

# Structure and Dynamics of Information Pathways in Online Media

Manuel  
Gomez-Rodriguez  
MPI for Intelligent Systems  
Stanford University  
manuelgr@stanford.edu

Jure Leskovec  
Stanford University  
jure@cs.stanford.edu

Bernhard Schölkopf  
MPI for Intelligent Systems  
bs@tuebingen.mpg.de

## ABSTRACT

Diffusion of information, spread of rumors and infectious diseases are all instances of stochastic processes that occur over the edges of an underlying network. Many times networks over which contagions spread are unobserved, and such networks are often dynamic and change over time. In this paper, we investigate the problem of inferring dynamic networks based on information diffusion data. We assume there is an unobserved dynamic network that changes over time, while we observe the results of a dynamic process spreading over the edges of the network. The task then is to infer the edges and the dynamics of the underlying network.

We develop an on-line algorithm that relies on stochastic convex optimization to efficiently solve the dynamic network inference problem. We apply our algorithm to information diffusion among 3.3 million mainstream media and blog sites and experiment with more than 179 million different pieces of information spreading over the network in a one year period. We study the evolution of information pathways in the online media space and find interesting insights. Information pathways for general recurrent topics are more stable across time than for on-going news events. Clusters of news media sites and blogs often emerge and vanish in matter of days for on-going news events. Major social movements and events involving civil population, such as the Libyan's civil war or Syria's uprising, lead to an increased amount of information pathways among blogs as well as in the overall increase in the network centrality of blogs and social media sites.

**Categories and Subject Descriptors:** H.2.8 [Database Management]: Database applications—*Data mining*

**General Terms:** Algorithms; Experimentation.

**Keywords:** Networks of diffusion, Information cascades, Blogs, News media, Meme-tracking, Social networks.

## 1. INTRODUCTION

Networks represent a fundamental medium for spreading and diffusion of various types of behavior, information, rumors and diseases [27]. A *contagion* appears at some node of a network and then spreads like an epidemic from node to node over the edges of

the underlying network. For example, in case of information diffusion, the contagion represents a piece of information [16, 18] and infection events correspond to times when nodes mention or copy the information from one of their neighbors in the network. Similarly, we can think about the spread of a new type of behavior or an action, e.g., purchasing a new cellphone [15], or the propagation of a contagious disease over the edges of the underlying social network [6].

In the context of network diffusion, we often observe the temporal traces of diffusion while the pathways over which contagion spreads remain hidden. In other words, we observe the times when each node gets infected by the contagion, but the edges of the network that gave rise to the diffusion remain unobservable. For example, we can often measure and observe the time when people decide to adopt a new behavior while we do not explicitly observe which neighbor in the social network influenced them to do so. In case of information diffusion, we often observe people (or media sites) talking about a new piece of information without explicitly observing the path it took in the information diffusion network to reach the particular node of interest. And, epidemiologists often observe when a person gets sick but usually cannot tell who infected her. In all these examples, one can observe the infection events themselves while not knowing over which edges of the network the contagions spread. Therefore, one of the fundamental research problems in the context of network diffusion is inferring the structure of networks over which various types of contagions spread [10]. Moreover, many times networks over which contagions diffuse are not static but change over time. Depending on the type of contagion, the time of the day, or death of the existing and birth of new nodes, the underlying network may dynamically change and shift over time.

In recent years, several network inference algorithms have been developed [9, 10, 12, 20, 24, 30]. Some approaches infer only the network structure [10, 30], while others infer not only the network structure but also the *strength* or the average latency of every edge in the network [9, 20]. However, to the best of our knowledge, previous work has always assumed networks to be static and contagion pathways to be constant over time. However, in most cases, networks are dynamic, and contagion pathways change over time, depending upon the contagions that propagate through them [22, 28]. For example, a blog can increase its popularity abruptly after one of its posts turns *viral*, this may create new edges in the information transmission network and so the content the blog produces in the future will likely spread to larger parts of the network. Similarly, at any given time a particular unexpected event may occur and a topic or piece of news may become very popular for a limited period of time. This again will lead to different emerging and vanishing information pathways, and thus to a time-varying underlying

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WSDM'13, February 4–8, 2013, Rome, Italy.

Copyright 2013 ACM 978-1-4503-1869-3/13/02 ...\$15.00.

| Model       | Edge $(j, i)$ transmission likelihood $f(t_i t_j; \alpha_{j,i})$  |
|-------------|---|
| EXponential | $\begin{cases} \alpha_{j,i} \cdot e^{-\alpha_{j,i}(t_i-t_j)} & \text{if } t_j < t_i \\ 0 & \text{otherwise} \end{cases}$  |
| POWER law   | $\begin{cases} \frac{\alpha_{j,i}}{\delta} \left(\frac{t_i-t_j}{\delta}\right)^{-1-\alpha_{j,i}} & \text{if } t_j + \delta < t_i \\ 0 & \text{otherwise} \end{cases}$ |
| RAYleigh    | $\begin{cases} \alpha_{j,i}(t_i-t_j)e^{-\frac{1}{2}\alpha_{j,i}(t_i-t_j)^2} & \text{if } t_j < t_i \\ 0 & \text{otherwise} \end{cases}$                               |

**Table 1: Various models of edge transmission likelihood.**

network. In order to better understand these temporal changes, one needs to reconstruct the time-varying structure and underlying temporal dynamics of these networks and then study the information pathways of real-world events, topics or content.

**Our approach to time-varying network inference.** In this paper we investigate the problem of inferring dynamic networks based on information diffusion data. We assume there is an unobserved dynamic network that changes over time, while we observe the node infection times of many different contagions spreading over the edges of the network. The task then is to infer the edges and the dynamics of the underlying network. We develop an efficient on-line dynamic network inference algorithm, INFOPATH, that allows us to infer daily networks of information diffusion between online media sites over a one year period using more than 179 million different contagions diffusing over the underlying media network.

We model diffusion processes as discrete networks of fully continuous temporal processes occurring at different rates building on our previous work [9, 11]. Our model allows information to propagate at different rates across different edges by adopting a data-driven approach, where only the recorded temporal diffusion events are used. The model considers the information which propagates through the network due only to diffusion, while ignoring any external sources [22]. However, our original diffusion model considered only static networks [9]. Here, we generalize the model and develop a new inference method to support dynamic networks. Our time-varying network inference algorithm, INFOPATH, uses stochastic gradient [26] to provide estimates of the time-varying structure and temporal dynamics of the inferred network. The framework enables us to study the temporal evolution of information pathways in the online media space.

We apply the INFOPATH algorithm to synthetic as well as real Web information propagation data. We study 179 million different information cascades spreading among 3.3 million blog and news media sites over a one year period, from March 2011 till February 2012.<sup>1</sup> Results on synthetic data show INFOPATH is able to track changes in the topology of dynamic networks and provides accurate on-line estimates of the time-varying transmission rates of the edges of the network. INFOPATH is also robust across network topologies, and temporal trends of edge transmission dynamics.

Experiments on large-scale real news and social media data lead to interesting insights and findings. For example, we find that the information pathways over which general recurrent topics propagate remain more stable over time, while unexpected events lead to dramatically changing information pathways. Clusters of mainstream news and blogs often emerge and vanish in a matter of days, and our on-line algorithm is able to uncover such structures. News events that involve large-scale social movements, as the Libyan civil war, Egypt’s revolution or Syria’s uprising, result in a greater increase in information transfer among blogs than among main-

<sup>1</sup>All data, code and additional results are available at the supporting website [1]: <http://snap.stanford.edu/infopath>.

stream media. Perhaps surprisingly, the amount of mainstream media and blogs among the most influential nodes for most topics or news events are comparable. However, we find that growing numbers of influential blogs on some topics or news events are often temporally correlated with large-scale social movements (*e.g.*, the Occupy Wall Street movement in Sept-Nov 2011).

**Further related work.** Previous methods for inferring diffusion networks [9, 10, 12, 20] also use a generative probabilistic model for modeling cascading processes over networks. NETINF [10] and MULTITREE [12] infer the network connectivity using submodular optimization. NETRATE [9] and CONNIE [20] infer not only the network connectivity but also transmission rates of infection or prior probabilities of infection using convex optimization. Moreover, there have been also attempts to model information diffusion without assuming the existence of an underlying network [33, 32].

