Proceedings of the Eleventh International AAAI Conference on Web and Social Media (ICWSM 2017)

Predicting Popular and Viral Image Cascades in Pinterest

Jinyoung Han,[1] Daejin Choi,[2] Jungseock Joo,[3] Chen-Nee Chuah[4]
[1]Hanyang University, [2]Seoul National Universtiy
[3]University of California, Los Angeles, [4]University of California, Davis
jinyounghan@hanyang.ac.kr, djchoi@mmlab.snu.ac.kr, jjoo@commstds.ucla.edu, chuah@ucdavis.edu

Abstract

The word-of-mouth diffusion has been regarded as an important mechanism to advertise a new idea, image, technology, or product in online social networks (OSNs). This paper studies the prediction of *popular* and *viral* image diffusion in Pinterest. We first characterize an image cascade from two perspectives: (i) volume – how large the cascade is, i.e., total number of users reached, and (ii) structural virality – how many users in the cascade are responsible for attracting other users. Our model predicts whether an image will be (a) popular in terms of the volume of its cascade, or (b) viral in terms of the structural virality. Our analysis reveals that a popular image is not necessarily viral, and vice versa. This motivates us to investigate whether there are distinctive features for accurately predicting popular or viral image cascades. To predict the popular or viral image cascades, we consider the following feature sets: (i) deep image features, (ii) image meta and poster's information, and (iii) initial propagation pattern. We find that using deep image features alone is not as effective in predicting popular or viral image cascades. We show that image meta and poster's information are strong predictors for predicting popular image cascades while image meta and initial propagation patterns are useful to predict viral image cascades. We believe our exploration can give an important insight for content providers, OSN operators, and marketers in predicting popular or viral image diffusion.

Introduction

Online social networks (OSNs) have seen phenomenal growth: over a billion users on Facebook, and hundreds of millions on Twitter, Instagram, and Pinterest. This new social ecosystem generates rich social data, which enables data-driven studies on human behavior as well as information diffusion patterns (e.g., political propaganda, product advertisement, or content distribution), which are often referred to as 'Computational Social Science' (Lazer et al. 2009).

The word-of-mouth information diffusion (Rodrigues et al. 2011; Rahman, Han, and Chuah 2015; Han et al. 2014) has been regarded as an important mechanism to advertise a new idea, technology, content, or product in OSNs. This in turn has spurred many data-driven computational studies on how information (e.g., a photo, news, URL, or product) spread in various OSNs such as Facebook, Twitter, or Pinterest (Bakshy et al. 2012; Kupavskii et al. 2012; Lerman and Ghosh 2010; Han et al. 2014). A news, idea, photo, URL, or product can be reshared multiple times (e.g., like retweet in Twitter or repin in Pinterest) in OSNs, hence generating a *cascade* that potentially reaches a large number of audiences.

These studies have revealed valuable insights into the macro-level information propagation patterns (e.g., cascade size). Despite the large theoretical and empirical literature on the information diffusion in OSNs, however, relatively little has been known about the micro-level dynamics of viral cascades, in part because requisite data for building a complete information cascade has not been available until very recently (Goel et al. 2015). Often time, 'virality' is used to refer to 'popularity' of diffusion (Khosla, Sarma, and Hamid 2014; Jenders, Kasneci, and Naumann 2013; Guerini, Staiano, and Albanese 2013; Totti et al. 2014; Khosla, Sarma, and Hamid 2014). However, viral diffusion is not only about popularity (macro-level view, i.e., cascade size) but also about how information propagates via person-to-person viral contagion (i.e., micro-level view). Such a person-to-person diffusion process is analogous to the spread of biological viruses. That is, unlike the 'broadcast' where a single source directly spreads information to the most of recipients like mass media, many individuals participate in spreading information in viral diffusion. Figure 1 shows an illustrative example of different 11-nodes diffusions where an information reaches 10 audiences through different propagation scenarios. The leftmost figure shows a typical broadcast scenario, while the rightmost showcases a viral diffusion through a chain.

Recently, with the availability of large-scale user interaction data over OSNs, there have been effort in investigating dynamics of viral cascades (Cheng et al. 2014; Goel et al. 2015; Deza and Parikh 2015). For example, the structural virality of retweet cascades was measured in Twitter (Goel et al. 2015), which showed that the correlation between the size of a cascade and its virality is low. Also, some work tried to use image features for predicting image virality (Deza and Parikh 2015; Guerini, Staiano, and Albanese 2013), revealing a possibility to use content information in predicting viral diffusion. Cheng *et al.* studied a cascade growth prediction problem, and tried to predict whether a given cascade (with size k) *grows* beyond the median size of all the cascades with at least k reshares (Cheng et al. 2014) in Facebook. These work have revealed valuable theoretical insight into viral diffusion in social media.

This paper demystifies the prediction of *popular* or *viral* image diffusion in Pinterest, an *interest-driven* OSN, from a practical and engineering standpoint. To this end, we first characterize an image cascade from two perspectives: (i) volume – how large the cascade is, i.e., total number of users reached, and (ii) structural virality – how many users in the cascade are responsible for attracting other users, which is measured by the 'Wiener Index (WI)' (Wiener 1947; Goel et al. 2015; Cheng et al. 2014; Choi et al. 2015). Here, the popular cascades are the ones that have high volumes while the viral cascades are defined as the ones that have high WIs. Note that 'popularity' does not imply 'virality', and vice versa. An image can be popular but not viral. Similarly, a viral diffusion may not end up in a large cascade.

We attempt to predict whether an image will be (a) popular in terms of cascade size/volume, and/or (b) viral in terms of WI. We investigate what factors have the strong predictive power of popular or viral image cascades. In particular, we explore two factors that contribute to popular or viral diffusion – *human social context* and *properties of content*. The ability to predict whether the diffusion of a content would go viral and reach a large user population is of interest to researchers, OSN operators, marketers, and content providers.

