

Differential Adaptive Diffusion: Understanding Diversity and Learning Whom to Trust in Viral Marketing

Hossam Sharara

Computer Science Department
University of Maryland, College Park

William Rand

Robert H. Smith School of Business
University of Maryland, College Park

Lise Getoor

Computer Science Department
University of Maryland, College Park

Abstract

Viral marketing mechanisms use the existing social network between customers to spread information about products and encourage product adoption. Existing viral marketing models focus on the dynamics of the diffusion process, however they typically: (a) only consider a single product campaign and (b) fail to model the evolution of the social network, as the trust between individuals changes over time, during the course of multiple campaigns. In this work, we propose an adaptive viral marketing model which captures: (1) multiple different product campaigns, (2) the diversity in customer preferences among different product categories, and (3) changing confidence in peers' recommendations over time. By applying our model to a real-world network extracted from the Digg social news website, we provide insights into the effects of network dynamics on the different products' adoption. Our experiments show that our proposed model outperforms earlier non-adaptive diffusion models in predicting future product adoptions. We also show how this model can be used to explore new viral marketing strategies that are more successful than classic strategies which ignore the dynamic nature of social networks.

Introduction

How information diffuses through social networks is a question that has attracted scholars from a wide variety of research disciplines. A richer understanding of the mechanism governing the spread of new ideas or trends in social media has implications for marketing, sociology, journalism, computer science and many other research areas. Models of network diffusion have been used to study phenomena as widespread as product recommendation systems (Leskovec, Singh, and Kleinberg 2006), viral marketing (Domingos 2005; Leskovec, Adamic, and Huberman 2007), disease transmission (Dodds and Watts 2005), herding behavior in financial markets (Drehmann, Oechssler, and Roeder 2005), and even the contagion properties of obesity (Christakis and Fowler 2007). This is in part because the widespread growth and use of online social networks has created a new opportunity to observe diffusion processes on a very large scale, and across different types of interactions from email to microblogging to the sharing of photos.

Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Viral marketing builds upon these network-based diffusion processes. The main goal of viral marketing is to exploit existing social networks among customers by encourage those customers to share product information with their friends. This goal is based on the premise that consumers' purchasing decisions are heavily influenced by recommendations and referrals from their family, friends, and colleagues; an assumption that has been supported by research since some of the earliest studies of diffusion (Ryan and Gross 1943). Recently, viral marketing has become more appealing to marketers as consumers have started to show an increasing resistance to traditional forms of advertising such as TV or newspaper ads.

One of the major early success stories of viral marketing was the introduction of "Hotmail", which was able to gain twelve million subscribers in just eighteen months by adding a simple promotional message with each outgoing email (Jurvetson 2000). Similarly, cell phone companies are another industry where providers take advantage of social network-based diffusion by offering highly discounted rates for customers talking to other customers within the same network. Thus, if a customer's social circle (family, friends, colleagues) is using a certain provider, there's an added incentive for her to use the same provider.

However, social networks are not static. In addition, as consumers continue to listen to their friends and family, they learn that some of their social connections have recommendations that are more appropriate for them and that other members of their social network simply do not have the same interests as they do. This is in part because different individuals are interested in different topics. For someone who is primarily interested in science, if their friend constantly talks to them about new sports developments, sending them emails, and links to promotions for sporting events, that friend is essentially acting as a spammer and the focal individual will eventually decrease their trust in her.

However, if another friend makes a recommendation and the focal individual adopts the product that they recommend then the trust of the focal individual in that friend will increase. As a result of these processes, the social network of confidence changes in time due to the recommendation and adoption process. Although the dynamics of social trust has attracted the attention of multiple researchers (Golbeck 2009), most current viral marketing models do not fully ad-

dress either the fact that social networks change in time, or the heterogeneity of preferences that individuals have for different topics.

In this paper, we present an adaptive model that addresses this shortcoming by allowing individuals to have different preferences for product categories, while adapting their confidence in other individuals' recommendations on the basis of history. This model is novel in that previous models assume the confidence that a user has in other individuals remains constant over time, and that preference for adoption is not dependent on product categories. By incorporating network-level dynamics into a standard diffusion model and allowing for heterogeneous preferences, our model provides a better prediction of expected users' adoption of a given product. We then build upon this model to examine whom to target using viral marketing.

