

A PREDICTIVE SELF-CONFIGURING SIMULATOR FOR ONLINE MEDIA

Shuochao Yao
Yifan Hao
Dongxin Liu
Shengzhong Liu

Department of Computer Science
University of Illinois
Urbana, IL 61801, USA

James Flamino

Department of Physics
Rensselaer Polytechnic Institute
Troy, NY 12180, USA

Huajie Shao
Jiahao Wu
Mouna Bamba
Tarek Abdelzaher

Department of Computer Science
University of Illinois
Urbana, IL 61801, USA

Boleslaw Szymanski

Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180, USA

ABSTRACT

This paper describes the design, implementation, and early experiences with a novel agent-based simulator of online media streams, developed under DARPA's SocialSim Program to extract and predict trends in information dissemination on online media. A hallmark of the simulator is its self-configuring property. Instead of requiring initial set-up, the input to the simulator constitutes data traces collected from the medium to be simulated. The simulator automatically learns from the data such elements as the number of agents involved, the number of objects involved, and the rate of introduction of new agents and objects. It also develops behavior models of simulated agents and objects, and their dependencies. These models are then used to run simulations allowing for future extrapolations and "what if" analyses. Results are presented on using this system to simulate GitHub transactions. They show good performance in terms of both simulation accuracy and overhead.

1 INTRODUCTION

In the 20th century, advances in computing enabled significant advances in the understanding and exploitation of physical systems due to faster modeling and simulation capabilities. The 21st century is widely believed to be the age of *information*. By interconnecting individuals and endowing them with unprecedented capabilities for sharing information at scale, a new global information fabric has emerged, wherein information signals are created, propagated, distorted, and consumed, not unlike signals on physical media. Yet, the properties of this new information fabric remain poorly understood.

Can advances in cyber-physical computing help uncover the properties of the new information fabric, the way they helped study laws of physical signals, thereby laying the analytic foundations for understanding online information dynamics in the 21st century? Information objects on online media are diverse. They include posts on social networks, such as Twitter, publication entries in databases, such as DBLP, and software updates on version control services, such as GitHub. Collectively, streams of these continually created and updated objects comprise a new reflection of reality and beliefs (e.g., on Twitter), innovation (e.g., on GitHub), and publications (e.g., on DBLP). What are the fundamental properties of this reflection? Does it reveal large-scale phenomena such as competition (e.g., between ideas or beliefs, software toolchains,

or publication pedigrees)? Does it predict the rise and fall in popularity of the underlying objects involved? Does it help understand synergies between ideologies, research directions, or innovations? How does the landscape of synergies and competitions mold the dynamics of online interactions?

A first step towards answering those questions is to develop good simulation tools of information dissemination on online media. Those tools will advance scientific discovery towards addressing the above questions, much the way wind tunnel simulators helped create more efficient aircrafts, and data center simulators helped inform energy management policies.

This paper describes a new simulator of online media. The simulator has two innovative features. First, it is highly retargetable. It can be adapted to simulate multiple different online media that share the abstractions of agents, (information) objects, and actions. Second, it is self-configuring. While a user can decide the configuration being simulated, it is possible to generate the configuration automatically from past traces of the target online medium.

The rest of the paper is organized as follows. Section 2 summarizes the overall architecture of the simulator. Section 3 describes the modules responsible for agent simulation and clustering. Section 4 presents initial evaluation results. Related work is reviewed in Section 5. The paper concludes with Section 6.

2 ARCHITECTURE

Our online media simulations rely on three main abstractions: *agents*, information *objects*, and *actions*. Agents perform actions on information objects. These actions may have dependencies on other actions, which is how signals spread through the medium. For example, on Twitter, tweets constitute the information objects. Users constitute agents. Individual tweets can be created or re-posted (retweeted), which constitutes actions. On GitHub, repositories are information objects. GitHub users are agents. Users can create a repository, fork it, commit content to it, or express interest in it (“star” it), among other capabilities that constitute actions.

2.1 Design Objectives

Our simulation architecture supports two novel features; namely, being retargetable and being self-configuring. These features are introduced below:

- **Retargetable simulation:** Our goal is to simulate multiple types of online media, centered around the abstractions of *agents*, *objects*, and *actions*. As described later in this paper, simulated configurations can be instantiated from data learned from media traces. To make the simulator retargetable, it is, therefore, important to convert media traces into a unified set of supported formats. This is the responsibility of a preprocessing step. Presently, it converts media-specific data structures into both a MySQL database and JSON files. Later modules can thus pick their preferred input format independently of the original source.
- **Self-configuring simulation:** Our online media simulator is novel in that it learns its own agent and object configuration and models from traces of the target online medium. An agent in the simulator is defined by a model that features several parameters. These parameters are learned by the *analysis library*. This library reads traces of online media (from either the MySQL database or the JSON files mentioned above) and computes statistics from which agent model parameters can be determined. Developers are free to implement and experiment with different kinds of machine learning or data mining algorithms, and integrate them into the analysis library to extract behaviors of different agents from the input data. These statistics constitute the *configuration file* needed for the simulator to operate.

