Effect of Persuasion on Information Diffusion in Social Networks

Abhimanyu Das
Microsoft Research
abhidas@microsoft.com

Sreenivas Gollapudi
Microsoft Research
sreenig@microsoft.com

Emre Kiciman
Microsoft Research
emrek@microsoft.com

ABSTRACT

One of the key factors guiding the act of communication between individuals in a social network is the desire to persuade or influence one another. In this paper, we study the interplay between a person writing (selecting) a message to send to another and the effect that the message has on its recipient. Using large-scale online user studies, we focus on a single effect (persuading or changing a recipient's opinion about a topic) and its relationship to various measurable properties of the written message often associated with persuasive techniques, namely the degree of emotional and logical content. We find that the persuasive efficacy of these properties varies by domain of discussion and by individual susceptibility, and that senders appear to strategically select their persuasion techniques. Based on these results, we develop a structural model of information diffusion and show through simulations that the emergent larger-scale behaviors are consistent with current models of information cascades and, moreover, are able to model as yet unexplained empirically observed variance in the structural virality of cascades.

1. INTRODUCTION

Persuasion is the act of convincing someone to take some action. For example, one individual might persuade another to buy a product, vote for a candidate, or share a message on social media. People wield a broad variety of persuasion techniques, from logical argument and emotional appeals to threats and bribes. These techniques and their appropriate applications have been widely studied from many perspectives in fields including psychology, rhetoric, marketing, and neurobiology [21, 13, 12]

While there is no single unified theory of persuasion, it is broadly clear from past research that the effectiveness of persuasion techniques is greatly dependent on the target individual as well as the requested action. While some individuals may be easy to persuade, there are many people who are challenging to convince to act. Moreover, persuasion techniques work well in some contexts, but not in others. Certain persuasion techniques are more likely to work for some individuals than others. In many ways, this is simply common sense. Individu-

als have different attitudes, backgrounds, assumptions, and these individual perspectives are entwined with the context of the action.

In this paper, we present a first investigation into the role of persuasion in information dissemination within a social network. We focus our study on the interplay between an individual who sends a message and the (at least one) person who receives it: can the sender persuade the recipient to adopt the message and repropragate it along the network?

In this setting, through constrained models and experiments, we study the implications of the most basic properties of persuasion—namely, that there are many different persuasion techniques, that the efficacy of a specific technique is contextually dependent, and that individuals vary in their susceptibility to specific techniques. First, through on-line user studies, we measure the contextual efficacy of two persuasion techniques; validate that individual susceptibility to such appeals varies; and that in aggregate, senders appear to strategically select their persuasion technique.

Secondly, based on the learnings from the user studies, we propose a simple structural model of persuasion's role in information dissemination based on the micro-behavior between a sender and recipient. The model associates with each node a d-dimensional vector representation of preferences (or susceptibilities) to different persuasion techniques. A message flowing through the network is also represented by a vector in the same latent space, and its propagation through a node is controlled by the distance between the message and node. Within this model, we also perform preliminary analysis for the problem of selecting a message for maximal propagation.

Finally, through simulations, we study the emergent properties of the large-scale information diffusions that arise from the micro-behaviors specified in our model. We find that, not only are the resultant diffusion trees broadly consistent with the previous models and empirical observations, they are also able to model empirically observe variances in structural virality that have not been captured by past models.

Contributions of this study

We make the following contributions in this paper:

- We perform carefully designed user studies to establish the role of persuasive properties of a message on senders and receivers of such messages across different topics. This lets us characterize the persuasive impact of different message properties on receivers, and the strategic behaviour of senders in choosing these message properties when trying to persuade receivers.

- Based on insights from our user studies, we posit a new analytical model of information diffusion that incorporates how the persuasive properties of a message plays a central role in its diffusion through a social network. We also present some preliminary analytical results for how a sender should optimize her choice of message in the context of our model to maximize its adoption in the network.

- We conduct extensive simulation studies based on our above information diffusion model to characterize properties related to the virality rate and structural virality [15] of information cascades in large scale networks. Our simulations reproduce empirically observed variances in virality of real-world cascades that could not be explained by earlier models.

2. RELATED WORK

There has been a huge body of work in the area of information diffusion through networks. Several early models for information diffusion were inspired from classical disease propogation models in epidemiology, such as SIR and SIS [3]. Other related diffusion models for product marketing included the Bass [4] model that is based on an S-shaped adoption curve.

There has also been extensive work on modeling the adoption or spread of an idea, content or product in a social network. Well known classes of models in this domain include Threshold [19] and Cascade models [17], that specify how a node adopts a particular idea or product based on the adoption pattern prevalent in its neighborhood. In a Threshold model, a node chooses a threshold uniformly at random, and adopts an idea if the (weighted) number of its neighbors who have adopted the idea is greater than the threshold. In the Cascade model, a node u adopts an idea with probability p_{uv} whenever a neighbor v of u also adopts the idea.

A paper that is quite relevant to our work is [15], which propose a formal measure (called structural virality) of the amount of viral diffusion in a cascade. Using this measure, the authors conduct a large scale empirical study of a billion diffusion events for news, videos, images and petitions on twitter, and observe a wide range of diverse cascading structures with varying structural virality, and show a low correlation between popularity and structural virality. The authors then show how a simple SIR model can capture several of the empirically-observed properties of the cascades. However, they note that their model could not explain the large variance in structural virality that they observed empirically.

Another recent related paper [29] also looks at the effect of a message's wording on its propogation on Twitter. The authors take a natural language processing approach to build classifiers based on textual and linguistic features that can predict the likelihood of propogation of a message. Unlike our approach, the authors do not address the problem in the context of information diffusion models.

Kempe et. al [22] analyzed the Threshold and Cascade models in the context of influence maximization: that is, the problem of selecting a set of influential nodes to seed with an idea such that it maximizes adoption of the idea in the entire social network. Several other papers [18, 26] studied various aspects of the phenomenon of diffusion of ideas or content in social networks such as the problem of tracing the path used for diffusion in the network, or of modeling external out-of-network sources in the diffusion process.

Goel et al. [16] empirically observed diffusion patterns over seven different domains including URL cascades corresponding to news stories, videos, and online games, and observed that the vast majority of cascades in practice are small and terminate within a few hops. Leskovec et al [23] studied information diffusion in the context of product recommendation for e-commerce, propose an algorithm for enumerating cascades, and observe that the distribution of cascade sizes follow a power-law model.

Several papers [20, 1] also analyzed and characterized information propogation (URLs, topics, etc) in blogs using well known models of information diffusion. In [24], the authors study cascades in linking propogation patterns for blogs and show that the SIS model is able to mimic the cascades quite well.

In social psychology, there has been a large body on work on persuasion and social influence [7, 10, 30] that talks about various cognitive theories and psychological processes behind how people convince and persuade each other. The book by [27] provides a comprehensive overview. Guadagno and Cialdini discuss persuasion and compliance in the context of Internet-mediated communications, especially textual messages [21]. However, to the best of our knowledge, not of these persuasion processes have been modeled in the context of information diffusion in social networks.

