

Exploring Image Virality in Google Plus

Marco Guerini

Trento RISE

Trento - Italy

Email: marco.guerini@trentorise.eu

Jacopo Staiano

University of Trento

Trento - Italy

Email: staiano@disi.unitn.it

Davide Albanese

Fondazione Edmund Mach, CRI-CBC

San Michele all'Adige (TN) - Italy

Email: davide.albanese@fmach.it

Abstract—Reactions to posts in an online social network show different dynamics depending on several textual features of the corresponding content. Do similar dynamics exist when images are posted? Exploiting a novel dataset of posts, gathered from the most popular Google+ users, we try to give an answer to such a question. We describe several virality phenomena that emerge when taking into account visual characteristics of images (such as orientation, mean saturation, etc.). We also provide hypotheses and potential explanations for the dynamics behind them, and include cases for which common-sense expectations do not hold true in our experiments.

I. INTRODUCTION

How do things become ‘viral’ on the Internet? And what exactly do we mean by ‘influence’? Since marketing and industry people want their messages to spread in the most effective and efficient way possible, these questions have received a great deal of attention, particularly in recent years, as we have seen a dramatic growth of social networking on the Web. Generally speaking, virality refers to the tendency of a content either to spread quickly within a community or to receive a great deal of attention by it. In studying the spreading process we will focus on the content and its characteristics, rather than on the structure of the network through which the information is moving. In particular, we will investigate the relationships between visual characteristics – of images enclosed in Google+ posts – and virality phenomena. We will use three virality metrics: pluserers, replies and resharers.

This exploratory work stems from the use people make of social networking websites such as Google+, Facebook and similar: we hypothesized that perceptual characteristics of an image could indeed affect the virality of the post embedding it, and that – for example – cartoons, panorama or self-portraits picture affect users’ reactions in different ways. The aim of this paper is to investigate whether signs of such “common-sense” intuition emerge from large-scale data made available on popular social networking websites like Google+ and, in such case, to open discussion on the associated phenomena.

The paper is structured as follows: first, we review previous works addressing the topic of virality in social networks, and particularly some focusing on content impact. Then, after describing the dataset collected and used for this work, we proceed with the study of virality of Google+ posts and the characteristics of their content. We also discuss the behavior of virality indexes in terms of their alternative use, arguing that pluseres and comments fulfill a similar purpose of followers’ “appreciation” while reshares have a different role of “self-representation”. Finally, we investigate possible interactions between image characteristics and users’ typology, in order to

understand to what extent results are generalizable or typical of a community, gathered around a common interest.

II. RELATED WORKS

Several researchers studied information flow, community building and similar processes using Social Networking sites as a reference [1], [2], [3], [4]. However, the great majority concentrates on network-related features without taking into account the actual content spreading within the network [5]. A hybrid approach focusing on both product characteristics and network related features is presented in [6]: the authors study the effect of passive-broadcast and active-personalized notifications embedded in an application to foster word of mouth.

Recently, the correlation between content characteristics and virality has begun to be investigated, especially with regard to textual content; in [7], for example, features derived from sentiment analysis of comments are used to predict the popularity of stories. The work presented in [8] uses *New York Times*’ articles to examine the relationship between emotions evoked by the content and virality, using semi-automated sentiment analysis to quantify the affectivity and emotionality of each article. Results suggest a strong relationship between affect and virality; still, the virality metric considered is interesting but very limited: it only consists of how many people emailed the article. The relevant work in [9] measures a different form of content spreading by analyzing which are the features of a movie quote that make it “memorable” online. Another approach to content virality, somehow complementary to the previous one, is presented in [10], trying to understand which modification dynamics make a meme spread from one person to another (while movie quotes spread remaining exactly the same).

More recently, some works tried to investigate how different textual contents give rise to different reactions in the audience: the work presented in [11] correlates several viral phenomena with the wording of a post, while [12] show that specific content features variations (like the readability level of an abstract) differentiate among virality level of downloads, bookmarking, and citations. Still, to our knowledge, no attempt has been made yet to investigate the relation between visual content characteristics and virality.