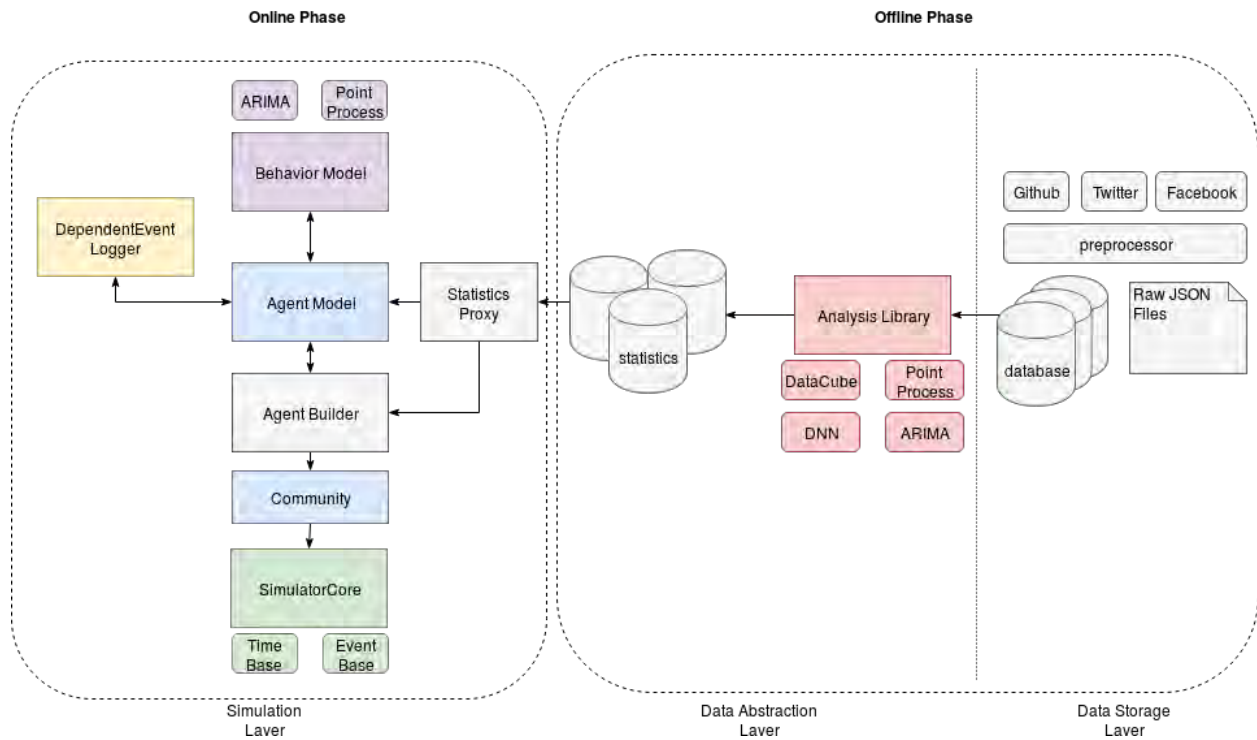


Figure 1: SocialCube simulator architecture.

When the simulator starts, an agent builder module reads the configuration files using a proxy that converts the stored configuration into a run-time instantiation of agent models. The simulator core module then runs the agent models and logs the observed behaviors.

2.2 Simulator Architecture

The simulator architecture is shown in Figure 1. Once a configuration file has been generated, the simulation can start. The *statistics proxy* in Figure 1 serves as a bridge between the offline phase and the online phase. The proxy needs to understand the format of different analysis outputs. A new statistics proxy module needs to be integrated whenever a new type of analysis is implemented in the library. The architecture makes it easy to extend the simulator with new types of analysis techniques.

The *behavior model* is a collection of algorithms describing how agent parameters described in the configuration file determine agent behavior. It is a place to simulate the behaviors and actions of an agent. A new behavior model is also needed whenever a new type of analysis (and hence new agent parameters) are introduced.

The *agent model* is a class that implements individual agents. Each agent will encapsulate a certain behavior model. Each agent (instance of the agent model class) also encapsulates the variables that are specific to the agent. It is essentially a stateful wrapper around the behavior model, whereas the behavior model describes the algorithm that interprets the variables.

The *agent builder* is a factory that instantiates agents during a simulation. When the simulator starts, the agent builder gets a list of agents from the statistics proxy. It then instantiates all the agents of the appropriate type and behavior model. The agent builder also creates new users on the fly according to a pattern dictated in the configuration file.

Some agents represent communities. Like normal agents, they have behavior models, except that the generated actions do not share a single user ID, but rather are distributed across multiple identities. Communities can be used to cluster similar agents that share the same behavior model parameters.

The *event logger* records key past events in the simulation. Developers can adjust the depth of the logger, which in turn will adjust the memory of the simulation. Agent behavior models can condition behaviors based on events logged.

The *simulator core* will hold all the agents and communities after the agent builder instantiates them. The simulator core manages the passage of time, making calls to agents (with reference to current time and logged events), so they can be concurrently simulated.

