

A Stochastic Differential Equation Framework for Guiding Online User Activities in Closed Loop

Yichen Wang
Georgia Tech

Evangelos Theodorou
Georgia Tech

Apurv Verma
Georgia Tech

Le Song
Georgia Tech & Ant Financial

Abstract

Recently, there is a surge of interest in using point processes to model continuous-time user activities. This framework has resulted in novel models and improved performance in diverse applications. However, most previous works focus on the “open loop” setting where learned models are used for predictive tasks. Typically, we are interested in the “closed loop” setting where a policy needs to be learned to incorporate user feedbacks and guide user activities to desirable states. Although point processes have good predictive performance, it is not clear how to use them for the challenging closed loop activity guiding task. In this paper, we propose a framework to reformulate point processes into stochastic differential equations, which allows us to extend methods from stochastic optimal control to address the activity guiding problem. We also design an efficient algorithm, and show that our method guides user activities to desired states more effectively than state-of-arts.

1 Introduction

Online social and information platforms have brought to the public a new style of social lives: people use these platforms to receive information, create content, and share opinions. The large-scale temporal event data generated by online users have created new research avenues and scientific questions at the intersection of social sciences and machine learning. These questions are directly related to the development of new models as well as learning and inference algorithms to understand and predict user activities from these data (Kempe et al., 2003; Zhao et al., 2015; Grover & Leskovec, 2016; Lian et al., 2015; Nowzari et al., 2016; Pan et al., 2016; Tan et al., 2016; He & Liu, 2017).

Recently, point processes have been widely applied to model

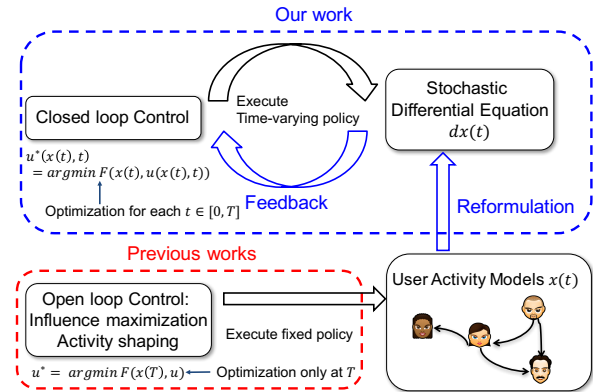


Figure 1: Comparison between our work and previous works in guiding user activities. Given a user activity model that is learned from observed behaviors, our “closed loop” guiding framework aims to find an optimal policy $u(x(t), t) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ that maps a user’s current state $x(t)$, e.g., opinion, to an action. On the contrary, previous works are “open loop”: they only optimize the objective function at terminal time and compute a fixed scalar policy $u \in \mathbb{R}$.

user activities. Instead of discretizing time into intervals, these models treat timestamps as random variables, and propose a mechanism to link model parameters with the observed timestamps. Such fine-grained modeling of temporal information has led to improved predictive performance in diverse applications, such as information diffusion (Du et al., 2013; Farajtabar et al., 2014; He et al., 2015; Wang et al., 2016a,b), recommender systems (Du et al., 2015; Wang et al., 2016a), and evolutionary networks (Farajtabar et al., 2015). However, most works deal with the “open loop” setting where models are used mainly for prediction, but user feedbacks are not incorporated into the model. Typically, we are interested in the “closed loop” setting where we want to design a policy to guide user activities and incorporate feedbacks timely. It is not clear how the current models can be used for such “closed loop” task.

For instance, a decision maker seeks to determine the best intervention policy such that the sentiment in user generated contents can be guided towards a positive state. This is significantly more challenging than traditional influence maximization or activity shaping tasks (Chen et al., 2009,

2010a,b, 2012; Farajtabar et al., 2014), where the policy is determined before the process unfolds, and they do not take into account the instantiation of the process. A framework for doing this is critically important for understanding the vulnerabilities of information networks, designing policies to suppress the spread of undesired contents, and exploring new ways of performing content recommendation. For instance, a network moderator may want to effectively contain the spread of rumors and misinformation, and an online platform may want to promote the level of long-term user activities rather than the short-term click-through rate.

In this paper, we provide a novel view of point process models, and reformulate them into stochastic differential equations (SDEs). This reformulation allows us to significantly generalize existing point process models to SDEs, and plays a critical role in connecting the task of guiding user activities to stochastic optimal control theory, which are often used in robotics. Hence, we can bring in and extend lots of tools from stochastic optimal control literature to address the “closed loop” sequential decision making problem. Figure 1 illustrates our framework.

Interestingly, these problems also introduce new technical challenges. Previous works in stochastic optimal control study SDEs driven by Wiener processes and/or Poisson processes (Boel & Varaiya, 1977; Pham, 1998; Oksendal & Sulem, 2005; Hanson, 2007). Online user activity modeling requires us to consider more advanced processes, such as (i) Hawkes processes for long term memory and mutual exciting phenomena in social interactions, (ii) survival processes (Aalen et al., 2008) for self-terminating behavior in influence propagation and link creation, and (iii) node birth processes for evolving networks (Farajtabar et al., 2015; Thalmeier et al., 2016). Thus, many technical results from control theory need to be extended for the activity guiding problem. In summary, we make the following contributions:

- We provide a general way to reformulate point process based user activity models into stochastic differential equations (SDEs), which allows us to bring in and extend tools from stochastic control literature to address the “closed loop” activity guiding problem.
- We extend technical results in stochastic optimal control theory and derive generalized Ito’s lemma and HJB equation for SDEs driven by more general point processes.
- We propose an algorithm that efficiently guides the dynamics of user activities towards a target. The algorithm can also deal with the challenging cases of time-varying networks with node births.

Finally, with synthetic and real networks, we showed our framework is robust, able to steer user activities to desired states with faster convergence rate and less variance than the state-of-art methods.

Further related work. Most previous works in influence maximization (Kempe et al., 2003; Chen et al., 2010b, 2012)

focus on selecting sources to maximize the spread of information in infinite time, which is an “open loop” setting. There is typically no quantitative prescription on how much incentive should be provided to each user. Moreover, in most cases, a finite time window must be considered. For example, a politician would like to maximize his support by a million people in one week instead of fifty years.

Another relevant area is optimal control for epidemic processes (Epanchin-Niell & Wilen, 2012; Nowzari et al., 2016; Ogura & Preciado, 2016; Pastor-Satorras et al., 2015). However, the epidemic processes are modeled by deterministic differential equations. Therefore, these works neither consider the influence of abrupt event nor the stochastic nature of user behaviors in social networks. Hawkes process models (Farajtabar et al., 2014; De et al., 2015) overcome this limitation by treating the user-generated event time as a random variable and model users’ stochastic behaviors. They address the problem of designing the base intensity of Hawkes process to achieve a desired steady behavior. However, these methods are open loop and do not consider user feedbacks, and the variance of the dynamics can still be very large. Recently, Thalmeier et al. (2016) studied the problem of controlling network growth in the macroscopic level, while we focus on guiding users’ microscopic temporal behaviors based on point processes and SDEs.