3. PERSUASION TECHNIQUES

Studies in the area of persuasion theory [27, 21, 14] recognize several models for effecting people's opinions. The Elaboration Likelihood Model argues that there are two fundamental methods (or "routes") of information processing and reasoning that occur when a subject considers a persuasive attempt—a central path of reasoning uses logic and deep reasoning; and a peripheral path relies on heuristics. According to the elaboration likelihood model, which path is taken depends on many factors, including the personal receiver's motivation and expertise. Learning theory views persuasion as a specific form of learning, where subjects learn the best response to a persuasive message through conditioning and incentives over time. Cognitive dissonance views persuasion in the context of our desire to reduce dissonance or discomfort due to holding conflicting views. In other words, successful persuasion depends on the elimination of conflict or congnitive dissonance within the subject's mental models about a given topic.

While they have many fundamental differences, substantially all of these theories and models about persuasion broadly recognize that the efficacy of persuasion techniques varies significantly based on topical and individual context. The importance and comprehensibility of the topic, the subjects familiarity, expertise and attitude, all play some role in how (and how much) a person considers and reacts to an attempt to persuade him or her.

In this paper, we focus on a taxonomy of persuasive techniques first proposed by Aristotle in Rhetoric, and still in use as a general framework today [11, 28, 6]. This taxonomy splits modes or strategies of persuasion into three categories, pathos, logos, and ethos:

Logos is a logical appeal or argument, attempting to convince through facts and other detailed information. Such messages often try to simplify the issue around a single key topic to minimize cognitive dissonance. Here are some examples of messages making logical appeals:

```
I like the Surface Pro because I can install
most any program that runs on my regular PC
and it adds in stylus/touch controls.

I don't like the Surface Pro because the
windows app store is not there yet -- some
of the apps are great, but you will find they
are higher priced than you are used to paying
in the google play or App store.}
```

Pathos is an emotional appeal or argument tries to convince through feelings. Such messages may directly mention emotions, like love, hate, or injustice; or indirectly by mentioning things people care about, like family, celebrities, memories.

Here are some examples of messages making emotional appeals:

```
I like the book ''The Valley of Amazement''
because Oprah loved this novel.

I love coke because it reminds me of the
holidays!  Always love their jingles during
the holidays.
```

Ethos is an ethical appeal that depends on the credibility or authority of the author. For example, a leader of a group may attempt to persuade people based on their credibility. Previous studies have shown that author identity is an important signal people use when judging the authoritativeness and credibility of messages on social networks [25].

In the user studies described in the next section, we will focus on measuring the efficacy of the Pathos and Logos modes of persuasion in various contexts. We focus on these two modes because of their ability to be expressed solely through the text of a message (whereas Ethos depends on contextual signals and prior knowledge about the author of a message). Note that we treat these two persuasion strategies as being orthogonal — for example, some messages make both logical and emotional appeals, while other appeals make use neither strategy. Furthermore, there are many persuasion techniques that fall outside of the scope of this taxonomy, such as the use of deception or force, and more complex multi-step strategies of persuasion, such as inoculation. We leave a more complete study of persuasion techniques for future work.

4. ONLINE USER STUDIES

We will now present experiments to establish the importance of the role of message properties on its eventual adoption by network users. Specifically, we confirm that there is a complex relationship between the measurable properties of a message, the domain of discussion, and the message's persuasive impact on recipients. Toward this goal, we ran two sets of experiments that center around the sender and recipient respectively. We choose the dimensions of *emotional* (pathos) and *logical* (logos) in a message to represent two of the three important persuasion techniques. We then measure the effect of these dimensions on the adoption of the message under different polarities (i.e., positive and negative sentiments). Further, we also vary the topics or entities about which the messages are drawn from to tease out the effect of the topic on the choice of dimensions of a message. The results are reported for three topics, *viz.*, Hyundai, Coca Cola, and organic food.

4.1 Effect of message type on the recipient

We begin with the creation of the data set used in this experiment. We manually selected a set of 200 messages, O, for each of the three different topics.

Message	Sentiment	Label
The hyundai has more crashes per vehicle than average. Its not worth the risk	Negative	Logical & Emotional
I had to replace the rear brakes on my hyundai after just 21k miles. Ridiculous!	Negative	Logical & Not Emotional
ugh. My hyundai has cost me so much in repairs. Stay away!	Negative	Not Logical & Emotional
I dont know why exactly, but I prefer honda and ford cars over hyundai	Negative	Not Logical & Not Emotional
if consumer reports says are safe and a good value, thats all I need to know	Positive	Logical & Emotional
the hyundai fits my 6'1" frame. Roomy and drives well	Positive	Logical & Not Emotional
My hyundai is like a member of the family	Positive	Not Logical & Emotional
I like how the hyundai looks	Positive	Not Logical & Not Emotional

Table 1: A sample of messages related to *Hyundai Cars*

In the first step, we gave the messages in $\mathcal{O}$ to judges on Amazon Mechanical Turk and asked them to label them along two dimensions: *logical* or *not logical* and *emotional* or *not emotional*. Each message was presented to 15 judges. We then slotted a message in one of the four quadrants resulting from the $\{logical, not\ logical\}$ and $\{emotional, not\ emotional\}$ labels based on what a majority of judges chose. The inter-annotator agreement for this step was 0.75 showing a large agreement between the judges. In addition, we also asked 15 judges to label each message as containing either a positive or negative sentiment. Table 1 shows a sample of the messages we presented to the user and their categorization into the label space.

In the second step, we aim to quantify the effect of a message's quadrant on the likelihood of its adoption by a user with a certain predisposition (positive or negative) to the topic. For this purpose, we carried out another study in which the user was initially asked to state her message on a given topic. Based on her response on a five-point scale (very negative=0 to very positive=4), she was presented 10 messages of the opposite sentiment, *all* chosen from a quadrant which was chosen uniformly at random. After being shown the 10 messages, she was asked for her message again. This experiment was repeated for all the aforementioned topics. Figure 1 illustrates the relative change in a user's message before and after she is exposed to the messages of the opposite sentiment is topic dependent. In fact, in the case of both *Hyundai cars* as well as *Coke*, there is a increase in the negative sentiment after the treatment and corresponding (although by not as much) decrease in positive sentiment. This suggests that users predisposed with a positive sentiment are more likely to change to a negative sentiment after seeing messages with negative sentiment than the other way around. In order to understand the same effect at a more granular level, we computed the average change from a particular sentiment (e.g., very negative) to another sentiment and measured the change as the difference between the two sentiment levels. For example, we compute the distance between *very negative* and *neutral* to be 2. In other words, we compute the average shift from each sentiment level for each topic and present the results of this experiment in Figure 2. In this figure, for *Organic*
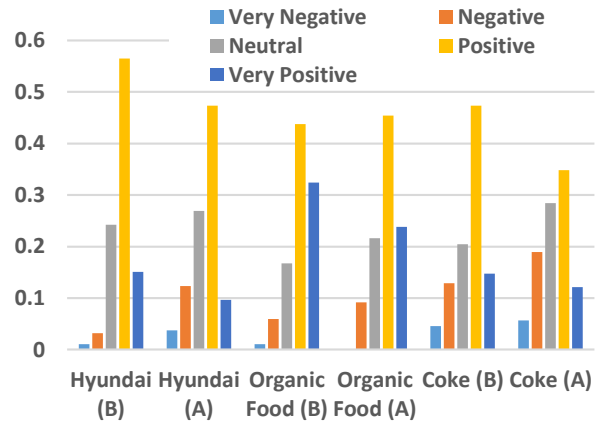