However, to the best of our knowledge, all previous approaches to network inference assume the network and the underlying dynamics of the edges to be constant, *i.e.*, the network structure and the transmission rates of each edge do not change over time. Therefore, they consider the pathways over which information propagates to be time-invariant. The main contribution of this paper is to combine stochastic gradient and the diffusion model introduced in [9] to develop an efficient on-line network inference algorithm that provides time-varying estimates of the edges of a network and the transmission rates of each edge. This allows us to detect how information pathways emerge and vanish over time, and identify when nodes produce highly viral content.

The remainder of the paper is organized as follows: in Sec. 2, we revisit the model of diffusion and state the dynamic network inference problem. Section 3 describes the proposed time-varying network inference method, called INFOPATH. Section 4 evaluates INFOPATH quantitatively and qualitatively using synthetic and real diffusion data. We conclude with a discussion of results in Section 5.

## 2. PROBLEM FORMULATION

In this section, we build on our fully continuous time model of diffusion [9, 11]. We start by briefly describing the generative model for the observed data. We then revisit how to compute the likelihood of a cascade using the model and state the continuous time network inference problem for both static and dynamic networks. Across the section, we explicitly point out which assumptions of the original model need to be extended in order to support dynamic networks.

**Observed data.** For now let’s consider a single static directed network. Over the edges of the network multiple contagions propagate. As the contagion spreads from infected to non-infected nodes over the edges of the network the contagion creates a *cascade*. For each contagion  $c$ , we observe a cascade  $\mathbf{t}^c$ , which is simply a record of observed node infection times during a time window of length  $T^c$ . In an information propagation setting, each cascade corresponds to a different piece of information and the infection time of a node is simply the time when the node first mentioned the piece of information  $c$ .

Cascade is a  $N$ -dimensional vector  $\mathbf{t}^c := (t_1^c, \dots, t_N^c)$  recording the times when each of  $N$  nodes got infected by the contagion  $c$ :  $t_k^c \in [t_0, t_0 + T^c] \cup \{\infty\}$ , where  $t_0$  is the infection time of the first node. Generally contagions do not infect all the nodes of the network, symbol  $\infty$  is used for nodes that were not infected by the contagion  $c$  during the observation window  $[t_0, t_0 + T^c]$ . Contagions often propagate simultaneously [21, 25] over the same net-

| Model | Log survival function                                     | Hazard function                          | Cascade gradient for uninfected             | Cascade gradient for infected   |
|-------|---|--|---|---|
|       | $\log S(t_i t_j; \alpha_{j,i})$                           | $H(t_i t_j; \alpha_{j,i})$               | $\nabla_{\alpha_{j,i}} L_c(\mathbf{A})$     | $\nabla_{\alpha_{j,i}} L_c(\mathbf{A})$   |
| EXP   | $-\alpha_{j,i}(t_i - t_j)$                                | $\alpha_{j,i}$                           | $T - t_j^c$                                 | $(t_i^c - t_j^c) - \frac{1}{\sum_{k:t_k^c < t_i^c} \alpha_{k,i}}$   |
| POW   | $-\alpha_{j,i} \log\left(\frac{t_i - t_j}{\delta}\right)$ | $\alpha_{j,i} \cdot \frac{1}{t_i - t_j}$ | $\log\left(\frac{T - t_j^c}{\delta}\right)$ | $\log\left(\frac{t_i^c - t_j^c}{\delta}\right) - \frac{(t_i^c - t_j^c)^{-1}}{\sum_{k:t_k^c < t_i^c} \alpha_{k,i} (t_i^c - t_k^c)^{-1}}$ |
| RAY   | $-\alpha_{j,i} \frac{(t_i - t_j)^2}{2}$                   | $\alpha_{j,i} \cdot (t_i - t_j)$         | $\frac{(T - t_j^c)^2}{2}$                   | $\frac{(t_i^c - t_j^c)^2}{2} - \frac{t_i^c - t_j^c}{\sum_{k:t_k^c < t_i^c} \alpha_{k,i} (t_i^c - t_k^c)}$                               |

**Table 2: Contagion transmission models for the three edge transmission likelihoods: Exponential, Power-law and Rayleigh.**

work but we assume each contagion to propagate independently of each other.

Given a set of node infection times of many different contagions, our goal is to infer the underlying dynamic network over which contagions propagated. We apply the Maximum Likelihood principle in order to infer the network that most likely generated the observed data. We proceed by assuming a static network and describe the generative model of information diffusion. We then generalize the model to dynamic networks.

**Pairwise transmission likelihood.** The first step in modeling diffusion dynamics is to consider pairwise node interaction. For every pair of nodes  $(j, i)$ , we define a pairwise transmission rate  $\alpha_{j,i}$  which models how frequently information spreads from node  $j$  to node  $i$ ; the *strength* of an edge  $(j, i)$ . We pay attention to the rather general case of heterogeneous pairwise transmission rates, *i.e.*, infections can occur at different transmission rates over different edges of the network. As  $\alpha_{j,i} \rightarrow 0$  the expected transmission time from node  $j$  to node  $i$  becomes arbitrarily long. In contrast with the original model [9], we will later allow transmission rates  $\alpha_{j,i}$  to change over time. In particular, we will allow the transmission rates  $\alpha_{j,i}$  to change across cascades but not within a cascade. Allowing edge transmission rates to dynamically increase and decay over time will enable us to infer time-varying diffusion networks.

Next, we define  $f(t_i|t_j; \alpha_{j,i})$  to be the conditional likelihood of transmission between node  $j$  and node  $i$  and assume it depends on the infection times  $(t_j, t_i)$  and the edge transmission rate  $\alpha_{j,i}$ . We allow information to only propagate forward in time, *i.e.*, node  $j$  that has been infected at time  $t_j$  may infect node  $i$  at time  $t_i$  only if  $t_j < t_i$ , otherwise  $f(t_i|t_j, \alpha_{j,i}) = 0$ . The shape of the conditional likelihood of transmission may depend on the particular setting (information, influence, diseases, etc.) in which propagation takes place. In some scenarios, it may be possible to estimate a non-parametric likelihood while in others, expert knowledge may be used to decide upon a parametric model. For simplicity, we consider three well-known parametric models of edge transmission rates: exponential, power-law, and Rayleigh, defined in Table 1. Exponential and power-law likelihoods have been used in modeling information propagation in social and information networks [9, 10, 11, 12, 20], while Rayleigh has been used in previous work in disease spread in epidemiology [31]. In all three models, as  $\alpha_{j,i} \rightarrow 0$ , the likelihood of infection tends to zero.

We recall some additional standard notation [14] that we will use in the remainder of the section. Given some node  $j$ , infected at time  $t_j$ , we define the *survival function* of edge  $j \rightarrow i$  as  $S(t_i|t_j; \alpha_{j,i}) = 1 - F(t_i|t_j; \alpha_{j,i})$  where  $F(t_i|t_j; \alpha_{j,i}) = \int_{t_j}^{t_i} f(t|t_j; \alpha_{j,i}) dt$  is the cumulative transmission density function, computed from the transmission likelihood. Finally, the *hazard function*, or instantaneous infection rate, of edge  $j \rightarrow i$  is the ratio  $H(t_i|t_j; \alpha_{j,i}) = f(t_i|t_j; \alpha_{j,i}) / S(t_i|t_j; \alpha_{j,i})$ . We derive the log survival and hazard functions for the three edge transmission models in Table 2.