2 Background and preliminaries

Point processes. A temporal point process (Aalen et al., 2008) is a random process whose realization is a list of events $\{t_i\}$, localized in time. Many types of data in social networks can be represented as point processes (Du et al., 2015; Farajtabar et al., 2014; Lian et al., 2015; Tan et al., 2016; Wang et al., 2016a,b, 2017a,b,c). The point process can be equivalently represented as a counting process, $N(t)$, which records number of events before time t . An important way to characterize point processes is via the conditional intensity function $\lambda(t)$, which is a stochastic model for the time of the next event given all previous events. Let $\mathcal{H}(t) = \{t_i | t_i < t\}$ be the history of events happened up to t . Formally, $\lambda(t)$ is the conditional probability of observing an event in a small window $[t, t + dt)$ given the history $\mathcal{H}(t)$: $\lambda(t)dt := \mathbb{P}\{\text{event in } [t, t + dt) | \mathcal{H}(t)\} = \mathbb{E}[dN(t) | \mathcal{H}(t)]$, where only one event can happen in a window of dt , i.e., $dN(t) \in \{0, 1\}$. The functional form of $\lambda(t)$ is often designed to capture the phenomena of interests. We will provide more examples later in Section 3.

Stochastic differential equations (SDEs). A jump diffusion SDE is a differential equation in which one or more terms is a stochastic process: $dx(t) = f(x)dt + g(x)dw(t) + h(x)dN(t)$, where $dx(t) := x(t + dt) - x(t)$ is the differential of $x(t)$. This SDE contains a drift, a diffusion and a jump term: the drift term $f(x)dt$ models the intrinsic evolution of $x(t)$; the diffusion Wiener process $w(t) \sim \mathcal{N}(0, t)$ captures the Gaussian noise; the jump point process $N(t)$ captures the influence of abrupt events.

3 Stochastic differential equations for user activity models

In this section, we establish a framework to reformulate many point process models into SDEs. This reformulation framework plays a critical role in connecting the task of guiding user activities to stochastic optimal control theory often used in robotics and designing closed loop policies.

3.1 User activity models

We first introduce the generic point process models for user activities, and present three examples.

Definition 1 (User Activity Model). *For a network with U users, the point process $N_i(t)$ models the generation of event times from user i , and its intensity f is defined as*

$$\lambda_i(t) = \eta_i(t) + \sum_{j=1}^U \beta_{ij} \sum_{t_j \in \mathcal{H}_j(t)} \kappa_{\omega_1}(t - t_j), \quad (1)$$

where $\eta_i(t)$ is the base intensity, $\beta_{ij} \geq 0$ models the strength of influence from user j to i , $\mathcal{H}_j(t)$ is the history of events for user j , and $\kappa_{\omega_1}(t) = \exp(-\omega_1 t)$ is the triggering kernel capturing the influence of past event. We also assume the additional event feature/content $x_i(t)$ follows the model

$$x_i(t) = b_i(t) + \sum_{j=1}^U \alpha_{ij} \sum_{t_j \in \mathcal{H}_j(t)} \kappa_{\omega_2}(t - t_j) h(x_j(t_j)) \quad (2)$$

where $b_i(t)$ is the base content, α_{ij} is the influence from user j to i , the function $h(\cdot)$ captures the influence of the activity content and typical forms include $h = 1$ and $h(x) = x$.

This model generalizes point processes since it not only models the on-off temporal behavior, but also the activity content $x_i(t)$, e.g., opinion or a vector of interested topics. It captures both exogenous and endogenous properties of networks. The exogenous term is the base intensity/activity content, and the endogenous term captures the fact that one's intensity/activity content is influenced by neighbors. α_{ij}, β_{ij} measure the strength of such influence, the kernel $\kappa_{\omega}(t)$ captures the decay of influence of past events/content over time, and the summation captures the influence of neighbors' events. Next we present three examples.

- **Continuous-time information propagation** (Du et al., 2013). The information propagation in social networks begins with a set of source nodes, and the contagion is transmitted from the sources along their out-going edges to their direct neighbors. We set $N_{ij}(t)$ to be the *survival process* capturing the infection on the edge $i \rightarrow j$, and $N_{ij}(t) = 1$ means node j is infected by node i at time t . Since there is only one event for an instantiation of this process, we set the infection intensity as follows:

$$\lambda_{ij}(t) = \eta_{ij}(1 - N_{ij}(t)) \quad (3)$$

where η_{ij} is intensity that i infects j , and $(1 - N_{ij}(t))$ ensures the infection happens only once.

- **Hawkes processes** (Hawkes, 1971). This model captures the mutual excitation phenomena between events, and has been successfully applied to analyze user behaviors

in social networks (Du et al., 2015; Farajtabar et al., 2014, 2015; Wang et al., 2016a,b, 2017b). The process $N_i(t)$ counts the number of events generated by user i up to time t and its intensity $\lambda_i(t)$ models mutual excitation between the events from a collection of U users as follows:

$$\lambda_i(t) = \eta_i + \sum_{j=1}^U \beta_{ij} \sum_{t_j \in \mathcal{H}_j(t)} \kappa_{\omega_1}(t - t_j)$$

Here, the occurrence of each historical event from one's neighbors increases the intensity by a certain amount determined by $\kappa_{\omega_1}(t)$ and α_{ij} .

- **Opinion dynamics** (De et al., 2015; He et al., 2015). This model considers both the timing and content of each event. It assigns each user i a Hawkes intensity $\lambda_i(t)$ to generate events and an opinion process $x_i(t)$, where $x_i(t) = 0$ corresponds to neutral opinion and large positive/negative values correspond to extreme opinions. The opinion of user i is modeled as a temporally discounted average of neighbors' opinion:

$$x_i(t) = b_i + \sum_j \alpha_{ij} \sum_{t_j \in \mathcal{H}_j(t)} \kappa_{\omega_2}(t - t_j) x_j(t_j)$$

This model has superior performance in predicting opinions. However, it is not clear whether it can be used to design feedback policies to guide the dynamics precisely to some target states.

3.2 Equivalent SDE reformulation

We are now ready to show the novel SDE reformulation of the user activity model in Definition 1.

Theorem 2 (Transformation Framework). *The equivalent SDE form of the user activity model is:*

$$d\lambda_i(t) = d\eta_i + \omega_1(\eta_i - \lambda_i)dt + \sum_j \beta_{ij} dN_j(t) \quad (4)$$

$$dx_i(t) = db_i + \omega_2(b_i - x_i(t))dt + \sum_j \alpha_{ij} h(x_j) dN_j(t)$$

Proof sketch. We first define a convolution operator and reformulate the user activity model in (1) and (2), next we apply a differential operator d to the reformulated equations and derive the differential form of $\lambda_i(t)$ and $x_i(t)$, which lead to two SDEs. Appendix A contains details.