III. DATA DESCRIPTION

Using the Google+ API¹, we harvested the public posts from the 979 top followed users in

¹<https://developers.google.com/+/api/>

Google+ (plus.google.com), as reported by the socialstatistics.com website on March 2nd 2012². The time span for the harvesting is one year, from June 28th 2011 (Google+ date of launch) to June 29th 2012.

We decided to focus on the most popular users for several reasons: (i) the dataset is uniform from the point of view of sample role, i.e. VIPs, (ii) the behavior of the followers is consistent – e.g. no friendship dynamics – and (iii) extraneous effects due to followers network is minimized, since top followed users’ network is vast enough to grant that, if a content is viral, a certain amount of reactions will be obtained.

We defined 3 subsets of our dataset, comprising respectively: (i) posts containing a static image, (ii) posts containing an animated image (usually, gif), (iii) posts without attachments (text-only). All other posts (containing as attachment videos, photo albums, links to external sources) were discarded. Statistics for our dataset are reported in Table I. For each post, we considered three virality metrics³:

- **Plusoners**: the number of people who +1’d;
- **Replies**: the number of comments;
- **Resharers**: the number of people who reshared.

TABLE I. AN OVERVIEW OF THE GOOGLE+ DATASET.

Global	
actors	979
posts	289434
published interval	6/28/11–6/29/12
Posts with static images	
actors	950
posts	173860
min/max/median posts per actor	1/3685/ 65.5
min/max/median plusoners per post	0/9703/33.0
min/max/median replies per post ^a	0/571/12.0
min/max/median resharers per post	0/6564/4.0
Posts with animated images	
actors	344
posts:	12577
min/max/median posts per actor	1/2262/3.0
min/max/median plusoners per post	0/5145/17.0
min/max/median replies per post	0/500/7.0
min/max/median resharers per post	0/6778/10.0
Posts without attachments, text-only	
actors	939
posts	102997
min/max/median posts per actor	1/1744/41.0
min/max/median plusoners per post	0/20299 /16.0
min/max/median replies per post	0/538/17.0
min/max/median resharers per post	0/13566/1.0

^aReplies count is cut around 500 by the API service.

In Figures 1 and 2 we display the evolution over time of the network underlying our dataset (using a week as temporal unit), and of the reactions to posts given by users, respectively. We notice that:

- 1) the average number of reactions per user shows quite different trends depending on the metric considered: while replies tend not to be affected by the growth of the network, reshares and, to a lesser degree plusones, show an ever-growing trend.
- 2) The temporal plot of the average number of followers per user (Figure 1) in our dataset (in Google+ terminology,

the number of people who *circled* them) shows a gradient increase around weeks 28/29. Interestingly, this is reflected in the plot of reactions over time (Figure 2): the gradient increases around the same weeks, for reshares and plusones; these effects are most probably due to Google+ transitioning from beta to public in late September 2011 (a similar phenomenon is reported also in [13]).

- 3) Finally, the orders of magnitude of such growths are very different: we notice that while reactions increase of a factor of 7 over the time period we took into account, the total number of followers increased of a factor of 25.

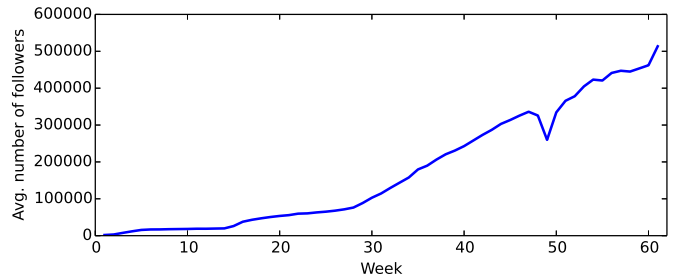


Fig. 1. Average number of followers per user, at 1-week temporal granularity.