**Likelihood of a cascade.** Consider some node  $i$  in a directed network. Node  $i$  can get infected by any of its parents (*i.e.*, nodes pointing to  $i$ ). Once infected, node  $i$  can then also spread the contagion to its children (*i.e.*, nodes  $i$  points to). As in the independent cascade model [13], we assume that node gets infected once the *first* parent infects it (*i.e.*, a node can get infected only once). Then, the likelihood of infection of node  $i$  at time  $t_i$  given a collection of previously infected nodes  $(t_1, \dots, t_N | t_N \leq t_i)$  results from summing over the likelihoods of the mutually disjoint events that each node is the first parent that generated the infection event of our node  $i$ :

$$f(t_i|t_1, \dots, t_N \setminus t_i; \mathbf{A}) = \sum_{j:t_j < t_i} f(t_i|t_j; \alpha_{j,i}) \times \prod_{j:j \neq k, t_k < t_i} S(t_i|t_k; \alpha_{k,i}), \quad (1)$$

where  $\mathbf{A} := \{\alpha_{j,i} | i, j = 1, \dots, n, i \neq j\}$ . If we assume that infections are conditionally independent given the parents of the infected nodes, the likelihood of the infections in a cascade is:

$$f(\mathbf{t}^{\leq T^c}; \mathbf{A}) = \prod_{t_i \leq T} \sum_{j:t_j < t_i} f(t_i|t_j; \alpha_{j,i}) \times \prod_{j:j \neq k, t_k < t_i} S(t_i|t_k; \alpha_{k,i}), \quad (2)$$

where  $\mathbf{t}^{\leq T^c}$  denotes the vector of infected nodes in the cascade up to  $T^c$ . Removing the condition  $k \neq j$  makes the product independent of  $j$ ,

$$f(\mathbf{t}^{\leq T^c}; \mathbf{A}) = \prod_{t_i \leq T} \prod_{k:t_k < t_i} S(t_i|t_k; \alpha_{k,i}) \times \sum_{j:t_j < t_i} \frac{f(t_i|t_j; \alpha_{j,i})}{S(t_i|t_j; \alpha_{j,i})}. \quad (3)$$

The fact that some nodes are *not* infected during the observation window is also informative. We therefore add multiplicative survival terms to Eq. 3 and rearrange with hazard functions:

$$f(\mathbf{t}; \mathbf{A}) = \prod_{i:t_i \leq T} \prod_{m:t_m > T} S(T|t_i; \alpha_{i,m}) \times \quad (4)$$

$$\prod_{k:t_k < t_i} S(t_i|t_k; \alpha_{k,i}) \left( \sum_{j:t_j < t_i} H(t_i|t_j; \alpha_{j,i}) \right). \quad (5)$$

Perhaps surprisingly, our continuous time model of diffusion is a particular case of Aalen's additive regression model, frequently used in survival theory analysis [3]. In Aalen's model, the hazard function, or instantaneous infection rate, of node  $i$  is parametrized as  $\alpha_{i,0}(t) + \alpha(t)^T \mathbf{s}_i(t)$ , where  $\alpha(t)$  is a vector that accounts for the effect of a collection of observable covariates  $\mathbf{s}(t)$  and  $\alpha_{i,0}(t)$  is a baseline. It is easy to show that the hazard function of node  $i$  at time  $t_i$  for the three pairwise transmission models: exponential, power-law and Rayleigh, has the following form:

$$H(t_i|t_1, \dots, t_N \setminus t_i; \mathbf{A}) = \alpha_i^T \mathbf{s}_i(t_i; t_1, \dots, t_N \setminus t_i), \quad (6)$$

---

**Algorithm 1** INFOPATH: the dynamic network inference algorithm

---

**Require:**  $C, t, K, T, \rho$ **while**  $k < K$  **do** $c_k \leftarrow$  cascade-sampling( $C, t, T$ );**for all**  $(j, i) : t_j^{c_k} < t_i^{c_k}$  **do**

$$\alpha_{j,i}^k = \left( \alpha_{j,i}^{k-1} - \gamma_k \nabla_{\alpha_{j,i}} L_{c_k}(\mathbf{A}^{k-1}) \right)^+;$$

**for all**  $(j, i) : \alpha_{j,i}^{k-1} > 0, t_j^{c_k} \rightarrow \infty$  **do**

$$\alpha_{j,i}^k = \rho \alpha_{j,i}^{k-1};$$

 $k = k+1;$  $\mathbf{A}^* \leftarrow \mathbf{A}^{K-1};$ **return**  $\mathbf{A}^*$ ;

---

where  $\alpha_i = (\alpha_{1,i}, \dots, \alpha_{N,i})$  accounts for the effect of a collection of observable covariates  $\mathbf{s}_i(t_i; t_1, \dots, t_N \setminus t_i)$ , the covariates depend on the pairwise transmission model (exponential, power-law or Rayleigh) and the previously infected nodes, and the baseline is zero.

**Dynamic network inference problem.** Given a static network with constant edge transmission rates  $\alpha_{j,i}$ , the network inference problem reduces to solving a maximum likelihood problem over set of recorded cascades  $C$  [9]:

$$\begin{aligned} & \text{maximize}_{\mathbf{A}} \quad \sum_{c \in C} \log f(\mathbf{t}^c; \mathbf{A}) \\ & \text{subject to} \quad \alpha_{j,i} \geq 0, i, j = 1, \dots, N, i \neq j, \end{aligned} \quad (7)$$

where  $\mathbf{A} := \{\alpha_{j,i} | i, j = 1, \dots, n, i \neq j\}$  are the edge transmission rates we aim to infer. The edges of the network are those pairs of nodes with transmission rates  $\alpha_{j,i} > 0$ .

Now we generalize the network inference problem to dynamic networks with edge transmission rates  $\alpha_{j,i}(t)$  that may change over time. To this aim, at any given time  $t$ , we solve a maximum weighted likelihood problem over the set of recorded cascades by time  $t$ ,  $C_t = \{\mathbf{t}^1, \dots, \mathbf{t}^{|C_t|}\}$ :

$$\begin{aligned} & \text{maximize}_{\mathbf{A}(t)} \quad \sum_{c \in C_t} w_c(t) \log f(\mathbf{t}^c; \mathbf{A}(t)) \\ & \text{subject to} \quad \alpha_{j,i}(t) \geq 0, i, j = 1, \dots, N, i \neq j, \end{aligned} \quad (8)$$

where  $w_c(t) \geq 0$  is a weight that penalizes the importance of cascade  $c$  based on how old it is at time  $t$  and  $\mathbf{A}(t) := \{\alpha_{j,i}(t) | i, j = 1, \dots, n, i \neq j\}$  are the variables. The intuition here is that diffusion network smoothly changes over time and that recent cascades have higher importance in determining current network structure than old cascades. Thus, at any point in time, we can solve the above optimization problem to obtain the structure of the diffusion network at that particular time. Next, we show how to efficiently solve the above optimization problem for all time points  $t$ .

### 3. THE INFOPATH ALGORITHM

The problem defined by equation Eq. 8 is convex for the three transmission models we consider. Therefore we can aim to find the unique optimal solution at any given time point  $t$ :

**THEOREM 1** ([9]). *Given log-concave survival functions and concave hazard functions in the parameter(s) of the pairwise transmission likelihoods, the network inference problem defined by equation Eq. 8 is convex in  $\mathbf{A}$ .*

Stochastic gradient (SG) methods have been shown to be extremely successful for taking advantage of the structure exhibited by the optimization problem stated in Eq. 8. They have received increasing attention in the machine learning literature [4, 5, 7, 8, 29]. Although many optimization methods based on stochastic gradient descent have been proposed, we have found that in practice

the basic projected stochastic gradient method [26] works well for our problem. Other more sophisticated methods, like the stochastic average gradient [29] or incremental average gradient [7] do not offer a significant advantage. Therefore, we proceed with the basic stochastic gradient method in the remainder of the paper.

**Projected Stochastic Gradient.** The projected stochastic gradient method [26] uses iterations of the form:

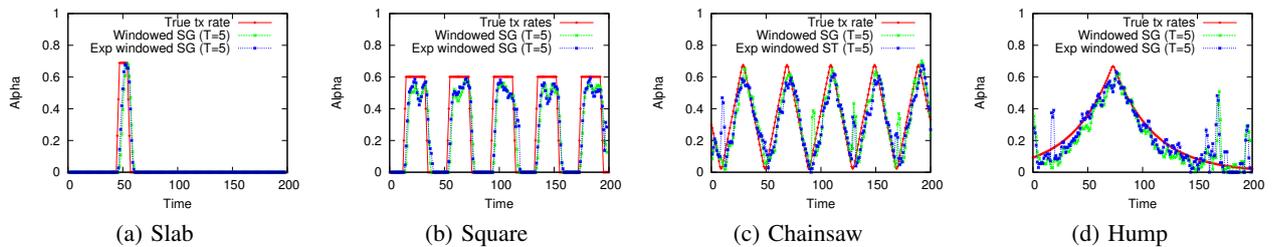
$$\alpha_{j,i}^k(t) = \left( \alpha_{j,i}^{k-1}(t) - \gamma_k \nabla_{\alpha_{j,i}} L_{c_k}(\mathbf{A}^{k-1}(t)) \right)^+, \quad (9)$$

where  $\nabla_{\alpha_{j,i}} L_{c_k}(\cdot)$  is the gradient of the log-likelihood  $L_c(\cdot)$  with respect to the edge transmission rate  $\alpha_{j,i}$ ,  $\gamma_k$  is a step-size,  $(z)^+ = \max(0, z)$ , and cascade  $c_k$  is sampled (with replacement) from  $C_t$ . The gradients for all three edge transmission models are given in Table 2.