We shed light on these issues by performing a data-driven analysis on a large-scale dataset that contains 337 K images shared by 1 M users in Pinterest. We had kept track of a propagation path of each image (i.e., pin in Pinterest) in each category (e.g., animal, kids, travel) in Pinterest for 44 days from June 5 to July 18, 2013. Our analysis reveals that there is overall a positive correlation between the volume and structural virality of an image cascade. However, we find that popular images are not always necessarily viral; we show that two popular images with the same popularity propagate through very different scenario, i.e., one by broadcast and the other by a person-to-person contagion process. This motivates us to study what different factors are useful for predicting popular or viral image cascades. In addition, we consider different scenarios in predicting the popular or viral image diffusion, e.g., only image information is available, user information is hidden due to privacy issues, etc. In particular, we seek answers to the following questions:

- **Q1 – If only an image is given (and available), can we predict whether it becomes popular or goes viral?** We apply deep learning techniques (Deng et al. 2009; Simonyan and Zisserman 2014) for extracting image features, which are used for predicting popular or viral image cascades. We find that using deep image features alone is not as effective in predicting popular or viral cascades, which is in line with previous work in other OSNs (Cheng et al. 2014; Totti et al. 2014).
- **Q2 – At the moment when an image is posted and its meta and/or poster (or pinner) information is available, can we predict whether the image becomes popular or goes viral?** We find that meta information of an image such as its category, source, or title, as well as its poster's information are useful in predicting popular or viral image cascades. We also find that combining image meta and pinner features improves the prediction performance, which signifies that image meta and pinner features are complementary to each other. Note that image meta and pinner features are the strongest predictor in predicting popular image cascades, implying that we can accurately forecast the image popularity using the image meta information (e.g., category, source, or title) and poster's information (e.g., his/her connectivity or activity), at the moment when the image is posted.
- **Q3 – If an initial image propagation pattern is observable, does it help to predict popular and/or viral image cascades?** We find that the initial propagation pattern of an image cascade is useful whether it will become popular or go viral. However, the information of users who initially participate in the cascade do not contribute much in predicting popular and viral image cascades. Note that the initial propagation pattern of an image and its meta info are the best predictors in predicting viral image cascades, implying that if we observe the *initial propagation pattern* of an image, we can accurately predict whether the image will go viral in the future. It is worth to note that image meta information is commonly useful for predicting future popular and viral image diffusion.

The rest of this paper is organized as follows. After reviewing the related work, we describe our measurement methodology. We then investigate the characteristics of image cascades in Pinterest. We finally propose models to predict popular or viral image cascades in Pinterest.

Related Work

Information cascades in OSNs: As OSNs have become one of the popular platform to spread information such as news, photo, URL, or product, there have been huge effort in studying information adoption and propagation in various OSNs (Aral and Walker 2012; Bakshy et al. 2012; Wang et al. 2011; Rahman, Han, and Chuah 2015; Choi et al. 2015). Bakshy *et al.* studied the role of social networks in information diffusion in Facebook, and showed that exposed users in the network are more likely to spread information (Bakshy et al. 2012). Aral *et al.* identified influential individual and susceptible users in adopting the product in Facebook (Aral and Walker 2012). Rahman *et al.* (Rahman, Han, and Chuah 2015) analyzed the adoptions and propagations of Facebook gifting applications, and showed that the evolutionary perspectives of cascades such as their initial growth rates are important factors for predicting the final population size of the application cascades. Choi *et al.* characterized online conversations in Reddit, and revealed how content properties and user participation behaviors are associated with successful conversation (Choi et al. 2015).

A few recent studies have shifted focus to micro-level dynamics of viral cascades (Cheng et al. 2014; Goel et al. 2015; Guerini, Staiano, and Albanese 2013; Deza and Parikh 2015; Khosla, Sarma, and Hamid 2014). The structural virality of cascades was measured based on the user dynamics information in Twitter (Goel et al. 2015). For predicting image virality, some work used image features (Deza and Parikh 2015; Guerini, Staiano, and Albanese 2013; Khosla, Sarma, and Hamid 2014; Cheng et al. 2014), revealing a possibility to use content information in predicting viral diffusion. Deza and Parikh studied the viral image prediction from a computer vision perspective (Deza and Parikh 2015). They evaluated several image features for predicting image virality. The most relevant work to this paper is that by Cheng *et al.* (Cheng et al. 2014), which studied models for predicting whether a given cascade (with size k) *grows* beyond the median size of all the cascades with at least k reshares, which is a *growth prediction problem*. They showed that temporal and structural features are key predictors of the photo reshare cascade growth in Facebook (Cheng et al. 2014). While the work by Cheng *et al.* (Cheng et al. 2014) provided an important theoretical insight into cascade prediction, this paper goes one step further from a practical and engineering standpoint; we focus on a *popularity or virality prediction problem* in Pinterest, based on the following feature sets which can be observed in different scenarios: (i) image features that can be obtained before posting, (ii) image meta and poster's information that can be obtained at the moment of posting, and (iii) initial propagation pattern. We explore which factors are strong predictors in predicting popular or viral image cascades, respectively.

Pinterest – an interest-driven OSN: Unlike other popular friendship-based OSNs such as Facebook, interests drive user activities or connectivities in Pinterest (Han et al. 2014; Gelley and John 2015). Han *et al.* revealed that pin propagation in Pinterest is mostly driven by content properties like its topic, not by users' characteristics (Han et al. 2014). Gelley and John also showed that 'following' is not significantly associated with content sharing in Pinterest (Gelley and John 2015). Zhong *et al.* proposed models to predict whether a user will be interested in repinning the given pin (Zhong, Karamshuk, and Sastry 2015). Han *et al.* (Han et al. 2015) proposed a method to predict which topics an individual user will be interested in. Totti *et al.* evaluated the predictive power of different features on image popularity (Totti et al. 2014), and showed that visual properties have a lower predictive power than social cues. This paper proposes models for predicting popular or viral images based on two factors – *human social context* and *properties of content*, which can give an important insight into resource allocation for content providers and marketers in Pinterest-like OSNs.

Methodology

In this section, we first characterize our image cascade model in terms of volume and structural virality. We then describe our dataset used in this paper.

Image Cascade Model

We first model an *image cascade* as an undirected tree, $T = (U, R)$, where U is the set of users including a pinner and follow-up repinners for *a given image* (or a pin) posted by the pinner, and R is the set of repinning activities, i.e., pin propagation. We characterize the image cascade T based on the following two metrics:

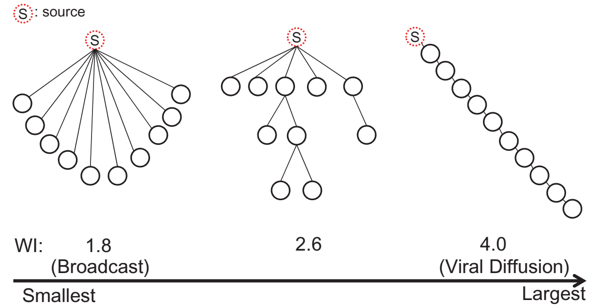

Figure 1: Illustrative examples of different 11-node diffusions: from a simple broadcast (on the far left) to a viral diffusion through a chain (on the far right). The bottom axis shows the Wiener Index (WI) values calculated for the 11-node cascades, $V_T = 11$.

