to setting all influence parameters in our model to zero. In all experiments involving the dynamic topic model, each data set was sliced into $K = 10$ or $K = 5$ equally-sized time slices (depending on the training–testing split), with the last slice taken to be the test set and utterances treated as documents. Each log probability reported for the dynamic topic model is the highest value obtained using either 5, 10, or 20 topics. Since inference for the dynamic topic model was performed using a variational inference algorithm,³ its log probabilities are also lower bounds and standard deviations are not available. For all data sets, the Bayesian Echo Chamber out-performed both the unigram language model and the dynamic topic model.

5.3 Influence Recovery

In this section, we demonstrate that the Bayesian Echo Chamber can recover known influence patterns in Supreme Court arguments and in the movie “12 Angry Men.” We also use these data sources to compare influence networks inferred by our model to those inferred by the model described in section 2. All reported influence parameters were obtained by averaging 3,000 posterior samples; posterior standard deviations are provided in the supplementary material.

5.3.1 US Supreme Court

As described previously, Supreme Court arguments are extremely formulaic: The attorneys representing the petitioner present their argument first, speaking for a total of 30 minutes before the respondent’s attorneys are allowed to present their argument. Justices routinely interrupt these presentations. We therefore anticipate that influence networks inferred from linguistic accommodation patterns will reveal significant influence exerted by the petitioner’s attorneys, simply because they speak first, establishing the language used in the rest of the discussion. We also anticipate

³Inference code obtained from <http://www.cs.princeton.edu/~blei/topicmodeling.html>

Table 1: Predictive Log Probabilities of Held-Out Data.

Data Set	10% Test Set			20% Test Set		
	Our Model	Unigram	DTM	Our Model	Unigram	DTM
Synthetic	-4292.97 ±0.02	-4297.92±0.04	-4364.81	-8702.92 ±0.04	-8717.77±0.08	-8948.07
DC v. Heller	-7383.45 ±0.12	-7794.25±0.21	-7533.58	-12404.21 ±0.15	-13126.73±0.26	-12744.73
L&G v. Texas	-6663.33 ±0.12	-6937.66±0.18	-6759.06	-10248.80 ±0.21	-10791.25±0.23	-10459.87
Citizens United v. FEC	-5770.12 ±0.14	-6120.67±0.18	-5851.224	-16370.7 ±0.95	-17157.21±0.40	-16400.46
“12 Angry Men”	-4667.47 ±0.24	-4920.21±0.14	-4691.11	-8722.97 ±0.27	-9222.99±0.25	-8787.35

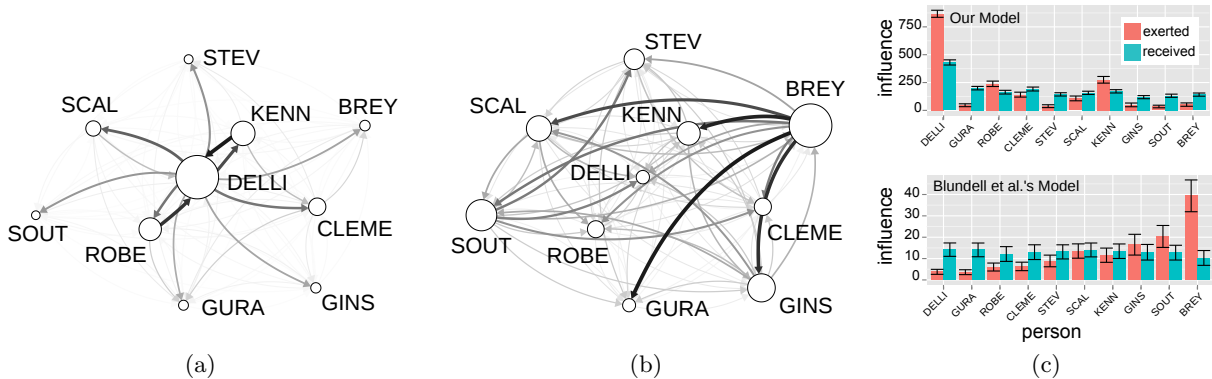


Figure 2: Influence Networks (Posterior Mean) Inferred from the DC v. Heller Case. (a) Network Inferred Using Our Model and (b) Inferred Using Blundell et al.’s Model. (c) Total Influence Exerted/Received by Each Person.