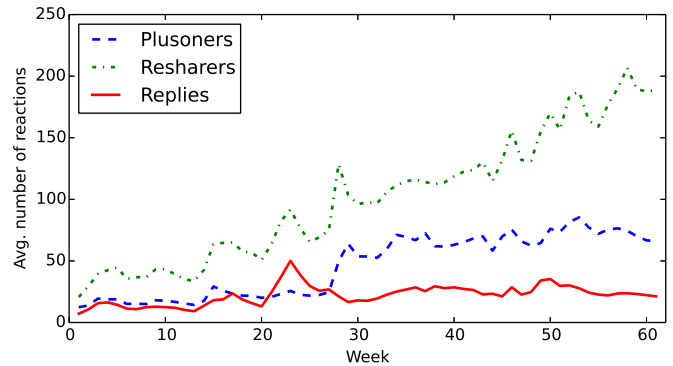


Fig. 2. Average number of reactions per user, at 1-week temporal granularity. This value represents the average number of reactions elicited by each user’s posts over 1-week time-slices.

The relative amount of followers’ reactions does not significantly increase as the network grows⁴. As detailed in the next section, our analyses are based on comparing probability distributions: e.g. we evaluate if grayscale images have a significantly higher or smaller probability of reaching a certain virality score than colored ones. In the following analyses, for the sake of clarity, our discussion will not take into account the normalization factor (i.e. the size of the audience when a content is posted). Indeed, we have run the same analyses normalizing the virality indexes of a given post against its *potential* audience: i) the effects are still visible, ii) the effects are consistent both in significance and sign with the not-normalized distributions, but iii) differences have lower magnitude (explained by the fact that virality indexes should be normalized using the *actual* audience – e.g. the

²The dataset presented and used in this work will be made available to the community for research purposes.

³Since the API provide only an aggregate number, we cannot make any temporal analysis of how reactions to a post were accumulated over time.

⁴It has been noted how (see, for instance, <http://on.wsj.com/zjRr06>), especially in the time frame we consider, users’ activity did not increase much in front of the exploding network size.

followers exposed to the content). Thus, since we are interested in comparing the virality of different image categories and our preliminary experiments showed that by normalizing the indexes their comparisons, their sign, and the derived interpretations still hold, we choose to report the non-normalized version of the results that are more intuitively readable.

In the following sections, after the analyses of text-only posts and of posts containing an animated image, we will consider the subset of static images as the reference dataset. Exemplar pictures taken from the dataset are shown in Figure 3, depicting some image categories that we will take into account in the following sections.

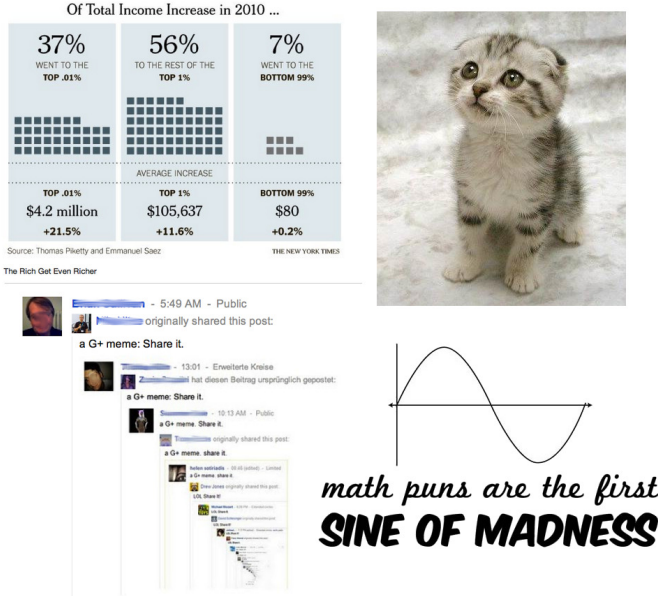


Fig. 3. Exemplar pictures from the dataset.