- **Volume ($|U|$)** of cascade T is the number of nodes in the tree. For example, $|U|$ of the cascade in Figure 1 is 11.
- **Structural virality (or WI)** of cascade T represents the average range of a node's effect in an image cascade. To quantify the structural virality, we adopt a well-known metric 'Wiener Index (WI)' (Wiener 1947; Goel et al. 2015; Cheng et al. 2014; Choi et al. 2015), which is defined as the average hop count between all pairs of nodes in a tree T. More specifically, WI of cascade T can be calculated as follows:

$$WI = \frac{1}{|U|(|U|-1)} \sum_{i,j \in U, i \neq j} d_{ij} \quad (1)$$

  where U is the set of users in T, and d_{ij} is the distance (or hop count) of the shortest path between users i and j in T. Figure 1 shows the three 11-node trees with their WIs. As shown in Figure 1, given the same number of nodes, i.e., an image reaches to 10 audiences through different propagation scenarios, the WI becomes the minimum if all the repinners directly get the image from the pinner (i.e., the leftmost scenario in Figure 1), and the maximum if T becomes a chain (i.e., the rightmost scenario in Figure 1).

Dataset

To explore image cascades in Pinterest, we analyze a dataset collected from Pinterest for 44 days from June 5 to July 18, 2013, which contains 337,345 (original) images (i.e., pins), their associated 1,190,220 repins, and 1,047,545 users. Among the 337,345 images, 144,080 images are shared by at least one repinners, hence we could obtain 144,080 image cascades whose volumes are higher than 1. During the measurement period, we collected all the newly posted pins (and their pinners) from the menu of each category (e.g., kids, education, history) in Pinterest. Note that there were 33 categories at the moment of data collection such as 'travel', 'kids', or 'women's fashion' in Pinterest. Table 1 summarizes the 33 Pinterest categories with their associated numbers of corresponding pins/repins. A Pinterest user can share a posted pin via repinning, and the shared pin also can be shared subsequently by other users. We kept track of this

1	diy & crafts(189176)	2	education(80910)	3	animals(75299)	4	food & drink(67730)	5	quotes(45849)	6	design(45758)	7	health & fitness(45225)
8	humor(42656)	9	art(41939)	10	womens' fashion(38698)	11	architecture(37557)	12	film, music & books(37403)	13	home & decor(34655)	14	products(34333)
15	men's fashion(28951)	16	science & nature(24324)	17	geek(20633)	18	technology(20306)	19	travel(19060)	20	outdoors(18149)	21	weddings(17927)
22	cars & motorcycles(17447)	23	hair & beauty(17074)	24	celebrities(15543)	25	gardening(14736)	26	tattoos(11644)	27	photography(10390)	28	history(7159)
29	kids(6508)	30	sports(6343)	31	holidays & events(5159)	32	shop(4075)	33	illustrations & posters(3916)				

Table 1: 33 Pinterest categories with their numbers of associated pins/repins.

1	fitsugar.com(8760)	2	designspiration.net(4503)	3	womenshealthmag.com(3410)	4	imdb.com(3131)	5	greatist.com(3008)
6	buzzfeed.com(2930)	7	saatchionline.com(2458)	8	fitnessmagazine.com(2445)	9	teacherspayteachers.com(2416)	10	archdaily.com(1934)
11	bhg.com(1717)	12	houzz.com(1595)	13	prevention.com(1466)	14	food.com(1442)	15	allrecipes.com(1440)
16	marthastewart.com(1439)	17	behance.net(1417)	18	ebay.com(1404)	19	wikipaintings.org(1297)	20	themetapicture.com(1263)

Table 2: A summary of top 20 sources and their numbers of original pins.

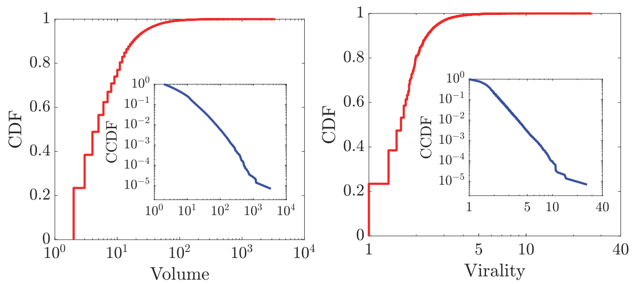

(a) Volume (b) Structural virality

Figure 2: Distributions of volume and structural virality of image cascades.

repinning process and collected all the corresponding repin and repinner information. Eventually, our dataset includes (i) pin and repin information, e.g., category, source[1], description, etc., and (ii) pinner and repinner information, e.g., number of followers/followees, number of pins, etc.

Image Cascade Analysis

In this section, we first investigate the characteristics of image cascades in Pinterest. We then explore whether popular images (i.e., cascades having high volume) are also viral.

Characteristics of image cascades

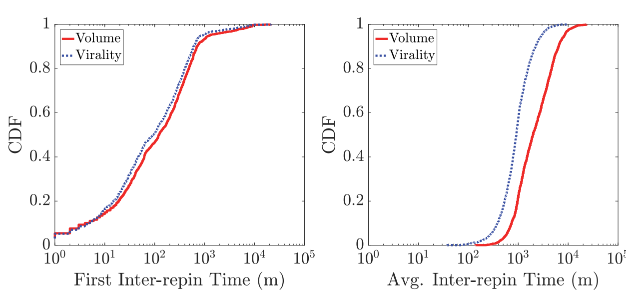

(a) First repin time (b) Average inter-repin time

Figure 3: Distributions of repin times of image cascades.

[1]A source indicates the URL the pin. Table 2 describes the top 20 sources and their numbers of original pins in Pinterest in our dataset.

We first investigate the distributions of volume and structural virality of image cascades in Figure 2. As shown in Figure 2(a), the volume of image cascades shows a heavy tail distribution; the range of volume spans four orders of magnitudes. While 77% of cascades have less than 10 users (including a pinner and 9 repinners), top 1% and 0.1% of cascades consist of more than 74 and 230 users, respectively, signifying a large deviation among cascades. Note that the average, median, and maximum volume is 9.26, 5, and 3,401, respectively. When we look at Figure 2(b), the structural virality also shows a heavy tail distribution, which only spans two orders of magnitudes. While around 81% of WI values are smaller than 2 and 99.7% of them are smaller than 5, top 0.1% of WIs are greater than 6.32. This implies that most of image cascades in Pinterest are not likely to span deep. Note that the average, median, and maximum structural virality is 1.66, 1.6, and 26.12, respectively.