Note that instead of using all historic data and then explicitly penalizing each cascade by a different weighting factor  $w_c(t)$ , we use a different, more scalable approach. We sample cascades with replacement, where the probability of a cascade being sampled decays with the age of the cascade. This way recent cascades get sampled more often and thus implicitly hold higher importance when inferring the network. In practice, we achieve a significant speed up by using this approach. Moreover, in our dynamic network inference problem, the edge transmission rates usually vary smoothly. This means that stochastic gradient descent is a perfect method as we can use the inferred network from the previous time step as initialization for the inference procedure in the current time step. We find that setting the starting point  $\alpha_{j,i}^0$  of each edge transmission rate  $\alpha_{j,i}$  to the last outputted estimate of the transmission rate allows us to further speed up the algorithm.

Importantly, in each iteration  $k$  of the projected stochastic gradient method, we only need to compute the gradients  $\nabla_{\alpha_{j,i}} L_{c_k}(\mathbf{A}^k)$  for edges  $(j, i)$  such that node  $j$  has been infected in cascade  $c_k$ , and the iteration cost and convergence rate are independent of  $|C|$  [5, 23]. Rigorous theoretical analysis of convergence turns out to be a challenging problem which we leave for future work. However, we point out that standard analyses [26] typically assume the gradients  $\nabla_{\mathbf{A}} L_c(\mathbf{A}^k)$  to be either bounded above by a constant  $M$ , where  $\|\nabla_{\mathbf{A}} L_c(\mathbf{A})\| \leq M$ , or Lipschitz-continuous with constant  $L$ ,  $\|\nabla_{\mathbf{A}} L_c(\mathbf{A}_2) - \nabla_{\mathbf{A}} L_c(\mathbf{A}_1)\| \leq L \|\mathbf{A}_2 - \mathbf{A}_1\|$ . In our case, these conditions are violated if at any iteration  $k$ , there is a node  $i$  infected in cascade  $c_k$  such that  $H(t_i^{c_k} | t_j^{c_k}, \alpha_{j,i}^{k-1}) = 0 \forall j : t_j^{c_k} < t_i^{c_k}$ , i.e., node  $i$  has no parents that *explain* the infection at  $t_i^{c_k}$ , and the objective function is positively unbounded. In practice, we obtain a good performance and avoid such scenario by bounding below each *feasible* transmission rate,  $\alpha_{j,i} \geq \epsilon$ . An edge transmission rate  $\alpha_{j,i}$  is *feasible* if there is at least one cascade in which both node  $j$  and  $i$  get infected. When outputting the final solution, we simply omit edges with transmission rates  $\epsilon$ .

**Aging edges.** Suppose we solve the dynamic network inference problem for a given time  $t$  using the projected SG method. In each iteration  $k$ , we only update edge transmission rate  $\alpha_{j,i}^k$  if node  $j$  has been infected in cascade  $c_k$ . Now, suppose a few edge transmission rates  $\alpha_{j,i}^0$  are greater than zero for a given node  $j$ , i.e., their last outputted estimates before  $t$  have been positive. Then, if node  $j$  turns to be infected in many sampled cascades at time  $t$  but it never transmits information to any of its neighbors  $i$  in the future, the edge transmission rates  $\alpha_{j,i}^k$  will eventually converge to zero for large  $k$ . However, if node  $j$  is never infected for any of the future cascades, then none of edge transmission rates  $\alpha_{j,i}$  will be updated, and they will remain positive. So, if node  $j$  is never infected in subsequent future cascades, the transmission rates  $\alpha_{j,i}$  will remain positive for-



**Figure 1: True and inferred edge transmission rates for edges with different 4 transmission rate evolution patterns: (a) Slab, (b) Square, (c) Chainsaw, (d) Hump. Results are for the Kronecker core-periphery with exponential edge transmission model for 200 time units with 1,000 cascades per time unit. Our INFOPATH method is able to track the evolving edge transmission rates over time. INFOPATH works better for continuously evolving edge transmission rates (c, d).**

ever. However, we would like these *unused* edges  $(j, i)$  to decay and eventually vanish, or equivalently the transmission rates  $\alpha_{j,i}$  to converge to zero. To achieve this, we multiply transmission rates of *unused* edges by *aging* factor  $\rho$  every time we solve the dynamic network inference problem. We use  $\rho = 0.95$  in all experiments.

**Cascade sampling.** In Eq. 9, instead of sampling cascades uniformly at random and explicitly penalizing each cascade by a different weighting factor  $w_c(t)$ , we achieve a significant speed up by sampling cascades using a sampling procedure that penalizes old cascades and considers  $w_c(t) = 1$  for all cascades. There are many possible sampling schemes. In practice, and for simplicity, we use windowed uniform sampling or windowed exponential sampling. Windowed means that when solving the network inference problem for time  $t$ , we only sample (uniformly or exponentially) cascades that have started in a sampling time window  $(t - T, T)$ . Here, we find an important tradeoff. The shorter is the sampling time window  $T$  in the stochastic gradient descent, the quicker our algorithm is tracking changes in the edge transmission rates. However, short sampling time window results in less reliable estimates because of sampling from a smaller universe of distinct cascades. Therefore, in order to be able to track changes quickly, we would need to observe many cascades over time.

## 4. EXPERIMENTAL EVALUATION

We evaluate the performance of INFOPATH on time-varying synthetic networks that mimic the structure of real networks as well as on a dataset of more than 179 million information cascades extracted from 300 million blogs and news articles from 3.3 million media sites over a period of one year, from March 2011 till February 2012. All the data, code and additional results are available at the supporting website [1].

### 4.1 Experiments on synthetic data

The goal of the experiments with synthetic data is to understand how temporal changes in a network affect the performance of our algorithm. We aim to detect not only when an edge appears (*i.e.*, its transmission rate becomes  $> 0$ ) or disappears (*i.e.*, its transmission rate becomes 0) but also provide instantaneous transmission rate estimates that track the true edge transmission rates over time.

**Experimental setup.** First, we generate synthetic networks using Kronecker graph models of directed real-world networks [17]. For all our experiments, we consider two different Kronecker networks, both with 1,024 nodes and 2,048 edges: A *core-periphery* Kronecker network with parameter matrix  $[0.9, 0.5; 0.5, 0.3]$  and a *hierarchical* Kronecker network with parameters  $[0.9, 0.1; 0.1, 0.9]$ .

The next step is to make each edge to follow a particular edge transmission rate evolution pattern. Our goal later will be to recover

the network as well as the evolution of the transmission rate of each individual edge.

We consider five edge evolution patterns: Slab, Square, Chainsaw, Hump and constant (see Figure 1). Slab and Hump patterns model outgoing connections of sites that become popular for a short period of time. Square and Chainsaw patterns model incoming connections to sites that perform updates periodically at specific times of the day or days of the week. Constant pattern represents connections between sites that interact at any time and during a long period of time, usually large media sites. We consider Chainsaw, Hump and Continuous to be examples of *Type I* pattern, without discontinuities, and Slab and Square to be examples of *Type II* pattern, with discontinuities.