IV. DATA ANALYSIS

Virality metrics in our dataset follow a power-law-like distribution thickening toward low virality score. In order to evaluate the “virality power” of the features taken into account, we compare the virality indexes in terms of empirical Complementary Cumulative Distribution Functions (CCDFs). These functions are commonly used to analyse online social networks in terms of growth in size and activity (see for example [14], [15], or the discussion presented in [16]) and also for measuring content diffusion, e.g. the number of retweets of a given content [17]. Basically, these functions account for the probability p that a virality index will be greater than n and are defined as follows:

$$\hat{F}(n) = \frac{\text{number of posts with virality index} > n}{\text{total number of posts}} \quad (1)$$

For example, the probability of having a post with more than 75 plusoners is indicated with $\hat{F}_{plus}(75) = P(\#\text{plusoners} > 75)$. In the following sections we use CCDFs to understand the relation between image characteristics and post virality; in order to assess whether the CCDFs of the

several types of posts we take into account show significant differences, we will use the Kolmogorov-Smirnov (K-S) goodness-of-fit test, which specifically targets cumulative distribution functions.

A. Image vs. text-only

First of all, we aim to understand what is the impact of “adding an image to a post” in Google+. Some studies [18] already show that posts containing an image are much more viral than simple plain-text posts, and that various characteristics of image based banners affect viewer’s recall and clicks [19]. This finding can be explained in light of a “rapid cognition” model [20], [21]. In this model, the user has to decide in a limited amount of time, and within a vast information flow of posts, whether to take an action on a particular post (e.g. to reply, reshare, give it a plusone). Thus, pictures, and the characteristics thereof analyzed in the following sections, might play a role of paramount importance in her decision-making process as she exploits visual cues that grab her attention. In some respects, the rapid cognition model is reminiscent of the mechanisms by which humans routinely make judgments about strangers’ personality and behavior from very short behavioral sequences and non-verbal cues [22], [23].

In order to investigate the general impact of images we compared posts containing a picture with posts containing only text. While our findings overall coincide with [18], some interesting phenomena emerged. First, we see that the probability for a post with an image to have a high number of resharers is almost three times greater ($\hat{F}_{resh}(10) = 0.28$ vs. 0.10, K-S test $p < 0.001$), see Figure 4.c. Still, the CCDFs for the other virality indexes show different trends:

- Posts containing images have lower probability of being viral when it comes to number of comments ($\hat{F}_{repl}(50) = 0.33$ vs. 0.22, K-S test $p < 0.001$), see Figure 4.b. This can be explained by the fact that text-only posts elicit more “linguistic-elaboration” than images (we also expect that the average length of comments is higher for text-only posts but we do not investigate this issue here).
- Also, if we focus on simple appreciation (plusoners in Figure 4.a), results are very intriguing: while up to about 75 plusoners the probability of having posts containing images is higher, after this threshold the situation capsizes. This finding can be of support to the hypothesis that, while it is easier to impress with images in the information flow — as argued with the aforementioned “rapid cognition” model — high quality textual content can impress more.

B. Static vs. Animated

Animated images add a further dimension to pictures expressivity. Having been around since the beginning of the Internet (the gif format was introduced in late 80’s), animated images have had alternate fortune, especially after the wide spread of services like youtube and the availability of broadband. Nonetheless, they are still extensively used to produce simple animations and short clips. Noticeably, the value of simple and short animations has been acknowledged by Twitter with the recently released *Vine* service.

