



OPEN

Temporal scaling in information propagation

SUBJECT AREAS:

MATHEMATICS AND
COMPUTING

COMPLEX NETWORKS

Junming Huang¹, Chao Li^{1,2}, Wen-Qiang Wang², Hua-Wei Shen¹, Guojie Li¹ & Xue-Qi Cheng¹

¹Institute of Computing Technology, Chinese Academy of Sciences, Beijing, People's Republic of China, ²Web Sciences Center, School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu, Sichuan, People's Republic of China.

Received

2 December 2013

Accepted

8 May 2014

Published

18 June 2014

Correspondence and requests for materials should be addressed to H.-W.S. (shenhuawei@ict.ac.cn)

For the study of information propagation, one fundamental problem is uncovering universal laws governing the dynamics of information propagation. This problem, from the microscopic perspective, is formulated as estimating the propagation probability that a piece of information propagates from one individual to another. Such a propagation probability generally depends on two major classes of factors: the intrinsic attractiveness of information and the interactions between individuals. Despite the fact that the temporal effect of attractiveness is widely studied, temporal laws underlying individual interactions remain unclear, causing inaccurate prediction of information propagation on evolving social networks. In this report, we empirically study the dynamics of information propagation, using the dataset from a population-scale social media website. We discover a temporal scaling in information propagation: the probability a message propagates between two individuals decays with the length of time latency since their latest interaction, obeying a power-law rule. Leveraging the scaling law, we further propose a temporal model to estimate future propagation probabilities between individuals, reducing the error rate of information propagation prediction from 6.7% to 2.6% and improving viral marketing with 9.7% incremental customers.

In recent years, information propagation on social networks has been attracting much attention from academia and industry^{1–9}. Understanding the mechanisms of information propagation, with or without exogenous and endogenous factors, is a fundamental task to uncover the universal laws governing the process of information propagation, which is important for better explaining the dynamics of information propagation¹⁰, predicting information popularity¹¹, and initiating viral marketing campaign^{12–16}. This task, from the microscopic perspective, is formulated as inferring and estimating the propagation probability that a piece of information propagates from one individual to another along social links connecting them.

The difficulty of estimating propagation probability lies in the complex interaction pattern between individuals and the co-existence of various confounding factors, such as the interplay between social selection and social influence. Previous studies empirically identified two classes of factors that drive information propagation: the attractiveness of information and the interactions between individuals. Existing studies on the first class mainly discussed three fundamental mechanisms with respect to message attractiveness¹⁷: the time-invariant intrinsic attractiveness or fitness^{18,19}, the Matthew effect in the popularity accumulation¹⁷, and the freshness of messages decaying in a power-law²⁰, exponential^{21,22}, Rayleigh^{23,24}, or log-normal¹⁷ manner with respect to the time span since the message is posted²⁵. In contrast, most conventional studies on the second class were limited to static or quasi-static scenarios, assuming time-invariant interactions between any pair of individuals. Researchers estimated a propagation probability by indifferently aggregating recent and long-ago interactions^{21,26}, or by learning a probability function with static features including structural characteristics of the underlying network^{11,27–29}, demographic features³⁰, and topical and contextual features^{31–33}. Few studies explored the possibility that individual interactions change with time. A recent study modeled social influence as a Markovian chain on temporally sliced snapshots of a social network, but did not reveal the intrinsic temporal scaling how social influence evolved³⁴.

Actually, most real-world social networks are far from static. On evolving social networks, whether a piece of information will be propagated is more related to instant frequency of individual interactions rather than average frequency indifferently aggregated over recent and long-ago interactions. Hence, it is problematic to neglect the dynamic nature of individual interactions and its crucial role at information propagation, leading to inaccurate predictions. A possible solution is working only on recent interactions based on temporally sliced snapshots of interactions. However, it is hard to determine the appropriate temporal scale of snapshots since the



frequency of interactions is scale-free³⁵. Therefore, we lack a full understanding about the temporal scaling of information propagation, which is crucial to grasp the propagation dynamics of information.

In this report, we study whether and how individual interactions vary temporally and their role at predicting the instant propagation probability. Intuitively, a high frequency of recent communication implies strong instant interaction and a high propagation probability. As the delegate of recency, *latency* is defined as the idle time since the latest communication between two individuals. A long latency generally reflects a low tendency of future interaction. Thus analyzing the interdependence between the latency and the trend of a propagation probability provides us a peculiar delegate for understanding the temporal effect of information propagation. With this delegate, we study on a population-scale social media dataset and conduct an empirical validation for the intuition that a longer latency indicates a relatively lower instant propagation probability.