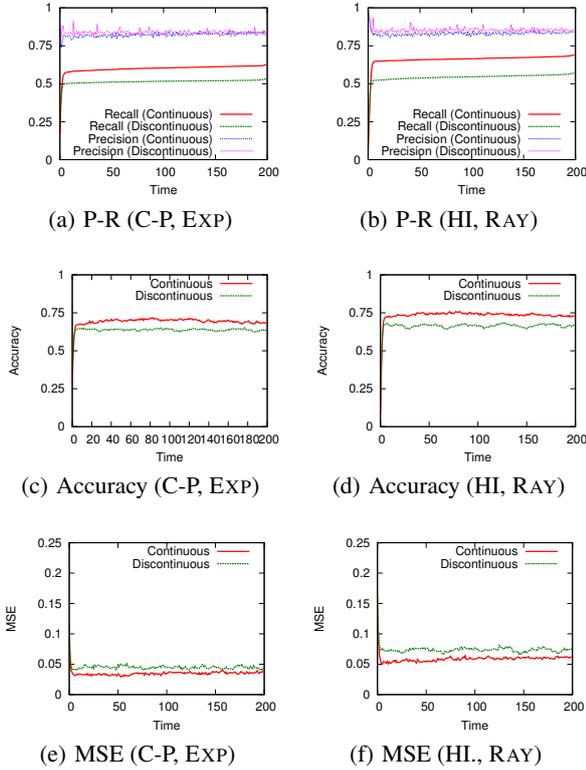
We assign to each edge in the network an evolution pattern chosen uniformly at random from the set of the above 5 patterns. Then, we generate transmission rate values  $\alpha_{j,i}^*(t)$  for each edge according to its chosen evolution pattern. The evolving edge transmission rate  $\alpha_{j,i}^*(t)$  models how quickly information spreads from one node to another. Finally, we generate 1,000 information cascades per time step. For each cascade we randomly pick the cascade initiator node.

Given the node infection times from the recorded cascades, our goal then is to find the true edges of the network and for each edge discover its transmission rate evolution pattern. In other words, inferring how each edge transmission rate  $\alpha(t)$  evolves over time. Figure 1 shows the true and inferred edge transmission rates for four different edges, each with a different evolution pattern: Slab, Square, Chainsaw and Hump. Observe that INFOPATH is able to track the evolving edge transmission rates over time for all evolution patterns. INFOPATH gives near perfect performance when edge transmission rate evolves continuously (Chainsaw, Hump). Interestingly, even when the edge transmission rate evolves discontinuously (Slab, Square), INFOPATH manages to track it.

**Accuracy of INFOPATH.** We evaluate the INFOPATH method quantitatively by computing four different measures for every time step: Precision, Recall and Accuracy of inferred edges as well as the Mean Squared Error (MSE) in the edge transmission rate. Precision at time  $t$  is the fraction of edges in the inferred network  $\hat{G}(t)$  present in the true network  $G^*(t)$ . Recall at time  $t$  is the fraction of edges of the true network  $G^*(t)$  present in the inferred network  $\hat{G}(t)$ . And accuracy at time  $t$  is defined as

$$1 - \frac{\sum_{i,j} |I(\alpha_{i,j}^*(t)) - I(\hat{\alpha}_{i,j}(t))|}{\sum_{i,j} I(\alpha_{i,j}^*(t)) + I(\hat{\alpha}_{i,j}(t))},$$

where  $\alpha^*(t)$  is the true transmission rate at time  $t$ ,  $\hat{\alpha}(t)$  is the estimated transmission rate at time  $t$ , and  $I(\alpha(t)) = 1$  if  $\alpha(t) > 0$  and  $I(\alpha(t)) = 0$  otherwise. Inferred networks with no edges or only false edges would have zero accuracy. Last, Mean Squared Error



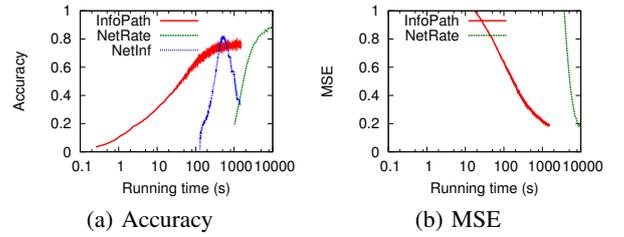
**Figure 2: Precision and Recall (P-R), Accuracy and Mean Squared Error (MSE) of our INFOPATH method against time. (a,c,e): Core-periphery (C-P) Kronecker network with exponential edge transmission model (b, d, f), and Hierarchical (HI) Kronecker network with Rayleigh edge transmission model. Performance on Type I (Chainsaw, Hump) and Type II (Slab, Square) edge transmission rate evolution patterns is plotted.**

(MSE) at time  $t$  is defined as  $E[|\alpha^*(t) - \hat{\alpha}(t)|^2]$ , where  $\alpha^*(t)$  is the true edge transmission rate at time  $t$  and  $\hat{\alpha}(t)$  is the estimated transmission rate.

Figure 2 shows Precision, Recall, Accuracy, and MSE over time for the time-varying core-periphery Kronecker network with exponential edge transmission model, and hierarchical Kronecker network with Rayleigh edge transmission model. Observe that the performance of our method is stable across time, and as mentioned before, continuous evolution patterns are easier to track and estimate than discontinuous ones.

**Accuracy vs. running time in static networks.** Our stochastic gradient descend based method, INFOPATH, can be also used to speed-up inference of static networks. In such scenario, stochastic gradient descend processes cascades in a random round-robin fashion. Here, we compare INFOPATH to the state of the art methods for inference of static networks: NETINF [10] and NETRATE [9]. First, we compare the methods by computing the accuracy against running time. Second, we compare INFOPATH to NETRATE in terms of mean squared error of the estimated transmission rates against the running time. We omit NETINF from this last comparison since it only infers the network structure (and no edge transmission rates). For the sake of fairness in the running time comparison we implemented all methods in C++. Our C++ implementation of NETRATE is much faster than the public Matlab implementation.

Figure 3 compares Accuracy and MSE against running time for the three network inference methods for a static core-periphery



**Figure 3: Accuracy and Mean Squared Error (MSE) against running time for a 1,024 node, 2,048 edge time-invariant core-periphery Kronecker network with power-law edge transmission model and 5,000 cascades. Longer running times mean the algorithms run for more iterations. INFOPATH and NETRATE improve accuracy until convergence. However, INFOPATH achieves the same level of performance 10-100 times faster.**

Kronecker network. INFOPATH is about 10 to 100 times faster than NETRATE and as fast as NETINF, while achieving the same accuracy as NETRATE. Importantly, INFOPATH and NETRATE always improve accuracy with the running time, until convergence. In terms of MSE, INFOPATH achieves lower MSE values much quicker than NETRATE.

## 4.2 Experiments on real data

The emergence of specific information pathways often depends on the information content of the news that is propagating [22, 28]. For example, a real world event may occur for a limited period of time and thus news related to the event spread quicker and to larger parts of the network around such time period. At any given time, there are many different real world events, topics, and content that propagates through the Web, leading to different emerging and vanishing information pathways, and thus an underlying time-varying network. In order to better understand these temporal changes, we aim to reconstruct time-varying networks and the information pathways for particular real world events and topics.

**Dataset description.** We experiment with more than 300 million blog posts and news articles collected from 3.3 million websites over a period of one year, from March 2011 till February 2012. We trace the flow of information using *memes* [16]. *Memes* are a short textual phrases (like, “lipstick on a pig”) that travel through the Web. We consider each meme  $m$  as a separate information cascade  $c_m$ . Since all documents which contain memes are time-stamped, a cascade  $c_m$  is simply a record of times when sites first mentioned meme  $m$ . We extracted more than 179 million memes, longer than four words. Out of these, 34 million distinct memes appeared at least twice, resulting in 34 million different information cascades.