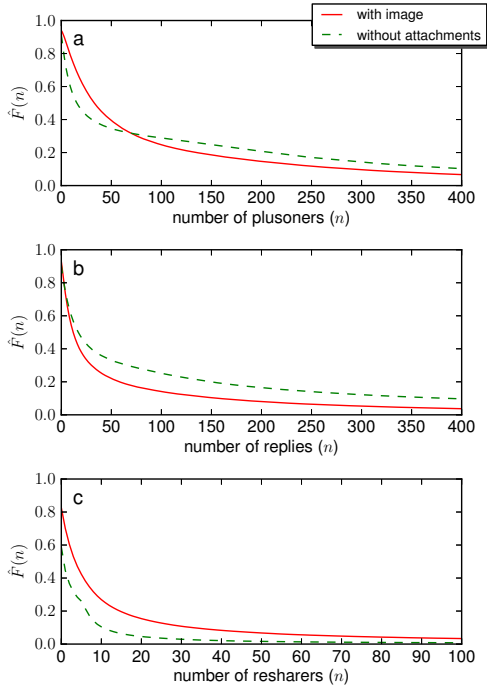


Fig. 4. Virality CCDFs for posts with image vs. text-only posts.

Whether a post contains a static or animated image has a strong discriminative impact on all virality indexes, see Figure 5. With respect to plusoners and replies, static images tend to show higher CCDFs (respectively two and three times more, $\hat{F}_{plus}(75) = 0.30$ vs. 0.17 , $\hat{F}_{repl}(50) = 0.22$ vs. 0.08 , K-S test $p < 0.001$), while on resharers the opposite holds.

The fact that $\hat{F}_{resh}(n)$ is two times higher for posts containing animated images ($\hat{F}_{resh}(10) = 0.48$ vs. 0.27 , K-S test $p < .001$) can be potentially explained by the fact that animated images are usually built to convey a small “memetic” clip - i.e. *funny*, *cute* or *quirky* situations as suggested in [24].

In order to verify this hypothesis we have annotated a small random subsample of 200 images. 81% of these animated images were found to be “memetic” (two annotators were used, positive example if the image score 1 at least on one of the aforementioned dimensions, annotator agreement is very high — Cohen’s kappa 0.78). These findings indicate that animated images are mainly a vehicle for amusement, at least on Google+.

C. Image Orientation

We then focused on the question whether image orientation (*landscape*, *portrait* and *squared*) has any impact on virality indexes. We included squared images in our analysis since they are typical of popular services a la *Instagram*. These services enable users to apply digital filters to the pictures they take and confine photos to a squared shape, similar to Kodak Instamatic and Polaroids, providing a so-called “vintage effect”.

We have annotated a small random subsample of 200 images. 55% of these images were found to be “Instagrammed” (two annotators were used, positive example if the image

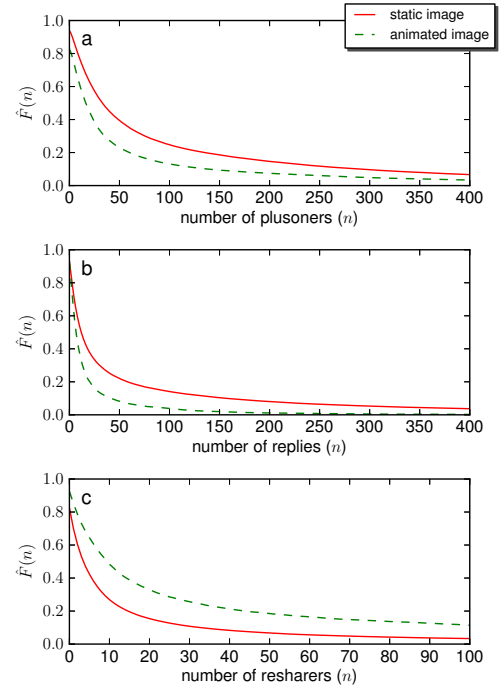


Fig. 5. Virality CCDFs for static vs. animated images.