To capture how quickly users propagate images in the top 1% image cascades in terms of volume and structural virality, we calculate the first inter-repin times (i.e., time difference between the original pinning and the first repinning) and the average inter-repin times of the top 1% cascades in Figure 3. In calculating the average inter-repin time, we only consider the time differences (in a cascade) within the range of $[\mu - 2\sigma,\ \mu + 2\sigma]$ for excluding outliers. As shown in Figure 3, the inter-repin times of the top popular cascades (with high volumes) are higher than those of the top viral cascades (with higher WIs), meaning that users in viral cascades tend to propagate images more quickly. This implies that the propagation speed can be used to predict popular or viral image cascades. For example, if we observe the initial propagation speed of a cascade, we may forecast whether the cascade goes viral in the future.

We next investigate the volume and structural virality of image cascades across 33 categories (Table 1) and the top 20 sources (Table 2) in Figures 4 and 5, respectively. As shown in Figure 4, the volume and structural virality are different across categories. This implies that category and source information of an image cascade can be one of the important factors in predicting whether the cascades grows much or goes viral. Interestingly, 'humor' (category index (CI) 8), 'quotes' (CI 5), and 'tatoos' (CI 26) show higher volume and structural virality than others while there are relatively less number of pins/repins in those categories. Also, the volume and structural virality are different across sources as shown in Figure 5. Interestingly, the images from 'themetapicture.com' (source index (SI) 20) are much more

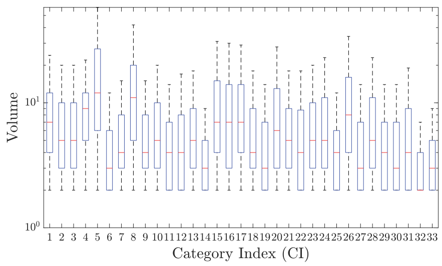

(a) Volume

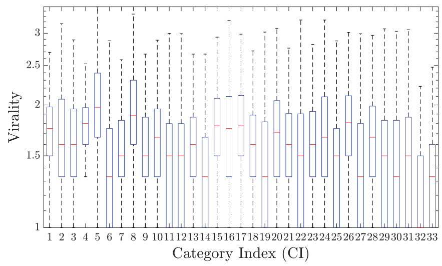

(b) Sturctural virality

Figure 4: Distributions of volume and structural virality of image cascades across 33 categories.

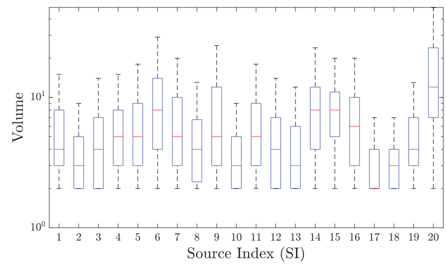

(a) Volume

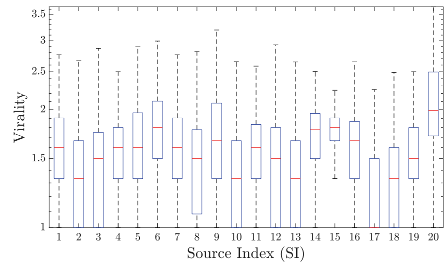

(b) Sturctural virality

Figure 5: Distributions of volume and structural virality of image cascades across the top 20 sources.

popular and viral than others; the website is not so popular in general (the Alexa rank is 20,328 as of October, 2016) and provides funny pictures. Also, food related sources such as ‘food.com’ (SI 14) and ‘allrecipes.com’ (SI 15) are likely to provide popular and viral images to Pinterest.

Are popular images also viral?

We now investigate whether the popular images (i.e., cascades having high volume) are also viral. To this end, we first plot the volume/virality of each image cascade in Figure 6(a). As shown in Figure 6(a), there is an overall positive correlation between the volume and virality, which means cascades with higher volumes tend to have higher structural viralities. However, the viralities of cascades with high volumes (e.g., over 100) tend to radiate. The Pearson correlation between the volume and virality of the top 1% cascades (whose volumes are higher than 74) is 0.42 ($p < 0.001$). This implies that popular images are not necessarily viral. Note that the Pearson correlation between the volume and virality of the bottom 99% cascades (whose volumes are 74 or smaller) is 0.8 ($p < 0.001$). When we look at the volume-based top 1% cascades and the virality-based top 1% cascades, only 17.4% of the cascades are overlapped, which signifies that top popular and viral cascades are disparate. This implies that different factors may be useful for predicting popular or viral image cascades, which will be discussed in the next section.

As an example, we illustrate two cascades with same volume ($N = 101$) but different viralities, $WI = 2.096$ for the red circle and $WI = 7.128$ for the blue diamond in Figures 6(b) and 6(c), respectively. As shown in Figures 6(b) and 6(c), two similarly popular images can be propagated through different scenarios: (i) broadcast where a pinner mostly spreads an image to the most of recipients and (ii) viral diffusion where an image propagates via the person-to-person contagion process. This confirms that an image can be popular but not viral.

To investigate whether popular images are viral in different categories and sources, we calculate the Pearson correlation coefficients between the volume and structural virality in each category and source. Figure 7 shows the two coefficient values for the top 1% (with high volumes, represented as redo dots) and bottom 99% (represented as bar) cascades, respectively, in each category and source. Overall, the Pearson correlation coefficients of the bottom 99% cascades are very high, i.e., mostly over 0.7. The coefficients of the bottom 99% cascades in ‘food & drink’ (CI 4) and ‘shop’ (CI 32) are even higher than 0.9 as shown in Figure 7(a). However, the coefficients of the top 1% cascades are substantially lower than those of the bottom 99%. Especially, the coefficients of the top 1% cascades in ‘men’s fashion’ (CI 15), ‘science & nature’ (CI 16), ‘sports’ (CI 30), and ‘diy & crafts’ (CI 1) are lower than 0.1, meaning that popular images in those categories are not necessarily viral. We observe a similar pattern in Figure 7(b) that shows popular images from particular sources (e.g., ‘imdb.com’ (SI 4), ‘greatist.com’ (SI 5), ‘houzz.com’ (SI 12), ‘allrecipes.com’ (SI 15), ‘wikipaintings.org’ (SI 19)) are not necessarily viral. Interestingly, top 1% popular images in ‘greatist.com’ (SI 5) and ‘houzz.com’ (SI 12) show even weak negative correlations, which is a significantly disparate pattern with the bottom 99% (unpopular) images from the sources.