**Experimental setup.** Our aim is to consider sites that actively spread memes over the Web. We achieve this by selecting top 5,000 sites in terms of the number of memes they mentioned. Moreover, we are interested in inferring dynamic networks related to particular topics or events. So, we assume we are also given a keyword query  $Q$  related to the event/topic of interest. When inferring a network for a given query  $Q$ , we only consider documents (and the memes they mention) that include keywords  $Q$ . Then, we build information cascades using only those memes and apply the INFOPATH algorithm to infer the edges and evolving edge transmission rates. The edge transmission rates explain the propagation of information related to a given topic or real world event  $Q$ . For each query  $Q$  we infer one network per day. Table 3 summarizes

| Topic or news event (Q) | # sites | # memes   |
|-------------------------|---------|-----------|
| Amy Winehouse           | 1,207   | 109,650   |
| Fukushima               | 1,666   | 383,745   |
| Gaddafi                 | 1,358   | 440,646   |
| Kate Middleton          | 1,427   | 191,777   |
| NBA                     | 2,087   | 1,543,630 |
| Occupy                  | 1,875   | 655,183   |
| Strauss-Kahn            | 1,263   | 204,238   |
| Syria                   | 1,565   | 615,176   |

**Table 3: Topic and news event statistics.**

the number of sites and meme cascades for several topics and real world events.<sup>2</sup>

**Implementation and scalability.** We developed an efficient distributed implementation of our INFOPATH algorithm in C++ based on the network analysis library SNAP [2]. We deployed the implementation in a cluster with 1,000 CPU cores and 6 TB of RAM. With this setup, we considered 38 different topics/events  $Q$ . For each topic, we inferred a time-varying network with a daily temporal resolution for a period of one year, from March 2011 to February 2012. Each network has thousands of nodes and is based on hundreds of thousands of cascades. Inferring 38 different time-varying networks took less than 4 hours on our cluster. Note that this is equivalent to solving Eq. 8 more than 13,000 times ( $38 \times 365$ ) for millions of pairwise transmission rates. We also tested our algorithm on larger datasets. For example, for “Occupy Wall Street movement”, we were able to infer a 43,415-node time-varying network over a period of 18 months, from January 2011 to June 2012, using 1,381,793 information cascades.

**Visualizing the information pathways.** Figure 4 plots diffusion networks for three different 2011 world events: Fukushima nuclear disaster, UK royal wedding, and civil uprising in Syria. Each network is shown at three different time points. Red nodes represent mainstream media sites, and blue nodes represent blogs [16].

Based on the figure, we draw several interesting observations. Most often, information propagates through a core-periphery network structure. Such structure emerges by few central media sites and blogs driving the adoption of memes across the Web [10]. However, the network structure often changes dramatically over time, and we find clusters that emerge and vanish in short periods of time. For example, the information networks for Syria’s uprising illustrated in Figures 4(g-h), do not have any clear clustering structure. However, on December 2, 2011 (Figure 4(i)) a cluster suddenly emerges in the network. Further investigation reveals that the cluster is composed of UK news sites and blogs that discuss recently implemented EU sanctions against Syria. Generally, it is common to observe sudden formation of clusters of sites from specific geographical areas. This is specially noticeable in the information network for Fukushima’s disaster, in Figures 4(a-c). Such clusters often form due to language boundaries, since such boundaries prevent memes to flow across countries or continents. Moreover, we often observe that such clusters are caused by a common external event [22], like in the case of UK discussion on EU sanctions against Syria. Inferred dynamic networks can thus be used to investigate the flow of information as well as to detect external events that cause sudden perturbations to the diffusion network structure.

<sup>2</sup>Additional time-varying diffusion networks for other topics and news events are available at the supporting website [1]

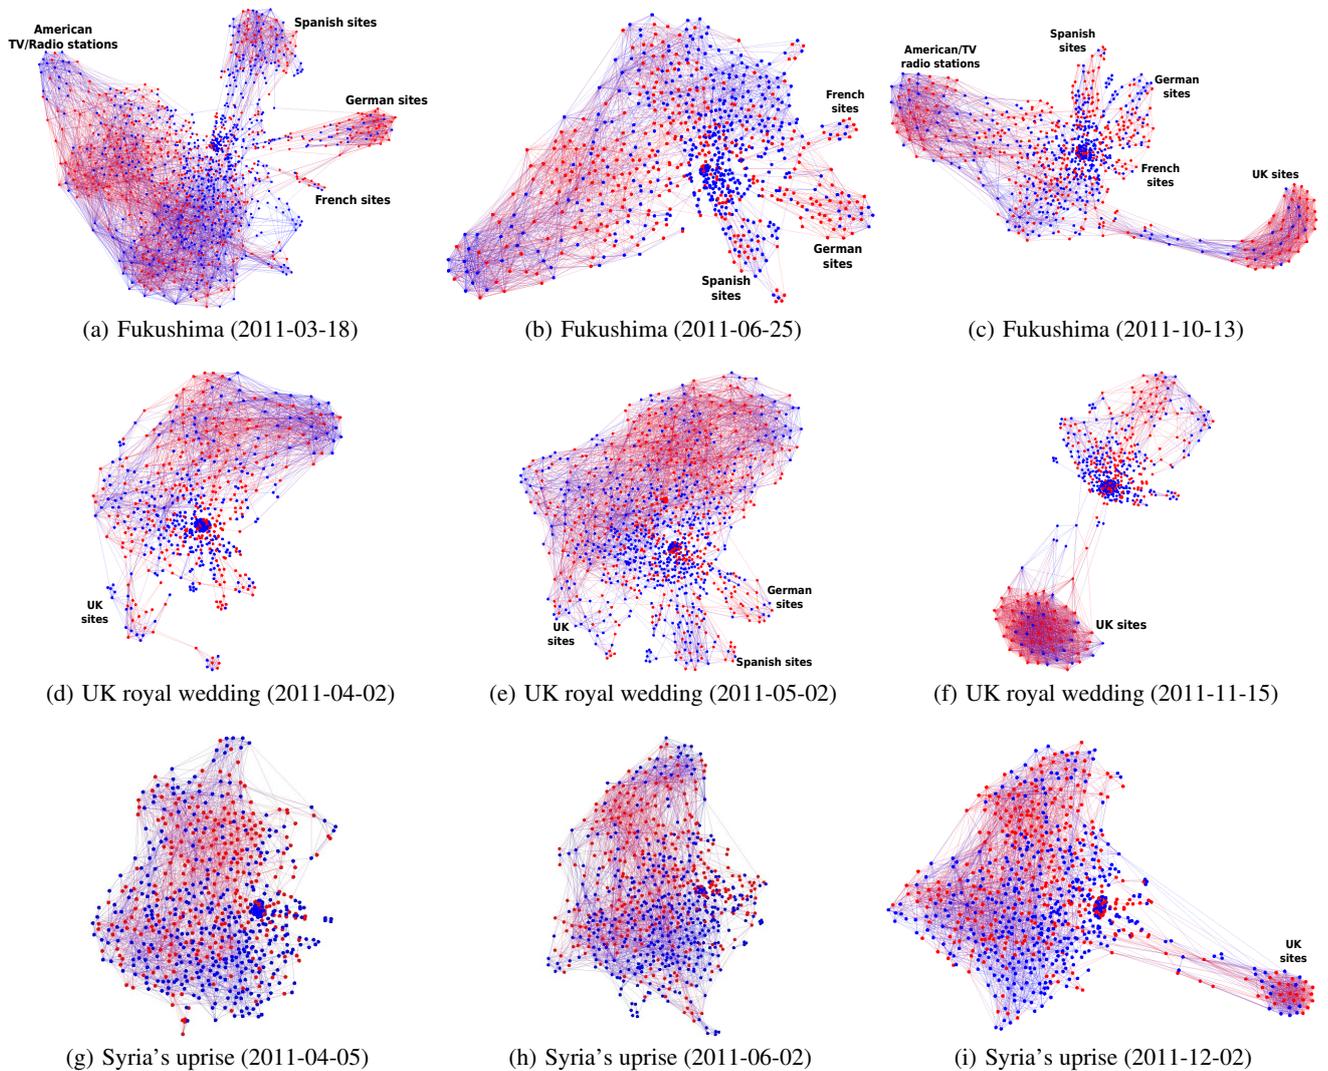
**Evolution of edge transmission rates.** Next, we aim to study the evolution of links among different types of sites. We label the nodes in our network as mainstream media and blog, and compute the number of links between different types of sites over time. Figure 5 gives the results for several inferred diffusion networks for different topics and world events. We note several interesting patterns.