Background

One of the first and most influential diffusion models was proposed by Bass (1969). This model of product diffusion predicts the number of people who will adopt an innovation over time, and though it does not explicitly account for the social network, it does assume that the rate of adoption is dependent on other members of the population, specifically the current proportion who have already adopted. The diffusion equation used by this model describes the cumulative proportion of adopters in the population at any time as a function of the intrinsic adoption rate, and a measure of social contagion. The model describes an S-shaped curve, where adoption is slow at first, takes off exponentially and flattens at the end. The Bass model has been shown to effectively model word-of-mouth product diffusion at the aggregate level (Mahajan, Muller, and Bass 1990), but does not explicitly model the decision of an individual consumer.

Though the Bass model can easily be generalized to address individual-level decisions (Stonedahl, Rand, and Wilensky 2010), most diffusion models that capture the process of adoption of an idea or a product at an individual level use a different mechanism and can generally be divided into two groups: *threshold models* and *cascade models*. Threshold models are based on the work performed by Granovetter (1978) and Schelling (1978) in the late 70's. Basically, each individual, v , in the network has a personal adoption threshold $\theta_v \in [0, 1]$, typically drawn from some probability distribution. A given individual v in the network adopts a new product if the sum of the connection weights of its neighboring peers that have already adopted the product $N(v)$ is greater than her personal threshold:

$$\sum_{u \in N(v)} w(u, v) \geq \theta_v.$$

Although the above model represents a *linear threshold model*, it can be easily generalized further with replacing the summation with an arbitrary function on the set of active neighbors of individual v . Dodds and Watts (2005) have also shown that a more general model than this can be used to describe both the Bass model and the threshold model.

Cascade models (Goldenberg, Libai, and Muller 2001) were originally inspired by research on interacting particle

systems. In these type of models, whenever a peer u of an individual v adopts a given product, then individual v also adopts with probability $p_{u,v}$. In other words, each individual has a single, probabilistic chance to activate each one of her currently inactive peers, after becoming active herself. A very common example is the *independent cascade model*, in which the probability that an individual is activated by a newly active peer is independent of the set of peers who have attempted to activate her in the past. Kempe et al. (2003) proposed a broader framework that simultaneously generalizes the linear threshold and independent cascade models, having equivalent formulations in both cases.

Regardless of the adoption model, one of the key aspects that affects information diffusion is the interaction structure. For instance, a model for product adoption in small-world networks was proposed by Centola et al. (2005), where an individual's probability of adopting a product is dependent on having more than one neighbor who has previously adopted the product. Wu et al. (2004) modeled opinion formation on different network topologies, and found that if highly connected nodes were seeded with a particular opinion, this would proportionally affect the long term distribution of opinions in the network. The work of Holme et al. (2006) focuses on coupling the evolution of both the social network and opinion formation, where both aspects adapt to each other during the evolution process.

Once a diffusion model and a network topology are specified, the next question is which set of individuals should be targeted to maximize the spread of information throughout the network. The problem of influence maximization was formalized by Domingos et al. (2001), who noticed that ordinary data mining techniques that reason about consumer behavior in independent settings, do not utilize network information. They proposed a probabilistic model of user-interaction to study influence propagation in networks, and then explored how to identify a group of individuals, who if they adopted a product, would maximize the speed and amount of adoption throughout the network. Even before Domingos et al. formalized this problem, one hypothesis as to how to maximize diffusion centered around the concept of *influentials*, who are individuals that have a disproportionate effect, compared to average individuals, on the amount and rate of information diffusion. In many information diffusion models, it has been shown that the most influential individuals in a network are the most central, where centrality is measured in a variety of different ways, including the most highly connected nodes, i.e. degree centrality (Wasserman and Faust 1994; Albert, Jeong, and Barabasi 2000). Other solutions have also been proposed, for instance, Stonedahl et al. (2010) show that not only is degree centrality important in maximizing diffusion, but in real social networks it is important to consider the clustering of a node's neighbors since tight clustering slows the diffusion process.

Case Study: Digg

Many popular online social network platforms allow for individuals to recommend items of interest and exchange knowledge. One such example is *Digg.com*, which is a popular social news website, where users can share and vote on




Figure 1: Topic distribution of stories in Digg dataset

different stories, referred to as “digging”, to elevate the ranking of the story on the website. Digg’s users form a social network by “following” other users in the network, which enables automatic tracking of their future diggs and submissions. Each news story on Digg belongs to one of ten topics; *Business*, *Entertainment*, *Gaming*, *Lifestyle*, *Offbeat*, *Politics*, *Science*, *Sports*, *Technology*, and *World News*. We constructed a sample from the Digg network which included both the diggs and follows for 11,942 users and the stories they submitted over a 6 months period (Jul - Dec 2010). The sample include 1.3 million follows relationships among the users, with over 1.9 million diggs, on 48,554 news stories.