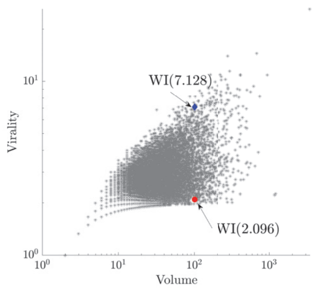

(a) Volume vs. Virality

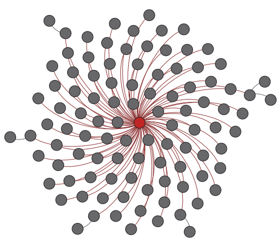

(b) Broadcast

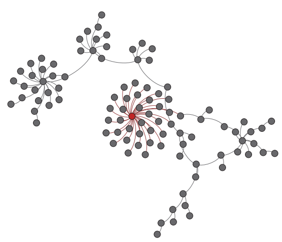

(c) Viral diffusion

Figure 6: Volume and virality of each image. Overall, there is a positive correlation between the volume and virality. However, popular images are not necessarily viral, e.g., two cascades with same volume ($N = 101$) have different viralities: $WI = 2.096$ for the red circle (b) and $WI = 7.128$ for the blue diamond (c).

Popular and Viral Image Prediction

We have revealed that popular images are not necessarily viral in Pinterest, which motivates us to study whether there are distinctive features to accurately predict popular and viral images, respectively. In this section, we aim to predict popular or viral image cascades. In particular, we identify popular cascades whose volumes are higher than 230 and 74, which account for the top 0.1% and 1% of all the image cascades (in terms of volume or cascade size), respectively. We also identify 'viral' cascades whose structural virality (or WIs) are higher than 6.32 and 3.85, which account for the top 0.1% and 1% of all the image cascades (in terms of WIs), respectively. Note that only a small portion of the top popular and viral cascades are overlapped as shown in the previous section; only 17.4% of the cascades are overlapped between the volume-based top 1% cascades and the virality-based top 1% cascades.

We cast this problem as a supervised learning problem, where we observe a set of features of a cascade and predict whether the given cascade belongs to the top popular or viral cascades. We build a learning model based on the Random Forest ensemble algorithm (Breiman 2001). We used

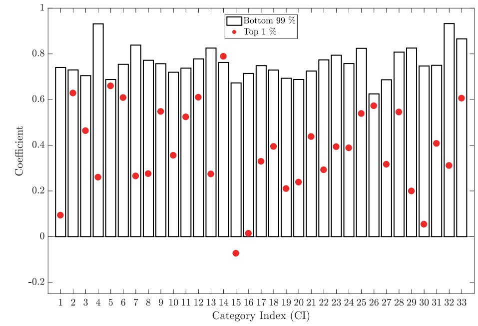

(a) Category

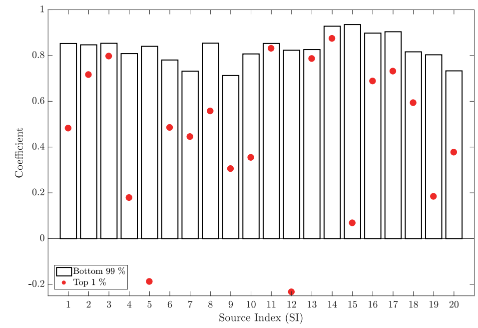

(b) Source

Figure 7: Pearson correlation coefficients between volume and virality of the top 1% (with high volumes, represented as red dots) and bottom 99% (represented as bars) cascades, respectively, in each category and source.

other classifiers including support vector machine or logistic regression, but we only report the results of the Random Forest ensemble classifier as it performs better than others. We report the following performance metrics: (i) accuracy ($ACC = \frac{TP+TN}{TP+FP+FN+TN}$), (ii) true positive rate or recall ($TPR = \frac{TP}{TP+FN}$), (iii) false positive rate ($FPR = \frac{FP}{FP+TN}$), and (iv) area under the receiver operating characteristics (ROC) curve (AUC)[2] (Fawcett 2006), where TP, FP, FN, and TN represents the true positive, false positive, false negative, and true negative, respectively. We perform a 10-fold cross validation.

Predicting the top popular or viral cascades can be suffered from the class imbalance problem, e.g., the ratio between the minority and majority classes for identifying the top 0.1% cascades is 1:999. To remedy this issue, we apply the Synthetic Minority Over-sampling TEchnique (SMOTE) (Chawla et al. 2002), which allows us to learn with over-sampled instances from the minority class (i.e., top cascades). We learn randomly under-sampled instances from the majority class (i.e., non-top cascades). We varied the sampling ratios of minority and majority classes, from 1:1 to 1:2 to 1:4 to 1:8, but we only report the results of 1:1 ratio as it shows a similar performance with others. Note that we apply this technique only in the learning set.

[2] AUC indicates the effectiveness of a model. A perfect model has an AUC of 1 while a random model generates an AUC of 0.5.

We consider different scenarios in predicting the popular or viral image cascades. In particular, we answer the following questions: (1) If only an image is given (and available), can we predict whether it becomes popular or goes viral?, (2) At the moment when an image is posted and its meta and/or poster (or pinner) information is available, can we predict whether the image becomes popular or goes viral?, and (3) If an initial image propagation pattern is observable, does it help to predict popular and/or viral image cascades? Answering these questions can give an important insight into predicting popular or viral image diffusion for content providers, OSN operators, and marketers.

Predictive power of image itself

We first study the role of image content in its popularity and virality prediction without using other features such as poster or posting information. This assumes the situation where (i) pinner or posting information is not available (e.g., due to privacy issues) or (ii) the image is not yet posted by anyone. To this end, we extract features from an image using the deep learning technique. We adopt a well-known visual categorization developed for the task of image classification and feature learning, ImageNet (Deng et al. 2009), which defines 1000 image classes (mostly object classes). We use a convolutional neural network (CNN) (Krizhevsky, Sutskever, and Hinton 2012), which is known as very effective for visual feature learning. Our model architecture is the VGG-16 (Simonyan and Zisserman 2014) and we use a publicly available pre-trained model on the Imagenet data. For each image, we extract the final image features at 1000 category level (referred to as 'IMAGE') as well as intermediate 4096 features at the last fully-connected layer ($FC7$) (referred to as 'IMAGE(FC7)'). In addition to high-level features ('IMAGE' and 'IMAGE(FC7)'), we further consider the following low-level features: (i) 512 'gist' image features (Oliva and Torralba 2001) which describe gradient-based (Gabor filters) scene features such as texture or edge (referred to as 'IMAGE(gist)') and (ii) the mean and standard deviation of image color in RGB (referred to as 'IMAGE(color)'). Based on the extracted image features, our classifier (i.e., the Random Forest ensemble) identifies whether the given image belongs to the top popular and/or viral cascades.