These SDEs describe how the intensity $\lambda_i(t)$ and content $x_i(t)$ change on $[t, t + dt)$; each consists of three terms:

- **Baseline change.** The differential $db_i(t)$ captures the infinitesimal change of base activity content.
- **Drift.** The change rate of the activity content, $dx_i(t)/dt$, is proportional to $-x_i(t)$, which means the activity content $x_i(t)$ tends to stabilize over time. In fact, if ignoring the jump term and $db_i(t)$, the expected activity content $\mathbb{E}[x_i(t)]$ will converge to $b_i(t)$ as t increases. This can be proved by setting $\mathbb{E}[dx_i(t)] = 0$. Similarly, equation (4) shows a user's intensity tends to stabilize over time.
- **Jump.** This term captures the influence of each event in the network and is a weighted summation of the neighbors' influence. The coefficient α_{ij} ensures that only neighbors' effect will be considered, and $dN_j(t) \in$

$\{0, 1\}$ models whether user j generates an event. Similarly, equation (4) shows that user i 's intensity increases by β_{ij} if his neighbor j generates an event.

Next, we present three applications of Theorem 2.

- **SDE for continuous-time information propagation.** The intensity in (3) is a simplified version of (1) without the historical influence term. Its SDE version is:

$$d\lambda_{ij}(t) = -\eta_{ij}dN_{ij}(t)$$

This SDE keeps the key property of survival process: before an infection happens, the intensity is constant, *i.e.*, $\lambda_{ij}(t) = \eta_{ij}$; after the infection happens, the intensity is 0, *i.e.*, $\lambda_{ij}(t + dt) = \lambda_{ij}(t) + d\lambda_{ij}(t) = 0$.

- **SDE for Hawkes process.** We set $\eta_i(t) = \eta_i$ and obtain:

$$d\lambda_i(t) = \omega_1(\eta_i - \lambda_i(t))dt + \sum_j \alpha_{ij}dN_j(t)$$

This SDE shows that the user's intensity tends to stabilize and its change is influenced by his neighbors' activities.

- **SDE for opinion dynamics.** We set $b_i(t) = b_i$, $h_j(x_j) = x_j$, and further generalize this model by adding a Wiener process term:

$$dx_i = \omega_2(b_i - x_i)dt + \beta dw_i + \sum_j \alpha_{ij}x_jdN_j$$

where the Wiener process $w_i(t)$ captures the Gaussian noise, such as fluctuations in the dynamics due to unobserved factors and activities outside the social platform. The jump term models the fact that the change of opinion is influenced by his neighbors' opinion.

3.3 Benefit of the SDE modeling framework

Our SDE formulation opens a new gate to extend tools from optimal control theory. Hence we can solve many important problems in social sciences, such as the least square activity guiding and activity maximization problem. Without this view, it is not easy to design algorithms with closed loop policies. Besides transforming an existing model to an SDE, we can also directly design an SDE to model many other factors. For example, one can model the Gaussian noise by adding a Wiener process to an SDE. Next, we show how to optimally control the SDE to guide user activities.

4 A convex activity guiding framework

In this section, we define the activity guiding problem. Let $\mathbf{x}(t) = (x_1(t), \dots, x_U(t))^T$ be the vector of each user's activity content, then we study the vector version of the SDE in (4) and reformulate this SDE with an extra control policy $\mathbf{u}(\mathbf{x}(t), t) : \mathcal{R}^U \times \mathcal{R} \rightarrow \mathcal{R}^U$ as follows.

$$d\mathbf{x} = (\mathbf{f}(\mathbf{x}) + \mathbf{u})dt + \mathbf{g}(\mathbf{x})d\mathbf{w}(t) + \mathbf{h}(\mathbf{x})d\mathbf{N}(t) \quad (5)$$

This policy \mathbf{u} maps the activity content \mathbf{x} to an action for all users. Next we present two examples.

- **Guiding Hawkes process.** To steer user activities to a desired level, we provide incentives to users and add a control policy $u_i(\lambda_i(t), t)$ to the SDE formulation of the

Hawkes intensity function as follows:

$$d\lambda_i(t) = (\eta_i + u_i(\lambda_i(t), t) - \lambda_i(t))dt + \sum_j \beta_{ij}dN_j(t),$$

where $u_i(\lambda_i(t), t)$ captures the amount of additional influence to change the baseline intensity η_i .

- **Guiding opinion dynamics.** We can guide the opinion SDE with a policy $u_i(x_i(t), t)$ as follows:

$$dx_i = (b_i + u_i - x_i)dt + \theta dw_i + \sum_j \alpha_{ij}x_jdN_j \quad (6)$$

where $u_i(x_i, t)$ determines how fast the opinion needs to be changed for user i . For example, a network moderator may request the user to change his opinion from -3 to 1 in one day, and this control policy quantifies the amount of change in unit time.

Next we present the objective for optimizing the SDE.

Definition 3 (Stochastic User Activity Guiding Problem).

For the SDE in (5), given the initial condition (\mathbf{x}_0, t_0) , we aim to find an optimal policy $\mathbf{u}^*(\mathbf{x}, t)$ for $t \in (t_0, T]$, which minimizes the convex objective function:

$$V(\mathbf{x}_0, t_0) := \min_{\mathbf{u}} \mathbb{E} \left[\phi(\mathbf{x}(T)) + \int_{t_0}^T \mathcal{L}(\mathbf{x}, \mathbf{u}, t)dt \right] \quad (7)$$

where V is called the value function that summarizes the optimal expected cost if \mathbf{u}^* is executed from t_0 to T . It is a function of the initial state. The expectation \mathbb{E} is over stochastic processes $\{\mathbf{w}(t), \mathbf{N}(t)\}$ for $t \in (t_0, T]$. ϕ is the terminal cost and \mathcal{L} is the instantaneous cost.

Terminal cost ϕ . It is the cost at final time T . We discuss several functional forms as follows:

- **Least square guiding.** The goal is to guide the expected state to a pre-specified target \mathbf{a} . For opinion dynamics, the goal can be to ensure nobody believes the rumor. Mathematically, we set $\phi = \|\mathbf{x}(T) - \mathbf{a}\|^2$. Moreover, to influence users' intensity of generating events, one can set the desired level of the intensity to be at a high level \mathbf{a} and conduct activity guiding: $\phi = \|\boldsymbol{\lambda}(T) - \mathbf{a}\|^2$.
- **Information/activity maximization.** The goal is to maximize the activity content of all users. For example, the goal for an educator is to maximize the students' recognition of the value of education, and we set $\phi = -\sum_u x_u(T)$ to maximize each user's positive opinion. Moreover, to improve the activity level in social platforms, one can also maximize the intensity: $\phi = -\sum_u \lambda_u(T)$.