3 AGENT MODELS AND SIMULATION ALGORITHMS

The simulator offers duality between agents and objects. In that sense, it can either run in an *agent-based* mode, or *object-based* mode. In the agent-based mode, the main abstraction of the simulator is the agent. The core simulation loop executes agents in (simulated) concurrency, one time-slot at a time. The model of each simulated agent is executed, producing the actions the agent will perform on objects during the simulated time interval. In the object-based mode, the main abstraction of the simulator is the object. The model of each simulated object is executed, one time slot at a time, producing the actions that agents will perform on the object during the simulated time interval.

The duality between agents and objects means that the two simulation modes and underlying data structures are symmetric. Below, we explain the data structures used in the agent-based mode. The data structures used in the object-based mode are symmetric and can be described by simply replacing “agent” with “object” and vice versa in the description below.

In the agent-based mode, an agent model is described by (i) total *activity rate* (i.e., rate of actions), (ii) *affinities* (set of objects it interacts with), and (iii) *interaction rates* (frequency of interaction with each object). The total activity rate can also be broken down by action type (e.g., tweets versus retweets). A Twitter user might, for example, have different rates for original tweets versus re-tweets. Users who represent official entities such as local traffic alerts might predominantly post original tweets. In contrast, propaganda publicists might predominantly curate and re-post previously published news that favor their viewpoints, thus primarily posting retweets.

Finally, the above parameters are recomputed at multiple time-scales. The current architecture supports two time-scales for parameter re-computation, called *global* and *local*, respectively. The global time-scale is slower. It adjusts aggregate parameters such as total activity rate and affinities by extrapolating trends described by auto-regressive time-series models. The local time-scale is shorter. It adjusts changes in probabilities of individual action types by extrapolating a more detailed point process model. Below, we describe these models in more detail.

3.1 Long-term Agent Trends

We predict the long-term trends (e.g., trends in total activity rate and affinities of agents) by applying time-series analysis models to extrapolate from observations. Here, we use the autoregressive integrated moving average (ARIMA) for the extrapolation.

We consider the aggregate parameters for each agent (or group of agents) a time series, and use X_t to denote the data vector at time slot t , where t is an integer. The ARIMA model can be expressed by “ARIMA(p, d, q)”, where p refers to the autoregressive terms, d refers to the order of differences in the model, and q refers to white noise terms. We first introduce the difference (Hamilton 1994) of a time series. Let X_t^d denote the d^{th} order difference of X_t :

$$X_t^d = \begin{cases} X_t, & \text{if } d = 0, \\ X_t^{d-1} - X_{t-1}^{d-1}, & \text{if } d \geq 1. \end{cases}$$

Then, the ARIMA(p, d, q) model can be written as:

$$X_t^d - \sum_{i=1}^p \phi_i X_{t-i}^d = c + \varepsilon_t - \sum_{j=1}^q \theta_j \varepsilon_{t-j}. \quad (1)$$

where c is a constant, the sequence $\{\varepsilon_t\}$ consists of independently identically distributed (iid) Gaussian noise terms with mean zero and variance σ^2 , and ϕ_i and θ_j refer to the coefficients of X_{t-i}^d and ε_{t-j} , respectively.

The parameters of ARIMA can be learned by maximizing likelihood. Assuming, without loss of generality, that $c = 0$, the likelihood function for the parameters $\phi_1, \phi_2, \dots, \phi_p, \theta_1, \theta_2, \dots, \theta_q$ and σ is:

$$L(\phi, \theta, \sigma | X) = \frac{1}{\sqrt{2\pi\sigma^2|\Sigma|^2}} \exp\left(-\frac{1}{2\sigma^2}x^T\Sigma^{-1}x\right),$$

where $\phi = [\phi_1, \phi_2, \dots, \phi_p]^T$, $\theta = [\theta_1, \theta_2, \dots, \theta_q]^T$, $x = [X_1^d, X_2^d, \dots, X_n^d]^T$, $\Sigma = \sigma^{-2}\text{cov}(x)$, and n is the number of the training data samples. The exact form of $L(\phi, \theta, \sigma | X)$ is too complicated to maximize directly. Hence, several approximations and modifications were proposed (Anderson 1977; Shaman 1975). After getting the parameters of the ARIMA model, we can predict the value of X_t^d using Equation (1).

For each agent (or group of agents), the analysis library builds an ARIMA model to predict the rate of each type of action and the interaction rates with individual objects over time.

3.2 Agent-Agent Influence

In addition to predicting long-term agent behavior trends, we also estimate the influence between different agents, expressed by directional dependencies between their actions on shared objects. Here, we use the temporal point process model, specifically the Hawkes process model, which is proposed by Hawkes (1971) to describe dependencies across different event types. Point processes provide the statistical language to describe the timing and properties of events.