To focus on analyzing the temporal scaling of propagation probabilities from the perspective of individual interactions, in this report we do not consider the factors of information attractiveness, and instead calculate a propagation probability between two individuals as the ratio of retweeted and neglected messages that are propagated from one to another. This methodology is reasonable when the number of messages is sufficient to largely average out information attractiveness. In this way the temporal scaling of information propagation fully reflects the temporal scaling of individual interactions.

Results

The studies are based on a publicly available dataset (WISE 2012 Challenge, <http://www.wise2012.cs.ucy.ac.cy/challenge.html>) collected from Sina Weibo, the largest Chinese micro blogging website, like Twitter. In the dataset with some simple preprocessing (see Section S1), half a million users created 1.2 million following relations among them, providing channels for propagation of 8 million messages. We denote with an edge (v_i, v_j) the relation that a user v_j (called the *follower*) follows another user v_i (called the *followee*). Each time v_j sees a message k posted or retweeted by v_i that v_j has not retweeted before, we say $\delta_{i,j,k} = 1$ if v_j retweets k , forming a *positive example* indicating v_j successfully activates v_j to retweet k ; otherwise $\delta_{i,j,k} = 0$ for a *negative example* if v_j neglects k . For each positive/negative example, we measure the *latency* $\tau_{i,j,k}$ as the time span since the latest time v_j retweets a message from v_i .

We start to explore the temporal scaling of information propagation by examining time stamps of positive examples on two randomly selected edges, a followee and two of his followers. Figure 1a and Figure 1b reveal a non-uniform density of positive examples that the followers frequently retweet messages from the followee in several short time periods, separated by long idle periods. This implies a *burst* phenomenon on individual interactions: short time frames of intense interactions are separated by long idle periods³⁵. To provide a solid evidence for the existence of burst in retweeting behaviors, we depict in Figure 1e the distribution of latency of all positive examples. The power-law distribution of latency, reflecting the emergence of bursty retweeting behaviors, exhibits the temporal nature of individual interactions. Note that static individual interactions lead to a time-invariant propagation probability on each edge in this scenario, which views retweeting behaviors as a homogeneous Poisson process, resulting in an exponential distribution of latency.

The temporal nature of individual interactions results in a necessity to assign a unique propagation probability to every retweeting/neglecting behavior even occurred on the same edge, reflecting the instant tendency that a follower retweets a followee's message at the time that message arrives. To uncover the temporal scaling of instant propagation probabilities, we investigate the interdependence between the propagation probability behind every retweeting/neglecting behavior and the latency associated with it. The interdependence

is suggested by the distribution of retweeting/neglecting behaviors on those two edges against associated latency, where most retweeting behaviors occur with short latency (Figure 1c and 1d). We calculate the ratio of retweeting and neglecting behaviors over all edges to estimate the invisible instant propagation probability given certain latency. The propagation probability decreases with the latency in a power-law manner (Figure 1f). Fitting the log-log curve in Figure 1f produces a consistently decaying speed of -0.71 slope, suggesting the temporal scaling between a propagation probability $Pr(\delta = 1)$ behind a retweeting/neglecting behavior and its associated latency τ as follows,

$$Pr(\delta_{i,j,k} = 1) \propto \tau_{i,j,k}^{-0.71}. \quad (1)$$

We further study whether retweeting behaviors on different edges share the same power exponent, governing the temporal scaling. As shown in Figure 1a–d, although the retweeting behaviors on the two edges both obey the power-law temporal scaling, the power exponents are quite different. Therefore, we need to assign an edge-specific exponent on each edge, in order to model the temporal scaling of information propagation on various edges of social networks.

Motivated by the observed temporal scaling, we propose a temporal model, namely *Decay model*, to predict propagation probability. We evaluate the performance of the model by applying it to predict retweeting behaviors and to launch a viral marketing strategy, compared with four mainstream baselines, namely MLE, EM²⁶, Static Bernoulli²¹, and Static PC Bernoulli²¹.

The first evaluation experiment measures the probability a model correctly predicts whether or not an individual will retweet an incoming message. Figure 2a reports AUC, the area under the Receiver Operating Characteristic (ROC) curve, equivalent to the probability that a classifier correctly distinguishes a positive example from a negative one. The Decay model outperforms all baselines, raising AUC from 93.3% to 97.4%. Intuitively speaking, when facing a randomly selected pair of a retweeting behavior and a neglecting behavior, the error rate to incorrectly distinguish them is reduced by a half by the Decay model over the best baseline. We then report the *perplexity* on the testing set against the training set ratio to obtain the probability that a model, trained with incomplete observations, correctly generates the testing examples. As shown in Figure 2b, the Decay model achieves the lowest (best) perplexity among all tested models. The priority of the Decay model is consistent in all examined training set ratios, with a more significant improvement on a relatively smaller training set. We also evaluate the Decay model with ROC curve, which is a metric appropriate for extremely imbalanced datasets such as the one we use in this report (as well as most real-world social media) where positive examples occupy less than 1%. ROC, measuring the sensitivity (true positive rate) against specificity (one minus false positive rate), is insensitive to the ratio between positive and negative examples. Figure 2c reports the ROC curves of the Decay model and baselines with 90% examples held out as the training set. Results of other training set ratios are similar. The figure shows that the Decay model achieves the best capability at distinguishing retweeting behaviors from neglecting behaviors with a significant improvement upon all baselines.