Figure 1: Distribution of the sentiment values before(B) and after(A)

Food, we observe that users with a negative sentiment are more likely to change to a positive sentiment (highlighted by the large average shift out of the *very negative* sentiment. On the other hand, users with an initial *positive* sentiment do not shift their sentiment by as much suggesting that they are more sticky or stubborn with their original message. In the next experiment, we will delve into more detail and try to tease out the actual message type that can likely cause the shifts we are observing in Figure 2.

Toward this end, we analyzed the effect of different dimensions on the change in sentiment. Thus, in each treatment, we presented messages belonging to only one quadrant to the user. We remind the reader that the messages were of the opposite sentiment as the user's. Table 2 summarizes the change in the distribution of sentiments before and after a particular treatment for the *Hyundai* cars topic. Clearly, there is a change in the sentiment distribution after the user sees the treatment. In fact, both *logical* and *emotional* messages induce bigger changes in the sentiments values corresponding to the treatments without such messages.

Finally, we computed the average shift *from* a prior sentiment which measures the overall drift away from that sentiment as a result of the treatment. For example, a change of 2 for a very *negative sentiment* implies that on the average, a user moved up two levels in her

Treatment	Initial					Final				
	(−−)%	(−)%	Neutral%	(+)%	(++)%	(−−)%	(−)%	Neutral%	(+)%	(++)%
Logical	1.0%	3.96%	24.75%	62.38%	7.92%	2.97%	17.8%	26.73%	45.54%	6.93%
Not Logical	1.18%	2.35%	23.53%	49.41%	23.53%	4.71%	5.88%	27.06%	49.41%	12.94%
Emotional	2.08%	2.08%	22.92%	60.42%	12.50%	5.21%	13.54%	30.21%	42.71%	8.33%
Not Emotional	0.00%	4.44%	25.56%	52.22%	17.78%	2.22%	11.11%	23.33%	52.22%	11.11%
Logical & Emotional	2.04%	2.04%	28.57%	65.31%	2.04%	6.12%	20.41%	26.53%	42.86%	4.08%
Neither	0%	2.63%	31.58%	42.11%	23.68%	5.26%	5.26%	18.42%	57.89%	13.16%
Logical Only	0%	5.77%	21.15%	59.62%	13.46%	0%	15.38%	26.92%	48.08%	9.62%
Emotional Only	2.13%	2.13%	17.02%	55.32%	23.4%	4.26%	6.38%	34.04%	42.55%	12.77%

Table 2: The effect of properties of the message on a user for the topic *Hyundai cars*
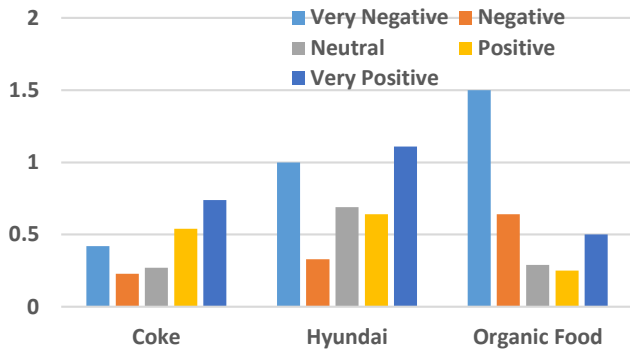

Figure 2: Distribution of the change in sentiment values

sentiment to around *positive*. We note that we also consider combinations of the quadrants in the set of treatments. For example, a treatment of *logical* implies all treatments that include *logical*, i.e., *logical* and *logical and emotional*. Figure 3 shows the average change from initial sentiment for all three topics.

Let us go through the results for each topic. In the case of *Hyundai cars*, we observe that there is movement in both the directions, i.e., users with a positive sentiment change their value downwards while users with negative sentiment move upwards. However, we qualify the moves as follows. While the positive sentiments move for all treatments, albeit in different measures, the presence of *logical* message in the treatment induced significant shifts in the negative sentiments (toward the positive). This becomes clear when one observes the shifts in the treatments *Not Logical*, *Not Emotional*, and *Neither Logical nor Emotional*. In the first and the third treatments where there are no *logical* messages in the treatment, there is *no* shift in the negative sentiments. Finally, we observe that the biggest shift (in both directions) comes when the user is presented with both *logical* and *emotional* messages and the smallest shift is associated when the user is presented with messages that are neither *logical* nor *emotional*.

In terms of the effect of *logical* messages on the shift

in negative sentiments, we observe the same behavior in the case of *Organic Food*. Indeed, all three cases, *viz.*, *Logical*, *Not Emotional*, and *Only Logical* produce the largest shift in the negative sentiment with *Only Logical* producing a shift of 2 levels - well into the positive levels! This suggests, that users with negative predispositions to organic food generally are more open to changing their sentiment when presented with logical explanations or messages while the change in the other direction is less pronounced. This could be attributed to the fact that people with positive predisposition to organic food are those that have already done their "research" and are therefore less likely to get swayed by any kind of message from others.

We observe a slightly different shifting behavior in the case of *Coke*. While the *logical* messages still have more effect on the shift than *emotional* messages, in this case the shift is more pronounced for positive sentiments compared to the topics.

Finally, we measure the relative standard deviation in the average shift in the receiver's sentiments. This measure gives us an insight into the variability in the underlying characteristics associated with a receiver. In particular, a high value of this measure corresponds to a large variability in the population. Figure 4 precisely illustrates such a phenomenon. A large variability allows us to infer different users are persuaded to different extent for a given type of message. Indeed, we will use this observation to motivate our model in Section 5.

4.2 Message selection by senders

In our second set of experiments, we ask people to write a message about a domain-specific item with the specific purpose of either persuading or charming the recipient. We measure the degree of emotional and logical content chosen by the sender and find that it does, indeed, vary by domain.

Then, we analyze and compare the distribution of message attributes chosen by our senders and find that, on average, senders do select the set of attributes that are best aligned for a specific domain and desired effect.