The Hawkes process is a natural extension of the well-known Poisson process. Instead of assuming that events happen independently as in a Poisson process, the probability of each coming event is dependent on every past event in the Hawkes process (with a decay function that effectively results in finite memory). A Hawkes process sequence can be completely specified by the distribution of its inter-event times:

$$f(t_1, \dots, t_n) = \prod_{k=1}^n f(t_k | t_1, \dots, t_{k-1}) = \prod_{k=1}^n f(t_k | \mathcal{H}_{t_k}). \quad (2)$$

where $\mathcal{H}(t)$ denotes the event history before time t and $f(t_k)$ denotes the probability density function (PDF) of the k th event instant. To derive $f(t | \mathcal{H}(t))$, we need to define the conditional intensity $\lambda(t)$ first, which is heuristically defined as the arrival rate of new events:

$$\lambda = \lim_{\Delta t \rightarrow 0} \frac{\mathbb{E}(N(t + \Delta t) - N(t) | \mathcal{H}_t)}{\Delta t} = \lim_{\Delta t \rightarrow 0} \frac{\mathbb{P}(N(t + \Delta t) - N(t) = 1 | \mathcal{H}_t)}{\Delta t}.$$

where $N(t)$ is the number of events up to time t . The relationship between the conditional intensity function $\lambda(t)$ and the probability density function $f(t | \mathcal{H}_t)$ is given by:

$$\lambda(t) = \frac{f(t | \mathcal{H}_t)}{1 - F(t)}.$$

where $F(t)$ is the cumulative distribution function (CDF) of the given event sequence up to time t . It is also the integral of $f(t | \mathcal{H}_t)$ from the start to time t . Therefore, the conditional intensity function is all we need in order to define a Hawkes process.

Now let's come back to our scenario. We want to use the point process model to describe the dependencies between events of different types. We model it from an object-centric viewpoint, because on online media the operations by different agents on the same object are usually more interrelated than operations by the same agent on different objects. For example, agents may influence other agents to retweet their text, cite

their paper, or view their repository update. The object-centric view-point, therefore, captures influence among agents. For each dimension (event type), i , the conditional intensity function is:

$$\lambda_i(t) = \mu_i + \sum_{j=1}^D \sum_{t_k^j < t} \phi_{ji}(t - t_k^j),$$

where the triggering kernel $\phi_{ji}(t)$ from event type j to event type i is an exponential decay function with respect to the time elapsed:

$$\phi_{ji}(t) = \alpha^{ji} \beta^i \exp(-\beta^i t) 1_{t>0}.$$

In this definition, for an object with L event types, we have three parts of parameters in total:

- μ : a vector of size L . It represents the base intensity of each dimension. It determines the arrival rate of events triggered by extraneous factors.
- α : a matrix of shape $L \times L$. It denotes the triggering coefficients within and across each dimension.
- β : a vector of size L . It represents the decay factor of the triggering effect. We have to set the triggering effect between dimensions as a time decay function, otherwise it will result in an infinite loop of event triggering.

With the definition of our object Hawkes process model, we are able to derive the log likelihood function of seeing the observed event time series. We maximize this function with respect to model parameters using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (Liu and Nocedal 1989), a widely used quasi-Newton method for unconstrained optimization problems without requiring the computation of a second-order Hessian matrix. The resulting mode is used to predict future events using Equation (2).

When both the ARIMA and the point process model are jointly used, we use the former to define long-term trends in total action budgets, and the latter to define action timing within those budgets.

3.3 Communities and Aggregation

Aggregation is a central component for agent-based simulations, as this allows for grouping of both agents and objects by similarity in behavior, which naturally reduces overhead at runtime and can even capture behavior of communities more accurately, if insufficient data are present on individual members.

One type of aggregation is the grouping of agents by similarity in temporal behavior. For example, clear daily cycles exist for a vast majority of users that are well-correlated with their time zones. Some networks, such as GitHub, also see a significant drop-off of activity on weekdays. The concept of agent clustering is not limited, of course, to time zone effects. The simulator leaves it open to define meta-agents that correspond to clusters of agents with similar behaviors. Ideally, such meta agents are used to model communities of agents when one does not have a sufficient number of data points to derive parameters of their individual agent model (e.g., per-agent ARIMA model parameters). If agents are grouped by predicted similarity, the group can collectively have enough data to be modeled accurately. The prediction of similarity is up to individual routines in the analysis library. For example, it may be based on organization, role, time-zone, or a combination thereof.

4 EVALUATION

In this section, we present some of the basic results of running components of the analysis library and simulation on data obtained from GitHub. The dataset used in this section is provided by Leidos. It covers records from Jan 2015 to Aug 2017. All the records have been anonymized to protect privacy of Github users. Each record specifies information of actor (user agent), payload, timestamp, action type, as well as repository id (object id).