The second evaluation measures the accuracy a model predicts propagation probabilities. Intuitively, predictions that are more accurate would help select a better initial seed set, triggering a larger fraction of individuals. We split all examples into 4 groups in a chronological order with respect to example time stamps. Each group contains examples in 30 weeks (see Section S6 for details). The Decay model and baselines train on examples in the earlier 205 days (training phase) and predict the propagation probabilities in the last 5 days (evaluation phase). Based on those predictions, a state-of-the-art influence maximization algorithm (CELF++¹⁵) is used to select an initial seed set maximizing the expected eventual influence spread. We then estimate the *pseudo actual* spread of such a seed set as the number of nodes reachable from the seed set on a propagation network, which is

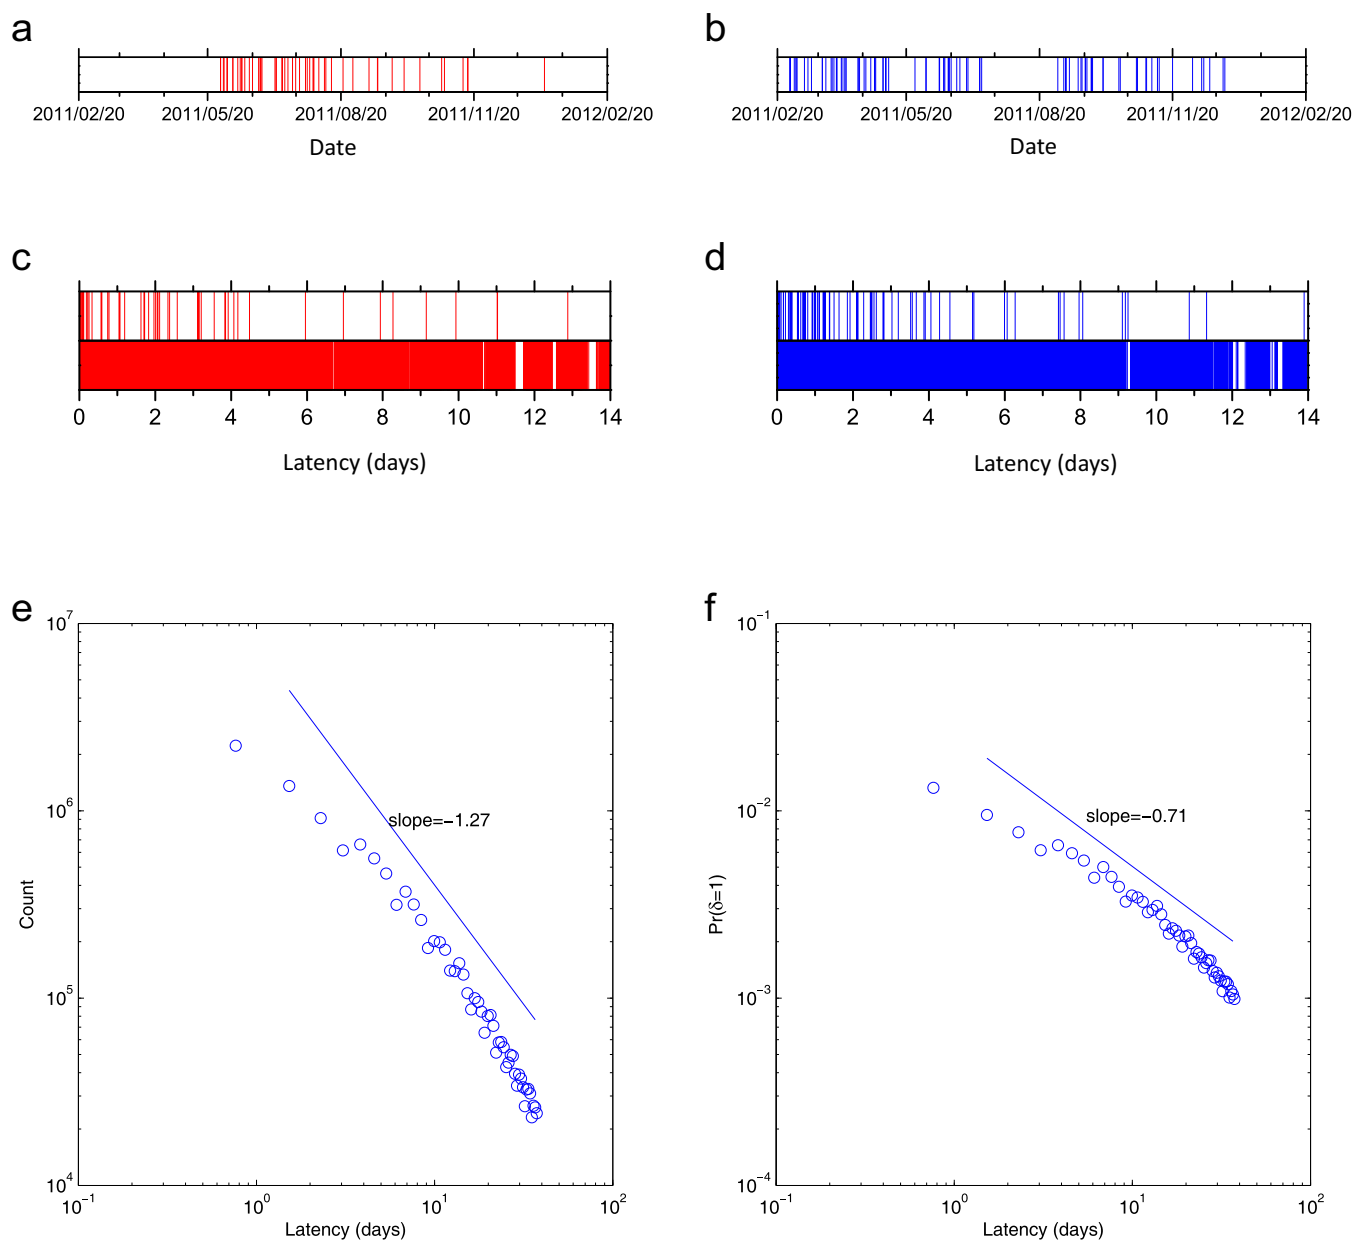


Figure 1 | Characterizing propagation probabilities. (a,b) Time stamps of positive examples (retweeting behaviors) on two random edges. Each vertical line represents a retweeting behaviors occurring with the time stamp marked on the horizontal axis. (c,d) Positive (retweeting) and negative (neglecting) examples on those two edges. Vertical lines in upper half represent positive examples, while those in lower half represent negative ones. It shows an obvious tendency that most positive examples are concentrated on the left zone, i.e., most retweeting behaviors occur with short latency. The tendency is stronger on (c) than that on (d). (e) Distribution of latency of retweeting behaviors over all edges. (f) Ratio of positive examples upon all examples on all edges with respect to the associated latency, demonstrating the power-law interdependence between the propagation probability and the latency.