Similar to the recipient experiments, we ran user studies by asking a set of 200 users to create messages on a particular topic (same as the three used in the previ-
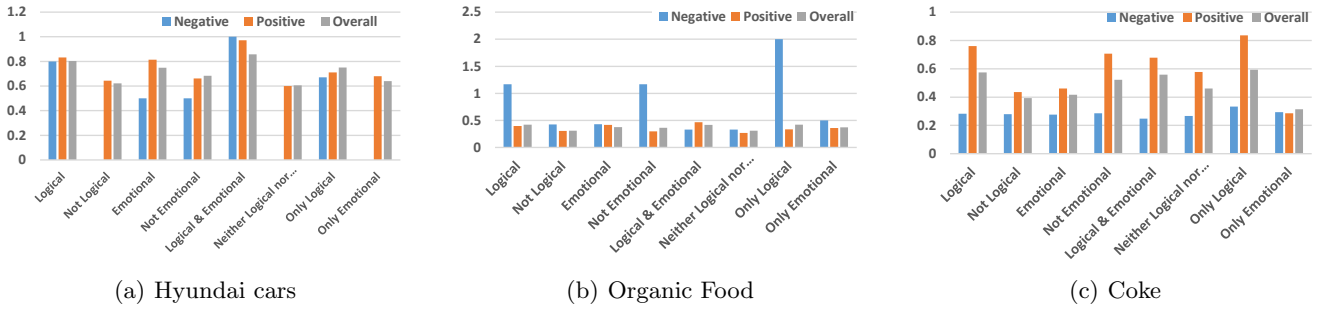
| (a) Hyundai cars | (b) Organic Food | (c) Coke |

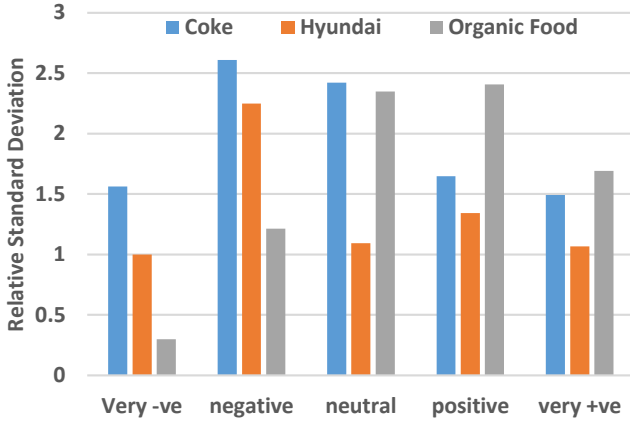**Figure 3: Average change in shift from a sentiment after a particular treatment**

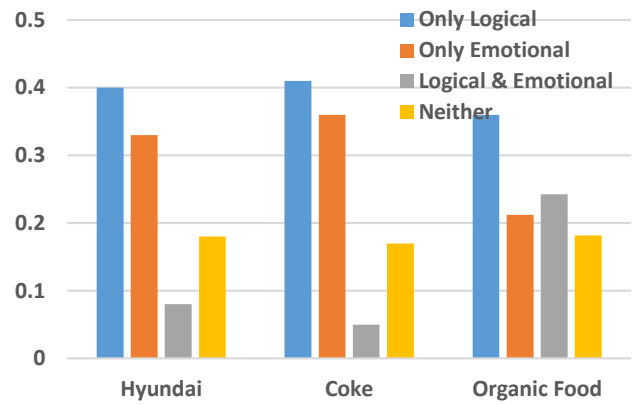**Figure 4: Relative standard deviation in the average shift in receiver sentiments**

Figure 5: Distribution of the message types chosen by the senders

ous set of experiments) with the goal of maximizing its reach. We created a HIT[1] on Amazon Mechanical Turk asking each judge to do the following -

"Your task is to write a short message (limited to one or two sentences) about why you like or dislike a car brand, such that you can convince as many of your friends on Facebook/Twitter who read your message to agree with your opinion. For example, if you have a reason to like Hyundai cars, your message should be written such that it also convinces your friends to like Hyundai cars. On the other hand, if you do not like Hyundai cars, your message should be written such that it persuades your friends to not prefer Hyundai cars."

We then went ahead and got each of the message labeled by 15 users in another user study, again using the Amazon Mechanical Turk platform. Table 3 shows a sample of the messages crafted by the sender and the subsequent categorization of the message. Finally, Figure 5 illustrates the distribution of message types that the senders choose to send in response to our HIT. We observed that the senders showed a clear preference for a specific message type based on the topic. Hence, sug-

---
[1]Human Intelligence Task

gesting that the senders are indeed strategic in their choice of the message. Further, they exhibited a general preference to *logical* messages over *emotional* messages across the three topics.

5. MODEL OF SOCIAL EFFECTS ON MESSAGES

In this section, we describe a simple model of the individual-level behaviors of senders and receivers that capture the two core observations from our crowdsourcing experiments. This model focuses on the interaction between a sender and a receiver: namely, how does the sender of a message reason about the persuasive effect that a message will have on recipients, and how does this influence what messages are actually sent and consequently adopted by users in a social network?

5.1 Motivation for a new model

We make the following observations from our user studies that existing models of information diffusion such as the Linear Threshold and Independent Cascade models [22] do not incorporate.

- Message selection by sender: Our first observation is that senders exhibit relatively consistent

Opinion	Sentiment	Label
Hyundai cars just suck. Mine broke down right after their guarantee period ended.	Negative	Logical & Emotional
Hyundai cars are not giving good milage compared to other car brands	Negative	Logical & Not Emotional
Hyundai should come up with some new designs for their cars	Negative	Not Logical & Emotional
Hyundai cars are unreliable and of cheap quality	Negative	Not Logical & Not Emotional
Hyundai offers a sweet, stylish ride for less money	Positive	Logical & Emotional
Hyundai cars are best comes with good mileage and pickup. Great value for our money.	Positive	Logical & Not Emotional
hyundai cars are very good to look and drive	Positive	Not Logical & Emotional
It looks good in design	Positive	Not Logical & Not Emotional

Table 3: A sample of opinions expressed by the senders related to *Hyundai Cars*

preferences for certain types of messages over others when trying to persuade the recipients of the message. In particular, we observed that across all the three domains in our user studies, most senders preferred composing messages with high-logical content and low-emotional content when asked to persuade their friends on Twitter or Facbook to "adopt" the message. This suggests that senders in a social networks are strategic in selecting their messages to maximize the persuasive effect of the messages, and consequently the adoption of the message in the social network.

- Effect of message type on persuasion of recipient: Our second observation is that the ability of a message containing a positive (c.f. negative) opinion towards an entity to persuade a recipient holding a negative (c.f. positive) opinion to change her mind varies greatly depending on the type of messsage. In particular, messages about an entity with high logical and emotional content were more likely to persuade people to change their opinion than other types of messages. This suggests that the extent to which a recepient of a message is persuaded to adopt the message's opinion depends to a large extent on the type of the message.

- Heterogeneous user interests resulting in different preferences and thresholds at which they are willing to be persuaded (i.e., adopt the sender's message and propagate it to their neighbors). This makes the propagation dynamics depend both the preferences of the recipient as well as the sender's message.