Figures 8 and 9 show the prediction results on popular and viral cascades, respectively, using image features. For a comparison purpose, we include the result of a null model, 'BASELINE'. Since the 'BASELINE' model predicts the popular or viral cascades according to their distributions, the $ACCs$ are 99.9% and 99% for predicting top 0.1% and 1%, respectively. Note that the 'BASELINE' has AUC of 0.5. As shown in Figures 8 and 9, the models using image features ('IMAGE', 'IMAGE(FC7)', 'IMAGE(gist)', and 'IMAGE(color)') perform slightly better than 'BASELINE', but their $AUCs$ are mostly lower than 0.55, implying that using image features alone is not as effective in predicting popular or viral image cascades. In other words, popular or viral image cascades are not predictable using only image features. Note that high-level features ('IMAGE' and 'IMAGE(FC7)') performs slightly better than low-level features

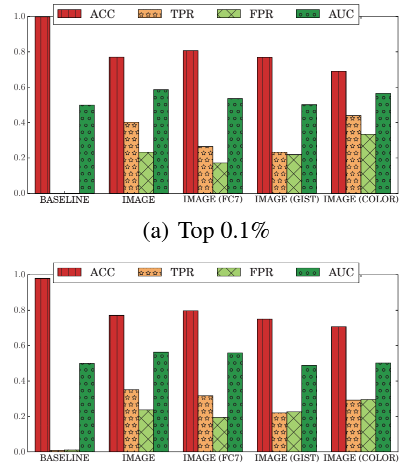

(a) Top 0.1%

(b) Top 1%

Figure 8: Prediction results on popular cascades using image features.

('IMAGE(gist)' and 'IMAGE(color)'). We find that 'IMAGE' mostly performs better than other image feature sets, hence we only consider 'IMAGE' features hereafter.

To identify the specific features that contribute most towards predicting top 1% popular and viral cascades based on 'IMAGE' features, respectively, we apply the *Chi-squared* (χ^2) statistic evaluation (Liu and Setiono 1995) to all of the 'IMAGE' features, which results in assigning a score to each feature. We rank the features according to the χ^2 values. The top 3 features for predicting top 1% popular cascades are 'menu', 'brassiere', and 'binoculars', while those for predicting top 1% viral cascades are 'menu', 'binoculars', and 'plate'.

Predictive power of image meta and pinner information

We next investigate whether image meta and/or poster (or pinner) information is useful in predicting popular or viral image cascades. For the meta information of a pin (referred to as 'META'), we consider the following features: (i) category popularity (i.e., number of pins) where the give pin belongs, (ii) source popularity (i.e., number of pins) where the given pin comes, (iii) maliciousness of the pin, and (iv) revealed sentiment from the pin's title and description. For detecting the maliciousness of a pin, we submit the source (i.e., URL where the pin comes) of the pin to a commercial URL scanner, VirusTotal (VirusTotal 2016), which scans a submitted URL over a corpus of over 60 website scanning engines. We identify each source as malicious if two or more security engines indicate it malicious. To calculate the revealed sentiment of a pin, we use LIWC (Linguistic Inquiry and Word Count), which counts words into psychologically meaningful categories (Pennebaker, Mehl, and Niederhoffer 2003). We calculate the positive, negative, cognitive, and social scores of each pin's title and description using LIWC.

We also consider the following characteristics of a pinner (referred to as 'PINNER'): (i) number of pins the pinner has,

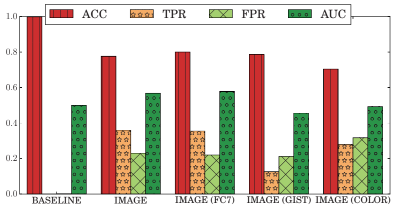

(a) Top 0.1%

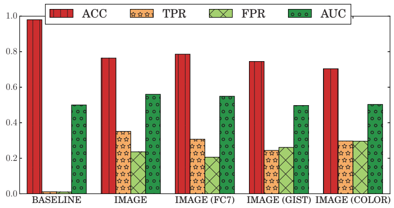

(b) Top 1%

Figure 9: Prediction results on viral cascades using image features.

(ii) number of followers who follow the pinner, (iii) number of followees the pinner follows, (iv) number of likes the pinner likes, (v) number of boards the pinner has, (vi) number of categories the pinner has, and (vii) category entropy of the pinner. A category entropy quantifies how a user's interest (pinning/repinning) is distributed across multiple categories. We calculate the category entropy of user u as follows:

$$H_{category}(u) = -\sum_{i=1}^{C_u} p_i^u ln(p_i^u) \qquad (2)$$

where C_u is the number of categories the user u has, and p_i^u is the portion of pins/repins in the category i by u.

Figures 10 and 11 show the prediction results on popular and viral cascades, respectively, using image meta and/or pinner information. To investigate whether there is a synergy among different feature sets, we also consider (i) image and image meta features ('IMAGE+META'), (ii) image and pinner features ('IMAGE+PINNER'), (iii) image meta and pinner features ('META+PINNER'), and (iv) image, image meta, and pinner features ('IMAGE+META+PINNER').

As shown in Figures 10 and 11, 'META' and 'PINNER' performs better than 'BASELINE', meaning that meta information of an image as well as its poster's information are useful in predicting both popular and viral image cascades. The image meta information shows a stronger predictive power than pinner's information in predicting popular image cascades, which implies that information about the image is more important than information of a user who posts the image. On the other hand, image meta information performs slightly better (or similarly) than pinner's information in predicting viral image cascades, implying that viral image cascades are similarly associated with image meta and pinner information. If we consider both of image meta and pinner features, i.e., 'META+PINNER', it performs better than others in Figures 10 and 11, which signifies that image meta and pinner features are complementary to each other.