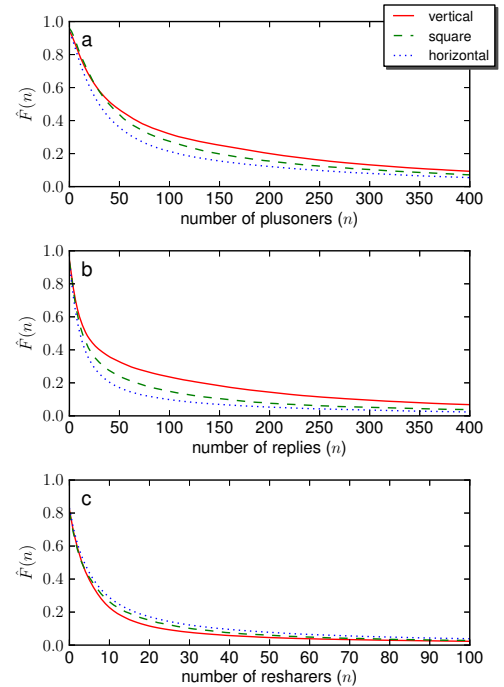


Fig. 6. Virality CCDFs for image orientation.

is clearly recognized as modified with a filter; annotator agreement is high – Cohen’s kappa 0.68). Note that, if we include also black and white squared pictures without any other particular filter applied (b/w is one of the “basic” filter provided by *Instagram*) the amount of Instagrammed pictures rises to 65%. Obviously, the ratio of pictures modified with this

and similar services could be higher; here, we rather wanted to identify those pictures that were clearly recognized as seeking for the aforementioned “vintage effect”.

While the orientation seems not to have strong impact on resharers, with a mild prevalence of horizontal pictures (see Figure 6.c), plusoners and replies tend to well discriminate among various image orientations. In particular, portrait images show higher probability of being viral than squared images than, in turn, landscapes (see Figure 6.a and 6.b).

Furthermore, CCDFs indicate that vertical images tend to be more viral than horizontal ones ($\hat{F}_{plus}(75) = 0.38$ vs. 0.26, $\hat{F}_{repl}(50) = 0.38$ vs. 0.17, K–S test $p < 0.001$). Hence, while squared images place themselves in the middle in any metric, landscape images have lower viral probability for plusones and replies but slightly higher probability for reshares.

This can be partially explained by the fact that we are analyzing “celebrities” posts. If the vertically-orientated image contains the portrait of a celebrity this is more likely to be appreciated rather than reshared, since the act of resharing can also be seen as a form of “self-representation” of the follower (we will analyze the impact of pictures containing faces in the following section). The opposite holds for landscapes, i.e. they are more likely to be reshared and used for self-representation.

D. Images containing one face

In traditional mono-directional media (e.g. tv, billboards, etc.) a widely used promotion strategy is the use of testimonials, especially celebrities endorsing a product. Is the same strategy applicable to Social Media? Understanding the effect of posting images with faces by most popular Google+ users (and hypothesizing that those are their faces) is a first step in the direction of finding an answer.

We computed how many faces are found in the images, along with the ratio of the area that include faces and the whole image area, using the Viola-Jones [25] face detection algorithm. We considered images containing one face vs. images containing no faces. We did not consider the surface of image occupied by the face (i.e. if it is a close-up portrait, or just a small face within a bigger picture). The discriminative effect of containing a face on virality is statistically significant but small. Still, the pictures containing faces tend to have mild effect on resharers (slightly higher replies and plusoners but lower resharers as compared to images with no faces).

In order to verify the hypothesis mentioned earlier, i.e. that self-portraits tend to be reshared less, we also focused on a subsample of images containing faces that cover at least 10% of the image surface (about 6400 instances). In this case, the differences among indexes polarize a little more (higher plusoners and comments, lower resharers), as we were expecting. Unfortunately, images with even higher face/surface ratio are too few to further verify the hypotheses.

E. Grayscale vs. Colored

The impact and meaning of black-and-white (i.e. grayscale) photographic images has been studied from different perspectives (e.g. semiotics and psychology) and with reference to different fields (from documentary to arts and advertising).