Existing diffusion models do not take into account the role played by the type of a message or opinion when estimating its spread in the social network. This motivates the following model.

5.2 Basic Model

We represent a message as a vector in d dimensional space, where every dimension corresponds to an attribute of the message such as logic, emotion etc. For simplicity, we assume that the range for each dimension lies in $[0, 1]$. Hence every message m is a vector in $[0, 1]^d$

We consider an online social network graph $G = (V, E)$ with nodes $\{v_1, v_2, \ldots, v_n\} \in V$. For ease of notation, we will frequently interchange node v_i and index i. The vertices v_i correspond to individuals, and edges $E = [e_{ij}]$, denote social ties between the individuals. Every vertex v_i has a d-dimensional unit-length preference vector $m_i \in \mathbb{R}^d$ and a persuasion threshold $t_i \in \mathbb{R}$. Vertices in the graph can propagate messages along their edges to their neighbors. For a node v_i, its set of neighbors is denoted by $N(i) = \{j : e_{ij} = 1\}$, and its degree is denoted by $d_i = |N_i|$. We denote the total number of edges in the graph as $|E|$.

The dynamics of message propagation in the graph works as follows: Every vertex in the graph can act as a sender or a receiver of messages. A sender in the graph can compose a d-dimensional message vector m and broadcasts it to its neighboring vertices. A receiving vertex is "persuaded" to adopt an incoming message only if the incoming message is "close" (in some L_p norm) to its preference vector. In particular, whenever a (receiver) vertex v_r receives a message m that it has not previously adopted, it adopts the message and propagates the same message to all other neighbors in N(r) iff $\|m - m_r\|_p < t_r$. Otherwise it drops the message, and does not adopt it. We typically assume p to be 1 (Manhattan distance) or 2 (Euclidean distance). Figure 6 illustrates an example of message propagation in a graph based on this model.

Note that unlike in the case of traditional information diffusion models such as Linear Threshold and Independent Cascade models, the adoption of a message composed by a sender node depends not only on the graph structure but also on the message vector itself. Indeed, as we saw in our user studies, most senders strategically compose their messages to have high logical and emotional content when given a task of persuading their friends to adopt the message. Based on this observation, and using the above model, a natural question is then to understand what message a given sender should compose in order to persuade as many nodes in the graph to adopt the message.

In the next section, we propose an optimization problem that arise from the above question.

5.3 Optimization Problem
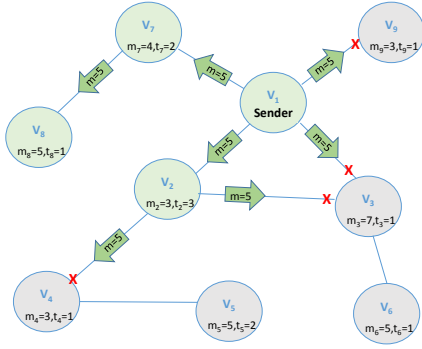
Figure 6: Message Propagation Model: Example of message propagation in a graph with 1-dimensional message vectors, where v_1 is a sender sending a message $m = 3$. Every vertex v_i has a preference vector m_i and persuasion threshold t_i. The green vertices are nodes that have been persuaded to adopt the message, while grey vertices are nodes that have not adopted the message (either they did not receive the message from any of their neighbors, or they received the message but dropped it because the message did not fall within a distance t_i of their preference vector m_i.

More formally, we define the following optimization problem: given a Graph $G = V, E$, preference vectors m_i and thresholds t_i of each node $v_i \in V$, and a sender node $v_s \in V$, choose a message vector m that v_s should compose that maximizes the number of vertices in G that adopt the message.

5.3.1 Complete Graph

We first consider the above problem for the special case of the complete graph. Since all nodes are directly reachable from a sender node, the problem can be reduced to the following version:

$$\operatorname*{argmax}_{m \in [0,1]^d} |v \in V : \|m - m_v\|_p \le t_v| \qquad (1)$$

This is equivalent to the following *depth* problem in computational geometry []: given n hyperspheres (or hypercubes, depending on the choice of p for the L_p norm) in d dimensions, find the point that has the maximum depth, i.e. lies inside the maximum number of hyperspheres (hypercubes). Figure 7 illustrates an example of the depth problem corresponding to the optimization problem for message selection in a complete graph with 8 vertices.

We have the following results for two special cases of this problem:

The first lemma corresponds to the case when $p = 1$.

Lemma 1 *For $p = 1$ and constant d, given a complete*
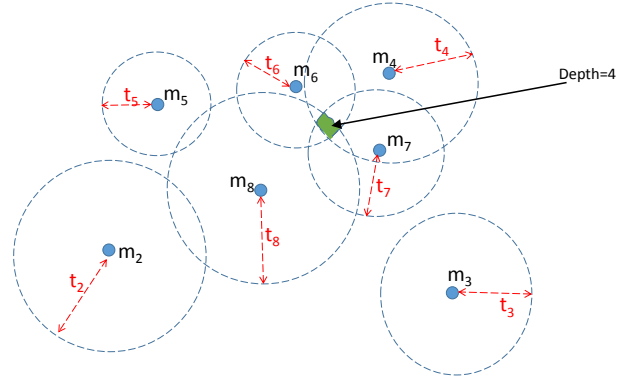
Figure 7: Optimization problem for message selection in a complete graph with vertices $\{v_1, v_2, \ldots, v_8\}$, where v_1 is the sender node seeking to choose a message that maximizes adoption. Every vertex v_i has a preference vector $m_i \in \mathbb{R}^2$ and persuasion threshold t_i, which corresponds to a hypersphere centered at m_i with radius t_i. The problem of selecting a message that is adopted by as many nodes as possible is equivalent to selecting a point in $\mathbb{R}^2$ that lies in the intersection of as many of the hyperspheres (corresponding to $\{v_2, v_3, \ldots, v_8\}$) as possible, or alternatively, a point with maximum depth. This corresponds to any point in the green region, which has a depth of 4.

graph $G = (V, E)$, there is a polynomial time exact algorithm for finding a message $m \in [0,1]^d$ that maximizes the number of vertices that adopt the message. The algorithm runs in $O(2n^d)$ time.

Proof. We note that each of the hypercubes centered around the preference vectors of the vertices are axis-aligned with each other. Furthermore, it is known [5] that the intersection graph of axis-aligned hypercubes satisfies the Helly property. A family of sets F satisfy the Helly property if whenever a subfamily $S \subseteq F$ has non-empty intersection for every pair of sets in S every pair of sets in S, then the whole subfamily has non-empty total intersection. Hence, finding the largest subset of hypercubes with a non-empty total intersection is equivalent to the problem of finding the maximum-clique in the intersection graph of the hypercubes. This is solvable exactly in time $O((2n)^d)$ [9, 8] ∎

An immediate corollary of this result gives us an approximation for the case when $p = 2$.