The connectivity changes tend to reflect the amount of attention that a news event or a topic triggers over time. Unexpected news events, like the sex scandal of the director of the International Monetary Fund Strauss-Kahn on May 14, 2011 in Fig. 5(g) or the death of British singer Amy Winehouse on July 23, 2011 in Fig. 5(a), result in a dramatic increase in the number of edges over a short period of time. More general topics, like the NBA in Fig. 5(e), result in a network with more stable connectivity over time. Certain types of news are sometimes spreading earlier among blogs than mainstream media. This is especially the case for population wide events like the Fukushima nuclear disaster, civil war in Libya and civil uprising in Syria (Fig. 5(b, c, h)). However, it happens more frequently that the largest amount of links are mainstream media-to-mainstream media and the fewest links point from blogs to mainstream media. These results are intuitive and consistent with previous work [10, 16] that observed most often information flows from mainstream media to blogs (and rarely the other way around). However, as we see here for population level events and social movements (like, in case of the civil unrest in the Middle East) social media plays crucial role in information dissemination and organization of civil movements.

**Evolution of node centrality.** Having studied the dynamics of edges in the network we now move towards investigating the network centrality of blogs and mainstream media sites over time for different topics and world events. To measure network centrality of node  $S$  in the network at time  $t$ , we first compute shortest path length from  $S$  to any other node  $R$  in the network. Then centrality of node  $S$  is defined as  $\sum_R 1/d(S, R)$ , where  $d(S, R)$  is the shortest path length from  $S$  to  $R$  (if  $R$  is not reachable from  $S$  then  $d(S, R) = \infty$ ). For networks with core-periphery structure, nodes with high centrality are typically located in the “central” core of the network.

Figure 6 plots the percentage of blogs among the top 100 most central sites over time for eight different topics/events of 2011. Perhaps surprisingly, we observe there is about the same number of mainstream media and blogs in the top-100 most central nodes for most networks – the number of blogs in the top-100 does not typically decrease below 30% or increase over 70%. For some topics, mainstream media are always more *central* (e.g., baseball and NBA in Figures 6(a, b)). In contrast, for other topics, blogs dominate mainstream media over a significant amounts of time (e.g., Gaddafi in Fig. 6(c)). Centrality of mainstream media and blogs can be relatively constant (Fig. 6(a,b)) or more time-varying (Fig. 6(c,h)). We find that a significant rise in the number of central blogs is often temporally correlated with an increasing social unrest (e.g., the Occupy Wall Street movement in Sept-Nov 2011 in Fig. 6(f)).

**Accuracy on real data.** So far, we have used *memes* to trace the flow of information over the Web and have made several qualitative observations about the structure and dynamics of information pathways in online media. We now proceed and attempt to also quantitatively evaluate INFOPATH on real data. In case of real data the ground-truth information diffusion network is impossible to obtain. However, we can use the temporal dynamics of hyperlinks created between news sites as a proxy for real information flow. Thus, by observing the times when sites create hyperlinks, our goal is to infer the ‘targets’ of the links (i.e., infer the hyperlink network from the hyperlinks times).



**Figure 4: Time-varying diffusion networks for three different major events of 2011. Red nodes are mainstream media, and blue nodes are blogs. Additional visualizations for other topics and events are available at the supporting website [1].**

We proceed as follows. First, we discretize the time in days, we generate one network  $G^*(t)$  per day  $t$ , in which we add an edge  $(u, v)$  if a document on a site  $u$  linked to a document on a site  $v$  within the last day. Then, we build a set of hyperlink cascades. A hyperlink cascade  $c_h$  starts when a site publishes a piece of information and then other sites use hyper-links to refer to it. Since all our documents/posts are time stamped, we can trace the hyper-links in the reverse direction and obtain information cascades. We extracted almost 0.5 million hyperlink cascades from 3.3 million websites from July 2011 till December 2012. Our aim is to use the hyperlink cascades to infer the time-varying network  $G^*(t)$ . We then evaluate how many edges INFOPATH estimates correctly by computing Accuracy, Precision and Recall for each day.

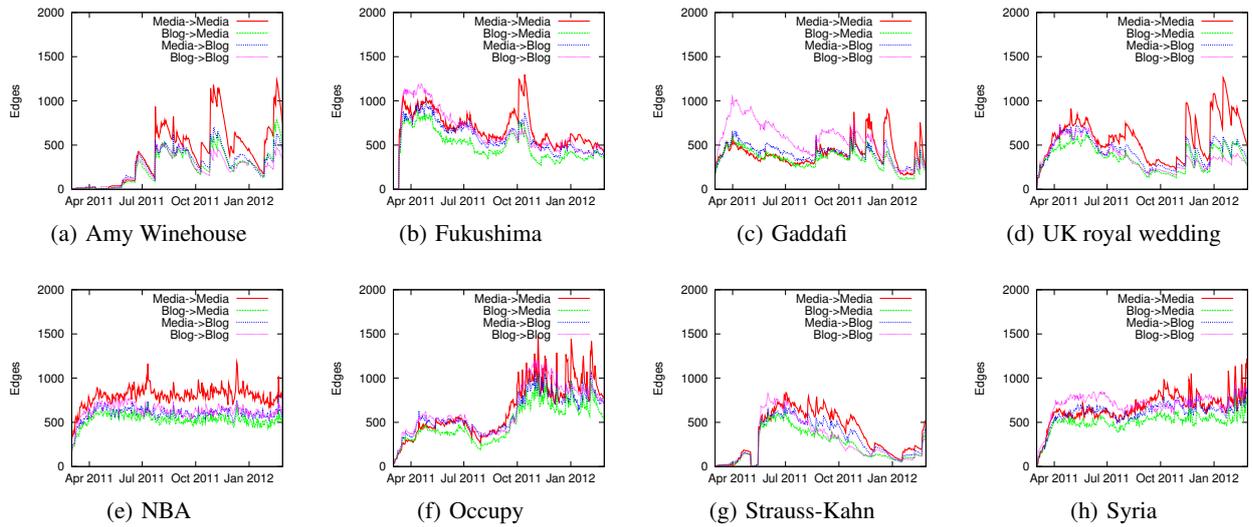
Figure 7 shows Precision, Recall, and Accuracy over time for a time-varying hyperlink network with 11,461 nodes and 19,915 edges created over time, using 495,655 hyperlink cascades from July 2011 to December 2011. We assume an exponential edge transmission model. We observe weekly periodicity and the overall encouraging performance of around 0.4 to 0.5 for all three performance metrics.

## 5. CONCLUSION

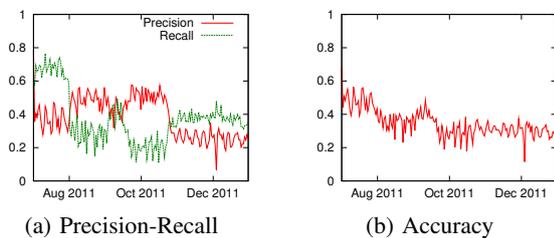
All previous network inference algorithms have assumed diffusion networks to be static. Therefore, they have considered the pathways over which information propagates to be static over time. In contrast, we developed an algorithm for time-varying network inference, INFOPATH. Our algorithm provides on-line time-varying estimates of the edges of the network as well as the dynamic edge transmission rates, which allows us to detect how information pathways emerge and vanish over time.

We evaluated our algorithm on synthetic data and demonstrated that INFOPATH successfully tracks changes in the topology of dynamic networks, provides accurate on-line estimates of the time-varying edge transmission rates and is also robust across network topologies, edge transmission models and patterns of evolution of edge transmission rates.

We also run INFOPATH on real data and investigated how real networks and information pathways evolve over time. We found that information pathways over which general recurrent topics propagate remain relatively stable across time. In contrast, major real-world events lead to dramatic changes and shifts in the information pathways. We observed that clusters of mainstream news and



**Figure 5: The number of links that point between different types of sites over time for eight different inferred diffusion networks. We split the sites into mainstream media and blogs and count the links among these two node types.**



**Figure 7: Precision, Recall, and Accuracy of INFOPATH for a time-varying hyperlink network with 11,461 nodes and 19,915 edges over time, using 495,655 hyperlink cascades from July 2011 to December 2011.**

blogs often emergence and vanish in matter of days. We discovered that there is an early greater increase in information transfer among blogs than among mainstream media for news involving general population and social unrest, such as the Libyan civil war, Egyptian revolution, Syria’s uprising and the Occupy Wall Street movement.