The network alone is not enough to describe the diffusion process in a network, it is also important to understand the mechanism by which a user provides recommendations to their peers. These mechanisms differ by platform and marketing strategy. For example, some mechanisms are based on broadcast techniques, where all the peers of a given user are informed when she adopts a product. In other settings, the user has to explicitly select peers to send her product recommendations to after adoption. *Digg.com* uses a broadcast mechanism, where connected users are able to see all the activities of their peers as soon as it is performed.

Analysis

We begin by analyzing the topic distribution of the news stories in the collected data. As shown in Figure 1, though there are differences, all ten topics are represented at comparable levels in our dataset, without a single topic dominating the others. *Technology*, *Entertainment*, and *Lifestyle* are among the topics with higher frequency, while *Gaming*, *Science*, and *Sports* are the ones with lowest number of submissions.

We use the topic distribution of individual user submissions (the actual stories / links they submitted), as opposed to their diggs, as an influence-independent source for determining a user’s topic preferences. Given this topic distribution, we then measure the correlation between the users’ topic preferences and their actual adoptions, i.e., their diggs. Figure 2 shows the Kullback-Leibler divergence between the topic distribution of the users’ submissions versus their diggs. For most users, there is very little divergence between their adoption behavior and their inferred preferences according to their submissions. However, in approximately

10% of the users, there is a quite significant difference between the topic distribution of the stories they digg and the ones they submit. One possible explanation is that while most people adopt only stories of interest to them, there are a smaller percentage of “imitators” who are easily influenced by their peers and do not weight their own preferences as highly. Similar results were obtained using normalized mutual information (NMI) between the topic distribution of users’ preferences and adoptions, with imitators appearing to be even more prominent (~16% of the users).

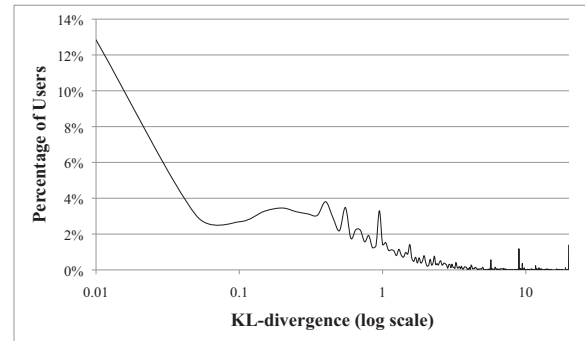


Figure 2: KL-divergence between the topic distribution of users’ submissions and diggs.

In order to characterize users’ topic preferences, we measure the KL-divergence between the topic distribution of each user’s submissions and a uniform distribution of topics. Lower values indicates that the user’s submission pattern is closer to uniform, while higher values indicate that the user is more interested in certain topics but not in others. From Figure 3, we can distinguish three different groups of users in the network: *Focused* users (~53% of the users) who are characterized by having highly skewed preferences towards one or two topics, *Biased* users (~32% of the users) who have less skewed preferences towards a larger set of topics, and *Balanced* users (~15% of the users) who have almost-uniform topic preferences in their submissions.

Finally, we analyze the dynamics of change in the nature of the social relationships between users, and how it affects peer influence over time. We hypothesize that as time passes, peers with similar preferences in topics start gaining confidence in each other’s recommendations, yielding higher levels of adoptions, while on the other hand, peers whose preferences are farther apart from each other become less confident in each other’s recommendations, resulting in lower adoption levels. To test our hypothesis, we measured, at different time points, the average number of diggs on the same story by different peers for different values of KL-divergence between their topic preferences. Figure 4 shows that peers with lower KL-divergence in their topic preferences increase their number of shared diggs over time, while the ones with higher levels of divergence have a decreasing pattern of adoptions over time.