that influence networks inferred from turn-taking behavior will reveal significant influence exerted by the justices over the attorneys. This is because the justices interrogate the attorneys’ during their presentations.

As an illustrative example, we present results obtained from the District of Columbia v. Heller case in figure 2. (The other two cases, Lawrence and Garner v. Texas and Citizens United v. Federal Election Commission exhibited remarkably similar influence networks.) The influence network⁴ inferred using the Bayesian Echo chamber is shown in 2(a), while the network inferred using Blundell et al.’s model is shown in 2(b). To illustrate posterior uncertainty, networks drawn with different posterior quantiles are provided in the supplementary material. The total influence exerted and received by each participant are shown for each model in figure 2(c). The error bars represent the posterior standard deviation. The justices present for this case were Alito, Breyer, Ginsburg, Kennedy, Roberts, Scalia, Stevens, Souter, and Thomas, while the attorneys were Dellinger (representing the petitioner), Gura (representing the respondent), and Clement (as amicae curae, supporting the petitioner). Ultimately, Alito, Kennedy, Roberts, Scalia, and Thomas (the majority) sided with the respondent, while Breyer, Ginsburg, and Stevens (the minority) sided with the petitioner. Neither Alito or Thomas spoke ten or more

utterances, so they were not included in our analyses.

The influence network inferred using our model is very sparse. As expected, Dellinger (who represented the petitioner and presented his argument first) is shown as exerting the most influence. The justices with the most influence are Kennedy and Roberts, both of whom ultimately supported the respondent and thus interrogated Dellinger much more the other justices.

The most striking pattern in the influence network inferred using Blundell et al.’s model is that the three attorneys received much more influence from the justices than vice versa. This pattern could be seen as reflecting the status difference between justices and attorneys or as reflecting the formulaic structure of the Supreme Court: attorneys present arguments, while justices interrupt to make comments or ask questions.

5.3.2 “12 Angry Men”

Unlike Supreme Court arguments, the dialog in “12 Angry Men” is informal and intended to seem natural. Since the focus of the movie is discussion-based consensus building in a group setting, we therefore anticipate that the narrative of the movie will be reflected in influence networks inferred from linguistic accommodation patterns and from turn-taking behavior.

The influence network inferred using our model is shown in figure 3(a), while the total influence exerted

⁴Plotted using `qgraph` [Epskamp et al., 2012].