a subgraph of the social network consisting of edges with at least one actual retweeting behavior in the last 5 days. As reported in Figure 2d (one group shown only), the largest pseudo actual spread comes from the seed set selected on propagation probabilities predicted by the Decay model, which eventually reaches 2,590 nodes, achieving a 9.7% increase upon what is reached by the best baseline, i.e., Static PC Bernoulli which reaches 2,361 nodes. The increase in pseudo actual spread demonstrates the advantage that the Decay model more accurately predicts the propagation probabilities, confirming our finding that individual interactions decay with latency.

Discussion

In this report, we uncovered the temporal scaling in information propagation from the perspective of individual interactions: a

propagation probability decays slowly in a power-law manner with the latency since their latest interaction. Such a dynamic nature was demonstrated by empirical studies on a large-scale public social media dataset, showing the power-law interdependence between a propagation probability and latency.

With the observed temporal scaling, a Decay model was proposed to predict future propagation probability among individuals, incorporating a time-invariant base probability and a time-decaying exponent on each edge. The model is applicable in scenarios where an underlying social network and tractable information propagation with time stamps are observed, such as micro blogging (Twitter and Sina Weibo), blog sites, book sharing sites and email promotion networks. Empirical evaluations supported that the Decay model outperformed mainstream baselines in predicting retweeting behaviors,