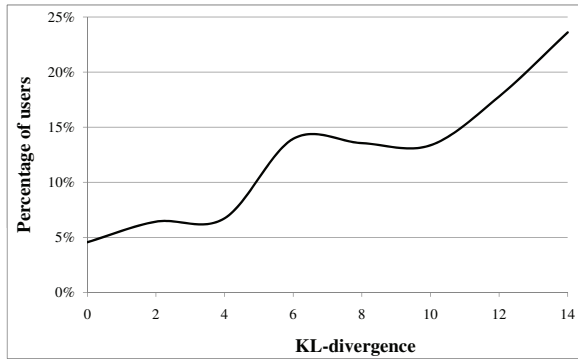


Figure 3: KL-divergence between uniform topic distribution and users’ submissions

Differential Adaptive Diffusion

We view our input social network as a directed weighted graph $G(V, E)$, where V represents the network users, and E represents the social relationships among them. Each edge $e(u, v) \in E$ is associated with a confidence value $w_i(u, v) \in [0, 1]$ representing the confidence user v has in the recommendations of her peer u during campaign i . This confidence value $w_i(u, v)$ is updated only once per campaign, and in general this update could take place either immediately after a recommendation or at the end of a campaign. In the model results presented here, we only update at the end of a campaign. Given a preference function $\mathcal{F}(v, c) : V \times C \rightarrow [0, 1]$ that quantifies user preferences for different product categories $c \in C$ for a given user v , we then define the probability of node v adopting a product of category $c \in C$ within campaign i as a result of node u adopting it as:

$$p(u, v) \triangleq w_i(u, v) \times \mathcal{F}(v, c)$$

To start a new campaign for a certain product x_c of category c , a marketing incentive is provided to a chosen set of seed nodes in the network to initiate the diffusion. As the diffusion process unfolds, the set of nodes who adopt the product at each time step, t , referred to as the “active” nodes, influence their peers through recommendations. These recommendations cause their neighbors to consider whether or not to adopt the product. The adoption function can take any form including any of the functions described in the background section, but throughout the following discussion we will assume an independent cascade process. Thus each active node u in time step t has a single chance of activating a peer v that has not already adopted the product where it succeeds with probability $p(u, v)$, which will result in v adopting the product. Once node u attempts to activate an inactive node v , it can never attempt to activate node v , in any future time step, i.e., node u will return to an inactive but adopted state after this time step. Given the set of active neighbors $N_t(v)$ of a given inactive node v at time t , the posterior probability of v adopting the product at time $t + 1$ can be defined as $p_{t+1}(v, x_c | N_t(v)) = 1 - \prod_{u \in N_t(v)} (1 - p(u, v))$. When a node adopts the product, it becomes active, and starts activating its currently inactive neighbors at future time points.




Figure 4: Heat map of the average number of diggs for different values of topic divergence between peers across time.

The diffusion process continues until no further adoptions occur for the current product.

At the end of each campaign, the confidence values among peers are updated according to the outcome of the product recommendation across the corresponding edge. We denote by $t_i^*(v)$ the time step within campaign i at which a node v adopts the product. If a given node u ends up not adopting the product by the end of campaign i , $t_i^*(u)$ is set to ∞ . Using a kernel function K , the change in confidence values at the end of campaign i for product x_c can be calculated as $\Delta W_{i+1} = K(W_i; \theta)$, where $\theta \in [0, 1]$ is a kernel parameter specifying the rate of change. For instance, a linear kernel can be defined as:

$$K_L(W_i; \theta) = \begin{cases} \theta \times \frac{1-w_i(u,v)}{t_i^*(v)-t_i^*(u)+1}, & t_i^*(u) < \infty \wedge t_i^*(v) < \infty \\ \theta \times \frac{-w_i(u,v)}{t_i^{max}(v)-t_i^*(u)+1}, & t_i^*(u) < \infty \wedge t_i^*(v) = \infty \end{cases}$$

where $t_i^{max}(v) = \max_{t_i^*} \{t_i^*(u) : (u, v) \in E \wedge t_i^*(u) < \infty\}$ represents the time of the last adoption by any of v ’s peers.

This linear kernel assigns credit to each peer u of a node v proportional to the elapsed time between that peer’s recommendation and node v adopting the product. The intuition is that the node u , that last recommended the product, has the highest impact for influencing node v to adopt the product, and thus should be assigned higher confidence in her future recommendations to v . If node v ends up not adopting the product by the end of the campaign, each peer u who recommended the product to node v is penalized relative to the time of the last recommendation. In this case, the last person to recommend the product, even though v still has not adopted it and will not adopt it, gets the maximum penalty for their recommendation.