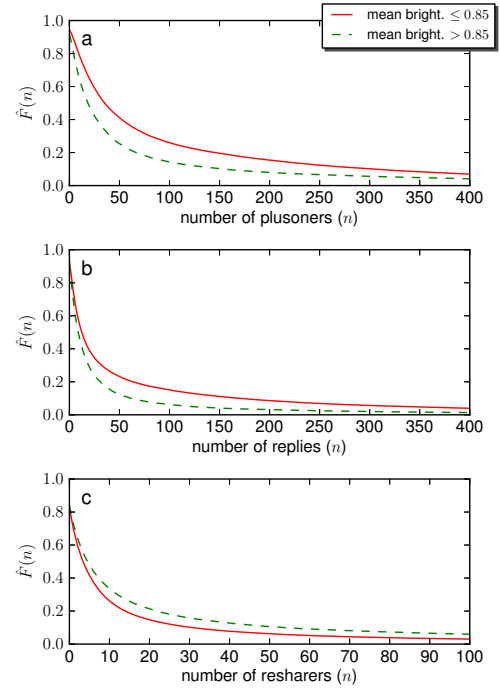


Fig. 7. Virality CCDFs for image Brightness.

Rudolf Arnheim, for example, argues that color produces essentially *emotional* experience, whereas shape corresponds to *intellectual* pleasure [26]. Hence, black-and-white photography, because of its absence of expressive colors, focuses on shapes that require intellectual reflection and brings to explore aesthetic possibilities. We want to understand if such functions and effects can be spotted in our virality indexes.

In order to have a “perceptual” grayscale (some images may contain highly desaturated colors and so perceived as shades of gray) we dichotomized the dataset according to the mean-saturation index of the images, using a very conservative threshold of 0.05 (on a 0-1 scale).

As can be seen in Figure 8.a and 8.b, colored images (with saturation higher than 0.05) have a higher probability of collecting more plusoners and replies as compared to images with lower saturation (grayscale). In particular the probability functions for replies is more than two times higher ($\hat{F}_{repl}(50)$ values are 0.26 vs. 0.10, K–S test $p < 0.001$). Instead, image saturation has no relevant impact on resharers.

F. Very Bright Images

After converting each image in our dataset to the HSB color space, we extracted its mean Saturation and Brightness. More in detail, the HSB (Hue/Saturation/Brightness) color space describes each pixel in an image as a point on a cylinder: the Hue dimension representing its color within the set of primary-secondary ones, while Saturation and Brightness describe respectively how close to the pure color (i.e. its Hue), and how bright it is. We split the dataset according to images mean brightness using a threshold of 0.85 (in a scale included between 0 and 1). Usually images with such a high mean brightness tend to be cartoon-like images rather than pictures.

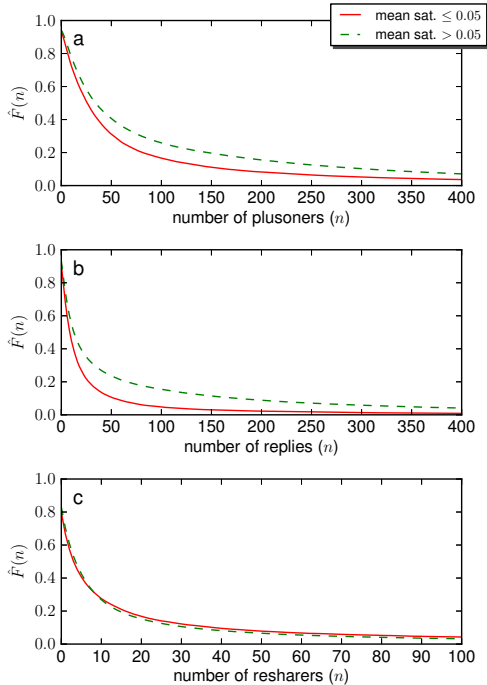


Fig. 8. Virality CCDFs for Grayscale vs. Colored images.