Corollary 2 *For $p = 2$ and constant d, given a complete graph $G = (V, E)$ and a threshold vector $\mathbf{t} = \{\mathbf{t_1}, \mathbf{t_2}, \ldots, \mathbf{t_n}\}$, let $m^*(G, \mathbf{t})$ be the optimal message in*

$[0,1]^d$ that maximizes the number of vertices adopting the message. Then, there is a polynomial time approximation algorithm that finds a message $m' \in [0,1]^d$ such that $N(m') \geq N(m^*(G, \mathbf{t} \cdot \sqrt{\mathbf{d}}))$ vertices, where $N(m)$ denote the number of vertices adopting a message m. The algorithm runs in $O((2n)^d)$ time.

The next lemma corresponds to the special case when $p = 2$ and $d = 2$.

Lemma 3 *For $p = 2$, given a complete graph $G = (V, E)$, there is a $1 + \epsilon$ approximation algorithm for finding a message $m \in [0,1]^2$ that maximizes the number of vertices that adopt the message. The algorithm runs in $O(n\epsilon^{-2} \log n)$ expected time.*

Proof. The proof follows directly from Theorem 4.6 of [2]. ∎

5.3.2 General Graph

In the case of a general graph, dependency constraints arise for each node which are of the form - "a node sees a message only if there exists at least one path from the sender to the node along which *all* the nodes adopt and propagate the message". Thus, for the general graph case with sender vertex v_s, the optimization problem can be formulated as follows:

$$\operatorname*{argmax}_{m \in [0,1]^d} |v_t \in V : \|m - m_{v_t}\|_p \leq t_v \text{ and } AdoptionPath(s, t) = 1|$$
(2)

where $AdoptionPath(s, t) = 1$ if there exists a path $P = \{v_s, v_{p_1}, \ldots, v_{p_k}, v_t\}$ in G such that for all $v_{p_i} \in P$, $\|m - m_{v_{p_i}}\|_p \leq t_{v_{p_i}}$

For this problem, we provide an approximation algorithm using epsilon-nets.

Lemma 4 *Given a general graph $G = (V, E)$, a sender $v_s \in V$, and a threshold vector $\mathbf{t} = \{\mathbf{t_1}, \mathbf{t_2}, \ldots, \mathbf{t_n}\}$, let $m^*(G, \mathbf{t})$ be the optimal message in $[0,1]^d$ that maximizes the number of vertices adopting the message. For constant d, there is a polynomial time approximation algorithm that finds a message $m' \in [0,1]^d$ such that $N(m') \geq N(m^*(G, \mathbf{t} + \epsilon\sqrt{\mathbf{d}} \cdot \mathbf{1}))$ vertices, where $N(m)$ denote the number of vertices adopting a message m, and $\mathbf{1}$ denotes the all ones vector. The algorithm runs in time $O(|E|\frac{1}{\epsilon}^d)$.*

Proof. We first uniformly discretize the hypercube $[0,1]^d$. It is easy to see that the space $[0,1]^d$ can be covered with $\frac{1}{\epsilon}^d$ grid hypercubes of size ϵ. We consider message vectors corresponding to the centers of each of the grid hypercubes. For every such message vector, we can compute the number of vertices in G that adopt the message by simulating the diffusion process over G in O($|E|$) steps. Thus, by exhaustively enumerating over the centers of all the grid hypercubes, we can find the optimal message vector m' among them in time $O(|E|\frac{1}{\epsilon}^d)$. Since each grid hypercube is of size ϵ, the true optimal message $m^*(G, \mathbf{t})$ has distance at most $\epsilon\sqrt{d}$ from m'. The result follows. ∎

Next, we study through simulations the emergent behavior implied by our core model and extensions.

6. SIMULATION EXPERIMENTS

In this section, we explore the large-scale cascades that emerge from the micro-behavior model of interplay between message senders and receivers. Using the model presented in Section 5.2, our simulation tracks the propagation of a message represented as a d-dimensional vector through a social network. Given the network of individuals, each associated with its own d-dimensional preference vector and threshold t, and an initial starting node, a message's propagation is deterministic.

In the rest of this section, we first present details of our simulation setup. Then, we characterize the relationships between key parameters of the model and the *virality rate* of information cascades. Finally, we study the *structural virality* of information cascades generated within our model, and find that our model captures properties of structural virality that have been empirically observed but not captured by previous models [15].

6.1 Simulation Details

In each of our simulation experiments, we build a 1 million node network and execute 100k information cascades. Our simulation model is governed by 2 parameters that control the creation of the network structure and the propagation of messages: d specifies the dimensionality of the latent message space and preference vectors; and t the maximum threshold acceptance distance for re-broadcasting messages.

There are two key stages to building and running an information dissemination model. The first stage is the initialization of the network of connections among nodes. The nodes of the network are each randomly assigned a desired message vector and a random threshold $\leq t$. The network is initialized as an asymmetric network. Every node's in-degree is assigned according to a power law distribution parameterized by $\alpha = 2.3$.

The association of a preference vector to nodes allows us to experiment with peer influence and homophily within our model. To do so, we add two additional parameters to control the network initialization. To represent the effect of homophily—that an individual is more likely to create a social tie to a user similar to herself, we replace the random selection of a neighbor with a preferential selection, where the closest neighbor in the d-space is selected from a random candidate pool of size
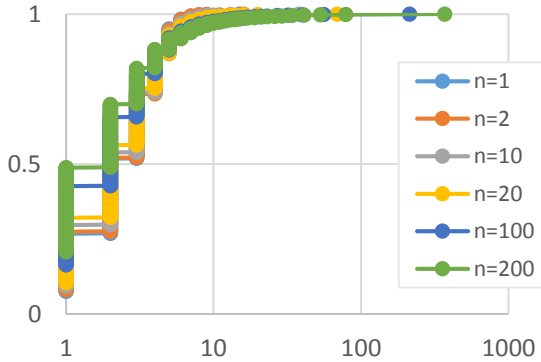
Figure 8: The out-bound degree distribution of nodes.

n, and the candidate pool is fully replaced for each selection. While the distribution of out-bound degrees is not directly controlled, this process does generate a heavy-tailed distribution, as shown in Figure 8.

To represent peer influence, s controls the smoothing of preference vectors between a vertex in the graph and its neighbors. More formally, each node's preference vector $m_i \leftarrow (1 - s) * m_i + s * avg(N_i)$ where N_i is the set of vertices neighboring i.

The second stage of the simulation—propagating a message through the nodes of this network—maps directly to the microbehavior of a sender and a recipient described in Section 5.2. We represent each message as a vector in the same latent message space associated with the nodes in our network. Following the basic model, whether or not a node will re-broadcast a message it has received is determined by whether the message lies within the threshold distance of the node's own desired message vector.

6.2 Virality Rate

In running our simulations, we validate that our model generates information cascades that are broadly consistent with past models and observations of cascade events. We find that for certain parameterizations, namely, $alpha \approx 2.3$, $t \approx 0.45$, $s \approx 0.45$, and $8 \leq k \leq 10$, our cascades fit previously empirically observed data. For example, we find that, consistent with past empirical observations, most cascades remain very small, and the rate of large-scale diffusion events (i.e., events reaching greater than 100 nodes) is roughly 1 in 1000 [15, 16].