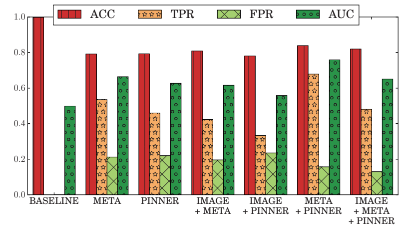

(a) Top 0.1%

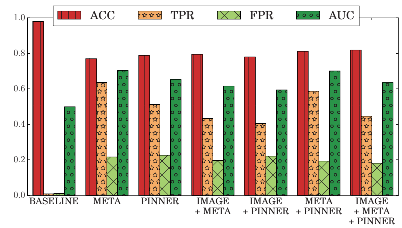

(b) Top 1%

Figure 10: Prediction results on popular cascades using image meta and/or pinner information.

Note that the AUC of 'META+PINNER' for predicting the top 0.1% popular cascades is 0.76, which is much higher than 'BASELINE'. The top 3 features (ranked by the χ^2 values) in 'META' for predicting top 1% popular cascades are 'social sentiment score', 'positive sentiment score', and 'maliciousness', while those for predicting top 1% viral cascades are 'maliciousness', 'source popularity', and 'cognitive sentiment score'. On the other hand, the top 3 features in 'PINNER' for predicting top 1% popular cascades are 'number of followers', 'social sentiment score', and 'number of categories', while those for predicting top 1% viral cascades are 'cognitive sentiment score', 'negative sentiment score', and 'positive sentiment score'. Interestingly, combining 'IMAGE' features (e.g., 'IMAGE+META', 'IMAGE+PINNER', 'IMAGE+META+PINNER') do not contribute much in predicting popular and viral image cascades. It is worth noting that prediction performance of popular cascades is higher than that of viral cascades, which signifies that image meta and pinner information are more useful in predicting popular image cascades than viral image cascades.

Predictive power of initial propagation pattern

We finally examine whether the observation of initial image propagation pattern helps to predict popular or viral image cascades. That is, we observe the first k repins of an image cascade, and predict whether the cascade will belong to the top popular or viral cascades in the future. Note that higher k shows better performance, but we only report $k = 5$ here since our goal is to observe the propagation pattern in the very early stage of a cascade. We consider two perspectives of initial propagation of a cascade: (i) how the cascade initially looks like (referred to as 'STRUCT') and (ii) who are the early adopters in the cascade (referred to as 'ADOPTER').

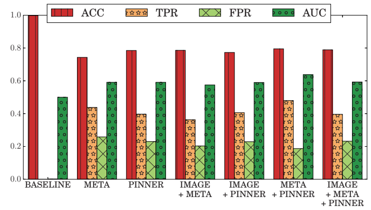

(a) Top 0.1%

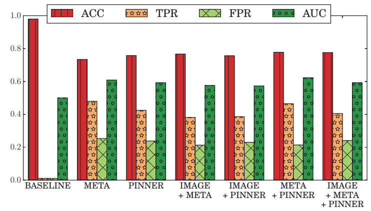

(b) Top 1%

Figure 11: Prediction results on viral cascades using image meta and/or pinner information.

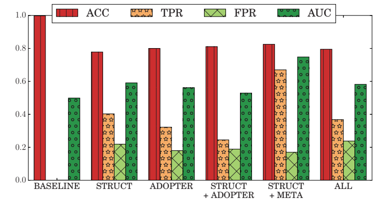

(a) Top 0.1%

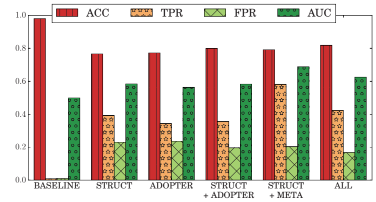

(b) Top 1%

Figure 12: Prediction results on popular cascades using initial propagation pattern.

The 'STRUCT' features consist of (i) max width, (ii) max depth, (iii) wiener index based on Equation 1 (when $N = 6$), (iv) width entropy quantifies how the distribution of widths in a cascade is even or skewed (calculated similarly with Equation 2), (v) inter-repin times for the five repins, and (vi) positive, negative, cognitive, and social sentiment scores for each repin's description using LIWC. The 'ADOPTER' consists of each repinner's following characteristics: (i) number of pins the repinner has, (ii) number of followers who follow the repinner, (iii) number of followees the repinner follows, (iv) number of likes the repinner likes, (v) number of boards the repinner has, (vi) number of categories the repinner has, (vii) category entropy of the repinner (Equation 2), and (viii) positive, negative, cognitive, and social sentiment scores of the repinner's introduction text.

Figures 12 and 13 show the prediction results on popular and viral cascades, respectively, using the initial propagation pattern. To investigate whether there is a synergy among different feature sets, we also consider (i) 'STRUCT+ADOPTER', (ii) 'STRUCT+META', and (iii) 'ALL' that includes all the features. As shown in Figures 12 and 13, 'STRUCT' performs better than 'ADOPTER', meaning that the initial propagation shape of the cascade is stronger predictor than the information of users who initially participate in the cascade. The 'STRUCT+ADOPTER' performs worse than 'STRUCT', meaning that 'ADOPTER' may not contribute much in predicting popular and viral image cascades. Note that the top 3 features (ranked by the χ^2 values) in 'STRUCT' for predicting top 1% popular cascades are 'wiener index', 'social sentiment score', and 'positive sentiment score', while those for predicting top 1% viral cascades are 'width entropy', 'max depth', and 'repin time'. On the other hand, if we combine the 'STRUCT' and 'META' features, it performs better than others in Figures 12 and 13, which implies that 'META' and 'STRUCT' features are complementary to each other. Note that 'STRUCT+META' are not mostly about *'human factors'* but more about *'content factors'*, implying that content factors are important predictors in predicting popular and viral image cascades.

In summary, 'META+PINNER' is the strongest predictor in predicting popular image cascades while 'STRUCT+META' is the strongest predictor in predicting viral image cascades. This implies that we can forecast popular image cascades using the image meta and pinner information at the moment when an image is posted. Also, if we observe initial propagation pattern of an image cascade, we can predict whether the image goes viral based on its meta information and initial propagation pattern. It is worth to note that 'META' information is commonly useful for predicting both popular and viral image cascades.