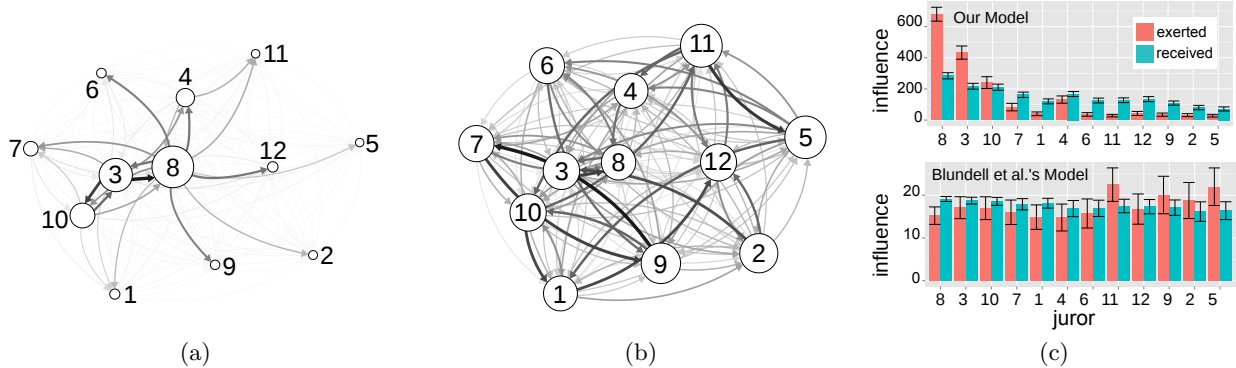


Figure 3: Influence Networks (Posterior Mean) Inferred from “12 Angry Men.” (a) Network Inferred Using Our Model and (b) Inferred Using Blundell et al.’s Model. (c) Total Influence Exerted/Received by Each Person.

and received by each juror are shown in the top of figure 3(c). The most significant pattern is that three individuals exert more influence over others than others do over them: Juror 8, Juror 3, and, to a lesser extent, Juror 10. Juror 8 is the protagonist of the movie, and initially casts the only “not guilty” vote. The other jurors ultimately change their votes to match his. Juror 3, the antagonist, is the last to change his vote. It therefore unsurprising that Juror 8, the first to vote “not guilty”, should dominate the discussion content. Similarly, Juror 3, the last to change his “guilty” vote, is most invested in discussing defendant’s supposed guilt. Juror 10 is one of the last three jurors, along with Jurors 3 and 4, to change his vote. However, unlike Juror 4 (who stands out marginally in figure 3(a) and, according to figure 3(c), has less influence over others than others do over him), Juror 10 is argumentative as he changes his mind. Overall, the consistency of the inferred influence network with the narrative of the movie confirms that the Bayesian Echo Chamber can indeed uncover substantive influence relationships.

The influence network inferred using Blundell et al.’s model and the total influence exerted and received by each juror are shown in figure 3(b) and the bottom of 3(c), respectively. The four jurors who exert more influence over others than others do over them (Juror 2, Juror 5, Juror 9, and Juror 11) are the first four jurors to change their votes. Jurors 5 and 11, who exert the most influence, are verbose, while Jurors 2 and 9 are comparatively taciturn. Jurors 8 exerts little influence because he must respond to questions and defend his position as he tries to persuade the others to agree with him, much like the attorneys in the Supreme Court.

5.4 Exploratory Analysis of FOMC Meetings

Finally, we performed an exploratory analysis of the relationships inferred from transcripts of 32 Federal

Open Market Committee meetings surrounding the 2007–2008 financial crisis, ranging from March 27, 2006 to December 15, 2008, inclusive. Since utterance durations are not available for these transcripts (also preventing the use of Blundell et al.’s model), we set the duration of each utterance to a value proportional to its length in tokens. We divided the meetings into three subsets: March 27, 2006 through June 28, 2006; August 8, 2006 through August 7, 2007; and August 10, 2007 through December 15, 2008. The first subset corresponds to meetings with a resultant policy of tightening; the second to meetings with a neutral outcome; and the third to meetings that resulted in easing. These meetings were all chaired by Bernanke.

Figure 4 depicts the influence network for each subset (aggregated by averaging over the meetings in that subset) inferred using our model. In the first network, corresponding to pre-crash meetings from March 27, 2006 through June 28, 2006, Bernanke, Fisher, and Lacker play the biggest roles with Bernanke, the chair, exerting the most influence over others. Given his role as chair, Bernanke’s involvement is arguably unsurprising, but Fisher and Lacker’s roles are notable. Unlike Bernanke, Fisher and Lacker are both “hawks” and thus generally in favor of tightening monetary policy; the meetings in this subset all resulted in an outcome of tightening. In the second network, corresponding to pre-crash meetings from August 8, 2006 through August 7, 2007, Bernanke, Fisher, and Lacker all continue to play significant roles, but the network is much less sparse, with both hawks and “doves” (those generally in favor of easing monetary policy) exerting influence over others. In contrast to the meetings in the previous subset, these meetings resulted in neutrality—i.e., neither tightening or easing. Finally, in the third network, corresponding to post-crash meetings from August 10, 2007 through December 15, 2008, there are fewer strong influence relation-