We can use different types of kernels to control the dynamics of the confidence levels in the network. For instance, this kernel could be exchanged with a kernel where only the last node to provide a recommendation is penalized or rewarded, as opposed to all nodes, or one where all nodes are punished or rewarded equally. Regardless, as a new campaign is initiated for a different product, the new, updated

confidence values are used to compute the influence probabilities, thus enabling the model to capture the dynamics of the diffusion process across different product types.

Experimental Evaluation

To test our proposed model, we used the first four months of interactions, i.e., diggs and submissions, on the Digg network as training data to learn the confidence values between different users, and we used the last two months for evaluation. We use the action of “digging” a story as a proxy for product adoption, and the topic distribution of users’ submissions to estimate their preferences. Starting from a uniform assignment of confidence values across all peers, we track the propagation of user diggs, and update the corresponding confidence values according to the proposed model. We use the learned values along with the user preferences to predict adoptions for new stories.

We compare our approach with two proposed approaches in (Goval, Bonchi, and Lakshmanan 2010) for learning the influence probabilities from training data. In the first approach (Bernoulli), they consider each recommendation a separate Bernoulli trial, and then estimate the confidence between two users as the maximum likelihood estimate (MLE) of the ratio of successful recommendations over the total number within a given contagion time. In the second proposed approach (Bernoulli-PC), they use the same Bernoulli representation but in this approach they give partial credit for each product adoption based to the set of peers who recommended the product within a given time frame. Although both approaches have comparable performance, Goval et al. show that introducing the notion of “contagion time” as a factor in estimating the influence probability outperforms static methods and yields more accurate results.

The above method utilizes a threshold adoption rule as opposed to the cascade rule that we utilize in our model (Adaptive). We can convert between these two models; as shown by Kempe et al. (Kempe, Kleinberg, and Tardos 2003), the independent cascade model is equivalent to a threshold model where the adoption threshold is set to the posterior probability of adoption; i.e. for a given user v , if we set $\theta_v = 1 - \prod_{u \in N(v)} (1 - p(u, v))$, the threshold model is equivalent to the independent cascade model. We use this conversion to facilitate in-depth evaluation of our model. We compare the different models by means of ROC curves, which are more appropriate than precision-recall curves in this setting (Provost, Fawcett, and Kohavi 1998). The ROC curve shows the relative trade-offs between the true positives (correctly identified adoptions) and the false positives (unrealized predicted adoptions) as the discrimination threshold is varied. Each point in the ROC curve corresponds to one possible value of activation threshold for the users.

Figure 5 illustrates the performance of all three models using ROC curves where the x-axis is the false positive rate (FPR) and the y-axis is the true positive rate (TPR). Our proposed model (Adaptive) outperforms both baselines (Bernoulli and Bernoulli-PC), yielding higher true positive rates at low values of false positives. We also experimented with using a predictor that ignores the peer-influence alto-

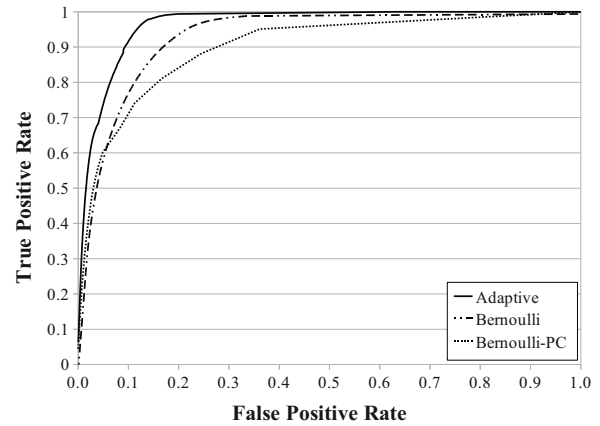


Figure 5: ROC performance of two comparison models (Bernoulli and Bernoulli-PC) and the proposed model (Adaptive) on the basis of the False Positive Rate (FPR) and True Positive Rate (TPR) for each model.

gether and relies only on the stories that were promoted to the “top stories” section in *Digg.com*. This popularity-based predictor yielded an accuracy of 45.7%, which is lower than random prediction. This indicates that individuals’ connections and interactions with their content preferences are more important factors than the overall popularity. Similar results were also confirmed by (Lerman 2007)

These results show that by modeling the dynamics of the diffusion process at a finer-grained level, taking into account the heterogeneity of users and the dynamics of the social network, it is possible to create a model which outperforms a more naïve model. This in turn leads to a better understanding of the whole diffusion process. In the next section we discuss the implications of our model for existing viral marketing strategies, and suggest a new strategy that better captures our findings.