Instantaneous cost \mathcal{L} . This is the cost at $t \in [t_0, T]$ and in the form of $\mathcal{L}(\mathbf{x}, \mathbf{u}, t) = q(\mathbf{x}(t)) + \rho c(\mathbf{u}(t))$. The state cost $q(\mathbf{x}(t))$ is optional and the control cost $c(\mathbf{u}(t)) = \|\mathbf{u}(t)\|^2$ is necessary. We set $q = 0$ if the cost only occurs at T ; otherwise q captures the cost at the intermediate time, *e.g.*, the cost incurred by maximizing students' positive recognition of the value of education over consecutive weeks. Its function form is the same as the terminal cost: $q = \phi$. The control cost captures the fact that the policy costs money or

human efforts. The scalar ρ is the trade-off between control cost and state cost.

Solving this activity guiding problem is challenging, since the objective function involves taking expectation over complex stochastic processes. Furthermore, it is a functional optimization problem since the optimal policy is a function of both state and time. Fortunately, the SDE formulations allow us to connect the problem to that of stochastic dynamic programming methods. As a result, we can extend lots of tools in stochastic optimal control to address sequential decision making problems.

5 Algorithm for optimal policy

In this section, we will find the optimal policy posed in (7). Prior works in control theory mostly study the SDE where the jump is a Poisson process (Boel & Varaiya, 1977; Hanson, 2007; Oksendal & Sulem, 2005; Pham, 1998). However, in our model, the jump process is a more complex point process with *stochastic intensity functions*, e.g., Hawkes process. Hence significant generalizations, both in theory and algorithms, are needed. We first derive the HJB partial differential equation (PDE) for a deterministic system, then generalize the procedure to the activity guiding problem. Further, we extend our framework to guide user activities in the challenging time-varying networks.

5.1 HJB equation for deterministic systems

To obtain the optimal policy, we need to compute the value function V in (7) subject to the constraint of the SDE. We will break down the complex optimization into simpler sub-problems. First, the initial condition $\mathbf{x}(t_0)$ needs to be replaced by an arbitrary start $\mathbf{x}(t)$, so that the start can be analytically manipulated and we obtain a time-varying objective $V(\mathbf{x}, t)$ amenable to analysis. Next, since the value function consists of an integral term, we break the integral into $[t, t + dt]$ and $[t + dt, T]$. If the system is deterministic, we can further split the value function as:

$$V(\mathbf{x}, t) = \min_{\mathbf{u}} \left[\underbrace{\phi + \int_{t+dt}^{T-} \mathcal{L} d\tau}_{V(\mathbf{x}(t+dt), t+dt)} + \underbrace{\int_t^{t+dt} \mathcal{L} d\tau}_{\text{cost } t \rightarrow t+dt} \right] \quad (8)$$

The first term is $V(\mathbf{x}(t + dt), t + dt)$ and the second term is the optimal cost on $[t, t + dt]$. Hence (8) follows the structure of a dynamic programming and we can solve it recursively: given $V(\mathbf{x}(t + dt), t + dt)$, we only need to proceed optimally on $[t, t + dt]$ to compute V backward.

To further simplify (8), we perform deterministic Taylor expansion up to second order of the first term on right-hand side as $V(\mathbf{x}(t + dt), t + dt) := V(\mathbf{x}, t) + dV(\mathbf{x}, t)$, where $dV = V_t dt + V_x^\top d\mathbf{x} + \frac{1}{2} d\mathbf{x}^\top V_{xx} d\mathbf{x} + d\mathbf{x}^\top V_{xt} dt + \frac{1}{2} V_{tt} dt^2$. Then we can cancel $V(\mathbf{x}, t)$ on both sides of (8), divide it by dt , and take the limit as $dt \rightarrow 0$. Since $d\mathbf{x} = \mathbf{x}'(t)dt$, all the second order term in dV goes to 0. Hence we obtain the Hamilton-Jacobi-Bellman (HJB) equation: $-V_t = \min_{\mathbf{u}} [\mathcal{L}(\mathbf{x}, \mathbf{u}, t) + V_x^\top \mathbf{x}']$. However, our

system is stochastic and the above procedure needs to be generalized significantly.

5.2 HJB equation for guiding user activities

To derive the HJB equation for our model, we need to address two challenges: (i) computing the stochastic Taylor expansion dV under our SDE, and (ii) taking expectation of the stochastic terms in (8) before minimization. We first derive Theorem 4 to compute dV , which generalizes Ito's Lemma and is applicable to SDEs driven by point processes with stochastic intensity functions. Then we compute the expectation and derive a HJB equation in Theorem 5.

Theorem 4 (Generalized Ito's Lemma for Jump Diffusion SDEs). *Given the SDE in (5), let $V(\mathbf{x}, t)$ be twice differentiable in \mathbf{x} and once in t , then we have*

$$dV = \left\{ V_t + \frac{1}{2} \text{tr}(V_{xx} \mathbf{g} \mathbf{g}^\top) + V_x^\top (\mathbf{f} + \mathbf{u}) \right\} dt + V_x^\top \mathbf{g} d\mathbf{w} + (V(\mathbf{x} + \mathbf{h}, t) - V(\mathbf{x}, t)) d\mathcal{N}(t) \quad (9)$$

Theorem 5 (HJB Equation). *Let $\mathbf{h}_j(\mathbf{x}, t)$ be the j -th column of $\mathbf{h}(\mathbf{x}, t)$. Then the HJB Partial Differential Equation for the user activity guiding problem is*

$$-V_t = \min_{\mathbf{u}} \left[\mathcal{L} + \frac{1}{2} \text{tr}(V_{xx} \mathbf{g} \mathbf{g}^\top) + V_x^\top (\mathbf{f} + \mathbf{u}) + \sum_j \lambda_j(t) (V(\mathbf{x} + \mathbf{h}_j(\mathbf{x}, t), t) - V(\mathbf{x}, t)) \right] \quad (10)$$

To prove Theorem 4, we derive a new set of stochastic calculus rules for general point processes and conduct stochastic Taylor expansion in the Ito's mean square limit sense. Appendix C contains details. To prove Theorem 5, we combine the Generalized Ito's Lemma with the property of point processes. Appendix D contains details. Next, we solve the HJB equation to obtain the value function V . We will show that under the optimal parameterizations of V , this HJB equation can be solved efficiently.

5.3 Parameterization of the value function

To solve the HJB equation, we first show the structure of the value function in the following proposition.

Proposition 6. *If the SDE in (5) is linear in \mathbf{x} , and the terminal and instantaneous cost are quadratic/linear in \mathbf{x} , then the value function $V(\mathbf{x}, t)$ in (7) must be quadratic/linear.*

This result is intuitive since the V is the optimal value of the summation of quadratic/linear functions. Appendix E contains the proof. This proposition is applicable to two important problems, including the least square activity guiding problem and the linear activity maximization problem.