Exploring the relationship between the four key parameters of our model, we confirm that as the dimensionality d factor for the latent space increases, the virality rate goes down (Figure 9(a)). The intuitive explanation is that as d increases, the latent space becomes more sparse, and the likelihood of a message satisfying the preference vector and threshold of a sufficient number of connected vertices decreases. Furthermore, as

the acceptance threshold t increases, the virality rate increases as well (Figure 9(b)). Increasing either the smoothing factor s or neighbor candidate pool size n, representing peer influence and homophily, we find that the virality rate increases (Figure 9(c) and 9(d)).

It is interesting to observe that in our simulations such relatively low values of d produce realistic cascades. An implication for the semantic interpretation of our model is that the variety of properties of messages that people care about and examine when deciding whether to post or re-post a message may be correspondingly small. While this may be intuitive, it is also possible that suitable parameterizations of our model (higher values of t for example, or more structured networks) might act as a counter-weight to enable higher values of d to produce realistic cascades as well. The manifold of semantically meaningful combinations of message properties within this space, and the resultant sparsity may also play a role. We leave further experimentation in this area for future work.

6.3 Structural virality

Structural virality is a measure proposed by Goel et al. [15] to measure and distinguish between information disseminations that occur primarily through broadcast mechanisms (one sender broadcasting a message to a large number of people, with relatively few or no independent decisions to rebroadcast or spread the message); and propagations where no one sender is responsible for most of the dissemination (i.e., multigenerational "viral" propagation). Specifically, structural virality $v(T)$ over a diffusion tree T is defined to be the average distance between all pairs of nodes:

$$v(T) = 1/n(n-1) \sum_{i=1}^{n} \sum_{j=1}^{n} d_{ij} \qquad (3)$$

where distance d_ij is the shortest path length through the tree T between nodes i and j.

One of the elements of empirically observed information cascades that previous models have not captured is the variance in structural virality observed in information cascades. In their large-scale study of information diffusion events in Twitter, Goel et al. find that structural virality is weakly correlated (0.36) with the size of a cascade— smaller cascades are more likely to have a broadcast-like spread, and larger cascades are more likely to have a viral-like spread. However, as the correlation is weak, there are many large cascades that have a broadcast-like spread, and many small cascades that have a viral-like spread.
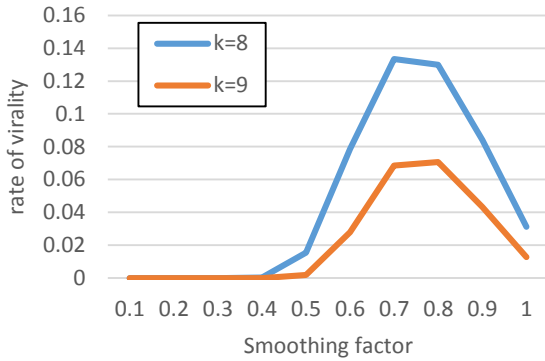
Figure 10(a) and 10(b) plots on a log-scale the structural virality of information cascades, binned by the size of the cascade under two information diffusion modes. Figure 10(a) shows structural virality under our model
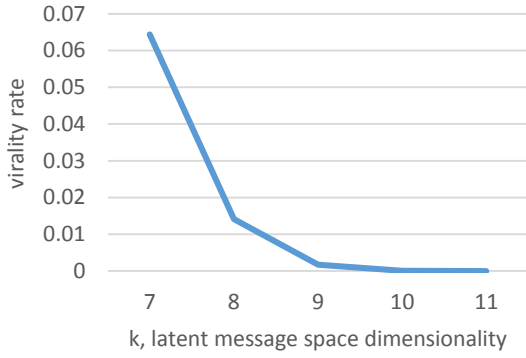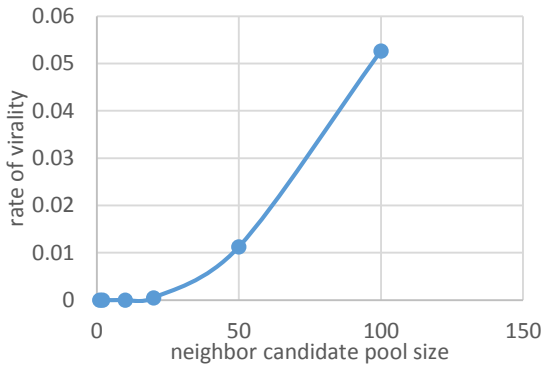
(a) Relationship between dimensionality d and virality rate. $n = 0.5, t = 0.5$
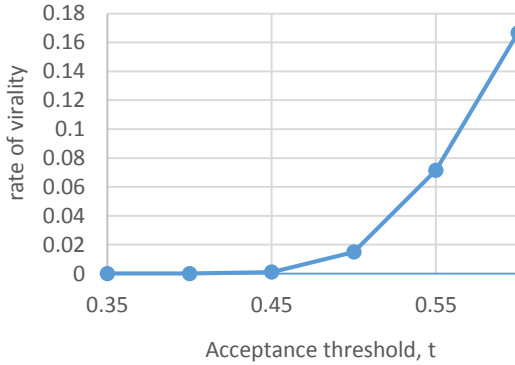
(b) Relationship between acceptance threshold and virality rate. $d = 8, t = 0.5$
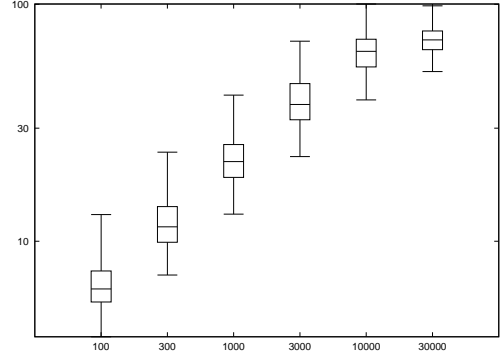

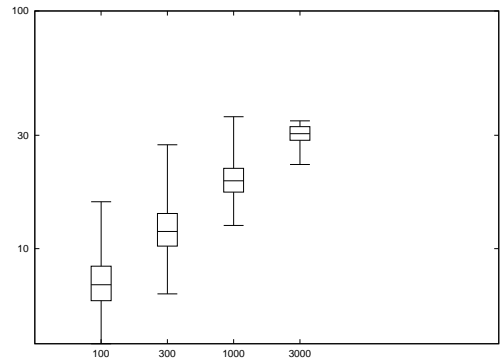(c) Relationship between smoothing factor and virality rate. $d = 8, t = 0.5$

(d) Relationship between candidate neighbor pool size virality rate. $d = 7, t = 0.5$

Figure 9: Relationship between virality rate and model parameters

(a) Structural virality under a persuasion model. $d = 7, s = 0.45, t = 0.45$

(b) Structural virality under a simple contagion model with uniform infection probability. $r = 0.7$

Figure 10: A boxplot of the structural virality of cascade size in simulations. The x-axis plots cascade size (number of message recipients) and the y-axis shows the structural virality. The box plot demarcates the quartiles of the structural virality.

of persuasion effects. Overall, we see that structural virality ranges from broadcast-like (low values) to very viral (high values). Within each bin of cascades, we also see a high variance in structural virality. Figure 10(b) shows structural virality under a simple contagion model, where the likelihood that any given node will propagate a message is β. We select β such that $r = \beta \bar{d}$; where $r = 0.7$, and $\bar{d}$ is the average degree of the network. Under this contagion model, we find that cascades are less likely to have a highly viral structure and, as cascade size increases, the variance of structural virality decreases significantly.