Adaptive Viral Marketing

One of the main implications of our model is a better understanding of the effects of existing viral marketing strategies on social networks in the long term. Our model suggests that user recommendations are most effective when recommended to the right subset of friends. If a user is very selective and makes each recommendation to only a few friends, then the chances of success are slim due to limited network exposure. On the other hand, recommending a product to everyone may have limited returns as well, due to the effect of irrelevant recommendations on the confidence levels between peers. From the perspective of a brand manager interested in maximizing the diffusion of recommendations, it is important to provide incentives to encourage the right balance between reaching as many users as possible and at the same time targeting the most appropriate consumers.

Given this dilemma, a natural question to ask is: what is the appropriate mechanism to maximize both spread and adoption of recommendations? We propose an “adaptive rewards” solution, where instead of rewarding an individual based only on successful recommendations, the reward is




Figure 6: Effect of varying the conservation parameter α

based on successful and unsuccessful recommendations.

Suppose a user v , with p_v peers in the social network, is chosen to start the campaign for a certain product. Assume only m_v of her peers have high preference for that specific product category. Then, whenever a recommendation is successful (a purchase based on a recommendation is carried out), user v gets rewarded ($\alpha \times r$), whereas if the recommendation is unsuccessful, v gets penalized ($(1 - \alpha) \times r$), where α is a conservation parameter, varying from 0 to 1, with 0 representing fully conservative behavior and 1 representing fully nonconservative behavior.

According to the classic viral marketing mechanism, where users only receive rewards for adoptions and no penalties for the lack of adoptions, there is no reason for a user v to be selective in the choice of whom to recommend the product to. In many cases a user will know which subset of her peers are the most probable ones to purchase a given product, based on their knowledge of their peers' preferences. However, there still exists a slight chance for any of v 's peers to purchase the product, including those that do not have a preference for the product, and there is no punishment for failed recommendations. Thus, the expected reward that user v will acquire through sending recommendations to all her peers is greater than or equal to the reward she would receive if she uses a more selective strategy under the classic viral marketing reward mechanism.

However, by utilizing the proposed adaptive rewards mechanism, there is an explicit penalty for unsuccessful recommendations. Following the same setup, if individual v chooses to be selective in recommending the product, thus sending the recommendations only to the interested m_v connections, her expected reward will be $(r \times \alpha \times m_v)$. However, if v chooses to follow a nonconservative strategy, the expected reward is decreased by a penalty relative to her unsuccessful recommendations and becomes $(r \times (\alpha \times m_v - (1 - \alpha) \times (p_v - m_v)))$. Tuning the conservation parameter α allows us to experiment with different mechanisms and their effect on product success and overall confidence levels.

Despite the fact that the main benefits of our proposed strategy appears on the network level through reducing the spamming behavior within the social network, it also carries an advantage for individuals by maximizing their re-

wards over time. While the users have different preferences for different product categories, their judgment in the confidence of their peers is evaluated on an aggregate level. So, if an individual chooses to engage in spamming behavior, this will lead to increased resistance by her peers to any future recommendation they receive from her, regardless of their preference for the product category, thus decreasing that individual's future rewards significantly. As a result, by using our proposed method, individuals must face the penalty of spamming behavior explicitly, and as a result they will be more likely to follow a strategy which will maintain their peers' confidences in them in the long run, and therefore increase their long term reward.

To test the proposed viral marketing strategy, we use an agent-based model to simulate the behavior of users in different settings. First, we generate a synthetic network using preferential attachment (Barabasi and Albert 1999). We use two modes of experiment where we allow the agents to either observe the preference values of their peers before making a recommendation, or learning these preference according to the peer's response to the recommended products. The main objective of each agent is to maximize its expected rewards according to the utilized strategy. Using our proposed adaptive diffusion model, we simulate the diffusion of 500 product campaigns for 5 different categories. We use a linear kernel for adjusting the confidence levels between peers.

Figure 6 shows that by decreasing the value of α , encouraging users to be more conservative in their decisions, the rate of decline in the average confidence level between peers decreases. However, as a side effect of being more conservative, the spread of the products decreases as well. This is illustrated by the fact that the adoption rate is always lower for lower levels of α in the early campaigns, and for very low values of α the adoption rate is always low, but for higher values of α , the adoption rate declines substantially in later campaigns due to the rapid decrease in confidence levels between peers. In fact, utilizing intermediate values for α (e.g. $\alpha = 0.5$, corresponding to equal chances of reward and penalty) consistently maintains high adoption rates and high overall confidence even over a large number of marketing campaigns. We tested the robustness of this result by varying the number of product categories and the size of the




Figure 7: Effect of allowing individuals to learn the preferences of their peers

initial seeding set. The same conclusion holds across all of these changes in the parameters of the systems.