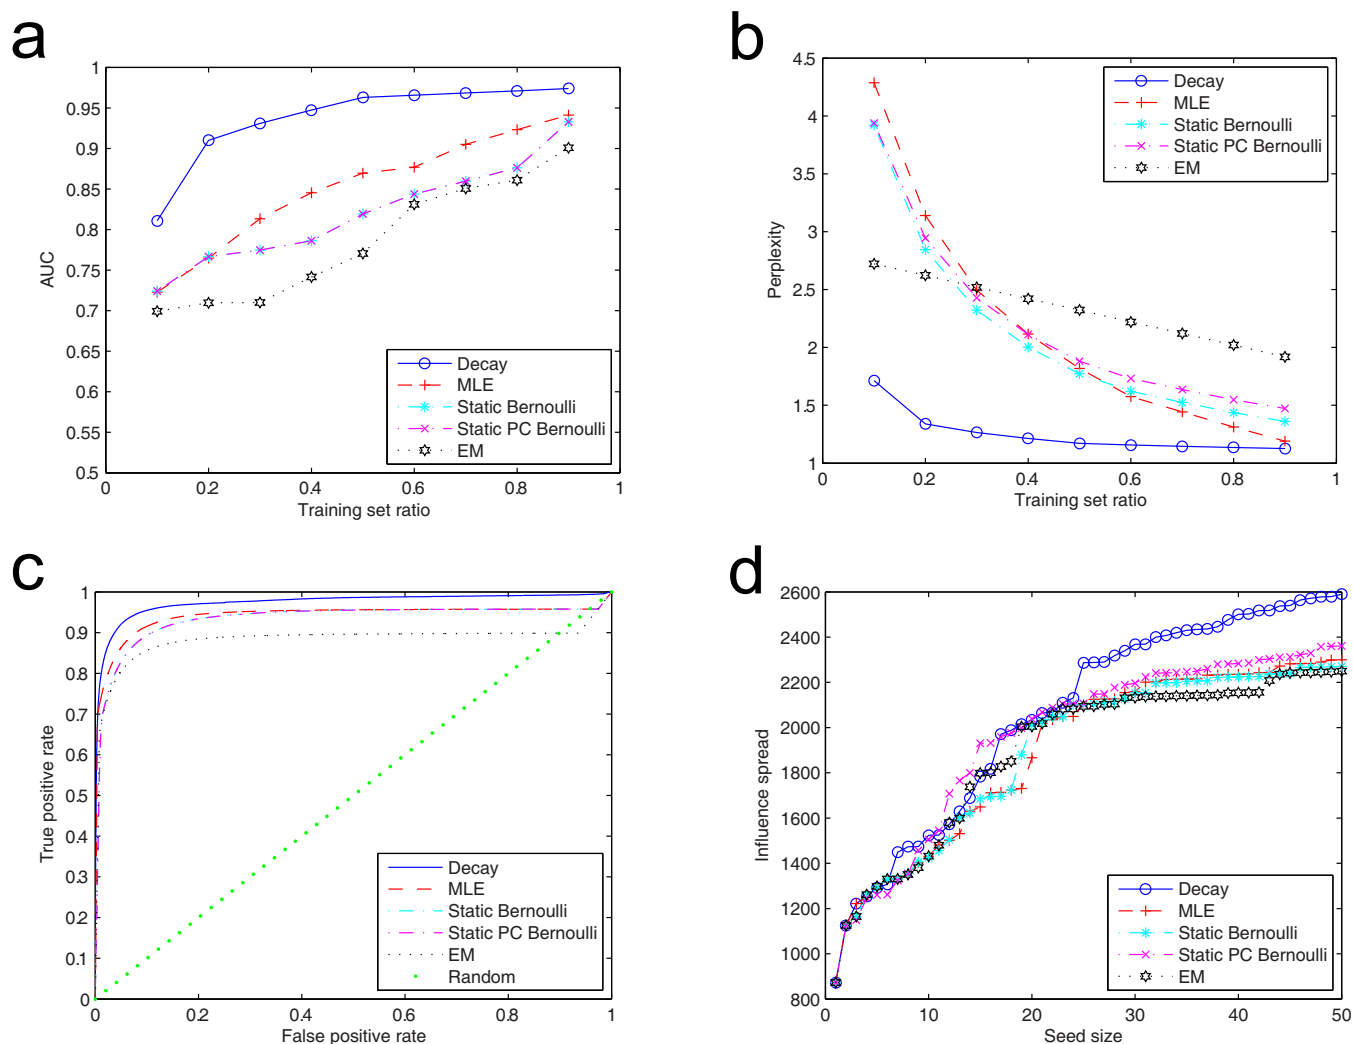


Figure 2 | Model evaluation. (a) AUC of the Decay model and baselines. AUC measures the area under the ROC curves, and thus is equivalent to the probability that a trained model correctly distinguish a randomly selected positive example from another randomly selected negative example. (b) Perplexity of the Decay model and baselines when predicting retweeting behaviors, against the training set ratio. A lower perplexity indicates a better prediction accuracy, meaning less extent a testing example surprises a trained model. (c) Receiver Operating Characteristic (ROC) curves with a training set of 90% examples. (d) Influence spreads of an initial seed set selected on propagation probabilities predicted by the Decay model and baselines.

significantly reducing by a half the expected error rate of incorrectly identifying a retweeting behavior.