In this section, we present derivations for the least square guiding problem, and Appendix H contains derivations for the activity maximization problem. Specifically, we set $V(\mathbf{x}, t)$ to be quadratic in \mathbf{x} with unknown coefficients $\mathbf{v}_1(t) \in \mathbb{R}^U$, $\mathbf{v}_{11}(t) \in \mathbb{R}^{U \times U}$ and $v_0(t) \in \mathbb{R}$:

$$V(\mathbf{x}, t) = v_0(t) + \mathbf{v}_1(t)^\top \mathbf{x} + \frac{1}{2} \mathbf{x}^\top \mathbf{v}_{11}(t) \mathbf{x}$$

Algorithm 1 OPTIMAL CONTROL POLICY

- 1: **Input:** network $\mathbf{A} = (\alpha_{ij})$, model parameters $\{\eta, \theta, \mathbf{b}\}$, timestamps $\{\tau_k\}_{k=1}^m$, events $\{t_i\}_{i=1}^n$, target \mathbf{a}
 - 2: **Output:** $\mathbf{v}_{11}(\tau_k), \mathbf{v}_1(\tau_k), k = 1, \dots, m$
 - 3: **for** $k = 1$ **to** m **do**
 - 4: Compute $\lambda_u(\tau_k) = \eta_u + \sum_{j:t_i < \tau_k} \alpha_{uu_i} \kappa(\tau_k - t_i)$,
 $\mathbf{\Lambda}(\tau_k) = \sum_{j=1}^U \lambda_j(\tau_k) \mathbf{B}^j$ for each user u
 - 5: **end for**
 - 6: Compute $\mathbf{v}_{11}(\tau_k), \mathbf{v}_1(\tau_k)$ using ODE45; then compute $\mathbf{u}(\tau_k)$ in (11).
-

To find the optimal control, we substitute $V(\mathbf{x}, t)$ to the HJB equation in (10) and set the gradient of its right-hand-side to $\mathbf{0}$. This yields the optimal feedback control policy:

$$\mathbf{u}^*(\mathbf{x}(t), t) = -V_{\mathbf{x}}/\rho = -(\mathbf{v}_1(t) + \mathbf{v}_{11}(t)\mathbf{x}(t))/\rho \quad (11)$$

This policy consists of two terms: the feedforward term $\mathbf{v}_1(t)$ controls the system as time goes by; the feedback term updates the policy based on the current state $\mathbf{x}(t)$. Moreover, ρ controls the tradeoff between control and state cost, and $\rho \rightarrow \infty$ means low budget; hence $\mathbf{u}^* \rightarrow \mathbf{0}$.

5.4 Stochastic optimal control algorithm

Given the form of optimal policy \mathbf{u}^* in (11), the final step is to compute its unknown coefficients $\{\mathbf{v}_1(t), \mathbf{v}_{11}(t)\}$. We substitute this optimal policy to the HJB equation in (10). Since the value function $V(\mathbf{x}, t)$ is quadratic, we can separate the HJB equation into terms that are scalar, linear, and quadratic in \mathbf{x} . Grouping coefficients of these terms leads to two Ordinary Differential Equations (ODEs) as follows:

$$-\mathbf{v}'_{11} = \mathbf{I} + 2\mathbf{v}_{11}(-1 + \mathbf{\Lambda}) - \frac{\mathbf{v}_{11}^2}{\rho} + \sum_j \lambda_j \mathbf{B}^{j\top} \mathbf{v}_{11} \mathbf{B}^j$$

$$-\mathbf{v}'_1(t) = -\mathbf{a} + (-1 + \mathbf{\Lambda}^\top - \mathbf{v}_{11}(t)/\rho)\mathbf{v}_1(t) + \mathbf{v}_{11}(t)\mathbf{b}$$

where $\mathbf{\Lambda}(t) = \sum_j \lambda_j(t) \mathbf{B}^j$, matrix \mathbf{B}^j has the j -th column as $(\alpha_{1j}, \dots, \alpha_{Uj})^\top$ and zero elsewhere. The terminal conditions are $\mathbf{v}_{11}(T) = \mathbf{I}$ and $\mathbf{v}_1(T) = -\mathbf{a}$. We use the Runge-Kutta method (Dormand & Prince, 1980) to solve the ODEs offline. Specifically, we partition $(t_0, T]$ to equally-spaced timestamps $\{\tau_k\}$ and obtain values of $\mathbf{v}_{11}(t), \mathbf{v}_1(t)$ at these timestamps. We use the ODE45 solver in MATLAB. Finally we update the policy online according to (11). Algorithm 1 summarizes the procedure.

5.5 Extensions to time-varying networks

We can extend our framework to networks with time-varying edges and node birth processes. Specifically, for a fixed network, the expectation in the objective function (7) is over Wiener processes and point processes. In stochastic networks, the network itself adds an extra stochasticity to the SDE and we need to take the expectation of the network topology $\mathbf{A}(t)$ to derive the HJB equation. Hence the input to Algorithm 1 is $\mathbb{E}[\mathbf{A}(t)]$ instead of \mathbf{A} . Appendix B contains details on computing this expectation.

6 Experiments

We focus on two tasks: least square opinion guiding (LSOG) and opinion influence maximization (OIM). We compare

with the following state-of-art stochastic optimization methods that are applicable to control SDEs and point processes.

- **Value Iteration (VI)** (Bertsekas, 2012): we directly formulate a discrete-time version of the opinion dynamics and compute the optimal policy using the value iteration algorithm. The discretized timestamps are the same as that for solving the HJB PDE.
- **Cross Entropy (CE)** (Stulp & Sigaud, 2012): it samples policies from a normal distribution, sorts them in ascending order w.r.t. the cost and recomputes the distribution parameters from the first K elite samples. This procedure repeats with new distribution until costs converge.
- **Finite Difference (FD)** (Peters & Schaal, 2006): it generates samples of perturbed policies and compute perturbed costs. Then it uses them to approximate the gradient of the cost w.r.t the policy. The cost in CE and FD is evaluated by executing the policy on the SDE.
- **Greedy:** It controls the SDE when the state cost is high. We divide time horizon into n timestamps. At each timestamp, we compute the state cost and control the system based on pre-specified rules if current cost is more than k times of the optimal cost of our method. We vary k from 1 to 5, n from 1 to 100 and report the best performance.
- **Slant** (De et al., 2015): It sets the open loop control policy only at the initial time.