In real settings, users do not necessarily know the preferences of their peers in advance, but rather learn them through peers’ responses to different recommendations. To account for this more realistic situation, we give agents the ability to learn the preferences of their peers instead of directly observing them. At each time step, if the agent decides to recommend a product to one of its peers, it stores whether or not that peer adopted the product. Then, when deciding whether or not to make a recommendation in the future, the agent uses the stored outcomes to estimate that peer’s preference toward different product categories.

The basic hypothesis is that the adoption rate will generally rise over the direct observation case, due to the fact that the agents inference of their peers’ preferences also takes into account the confidence levels, since the peers’ response to recommendations account for both factors. This information is not contained in the direct observation of peer preferences’ and since it is the composite of confidence and preference that determines actual adoption, the agents are better able to predict their peers’ adoptions. This indicates that the adaptive rewards mechanism may work even better in contexts when individuals do not have perfect knowledge of their peers’ preference but must instead learn both the preference and confidence levels from observing past behavior. Moreover, as shown in Figure 7, for moderate values of α , the performance of the proposed strategy is remarkably better than low and high levels of α , in terms of both product adoption and maintaining confidence levels in the network, which indicates that encouraging agents to target a small subset of their peers is the optimal strategy.

In order to analyze our model, we carried out another experiment where we manually inserted a set of spammers into the network. A spammer in our model forwards recommendations for any product it adopts to all its peers, regardless of their preferences. We set ($\alpha = 0.5$) for the rest of the users, and examined various numbers of seeded spammers.

As illustrated in Figure 8, the agents in the network were able to identify the spamming agents after a relatively small number of campaigns, dropping their confidence in them. The effect of spamming behavior is obvious in this figure through the decreased adoption rate as the percentage of

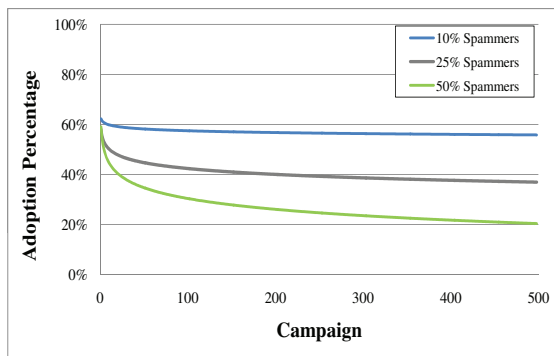
spammers present in the network is increased, but the collective behavior of the non-spammer agents maintains the confidence level among trusted peers, while removing any confidence in spammers. This minimizes the effect of the spamming behavior on the adoption rates over time.

Conclusion and Future Work

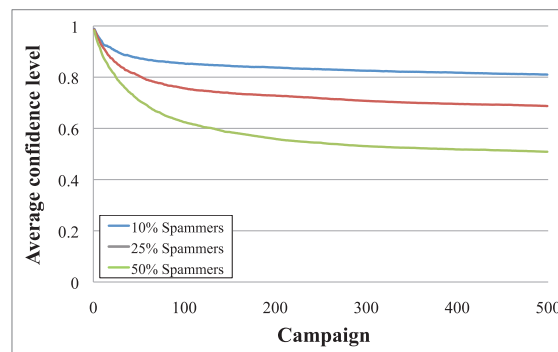
In this work, we provided insight into the effect of network-level dynamics and individual heterogeneity on the diffusion process in real-world networks. Utilizing a sample of users’ interactions on the *Digg.com* social news website, we analyzed the effect of peers’ confidence in each other’s recommendations on the adoption of different products over time. We presented an adaptive diffusion model that is able to capture the observed properties, and showed that it outperforms earlier non-adaptive models in predicting future adoptions.

By analyzing the implications of our proposed model for existing viral marketing strategies, we illustrated that most existing strategies focus on maximizing the product spread within each campaign, but fail to account for the long-term effects that spamming behavior can have on the underlying social network across campaigns. We then introduced a new viral marketing strategy based on our proposed adaptive diffusion model, that accounts for the social network dynamics across different product campaigns. Our experiments have shown that the proposed adaptive viral marketing strategy is able to account for the changes in peers’ confidence across multiple campaigns, maintaining higher levels of product adoptions than those attained by classic strategies in the long term. We also showed that the proposed adaptive strategy is less prone to spamming behavior.