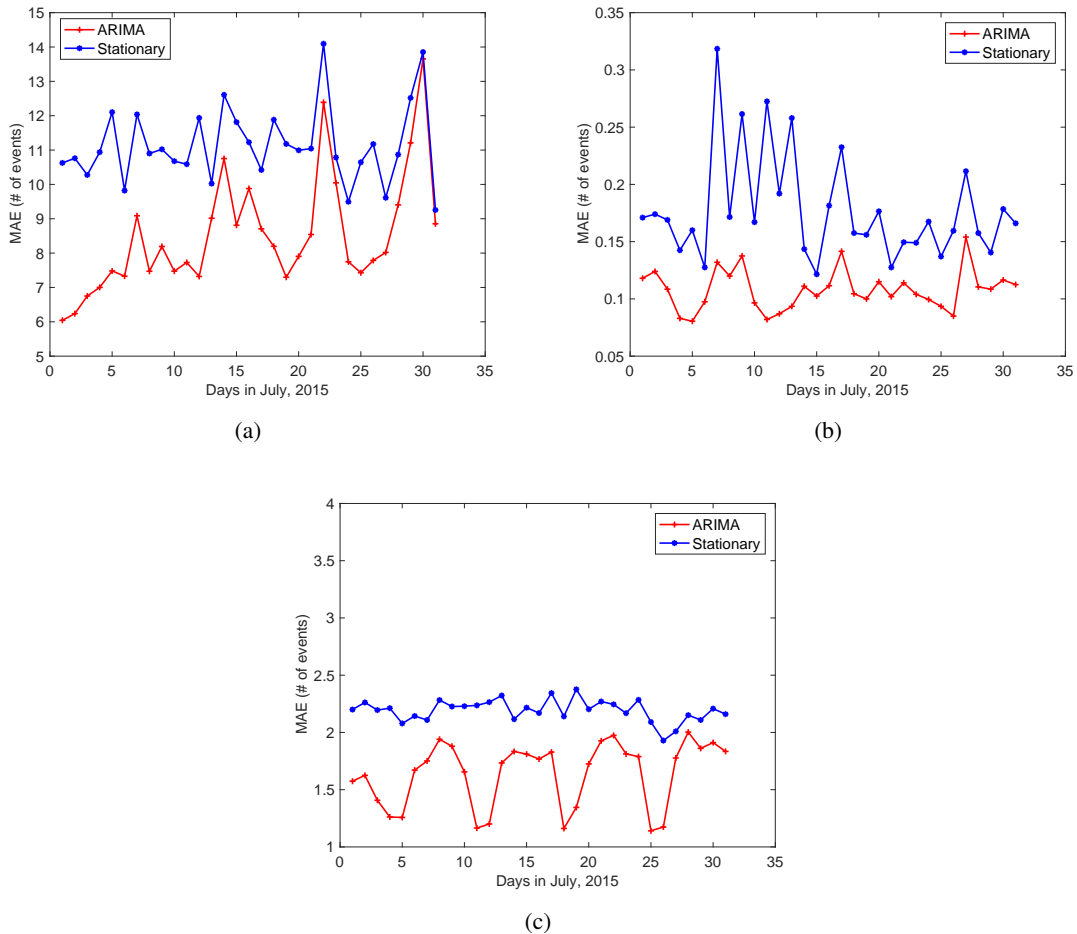


Figure 2: Comparison of MAE for ARIMA and stationary on (a) Push events, (b) Fork events, and (c) PullRequest events.

4.1 Predicting Activity Rate and Type

In this experiment, we are going to predict the number of different event types for a test period (month) for each user on Github. On Github, there are about 10 different event types such as PushEvent, ForkEvent, DeleteEvent and PullRequest. Due to space limitations, we use only the three more common event types, PushEvent, ForkEvent and PullRequest, as representatives. We compare ARIMA and a stationary method (where the future prediction is set identical to the current rates).

We predict the number of event types for each user using the ARIMA model. It took about two hours to train the analysis library to generate models of 2,000 agents based on six months of training data. Figure 2 shows the mean absolute error (MAE) in predicting the rate of different event types of the top 2,000 users using ARIMA and the stationary method (where future is equal to the present). As expected, we observe that the ARIMA method outperforms the stationary method.

After predicting the number of different event types for each user, we can compute the probability distribution of each event type. We take one user as an example in Table 1. The second row is the number of events predicted by the ARIMA model for this user while the last row is the corresponding probability distribution. This probability distribution would then be part of the configuration file for simulating the desired epoch of the user.

Table 1: Probability distribution of different events for a certain user.

Issues	PullReviewComment	Push	PullRequest	IssueComment	Create	Watch	Fork	Delete	CommitComment
95	197	3796	45	893	25	0	0	4	0
.0188	.0389	.7509	.0089	.1766	.0049	0	0	.0008	0

4.2 Predicting Affinity

In this part, we evaluate our ARIMA model for predicting affinities of the agents, and then compare the performance with the stationary agent model. We use all the Github events from January to June in 2015 to train our ARIMA model and use the data in July as the test set. Per affinity definition in Section 3, for Github, affinities refer to the set of repositories a user contributes to. In the evaluation, we build the agent model for 10^5 users who are active in June 2015, and predict their interaction rates with their repositories in July 2015 based on our ARIMA model. We then compare to the stationary model. We use the difference between the ground truth and the prediction values as the absolute errors, and define the accumulated error as the accumulation of the absolute errors over users.




Figure 3: Comparison between ARIMA and stationary on (a) average absolute error, (b) ratio of smaller absolute errors, and (c) accumulative absolute error.

Figure 3(a) shows the average absolute error of ARIMA and the Stationary model over the 10^5 users. We observe that our ARIMA model generates a smaller absolute error overall. Breaking it down by user, Figure 3(b) shows that ARIMA gives a better error for 78.6% users, while the Stationary model performs better for the remaining 21.4%. Figure 3(c) shows the cumulative absolute errors, verifying the advantages of ARIMA-based prediction.