7. CONCLUSIONS AND FUTURE WORK

11

We present a first study into analyzing the effects of persuasion on well-studied social network phenomena such as information diffusion and propagation. We performed large-scale online user studies to motivate the need for a model of persuasion that we then use to explain structural properties of information diffusion. We leave the analysis of the model both in terms of the characterization of the resulting dynamics as well as the performance of algorithms for message selection (e.g., see Section 5.3) to future work.

In addition to the analysis of the current model, there are rich set of open questions we leave to future work. One deals with making the model richer to distinguish among individual nodes or messages based on their vector representation in the latent space allows us to ask new kinds of questions. These need to be asked and validated against empirical observations we make in Section 4. One direction would be to better explore implications of homophily and peer influence; for example, model whether the same message in a different network location would likely generate the same cascade. (i.e., to what degree are cascades caused by the message, the user or the network structure? or all of the above! Another example would be to broadly validate the model through more and deeper user studies (e.g., user studies across more domains); explore model susceptibility to specific persuasion techniques such as the authority of the sender, message repetition etc.

Finally, understanding the synergies between the sender and the recipient in terms of message content has wider implications for for social systems engineering. For example, how do individuals learn what messages (i.e., persuasion strategies) are effective? how does the feedback they receive through social systems features (newsfeeds, retweets, likes, comments, etc.) influence their learning?

8. REFERENCES

[1] E. Adar and L. A. Adamic. Tracking information epidemics in blogspace. In *Web intelligence, 2005. Proceedings. The 2005 IEEE/WIC/ACM international conference on*, pages 207–214. IEEE, 2005.

[2] B. Aronov and S. Har-Peled. On approximating the depth and related problems. *SIAM Journal on Computing*, 38(3), 2008.

[3] N. T. Bailey et al. *The mathematical theory of infectious diseases and its applications*. Charles Griffin & Company Ltd, 5a Crendon Street, High Wycombe, Bucks HP13 6LE., 1975.

[4] F. M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5), 1969.

[5] A. Bohra, L. S. Chandran, and J. K. Raju. Boxicity of series-parallel graphs. *Discrete mathematics*, 306(18):2219–2221, 2006.

[6] A. C. Braet. Ethos, pathos and logos in aristotle's rhetoric: A re-examination. *Argumentation*, 6(3):307–320, 1992.

[7] S. Chaiken, W. Wood, and A. H. Eagly. Principles of persuasion. 1996.

[8] L. S. Chandran, M. C. Francis, and N. Sivadasan. Geometric representation of graphs in low dimension using axis parallel boxes. *Algorithmica*, 56(2):129–140, 2010.

[9] M. Chlebík and J. Chlebíková. Approximation hardness of optimization problems in intersection graphs of d-dimensional boxes. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 267–276. Society for Industrial and Applied Mathematics, 2005.

[10] R. B. Cialdini. Influence: The psychology of persuasion. 1993.

[11] Ü. D. Demırdöğen. The roots of research in (political) persuasion: Ethos, pathos, logos and the yale studies of persuasive communications. *International Journal of Social Inquiry*, 3(1), 2010.

[12] E. B. Falk, L. Rameson, E. T. Berkman, B. Liao, Y. Kang, T. K. Inagaki, and M. D. Lieberman. The neural correlates of persuasion: A common network across cultures and media. *Journal of cognitive neuroscience*, 22(11):2447–2459, 2010.

[13] M. Friestad and P. Wright. Persuasion knowledge: Lay people's and researchers' beliefs about the psychology of advertising. *Journal of consumer research*, pages 62–74, 1995.

[14] R. H. Gass and J. S. Seiter. Persuasion, social influence, and compliance gaining. *Utah State University Faculty Monographs*, 2011.

[15] S. Goel, A. Anderson, J. Hofman, and D. J. Watts. The structural virality of online diffusion. In *Preprint*, 2013.

[16] S. Goel, D. J. Watts, and D. G. Goldstein. The structure of online diffusion networks. In *Proc. of the 13th ACM EC*, pages 623–638. ACM, 2012.

[17] J. Goldenberg, B. Libai, and E. Muller. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters*, 12(3):211–223, 2001.

[18] M. Gomez Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *Proc. of the 16th ACM SIGKDD*, pages 1019–1028. ACM, 2010.

[19] M. Granovetter. Threshold models of collective behavior. *American journal of sociology*, pages 1420–1443, 1978.

[20] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins. Information diffusion through blogspace. In *Proceedings of the 13th international*

conference on World Wide Web, pages 491–501. ACM, 2004.

[21] R. Guadagno and R. Cialdini. Online persuasion and compliance: Social influence on the internet and beyond. *The social net: Human behavior in cyberspace*, pages 91–113, 2005.

[22] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proc. of the ninth ACM SIGKDD*, pages 137–146. ACM, 2003.

[23] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Trans. Web*, 1(1), May 2007.

[24] J. Leskovec, M. McGlohon, C. Faloutsos, N. S. Glance, and M. Hurst. Patterns of cascading behavior in large blog graphs. In *SDM*, volume 7, pages 551–556. SIAM, 2007.

[25] M. R. Morris, S. Counts, A. Roseway, A. Hoff, and J. Schwarz. Tweeting is believing?: understanding microblog credibility perceptions. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, pages 441–450. ACM, 2012.

[26] S. A. Myers, C. Zhu, and J. Leskovec. Information diffusion and external influence in networks. In *Proc. of the 18th ACM SIGKDD*, pages 33–41. ACM, 2012.

[27] D. J. O'keefe. *Persuasion: Theory and research*, volume 2. Sage, 2002.

[28] Y. Shao. Ethos, logos, pathos: Strategies of persuasion in social/environmental reports. *Social and Environmental Accountability Journal*, 33(3):179–180, 2013.

[29] C. Tan, L. Lee, and B. Pang. The effect of wording on message propagation: Topic- and author-controlled natural experiments on twitter. In *Proc. of the 52nd Annual Meeting of the Association for Computational Linguistics*. ACL, 2014.

[30] W. Wood. Attitude change: Persuasion and social influence. *Annual review of psychology*, 51(1):539–570, 2000.