From the perspective of machine learning, the discovered temporal scaling provides an additional feature to estimate propagation probability. While traditional models assume static propagation probability, the proposed Decay model additionally explores the temporal effect of a propagation probability, explaining the increased accuracy. Generally speaking, a model with more features requiring more data for training suffers severe over-fitting problem on sparse data. This partly explains why traditional models do not consider temporal features. In order to reduce the pain of sparsity, the Decay model introduces a prior distribution of the decaying exponent $p(x)$, suggested by the global decaying exponent in empirical study results. The prior distribution successfully reduces the pain of sparsity: the improvement of the Decay model upon baselines is even more significant with a relatively smaller training set (Figure 2a and 2b). Note that typically only several retweeting behaviors are observed on an edge in a real-world scenario, the outstanding performance of the Decay model on sparse data is of great importance in practice.

It is worth noting that the viral marketing evaluation is not conducted using Monte Carlo simulations, as done in most influence maximization studies. That is because what we compare is the

configurations of propagation probabilities estimated with various model, and thus it is unfair to run Monte Carlo simulations with any estimated configuration, otherwise estimating all probabilities equal to one will surely win. Instead, we estimate the propagation spread in a pseudo-actual way. We build a propagation network, a subgraph of the social network, with edges where at least one retweeting behavior occurs in the 5-day evaluation phase. Therefore the reachability of a node on the propagation network measures its pseudo actual influence spread during that 5 days. It is equivalent to one Monte Carlo simulation that is produced from the (unknown) actual individuals and observed by actual retweeting behaviors. The estimated propagation spread is deterministic without any random deviation.

In the Decay model, the base probability q is considered as a free variable whose value is fully determined by maximum-a-posteriori inference with a prior distribution. In fact, the Decay model can certainly incorporate any endogenous or exogenous factors through rewriting q as a function of those factors, such as demographical, structural, content and context features. Parameters of such a function could also be estimated in maximum-a-posteriori inference.

In the first testing evaluation experiment, the Decay model is tested with only one testing example on each edge, for the ease of calculating latency. When facing multiple testing examples (e.g., predicting



whether an individual will retweet a series of messages in a month), one should predict those examples one by one in a chronological order and calculate the expected latency of a later example over the joint probability distribution of predicted results of all previous testing examples.

Choosing the latency as a delegate of recency is equivalent to approximating the information propagation occurrences as a first order Markov process, i.e., only the idle time since the latest interaction, instead of all historical interactions, affects the current decision. Such an approximation, effectively avoiding expensive calculation with a nondeterministic number of parameters required to build a complicated function defined on all historical interactions, succeeds in revealing strong evidence of interdependence between propagation probabilities and latency and in building an outperforming prediction model. That supports the important role that the temporal scaling plays in characterizing a propagation probability.

As an open question in future, it would be attractive to characterizing influential nodes identified with high propagation probabilities estimated by the Decay model, and to demonstrate the evolving distribution of instant influential nodes on a social network.

Methods

The proposed *Decay model* describes the propagation probability $P(\delta_{i,j,k} = 1)$, that an individual v_i will successfully activate another individual v_j to retweet a message k , which is believed to be determined by two factors:

- $q_{i,j} \in [0, 1]$: the base probability associated with the edge (v_i, v_j) ;
- $\tau_{i,j,k} \in [1, +\infty)$: latency, the time span since the latest time v_i activated v_j , i.e., $\tau_{i,j,k} = t_{k,i} - t_{k',j}$, where $t_{k,i}$ is the time stamp when v_i posts or retweets k , and k' is the latest message before k that v_i activates v_j to retweet.

Specifically, the propagation probability is as follows,

$$P(\delta_{i,j,k} = 1) = q_{i,j} \tau_{i,j,k}^{-\alpha_{i,j}}, \quad (2)$$

where $\alpha_{i,j} > 0$ is a decaying exponent associated with the edge (v_i, v_j) . The decaying exponent is edge-specific, with a prior distribution $p(x)$ reflecting the global decaying exponent. Traditional models without temporal scaling of propagation probabilities can be viewed as special cases of the Decay model with constant $\alpha = 0$.

Latency is required to be bounded, i.e., $\tau \geq 1$, to guarantee $P(\delta_{i,j,k} = 1) \in [0, 1]$. Specifically, $\tau_{i,j,k} = 1$ results in that $q_{i,j} \tau_{i,j,k}^{-\alpha_{i,j}} = q_{i,j}$, revealing the intuitive meaning of the base probability that $q_{i,j}$ equals to the probability v_i successfully activate v_j to retweet a message k which arrives immediately after a previous successful activation.

The hidden parameters q and α are inferred with a maximum-a-posteriori estimate with prior distributions $p(q)$ and $p(\alpha)$. See Section S3 for details.

To demonstrate the performance of the Decay model, four mainstream baselines are implemented to estimate and predict propagation probabilities on all edges, including MLE, EM²⁶, Static Bernoulli²¹, and Static PC Bernoulli²¹ (see Section S4). Some other widely used models are not compared because those models require user profiles or message content that are absent in this scenario.