4.3 Understanding Short-term Dependencies

In this subsection, we briefly present the evaluation results for simulated events provided by the short-term point process model we introduced before. First, we describe the performance of single repository simulation results. Second, we enlarge the scope of comparison and evaluate aggregate performance on a set of popular repositories.

In the first experiment, we randomly choose one repository out of the January 2015 event set and estimate its point process parameters. Then, we simulate this model in February 2015 and compare it with ground truth. The estimated triggering kernels and the cumulative distributions of number of events over time are shown in Figure 4. From the triggering kernel map, we can see that the different event types on this repository are relatively independent because the triggering coefficients are higher in diagonal elements except for the IssuesEvent. As presented in this figure, our point process model maintains the daily activity level (i.e., slope of curve) and the burstiness (i.e., smoothness of curve) well in the simulation.

In addition, we compare other statistical metrics in Table 2. In this table, the diffusion delay is defined as the time difference between the time of the first event on this repository and the time of each following event. It represents the average offset of all the events in this repository. Our simulation is within 10% error

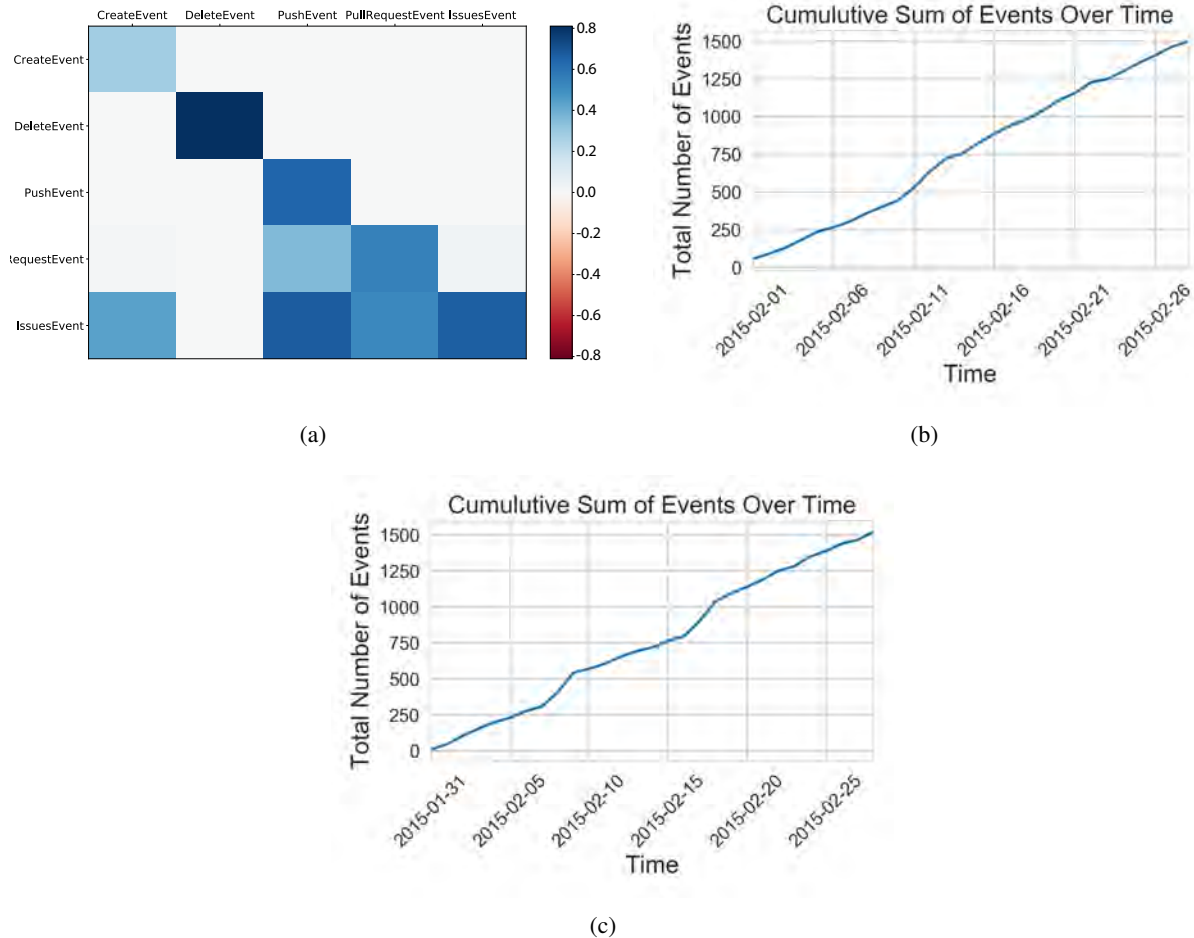


Figure 4: Estimation and simulation results on repository "JwcbE-owDa9ymJm3D4xhsw" based on (a) triggering kernels, (b) groundtruth events, and (c) simulation event.

of the groundtruth. The simulation also performs well on the average inter-event time and total number of events because it is designed to model the triggering effects on inter-event times.

Table 2: Comparison on single repository.