6.1 Experiments on synthetic data

We generate a synthetic a network with 1000 users, where the topology matrix is randomly generated with a sparsity of 0.001. We simulate the opinion SDE on $[0, 10]$ using the Euler method (Hanson, 2007) to compute its difference form. The time window is divided into 100 equally spaced intervals. We set the base opinion uniformly at random, $b_i \sim \mathcal{U}[-1, 1]$, $\omega = 1$, noise level $\theta = 0.2$, $\alpha_{ij} \sim \mathcal{U}[0, 0.01]$, and $x_i(0) = -10$. The Wiener process is simulated from the Gaussian distribution and the Hawkes process is simulated using the Thinning algorithm (Ogata, 1981). We set the budget level parameter $\rho = 10$, and our results generalize beyond this value. We repeat simulations of the SDE for ten times and report the averaged performance. Appendix F contains details on this experimental setup.

Total cost. For LSOG, we set the target $a_i = 1$. The total cost per user is computed by dividing the value function by # users. Since OIM aims to maximize positive opinions, the total opinion per user is computed by dividing the negative value function by # users. Figure 2(a) shows that our method has around $2.5 \times$ improvement over CE and $4 \times$ improvement over FD. CE assumes the policy is sampled from a Gaussian distribution, and FD approximates the gradient. However, our method does not have such restrictions or approximations, and it exactly minimizes the cost.

Importance of the SDE formulation. Our framework significantly outperforms VI, since our SDE reformulation of user activity models preserves the *continuous time* property

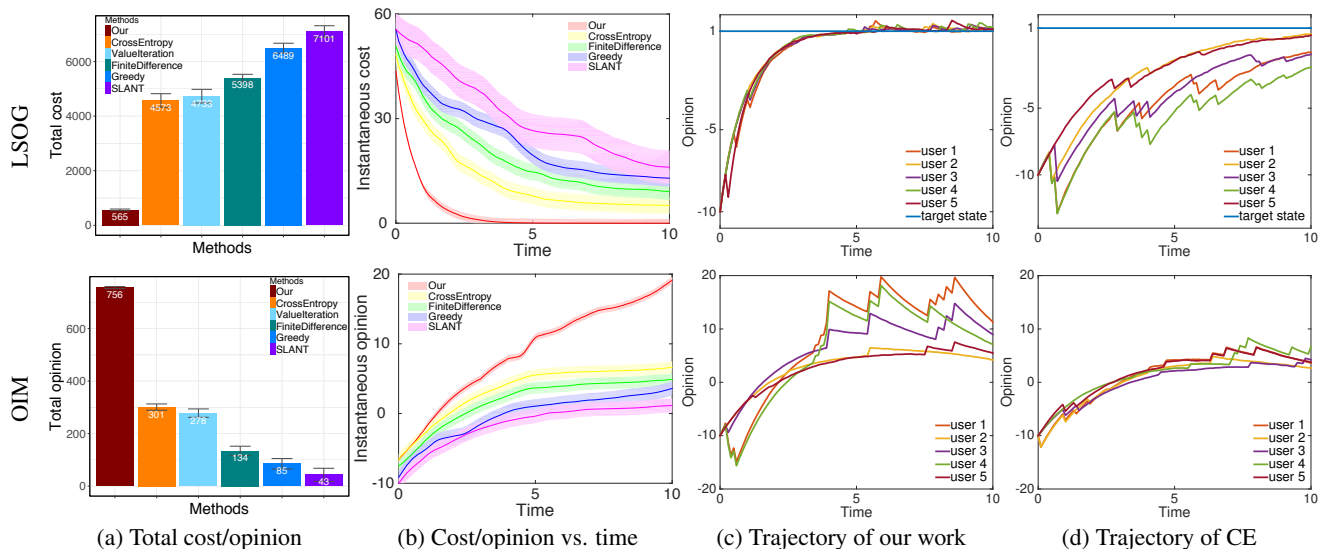


Figure 2: Results in least square guiding (LSOG) and opinion maximization (OIM). (a) total cost for LSOG and total opinion for OIM per user. Error bar is the variance; (b) instantaneous cost/opinion per user. Line is the mean and pale region is the variance; (c,d) sample trajectories of five users.

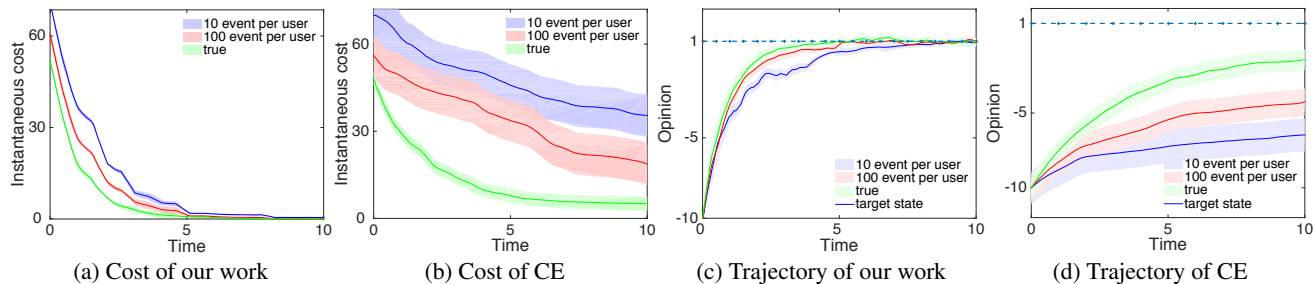


Figure 3: Robustness analysis when parameters are learned with different sizes of training data. (a) and (b) instantaneous cost; (c) and (d) opinion trajectory for one randomly sampled user.

and our policy exactly optimizes the objective function using the HJB equation. In contrast, discretizing the original opinion model introduces approximation errors, which further influences the policies in VI. Hence it is important to directly reformulating these user activity models into SDEs and study the *continuous time* control problem.

Instantaneous cost & trajectory. Figure 2(b) shows the instantaneous cost per user over time. Our method has the fastest convergence rate to the optimal cost and the cost is much smaller than competitors. Moreover, our method has the smallest variance and is stable despite multiple runs and the stochasticity in the SDE. Figure 2(c,d) compare opinion trajectories. The jumps in the opinion trajectories correspond to the opinion posting events that are modulated by the Hawkes process. For LSOG, the opinion converges to the target the fastest. For OIM, our method maximizes the opinion from the negative initial value quickly: around time 2.5, all users’ opinion are positive. Moreover, our method achieves the largest opinion value, *e.g.*, 20, while that of CrossEntropy is smaller than 10.

Robustness. Since error exists between estimated parameters and ground truth, we investigate how our method performs with this discrepancy. We generate data with 10 and 100 events per user, and learn parameters by maximum likelihood estimation (Iacus, 2009). Figure 3(a,c) show that learned parameters are close to ground truth as the training data increases. Moreover, even with less accurate parameters, our cost and trajectories are close to ground-truth, while CE has high variance.