In the retweeting prediction experiment, we apply a next-one strategy to split a training set and a testing set. On each edge, we sort all examples in a chronological order, take the earliest $N\%$ examples as the training set, and leave the next one example as the testing set. Thus the size of the training set increases with $N\%$, the *training set ratio*, while the size of the testing set is a constant equal to the number of edges. With parameters trained on the training set, the Decay model predicts the label δ of examples in the testing set.

The evaluation metrics include *perplexity*, *ROC curve* and *AUC*. The perplexity measures how the testing examples surprise a trained model. A lower perplexity demonstrates better prediction ability.

$$\text{perplexity} = e^{-\frac{\sum_{\{v_i, v_j, k\} \in D_{test}} \delta_{i,j,k} \ln \hat{P}(\delta_{i,j,k} = 1) + (1 - \delta_{i,j,k}) \ln (1 - \hat{P}(\delta_{i,j,k} = 1))}{|D_{test}|}}. \quad (3)$$

where D_{test} represents the testing set, and $\hat{P}(\delta_{i,j,k} = 1)$ is the estimated propagation probability. The Receiver Operating Characteristic (ROC) curve plots sensitivity (true positive rate) against specificity (one minus false positive rate). AUC measures the area under the Receiver Operating Characteristic curve, which is equivalent to the probability that a model correctly distinguishes a randomly selected positive example from a randomly selected negative example. A higher AUC indicates a better distinguishing ability. See Section S5 for details.

1. Leskovec, J., McGlohan, M., Faloutsos, C., Glance, N. & Hurst, M. Patterns of cascading behavior in large blog graphs. *Proceedings of 2007 SIAM International Conference on Data Mining* Minneapolis, USA. Philadelphia: SIAM. (2007).

2. Bakshy, E., Hofman, J. M., Mason, W. A. & Watts, D. J. Everyone's an influencer: quantifying influence on twitter. *Proceedings of the fourth ACM international conference on Web search and data mining* Hong Kong, China. New York, NY, USA: ACM. (2011).
3. Huang, J., Cheng, X.-Q., Guo, J., Shen, H.-W. & Yang, K. Social recommendation with interpersonal influence. *Proceedings of the 19th European Conference on Artificial Intelligence* Lisbon, Portugal. Amsterdam, The Netherlands: IOS Press. (2010).
4. Huang, J., Cheng, X.-Q., Shen, H.-W., Zhou, T. & Jin, X. Exploring social influence via posterior effect of word-of-mouth recommendations. *Proceedings of the fifth ACM international conference on Web search and data mining* Seattle, USA. New York, NY, USA: ACM. (2012).
5. Watts, D. J. & Dodds, P. S. Influentials, networks, and public opinion formation. *Journal of Consumer Research* **34**, 441–458 (2007).
6. Tang, J., Sun, J., Wang, C. & Yang, Z. Social influence analysis in large-scale networks. *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining* Paris, France. New York, NY, USA: ACM. (2009).
7. Muchnik, L., Aral, S. & Taylor, S. J. Social influence bias: A randomized experiment. *Science* **341**, 647–651 (2013).
8. Aral, S., Muchnik, L. & Sundararajan, A. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences* **106**, 21544–21549 (2009).
9. Centola, D. An experimental study of homophily in the adoption of health behavior. *Science* **334**, 1269–1272 (2011).
10. Leskovec, J., Adamic, L. A. & Huberman, B. A. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* **1** (2007).
11. Bao, P., Shen, H.-W., Huang, J. & Cheng, X.-Q. Popularity prediction in microblogging network: a case study on sina weibo. *Proceedings of the 22nd international conference on World Wide Web companion*, Rio de Janeiro, Brazil. New York, NY, USA: ACM. (2013).
12. Kempe, D., Kleinberg, J. & Tardos, E. Maximizing the spread of influence through a social network. *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining* Washington, D.C., USA. New York, NY, USA: ACM. (2003).
13. Leskovec, J. et al. Cost-effective outbreak detection in networks. *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining* San Jose, California, USA. New York, NY, USA: ACM. (2007).
14. Chen, W., Wang, Y. & Yang, S. Efficient influence maximization in social networks. *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining* Paris, France. New York, NY, USA: ACM. (2009).
15. Goyal, A., Lu, W. & Lakshmanan, L. V. Celf++: optimizing the greedy algorithm for influence maximization in social networks. *Proceedings of the 20th international conference companion on World wide web* Hyderabad, India. New York, NY, USA: ACM. (2011).
16. Cheng, S., Shen, H.-W., Huang, J. & Cheng, X.-Q. Staticgreedy: solving the scalability-accuracy dilemma in influence maximization. *Proceedings of the 22nd ACM international conference on Information and knowledge management* (ACM, New York, NY, USA, 2013).
17. Wang, D., Song, C. & Barabási, A.-L. Quantifying Long-Term Scientific Impact. *Science* **342**, 127–132 (2013).
18. Lerner, K. & Hogg, T. Using a model of social dynamics to predict popularity of news. *Proceedings of the 19th international conference on World wide web* Raleigh, North Carolina, USA. New York, NY, USA: ACM. (2010).
19. Ye, M., Liu, X. & Lee, W.-C. Exploring social influence for recommendation: a generative model approach. *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval* Portland, Oregon, USA. New York, NY, USA: ACM. (2012).
20. Gomez-Rodriguez, M., Leskovec, J. & Krause, A. Inferring networks of diffusion and influence. *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining* Washington, D.C., USA. New York, NY, USA: ACM. (2010).
21. Goyal, A., Bonchi, F. & Lakshmanan, L. V. Learning influence probabilities in social networks. *Proceedings of the third ACM international conference on Web search and data mining* New York, NY, USA. New York, NY, USA: ACM. (2010).
22. Song, X., Chi, Y., Hino, K. & Tseng, B. L. Information flow modeling based on diffusion rate for prediction and ranking. *Proceedings of the 16th international conference on World Wide Web* Banff, Alberta, Canada. New York, NY, USA: ACM. (2007).
23. Gomez-Rodriguez, M., Balduzzi, D. & Schölkopf, B. Uncovering the temporal dynamics of diffusion networks. *Proceedings of the 28th International Conference on Machine Learning* Bellevue, Washington, USA. International Machine Learning Society. (2011).
24. Manuel Gomez, R., Jure, L., Bernhard, S. & Ipkopf, B. Structure and dynamics of information pathways in online media. *Proceedings of the sixth ACM international conference on Web search and data mining* Rome, Italy. New York, NY, USA: ACM. (2013).
25. Wu, F. & Huberman, B. A. Novelty and collective attention. *Proceedings of the National Academy of Sciences* **104**, 17599–17601 (2007).