Metric	Groundtruth	Simulation
Average Diffusion Delay	325.599 h	348.077 h
Average Inter-event Time	1575.643 s	1579.644 s
Total Number of Events	1497	1532

Next, to test the performance of our model on more data, we choose the 1,000 most popular repositories in the first half year of 2015 to perform simulation. In this experiment, we only show the WatchEvent and ForkEvent, which reflect the popularity of these repositories. The training period is January 2015 to June 2015, and the simulation is conducted in July 2015. We first independently estimate the parameters for each repository, then perform separate simulations on them, and finally combine the simulated events together to compare with the groundtruth events on these repositories. The comparison results of accuracy are given in Table 3.

Table 3: Comparison on top 1000 repositories.

Metric	Groundtruth	Simulation
Average Diffusion Delay	372.401191 h	373.485502 h
Average Inter-event Time	83815.010660 s	80263.662305 s
Total Number of Events	242247	275133

Table 4: Evaluation results for temporal preference clustering.

	With Activity Level Prediction	Full Simulation
Runtime (sec)	42.5	203.3
RMSE	305.4	273.1

In Table 3, since we include more repositories in the comparison, the diffusion delay distribution is more balanced, resulting in an average delay very close to the middle point of a month. We still obtain good accuracy on average inter-event time (relative error $< 5\%$) and total number of events (relative error $\simeq 12\%$), which proves the stability of the point process model.

4.4 Trend Aggregation

Here, we evaluate both the efficiency and accuracy of clustering by temporal behavior patterns. In the simulation, the offline analysis library is fed with a data set with a date range of 2016-03-01 to 2016-04-01, comprising 27803145 individual events. We compute regional predictions of temporal activity patterns and use them to eliminate agents that are (predicted to be) asleep from simulation at every time step. The results are detailed in Table 4. The table shows that eliminating users predicted to be inactive based on time zone saves 79.1% of simulation time, while the difference in simulation error (RMSE) in predicting hourly PushEvents for the simulated month is only 11.2%. Better clustering remains an open challenge and an active area of development of analysis libraries for our simulator. Regardless, the decrease in overhead is substantial and the speed-up scales well with the number of agents in the training data.

5 RELATED WORK

A key direction in simulation literature is to empower diverse application domains with agent-based simulators for in-depth understanding of real-world complexities, and forecasting quantities of interest. Railsback et al. (2006) proposed an agent-based simulation platform for development recommendations. Drogoul et al. (2002) design a multi-agent-based simulator for distributed artificial intelligence research.

With the proliferation of interconnected users, many studies were made on modelling user behaviour. Wang et al. (2015) jointly model user and object interactions with latent credible status and user dependencies. Yao et al. (2016a) incorporate time information. Yao et al. (2016b) compute a fundamental error bound. Shao et al. (2017) formulate a nonlinear integer programming problem to select optimal sources to solicit in order to minimize the expected fusion error. To the best of our knowledge, this paper develops a novel modular framework for integrating models at different scales to address difference micro or macro concerns while allowing for self-configuration and easy retargeting.

6 CONCLUSIONS

In this paper, we described a new simulator of social media streams. The simulator is novel in that it does not require manual configuration. Rather, the simulated medium is automatically instantiated from analysis of past data traces from the target medium. Prediction algorithms allow for extrapolating from the past trace to a predicted future. Preliminary evaluation results show that the predictions are accurate and scalable. The simulator is extensible. As better prediction and clustering algorithms are developed over time, they can be integrated as modules in the analysis library and as agent behavior models. Similarly,

as new social media are considered, they can be converted into an agent/object/action representation by new thin pre-processing modules, allowing existing analysis algorithms to extract parameters of agents and objects for the new medium. Thus, analysis and simulation code can be reused across different online media. Future work of the authors will address performance optimization of the simulator, extensions of analysis libraries and agent behavior models, and support for additional input media types.

ACKNOWLEDGMENTS

This work was sponsored in part by DARPA under contract W911NF-17-C-0099.