We believe one major application of our work is in identifying influentials. Our model suggests that using only structural-based measures for determining influentials ignore individual behavior, and may lead to decreased efficacy of these strategies in the long run if the chosen individual turn out to be engaged in spamming behavior. One direction for future work is incorporating peer-confidence, by analyzing past interactions, into the process of identifying influentials. Other directions for future work include analyzing the dynamics of change in individual-level preferences, and whether these changes result from peer influence (contagion) or other external factors.



(a) Adoption Rate



(b) Confidence Level

Figure 8: Varying the percentage of spammers at ($\alpha = 0.5$)

Acknowledgments

The authors would like to thank the anonymous reviewers for their valuable feedback. This work was supported by NSF under Grants # IIS-0746930 and IIS-1018361

References

- Albert, R.; Jeong, H.; and Barabasi, A. 2000. Error and attack tolerance of complex networks. *Nature* 406:378 – 382.
- Barabasi, A. L., and Albert, R. 1999. Emergence of scaling in random networks. *Science* 285(5439):509 – 512.
- Bass, F. 1969. A new product growth for model consumer durables. *Management Science* 15(5):215 – 227.
- Centola, D., and Macy, M. 2005. Complex contagion and the weakness of long ties. *American Journal of Sociology*.
- Christakis, N. A., and Fowler, J. H. 2007. The spread of obesity in a large social network over 32 years. *New England Journal of Medicine* 357:370 – 379.
- Dodds, P. S., and Watts, D. J. 2005. A generalized model of social and biological contagion. *Journal of Theoretical Biology* 232:587 – 604.
- Domingos, P., and Richardson, M. 2001. Mining the network value of customers. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*.
- Domingos, P. 2005. Mining social networks for viral marketing. *IEEE Intelligent Systems* 20(1):80 – 82.
- Drehmann, M.; Oechssler, J.; and Roeder, A. 2005. Herding and contrarian behavior in financial markets : An internet experiment. *American Economic Review* 95(5):1403 – 1426.
- Golbeck, J., ed. 2009. *Computing with Social Trust*. Springer.
- Goldenberg, J.; Libai, B.; and Muller, E. 2001. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Letters* 3(12):211 – 223.
- Goval, A.; Bonchi, F.; and Lakshmanan, L. V. S. 2010. Learning influence probabilities in social networks. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining (WSDM)*.

- Granovetter, M. 1978. Threshold models of collective behavior. *American Journal of Sociology* 83(6):1420 – 1443.
- Holme, P., and Newman, M. E. J. 2006. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E* 74.
- Jurvetson, S. 2000. What exactly is viral marketing? *Red Herring* 110 – 111.
- Kempe, D.; Kleinberg, J.; and Tardos, E. 2003. Maximizing the spread of influence in a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*.
- Lerman, K. 2007. Social networks and social information filtering on digg. In *Proceedings of 1st International Conference on Weblogs and Social Media (ICWSM)*.
- Leskovec, J.; Adamic, L.; and Huberman, B. 2007. The dynamics of viral marketing. *ACM Transactions on the Web*.
- Leskovec, J.; Singh, A.; and Kleinberg, J. 2006. Patterns of influence in a recommendation network. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*.
- Mahajan, V.; Muller, E.; and Bass, F. 1990. New product diffusion models in marketing: A review and directions for research. *The Journal of Marketing* 54(1):1–26.
- Provost, F. J.; Fawcett, T.; and Kohavi, R. 1998. The case against accuracy estimation for comparing induction algorithms. In *Proceedings of the Fifteenth International Conference on Machine Learning (ICML)*.
- Ryan, B., and Gross, N. 1943. The diffusion of hybrid seed corn in two Iowa communities. *Rural sociology* 8(1):15–24.
- Schelling, T. C. 1978. *Micromotives and macrobehavior*. Norton.
- Stonedahl, F.; Rand, W.; and Wilensky, U. 2010. Evolving viral marketing strategies. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*.
- Wasserman, A., and Faust, K. 1994. *Social Network Analysis*. Cambridge University Press.
- Wu, F., and Huberman, B. A. 2004. Social structure and opinion formation. *Computational Economics*.