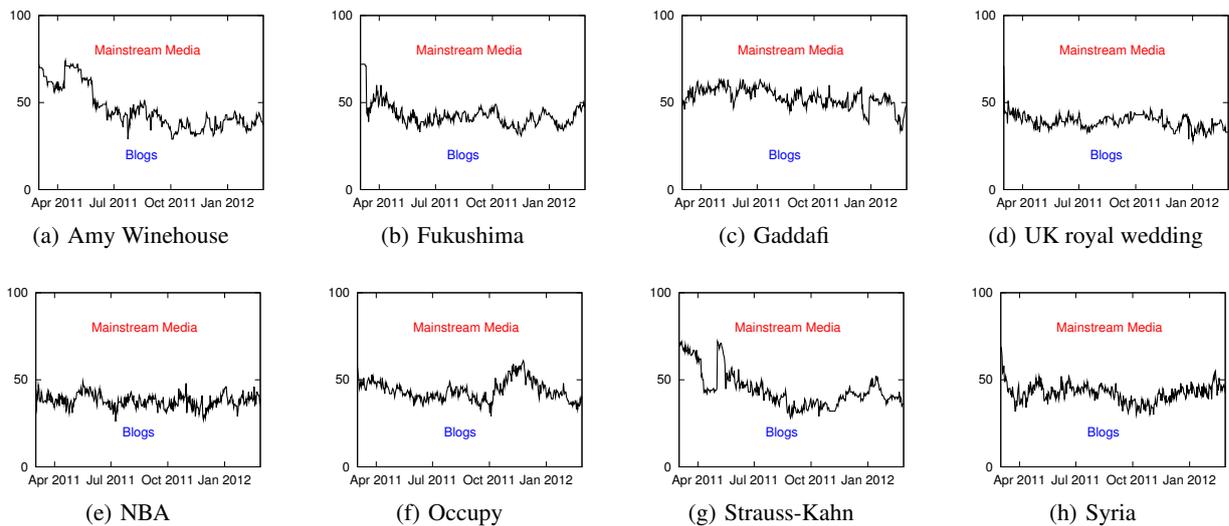
Our work also opens various venues for future work. For example, rigorous theoretical analysis of the convergence of our stochastic gradient descent method would provide further insights for its performance. Moreover, we notice that many times the changes in the inferred network structure could be attributed to sudden external real-world events. This opens two interesting questions. How can diffusion network inference be combined with methods for detecting external influence in networks [22]? And also, how can dynamic network inference be extended for detecting unexpected real-world events based on a stream of documents? Last, many times not only information but also sentiment attached to a piece of information spreads through the network [19]. It would be interesting to think about inference of signed networks, where a positive/negative valence of an edge models sentiment relationship between a pair of nodes. Overall, such methods would allow us to improve our understanding of the current landscape of news coverage, the role that news media plays in framing the discussion of important topics, and the evolving ecosystem that news media occupies.

**Acknowledgements.** This research has been supported in part by NSF IIS-1016909, CNS-1010921, IIS-1159679, IIS-1149837, DAR-

PA SMISC, ARO MURI, Brown Institute for Media Innovation, Albert Yu & Mary Bechmann Foundation, Allyes, Boeing, Docomo, Intel, Samsung, Alfred P. Sloan Fellowship, the Microsoft Faculty Fellowship, and Barrie de la Maza Graduate Fellowship.

## 6. REFERENCES

- [1] INFOPATH supporting website: <http://snap.stanford.edu/infopath>.
- [2] SNAP: Stanford Network Analysis Platform. <http://snap.stanford.edu>.
- [3] O. Aalen, Ø. Borgan, and H. Gjessing. *Survival and event history analysis: a process point of view*. Springer, 2008.
- [4] A. Agarwal and J. Duchi. Distributed delayed stochastic optimization. 2011.
- [5] F. Bach, E. Moulines, et al. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing*, 2011.
- [6] N. T. J. Bailey. *The Mathematical Theory of Infectious Diseases and its Applications*. Hafner Press, 1975.
- [7] D. Blatt, A. Hero, and H. Gauchman. A convergent incremental gradient method with a constant step size. *SIAM Journal on Optimization*, 18(1):29–51, 2008.
- [8] J. Duchi, A. Agarwal, M. Johansson, and M. Jordan. Ergodic subgradient descent. In *Allerton ’11: Proc. of the 50th Annual Allerton Conference on Communication, Control, and Computing*, 2011.
- [9] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the Temporal Dynamics of Diffusion Networks. In *ICML ’11: Proc. of the 28th International Conference on Machine Learning*, 2011.
- [10] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring Networks of Diffusion and Influence. In *KDD ’10: Proc. of the 16th ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, 2010.
- [11] M. Gomez-Rodriguez and B. Schölkopf. Influence Maximization in Continuous Time Diffusion Networks. In *ICML ’12: Proc. of the 29th International Conference on Machine Learning*, 2012.



**Figure 6: Percentage of blogs and mainstream media in top-100 most central sites for eight different inferred diffusion networks. Mainstream media are represented in red, and blogs in blue.**

- [12] M. Gomez-Rodriguez and B. Schölkopf. Submodular Inference of Diffusion Networks from Multiple Trees. In *ICML '12: Proc. of the 29th International Conference on Machine Learning*, 2012.
- [13] D. Kempe, J. M. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD '03: Proc. of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003.
- [14] J. Lawless. *Statistical models and methods for lifetime data*. Wiley New York, 1982.
- [15] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. In *EC '06: Proc. of the 7th ACM conference on Electronic commerce*, 2006.
- [16] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *KDD '09: Proc. of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2009.
- [17] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *The Journal of Machine Learning Research*, 11:985–1042, 2010.
- [18] D. Liben-Nowell and J. Kleinberg. Tracing the flow of information on a global scale using Internet chain-letter data. *Proc. of the National Academy of Sciences*, 105(12):4633–4638, 2008.
- [19] M. Miller, C. Sathi, D. Wiesensthal, J. Leskovec, and C. Potts. Sentiment flow through hyperlink networks. In *ICWSM '11: Proc. of the AAAI International Conference on Weblogs and Social Media*, 2011.
- [20] S. Myers and J. Leskovec. On the Convexity of Latent Social Network Inference. In *NIPS '10: Advances in Neural Information Processing Systems*, 2010.
- [21] S. Myers and J. Leskovec. Clash of the contagions: Cooperation and competition in information diffusion. In *ICDM '12: Proc. of the 12th IEEE International Conference on Data Mining*, 2012.
- [22] S. Myers, C. Zhu, and J. Leskovec. Information diffusion and external influence in networks. In *KDD '12: Proc. of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2012.
- [23] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574, 2009.
- [24] P. Netrapalli and S. Sanghavi. Finding the graph of epidemic cascades. In *Proc. of the 12th ACM SIGMETRICS/Performance*, 2012.
- [25] B. Prakash, A. Beutel, R. Rosenfeld, and C. Faloutsos. Winner takes all: competing viruses or ideas on fair-play networks. In *Proc. of the 21st International Conference on World Wide Web*, pages 1037–1046, 2012.
- [26] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- [27] E. M. Rogers. *Diffusion of Innovations*. Free Press, New York, fourth edition, 1995.
- [28] D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *WWW '11: Proc. of the 20th International Conference on World Wide Web*, 2011.
- [29] N. Roux, M. Schmidt, and F. Bach. A stochastic gradient method with an exponential convergence rate for strongly-convex optimization with finite training sets. *Advances in Neural Information Processing*, 2012.
- [30] T. Snowsill, N. Fyson, T. De Bie, and N. Cristianini. Refining causality: who copied from whom? In *KDD '11: Proc. of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2011.
- [31] J. Wallinga and P. Teunis. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of Epidemiology*, 160(6):509–516, 2004.
- [32] J. Yang and J. Leskovec. Modeling information diffusion in implicit networks. In *ICDM '10: Proc. of the IEEE International Conference on Data Mining*, 2010.
- [33] J. Yang and J. Leskovec. Patterns of Temporal Variation in Online Media. In *WSDM '11: ACM International Conference on Web Search and Data Mining*, 2011.