26. Saito, K., Nakano, R. & Kimura, M. Prediction of information diffusion probabilities for independent cascade model. *Lecture Notes in Computer Science* **5179**, 67–75 (2008).
27. Zhang, J., Liu, B., Tang, J., Chen, T. & Li, J. Social influence locality for modeling retweeting behaviors. *Proceedings of the Twenty-Three International Joint Conference on Artificial Intelligence* Beijing, China. Menlo Park, California, USA: AAAI Press. (2013).
28. Lü, L., Chen, D.-B. & Zhou, T. The small world yields the most effective information spreading. *New Journal of Physics* **13** (2011).
29. Fang, X., Hu, P. J.-H., Li, Z. & Tsai, W. Predicting adoption probabilities in social networks. *Information Systems Research* **24**, 128–145 (2013).
30. Aral, S. & Walker, D. Identifying influential and susceptible members of social networks. *Science* **337**, 337–341 (2012).
31. Suh, B., Hong, L., Pirolli, P. & Chi, E. H. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. *IEEE Second International Conference on Social Computing* Minneapolis, MN, USA. New York City, USA: Institute of Electrical and Electronics Engineers. (2010).
32. Feng, W. & Wang, J. Retweet or not?: personalized tweet re-ranking. *Proceedings of the sixth ACM international conference on Web search and data mining* Rome, Italy. New York, NY, USA: ACM. (2013).
33. Yang, Z. *et al.* Understanding retweeting behaviors in social networks. *Proceedings of the 19th ACM international conference on Information and knowledge management* Toronto, ON, Canada. New York, NY, USA: ACM. (2010).
34. Wang, C., Tang, J., Sun, J. & Han, J. Dynamic social influence analysis through time-dependent factor graphs. *2011 International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 239–246 (2011).
35. Barabási, A.-L. The origin of bursts and heavy tails in human dynamics. *Nature* **435**, 207–211 (2005).

Acknowledgments

We thank Tao Zhou, Yanyan Lan, Suqi Cheng and Ming Tang for valuable discussions. This work was funded by the National Basic Research Program of China (973 Program) under grant number 2014CB340401, the National High-tech R&D Program of China (863 Program) under grant number 2014AA015103, and the National Natural Science Foundation of China under grant numbers 61232010, 61202215, 61272536.

Author contributions

J.H. and H.-W.S. designed research. J.H., C.L. and W.-Q.W. performed experiments. J.H., C.L., W.-Q.W., H.-W.S., G.L. and X.-Q.C. wrote and reviewed the manuscript.

Additional information

Supplementary information accompanies this paper at <http://www.nature.com/scientificreports>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Huang, J. *et al.* Temporal scaling in information propagation. *Sci. Rep.* **4**, 5334; DOI:10.1038/srep05334 (2014).



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder in order to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>