REFERENCES

- Anderson, T. 1977. "Estimation for Autoregressive Moving Average Models in the Time and Frequency Domains". *The Annals of Statistics* 5(5):842–865.
- Drogoul, A., D. Vanbergue, and T. Meurisse. 2002. "Multi-agent Based Simulation: Where Are the Agents?". In *International Workshop on Multi-Agent Systems and Agent-Based Simulation*, edited by J. S. Sichman et al., 1–15. Berlin, Heidelberg: Springer.
- Hamilton, J. D. 1994. *Time Series Analysis*, Volume 2. Princeton: Princeton university press.
- Hawkes, A. G. 1971. "Spectra of some Self-exciting and Mutually Exciting Point Processes". *Biometrika* 58(1):83–90.
- Liu, D. C., and J. Nocedal. 1989. "On the Limited Memory BFGS Method for Large Scale Optimization". *Mathematical programming* 45(1-3):503–528.
- Railsback, S. F., S. L. Lytinen, and S. K. Jackson. 2006. "Agent-based Simulation Platforms: Review and Development Recommendations". *Simulation* 82(9):609–623.
- Shaman, P. 1975. "An Approximate Inverse for the Covariance Matrix of Moving Average and Autoregressive Processes". *The Annals of Statistics* 3(2):532–538.
- Shao, H., S. Wang, S. Li, S. Yao, Y. Zhao, T. Amin, T. Abdelzaher, and L. Kaplan. 2017. "Optimizing Source Selection in Social Sensing in the Presence of Influence Graphs". In *2017 IEEE 37th Int. Conf. on Distributed Computing Systems*, edited by K. Lee and L. Liu, 1157–1167. Piscataway, NJ: IEEE.
- Wang, S., L. Su, S. Li, S. Hu, M. T. A. Amin, H. Wang, S. Yao, L. M. Kaplan, and T. F. Abdelzaher. 2015. "Scalable Social Sensing of Interdependent Phenomena". In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, April 14th–16th, Seattle, USA, 202–213.
- Yao, S., M. T. Amin, L. Su, S. Hu, S. Li, S. Wang, Y. Zhao, T. Abdelzaher, L. Kaplan, C. Aggarwal et al. 2016a. "Recursive Ground Truth Estimator for Social Data Streams". In *Proceedings of the 15th International Conference on Information Processing in Sensor Networks*, April 11th–14th, Vienna, Austria, article 14. IEEE Press.
- Yao, S., S. Hu, S. Li, Y. Zhao, L. Su, L. M. Kaplan, A. Yener, and T. F. Abdelzaher. 2016b. "On Source Dependency Models for Reliable Social Sensing: Algorithms and Fundamental Error Bounds". *36th International Conference on Distributed Computing Systems (ICDCS)*, 467–476. Piscataway, NJ: IEEE.

AUTHOR BIOGRAPHIES

SHUOCHAO YAO is a PhD student in the Department of Computer Science at the University of Illinois Urbana-Champaign, USA (UIUC). His research interests include deep learning on Internet of Things (IoT), cyber-physical systems, and crowd and social sensing. Yao received a BS in Information Engineering from Shanghai Jiao Tong University. His email address is syao9@illinois.edu.

YIFAN HAO is a master student from the UIUC supervised by Prof. Tarek Abdelzaher. His research interests cover the full stack of Internet of Things, including networking infrastructure, low power design for IoT systems, and applications. His email address is yifanh5@illinois.edu.

DONGXIN LIU received the BS and MS degrees in Computer Science from the Shanghai Jiao Tong University, China, in 2014 and 2017. He is currently a PhD student in the Computer Science department of the UIUC. His research interests include social networks, Internet of things, and deep learning. His email address is dongxin3@illinois.edu.

SHENGZHONG LIU is a PhD student at the Department of Computer Science of the UIUC, supervised by Professor Tarek Abdelzaher. His research broadly lies in the intersection of social network analysis, social sensing, and machine learning. He received his BS in Computer Science from the Shanghai Jiao Tong University. His email address is sl29@illinois.edu.

HUAJIE SHAO is a PhD student in the Department of Computer Science of the UIUC. His research interests include data analysis in social networks, applied machine learning, sensor networks, and distributed data centers. Shao received an MS from Zhejiang University. His email address is dongxin3@illinois.edu.

JIAHAO WU is a senior undergraduate student in the Department of Computer Science of the UIUC. His research interests include generative models, machine learning in signal processing, and social sensing. His email address is jwu108@illinois.edu.

MOUNA BAMBA is as Senior Research Programmer in the Cyber Physical Computing Group in the Department of Computer Science of the UIUC. She is also a co-founder of Network Perception, a startup company developing network analysis software and providing services to protect critical infrastructures. Mouna holds a MS degree in Electronics and Communications Engineering from the EPF Engineering School and a master specializing in networks and multimedia information systems from Telecom Bretagne in France. She cumulates over 20 years of professional experience and has a special interest in social networks as well as dependable and intrusion-tolerant distributed systems. Her email address is mbamba@illinois.edu.

TAREK ABDELZAHER is a professor and Willett Faculty Scholar in the Department of Computer Science of the UIUC. His research interests include understanding and influencing performance and temporal properties of networked-embedded social and software systems in the face of increasing complexity, distribution, and degree of interaction with an external physical environment. Abdelzaher received a PhD in quality of service adaptation in real-time systems from the University of Michigan. He has authored or coauthored more than 200 refereed publications in real-time computing, distributed systems, sensor networks, and control. He is a member of IEEE and ACM. His email address is zaher@illinois.edu.

JAMES FLAMINO is a PhD student in Physics at the Rensselaer Polytechnic Institute under Boleslaw Szymanski. He first received his BS in Physics at the same institution. His research fields cover the analysis of social systems and networks primarily through the use of statistical physics, as well as predictive simulation of data trends. James is currently researching spectral embedding for linked social networks and the nature of cascades in these networks when triggered by exogenous events. His email address is flamij@rpi.edu.

BOLESŁAW SZYMANSKI is the Claire and Roland Schmitt Distinguished Professor of Computer Science and the Founding Director of the Network Science and Technology Center at Rensselaer Polytechnic Institute, Troy, NY, USA. He received his PhD in Computer Science from the Institute of Informatics in Warsaw, Poland. He is author or co-author of over four hundred publications and editor of six books. Dr. Szymanski is an IEEE Fellow and a foreign member of the National Academy of Sciences in Poland. His interests focus on network science, especially social networks and computer networks. His email address is szymab@rpi.edu.