Conclusion

This paper studied the prediction of popular or viral image cascades in Pinterest. To predict popular or viral image cascades, we considered the following feature sets: (i) image features, (ii) image meta and poster's information, and (iii) initial propagation patterns. We summarize three main contributions as follows. First, we found that image meta and poster's information are strong predictors for predicting popular image cascades, which implies that image popularity is predictable using the image meta and pinner information at the moment when the image is posted. Second, initial propagation pattern of a cascade and image meta information are useful for predicting whether the cascade will go viral in the future. Third, we revealed that using deep image features alone is not as effective in predicting popular or viral cascades, which suggests that a more effective image representation is needed to capture subtle visual traits exhibited in popular or viral images. We believe this work can provide an important insight for content providers, OSN operators, and marketers.

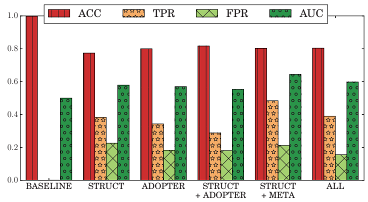

(a) Top 0.1%

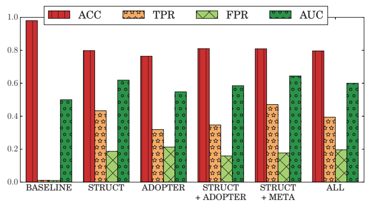

(b) Top 1%

Figure 13: Prediction results on viral cascades using initial propagation pattern.

Acknowledgments

This work is supported in part by National Science Foundation CNS-1302691 grant and the research fund of Hanyang University (HY-2017-N). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research.

References

Aral, S., and Walker, D. 2012. Identifying Influential and Susceptible Members of Social Networks. *Science* 337(6092):337–341.

Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012. The Role of Social Networks in Information Diffusion. In *WWW*.

Breiman, L. 2001. Random Forests. *Springer Machine learning* 45(1):5–32.

Chawla, N. V.; Bowyer, K. W.; Hall, L. O.; and Kegelmeyer, W. P. 2002. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research* 16(1):321–357.

Cheng, J.; Adamic, L.; Dow, P. A.; Kleinberg, J. M.; and Leskovec, J. 2014. Can Cascades Be Predicted? In *WWW*.

Choi, D.; Han, J.; Chung, T.; Ahn, Y.-Y.; Chun, B.-G.; and Kwon, T. T. 2015. Characterizing Conversation Patterns in Reddit: From the Perspectives of Content Properties and User Participation Behaviors. In *ACM Conference on Online Social Networks (COSN)*.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE CVPR*.

Deza, A., and Parikh, D. 2015. Understanding Image Virality. In *IEEE CVPR*.

Fawcett, T. 2006. An Introduction to ROC Analysis. *Pattern Recognition Letters* 27(8):861–874.

Gelley, B., and John, A. 2015. Do I Need To Follow You?: Examining the Utility of The Pinterest Follow Mechanism. In *ACM CSCW*.

Goel, S.; Anderson, A.; Hofman, J.; and Watts, D. J. 2015. The Structural Virality of Online Diffusion. *Management Science* 1–17.

Guerini, M.; Staiano, J.; and Albanese, D. 2013. Exploring Image Virality in Google Plus. In *IEEE International Conference on Social Computing*.

Han, J.; Choi, D.; Chun, B.-G.; Kwon, T. T.; Kim, H.-C.; and Choi, Y. 2014. Collecting, Organizing, and Sharing Pins in Pinterest: Interest-driven or Social-driven? In *ACM SIGMETRICS*.

Han, J.; Choi, D.; Choi, A.-Y.; Choi, J.; Chung, T.; Kwon, T. T.; Rha, J.-Y.; and Chuah, C.-N. 2015. Sharing Topics in Pinterest: Understanding Content Creation and Diffusion Behaviors. In *ACM Conference on Online Social Networks (COSN)*.

Jenders, M.; Kasneci, G.; and Naumann, F. 2013. Analyzing and Predicting Viral Tweets. In *WWW Companion*.

Khosla, A.; Sarma, A. D.; and Hamid, R. 2014. What Makes an Image Popular? In *WWW*.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Neural Information Processing Systems (NIPS)*.

Kupavskii, A.; Ostroumova, L.; Umnov, A.; Usachev, S.; Serdyukov, P.; Gusev, G.; and Kustarev, A. 2012. Prediction of Retweet Cascade Size over Time. In *ACM CIKM*.

Lazer, D.; Pentland, A.; Adamic, L.; Aral, S.; Barabasi, A.-L.; Brewer, D.; Christakis, N.; Contractor, N.; Fowler, J.; Gutmann, M.; Jebara, T.; King, G.; Macy, M.; Roy, D.; and Alstyne, M. V. 2009. Computational Social Science. *Science* 323(5915):721–723.

Lerman, K., and Ghosh, R. 2010. Information Contagion: An Empirical Study of the Spread of News on Digg and Twitter Social Networks. In *ICWSM*.

Liu, H., and Setiono, R. 1995. Chi2: Feature Selection and Discretization of Numeric Attributes. In *IEEE International Conference on Tools with Artificial Intelligence*.

Oliva, A., and Torralba, A. 2001. Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope. *International Journal of Computer Vision* 42(3):145–175.

Pennebaker, J. W.; Mehl, M. R.; and Niederhoffer, K. 2003. Psychological Aspects of Natural Language Use: Our Words, Ourselves. *Annual Review of Psychology* 54:547–577.

Rahman, M. R.; Han, J.; and Chuah, C.-N. 2015. Unveiling the Adoption and Cascading Process of OSN-based Gifting Applications. In *IEEE INFOCOM*.

Rodrigues, T.; Benevenuto, F.; Cha, M.; Gummadi, K.; and Almeida, V. 2011. On Word-of-mouth Based Discovery of the Web. In *ACM IMC*.

Simonyan, K., and Zisserman, A. 2014. Very Deep Convolutional Networks for Large-scale Image Recognition. *arXiv preprint arXiv:1409.1556*.

Totti, L. C.; Costa, F. A.; Avila, S.; Valle, E.; Meira, Jr., W.; and Almeida, V. 2014. The Impact of Visual Attributes on Online Image Diffusion. In *ACM Conference on Web Science*.

VirusTotal. 2016. VirusTotal - Free Online Virus, Malware and URL Scanner. https://www.virustotal.com.

Wang, D.; Wen, Z.; Tong, H.; Lin, C.-Y.; Song, C.; and Barabási, A.-L. 2011. Information Spreading in Context. In *WWW*.

Wiener, H. 1947. Structural Determination of Paraffin Boiling Points. *Journal of the American Chemical Society* 69(1):17–20.

Zhong, C.; Karamshuk, D.; and Sastry, N. 2015. Predicting Pinterest: Automating a Distributed Human Computation. In *WWW*.