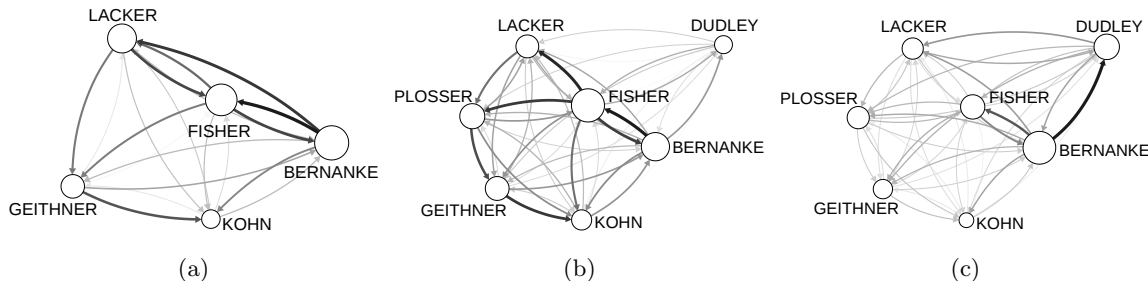


Figure 4: Influence Networks (Posterior Mean) Inferred from FOMC Meetings Using Our Model. (a) March 27, 2006–June 28, 2006. (b) August 8, 2006–August 7, 2007. (c) August 10, 2007–December 15, 2008.

ships. Bernanke (the chair and a dove) still plays a major role, while Fisher and Lacker’s roles are significantly diminished. Instead, Dudley, also a dove and a close ally of Bernanke, plays a much greater role, especially in his relationship with Bernanke. These meetings all resulted in monetary policy easing, a strategy generally favored by doves and opposed by hawks.

There has been little work in political science, economics, or computer science on analyzing these meeting transcripts. As a result, the inferred networks not only showcase our model’s ability to discover latent influence relationships from linguistic accommodation, but also constitute a research contribution of substantive interest to political scientists, economists, and other social scientists studying the financial crisis.

5.5 Model Combination

Since influence can be inferred from both turn-taking behavior and linguistic accommodation, we explored the possibility of combining the Bayesian Echo Chamber and Blundell et al.’s model to form a “supermodel” with a single set of shared influence parameters. The simplest way to share these parameters is to tie them together as $\rho^{(qp)} = r\nu^{(qp)}$, where r is a scaling factor and $\rho^{(qp)}$ and $\nu^{(qp)}$ correspond to the influence from person q to person p in our model and Blundell et al.’s model, respectively. Tying the influence parameters in this way provides the model with the capacity to capture a global notion of influence that is based upon both turn-taking and linguistic accommodation.

This tied model, whose likelihood is the product of the Bayesian Echo Chamber’s likelihood and that of Blundell et al.’s model but with shared influence parameters, assigned lower probabilities to held-out data than the fully factorized model (i.e., separate influence parameters). Log probabilities, obtained using a 90%–10% training–testing split and a vocabulary of $V = 300$ word types in order to reduce computation time, are provided in the supplementary material.

Interestingly, the networks inferred by the model with

tied parameters are extremely similar to those inferred using the Bayesian Echo Chamber. These results suggest that linguistic accommodation reflects a more informative notion of influence than that evidenced via turn-taking. We expect that investigating other ways of combining turn-taking-based models with ours will be a promising direction for future exploration.