6.2 Experiments on real-world data

We study the least square opinion guiding problem over two node birth networks. *Twitter* (Farajtabar et al., 2015) contains nearly 550,000 tweet, retweet and link creation events from around 280,000 users. We use events from Sep. 21-30, 2012 and use the data before Sep. 21 to construct the initial social network. We consider the links created in the second 10-day period to be the node birth. *MemeTracker* (Leskovec et al., 2009) contains online social media activities from August 2008 to April 2009. Users track the posts of others and the network growth is captured by hyperlinks of comments on one site to others. We extract 11,321,362 posts among

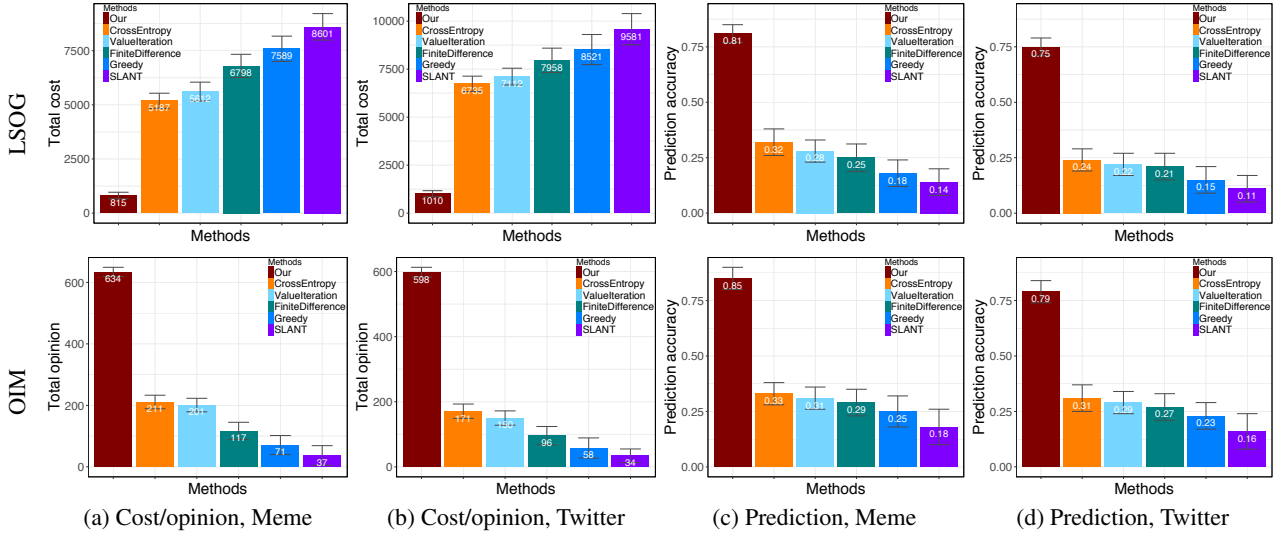


Figure 4: Results in LSOG and OIM over real-world networks with node birth processes. (a,b) total cost (for LSOG) and opinion (for OIM) in two datasets; (c,d) prediction accuracy in two datasets.

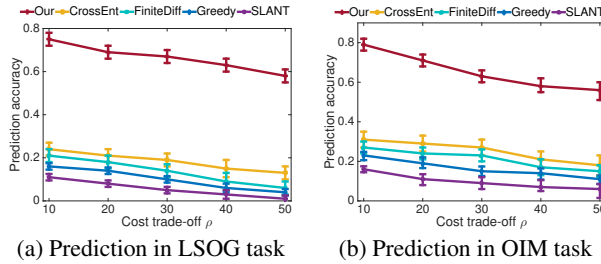


Figure 5: Budget sensitivity analysis: prediction accuracy as a function of ρ on Twitter.

5000 nodes. We use the data in Aug. 2008 to construct the initial network and use the LIWC (Pennebaker et al., 2015) toolbox to extract opinions from posts. We learn parameters of the opinion dynamics and the link creation process by maximizing the likelihood (Iacus, 2009).

We use two evaluation procedures. First, we have a real network and learned parameters; hence we simulate user behaviors, control simulated behaviors, and evaluate the total cost of different policies. The second and more interesting evaluation scheme would entail carrying real policy in a social platform. Since it is very challenging to evaluate on a real platform, we mimic such procedure using held-out data. The key idea is to predict which real trajectory reaches the objective better (has lower cost), by comparing it to the optimal trajectory x^* . Different methods yield different x^* , and the prediction accuracy depends on how optimal x^* is. If it is optimal, it is accurate if we use it to order the real trajectories, and the predicted list should be similar to the ground truth, which is close to the accuracy of 1.

Total cost. Figure 4(a,b) show that our method performs the best for the two time-varying networks. Compared with CrossEntropy, it achieves around $6\times$ improvement on LSOG and $3\times$ on OIM. This result suggests that controlling

the SDE over time-varying networks is a challenging problem for traditional stochastic optimal control algorithms. Moreover, the total costs of all methods for *Twitter* are higher than that of *Memetracker*. This is because *Twitter* has a much higher frequency of node birth, *i.e.*, users join the network in the timescale of minute-to-minute rather than day-to-day in *Memetracker*. Hence it is more challenging to control due to the high stochasticity in the network.

Prediction accuracy & budget sensitivity. Figure 4(c,d) show that our method achieves more than 0.4+ improvement over CrossEntropy. It means that our method accommodates 40% more of the total realizations correctly. An accurate prediction means that if applying our control policy, we will achieve the objective better than alternatives. Figure 5 shows that our method performs the best as the budget level decreases. Large value of the cost tradeoff parameter ρ means small budget.

7 Conclusions

We have proposed a novel SDE reformulation for user activity models and presented the activity guiding problem, which builds a new bridge between the problem of guiding of user activities in “closed loop” and stochastic optimal control theory. Moreover, we have shown that it is important to incorporate the system status information to design a feedback control policy, which will achieve a lower cost with faster speed. Our method also provides an efficient way to guide user activities over time-varying networks with link creation and node birth processes. There are many extensions we can imagine along this direction. For example, the jump process in the SDE model can be extended to marked point processes (Jacobsen, 2006), and we can also use nonlinear SDEs to model user activities.

Acknowledgements. This project was supported in part by NSF IIS-1218749, NIH BIGDATA 1R01GM108341, NSF CAREER IIS-1350983, NSF IIS-1639792 EAGER, NSF CNS-1704701, ONR N00014-15-1-2340, DMS-1620342, CMMI-1745382, IIS-1639792, IIS-1717916, NVIDIA, Intel ISTC and Amazon AWS.

References

- Aalen, Odd, Borgan, Ornulf, and Gjessing, Hakon. *Survival and event history analysis: a process point of view*. Springer, 2008.
- Bertsekas, Dimitri P. *Dynamic programming and optimal control*. Athena Scientific, 2012.
- Boel, R and Varaiya, P. Optimal control of jump processes. *SIAM Journal on Control and Optimization*, 15(1):92–119, 1977.
- Chen, Wei, Wang, Yajun, and Yang, Siyu. Efficient influence maximization in social networks. In *KDD*, 2009.
- Chen, Wei, Wang, Chi, and Wang, Yajun. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *KDD*, 2010a.
- Chen, Wei, Yuan, Yifei, and Zhang, Li. Scalable influence maximization in social networks under the linear threshold model. In *ICDM*, 2010b.
- Chen, Wei, Collins, Alex, Cummings, Rachel, Ke, Te, Liu, Zhenming, Rincon, David, Sun, Xiaorui, Wang, Yajun, Wei, Wei, and Yuan, Yifei. Influence maximization in social networks when negative opinions may emerge and propagate. In *SDM*, 2012.
- De, Abir, Valera, Isabel, Ganguly, Niloy, Bhattacharya, Sourangshu, and Rodriguez, Manuel Gomez. Learning opinion dynamics in social networks. *arXiv preprint arXiv:1506.05474*, 2015.
- Dormand, John R and Prince, Peter J. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.
- Du, Nan, Song, Le, Gomez-Rodriguez, Manuel, and Zha, Hongyuan. Scalable influence estimation in continuous-time diffusion networks. In *NIPS*, 2013.
- Du, Nan, Wang, Yichen, He, Niao, and Song, Le. Time sensitive recommendation from recurrent user activities. In *NIPS*, pp. 3492–3500, 2015.
- Epanchin-Niell, Rebecca S and Wilen, James E. Optimal spatial control of biological invasions. *Journal of Environmental Economics and Management*, 63(2):260–270, 2012.
- Farajtabar, Mehrdad, Du, Nan, Gomez-Rodriguez, Manuel, Valera, Isabel, Zha, Hongyuan, and Song, Le. Shaping social activity by incentivizing users. In *NIPS*, pp. 2474–2482, 2014.
- Farajtabar, Mehrdad, Wang, Yichen, Gomez-Rodriguez, Manuel, Li, Shuang, Zha, Hongyuan, and Song, Le. Co-evolve: A joint point process model for information diffusion and network co-evolution. In *NIPS*, pp. 1954–1962, 2015.
- Grover, Aditya and Leskovec, Jure. node2vec: Scalable feature learning for networks. In *KDD*, 2016.
- Hanson, Floyd B. *Applied stochastic processes and control for Jump-diffusions: modeling, analysis, and computation*, volume 13. Siam, 2007.
- Hawkes, Alan G. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- He, Xinran and Liu, Yan. Not enough data?: Joint inferring multiple diffusion networks via network generation priors. In *WSDM*, 2017.
- He, Xinran, Rekatsinas, Theodoros, Foulds, James, Getoor, Lise, and Liu, Yan. Hawkestopic: A joint model for network inference and topic modeling from text-based cascades. In *ICML*, pp. 871–880, 2015.
- Iacus, Stefano M. *Simulation and inference for stochastic differential equations: with R examples*, volume 1. Springer Science & Business Media, 2009.
- Jacobsen, Martin. *Point process theory and applications*. Springer, 2006.
- Kempe, David, Kleinberg, Jon, and Tardos, Éva. Maximizing the spread of influence through a social network. In *KDD*, pp. 137–146. ACM, 2003.
- Leskovec, J., Backstrom, L., and Kleinberg, J. Memetracking and the dynamics of the news cycle. In *KDD*, 2009.
- Lian, Wenzhao, Henao, Ricardo, Rao, Vinayak, Lucas, Joseph E, and Carin, Lawrence. A multitask point process predictive model. In *ICML*, pp. 2030–2038, 2015.
- Nowzari, Cameron, Preciado, Victor M, and Pappas, George J. Analysis and control of epidemics: A survey of spreading processes on complex networks. *IEEE Control Systems*, 2016.
- Ogata, Yoshihiko. On lewis’ simulation method for point processes. *IEEE Transactions on Information Theory*, 27(1):23–31, 1981.
- Ogura, Masaki and Preciado, Victor M. Stability of spreading processes over time-varying large-scale networks. *IEEE Transactions on Network Science and Engineering*, 3(1):44–57, 2016.
- Oksendal, Bernt Karsten and Sulem, Agnes. *Applied stochastic control of jump diffusions*, volume 498. Springer, 2005.
- Pan, Jiangwei, Rao, Vinayak, Agarwal, Pankaj, and Gelfand, Alan. Markov-modulated marked poisson processes for check-in data. In *ICML*, pp. 2244–2253, 2016.

- Pastor-Satorras, Romualdo, Castellano, Claudio, Van Mieghem, Piet, and Vespignani, Alessandro. Epidemic processes in complex networks. *Reviews of modern physics*, 87(3):925, 2015.
- Pennebaker, James W, Boyd, Ryan L, Jordan, Kayla, and Blackburn, Kate. The development and psychometric properties of liwc2015. *UT Faculty/Researcher Works*, 2015.
- Peters, Jan and Schaal, Stefan. Policy gradient methods for robotics. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*. IEEE, 2006.
- Pham, Huy en. Optimal stopping of controlled jump diffusion processes: a viscosity solution approach. In *Journal of Mathematical Systems, Estimation and Control*. Cite-seer, 1998.
- Stulp, Freek and Sigaud, Olivier. Path integral policy improvement with covariance matrix adaptation. In *ICML*, 2012.
- Tan, Xi, Naqvi, Syed AZ, Qi, Alan Yuan, Heller, Katherine A, and Rao, Vinayak. Content-based modeling of reciprocal relationships using hawkes and gaussian processes. In *UAI*, pp. 726–734, 2016.
- Thalmeier, Dominik, G omez, Vicen , and Kappen, Hilbert J. Action selection in growing state spaces: Control of network structure growth. *Journal of Physics A: Mathematical and Theoretical*, 50(3):034006, 2016.
- Wang, Yichen, Du, Nan, Trivedi, Rakshit, and Song, Le. Coevolutionary latent feature processes for continuous-time user-item interactions. In *NIPS*, pp. 4547–4555, 2016a.
- Wang, Yichen, Xie, Bo, Du, Nan, and Song, Le. Isotonic hawkes processes. In *ICML*, pp. 2226–2234, 2016b.
- Wang, Yichen, Williams, Grady, Theodorou, Evangelos, and Song, Le. Variational policy for guiding point processes. In *ICML*, 2017a.
- Wang, Yichen, Ye, Xiaojing, Zha, Hongyuan, and Song, Le. Predicting user activity level in point processes with mass transport equation. In *NIPS*, 2017b.
- Wang, Yichen, Ye, Xiaojing, Zhou, Haomin, Zha, Hongyuan, and Song, Le. Linking micro event history to macro prediction in point process models. In *AISTAT*, pp. 1375–1384, 2017c.
- Zhao, Qingyuan, Erdogdu, Murat A., He, Hera Y., Rajaraman, Anand, and Leskovec, Jure. Seismic: A self-exciting point process model for predicting tweet popularity. In *KDD*, 2015.