6 DISCUSSION

The Bayesian Echo Chamber is a new generative model for discovering latent influence networks via linguistic accommodation patterns. We demonstrated that our model can recover known influence patterns in synthetic data, arguments heard by the US Supreme Court, and in the movie “12 Angry Men.” We compared influence networks inferred using our model to those inferred using a variant of Blundell et al.’s turn-taking-based model and showed that by modeling linguistic accommodation patterns, our model infers different, and often more meaningful, influence networks. Finally, we showcased our model’s potential as an exploratory analysis tool for social scientists by inferring latent influence relationships between members of the Federal Reserve’s Federal Open Market Committee.

Promising avenues for future work include (1) modeling linguistic accommodation separately for function and content words and (2) explicitly modeling the dynamic evolution of influence networks over time.

Acknowledgements

Thanks to Juston Moore and Aaron Schein for their work on early stages of this project, and to Aaron for the “Bayesian Echo Chamber” model name. This work was supported in part by the Center for Intelligent Information Retrieval, in part by NSF grant #IIS-1320219, and in part by NSF grant #SBE-0965436. Any opinions, findings and conclusions or recommendations expressed in this material are the authors’ and do not necessarily reflect those of the sponsor.

References

- Backstrom, L., Huttenlocher, D., Kleinberg, J., and Lan, X. (2006). Group formation in large social networks: Membership, growth, and evolution. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Blei, D. M. and Lafferty, J. D. (2006). Dynamic topic models. In *Proceedings of the 23rd International Conference on Machine Learning*.
- Blundell, C., Heller, K. A., and Beck, J. (2012). Modelling reciprocating relationships with Hawkes processes. In *Advances In Neural Information Processing Systems*.
- Danescu-Niculescu-Mizil, C., Lee, L., Pang, B., and Kleinberg, J. (2012). Echoes of power: Language effects and power differences in social interaction. In *Proceedings of the 21st International Conference on World Wide Web*, pages 699–708.
- de Solla Price, D. J. (1965). Networks of scientific papers. *Science*, 149(3683):510–515.
- DuBois, C., Butts, C., and Smyth, P. (2013). Stochastic blockmodeling of relational event dynamics. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, pages 238–246.
- Epskamp, S., Cramer, A. O., Waldorp, L. J., Schmittmann, V. D., and Borsboom, D. (2012). **qgraph**: Network visualizations of relationships in psychometric data. *Journal of Statistical Software*, 48(4):1–18.
- Fowler, J. H. (2006). Legislative cosponsorship networks in the US House and Senate. *Social Networks*, 28:454–465.
- Gerrish, S. M. and Blei, D. M. (2010). A language-based approach to measuring scholarly impact. In *Proceedings of the 26th International Conference on Machine Learning*.
- Hawkes, A. G. (1971). Point spectra of some self-exciting and mutually exciting point processes. *Journal of the Royal Statistical Society: Series B (Methodology)*, 58:83–90.
- Iwata, T., Shah, A., and Ghahramani, Z. (2013). Discovering latent influence in online social activities via shared cascade Poisson processes. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 266–274.
- Linderman, S. W. and Adams, R. P. (2014). Discovering latent network structure in point process data. *International Conference on Machine Learning (ICML)*.
- MacWhinney, B. (2007). The TalkBank project. *Creating and Digitizing Language Corpora: Synchronic Databases*.
- Neal, R. M. (2003). Slice sampling. *Annals of Statistics*, 31(3):705–767.
- Perry, P. O. and Wolfe, P. J. (2013). Point process modelling for directed interaction networks. *Journal of the Royal Statistical Society: Series B (Methodology)*, 75(5):821–849.
- Simma, A. and Jordan, M. I. (2010). Modeling events with cascades of Poisson processes. *Proceedings of Uncertainty in Artificial Intelligence*.
- West, R. and Turner, L. (2010). *Introducing Communication Theory: Analysis and Applications*. McGraw Hill.
- Zhou, K., Zha, H., and Song, L. (2013). Learning triggering kernels for multi-dimensional Hawkes processes. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1301–1309.