Previous research [27] has shown that pixel brightness is expected to be higher in cartoon-like (or significantly “photoshopped”) than in natural images.

Image brightness level has a strong impact on plusoners and replies, and a milder one on resharers. Brighter images have a lower probability of being viral on the first two indexes (Figure 7.a and 7.b) and a higher probability on the latter (Figure 7.c). In particular, lower brightness images have a plusone and reshare probability almost two times higher ($\hat{F}_{plus}(75) = 0.31$ vs. 0.18 , $\hat{F}_{repl}(50) = 0.23$ vs. 0.12 , K–S test $p < 0.001$), while for resharers it is 27% higher in favor of high brightness images ($\hat{F}_{resh}(10) = 0.33$ vs. 0.26). Surprisingly, analyzing a small random subsample of 200 very bright images, we found that while 88% of these images contained some text, as we would have expected, only 13% were cartoon/comics and only 13% contained the real picture of an object as subject, even if highly “photoshopped”. Above all, only a small amount of these images (21%) was considered funny or memetic⁵. The great majority comprised pictures containing infographics, screenshots of software programs, screenshots of social-networks posts and similar. In this respect we are analyzing a content that is meant to be mainly informative, and is somehow complementary to the content of animated pictures (mainly intended for amusement, see IV-B).

G. Vertical and Horizontal edges

Finally, we want to report on an explorative investigation we made. We focused on the impact of edges intensity on posts virality. The intensity of vertical/horizontal/diagonal edges was

⁵Two annotators were used, four binary categories were provided (contain-text/comics/real-picture-obj/funny). The overall inter-annotator agreement on these categories is high, Cohen’s kappa 0.74.

computed using Gaussian filters, based on code used in [28] in the context of real-time visual concept classification. The probability density of the average edges intensity follows a gaussian-like distribution, with mean of about 0.08 (both for horizontal and vertical edges). We divided images into two groups: those having an average edge intensity below the sample mean, and those having an average edge intensity above the mean. Results showed that images with horizontal edge intensity below the sample mean are far more viral on the plusoners and replies indexes, while vertical are less discriminative. Results for horizontal hedges are as follows: $\hat{F}_{plus}(75) = 0.36$ vs. 0.22 , $\hat{F}_{repl}(50) = 0.27$ vs. 0.14 , $\hat{F}_{resh}(10) = 0.25$ vs. 0.29 , K–S test $p < 0.001$. While these results do not have an intuitive explanation, they clearly show that there is room for further investigating the impact of edges.

H. Virality Indexes Correlation

From the analyses above, virality indexes seem to “move together” (in particular plusoners and replies) while resharers appear to indicate a different phenomenon. We hypothesize that plusoners and replies can be considered as a form of endorsement, while reshares are a form of self-representation. This explains why, for example, pictures containing faces are endorsed but not used for self-representation by VIPs’ followers. On the contrary, animated images that usually contain funny material are more likely to provoke reshares for followers’ self-representation. In fact, people usually tend to represent themselves with positive feelings rather than negative ones (especially popular users, see [29]), and positive moods appear to be associated with social interactions [30], [31].

TABLE II. VIRALITY INDEXES CORRELATION ON THE DATASETS

	Pearson	MIC
Static images		
plusoners vs. replies	0.723	0.433
plusoners vs. resharers	0.550	0.217
replies vs. resharers	0.220	0.126
Animated Images		
plusoners vs. replies	0.702	0.304
plusoners vs. resharers	0.787	0.396
replies vs. resharers	0.554	0.205
Text Only		
plusoners vs. replies	0.802	0.529
plusoners vs. resharers	0.285	0.273
replies vs. resharers	0.172	0.185

This is supported also by the correlation analysis of the three virality indexes, reported in Table II, made on the various datasets we exploited. In this analysis we used both the Pearson coefficient and the recent Maximal Information Coefficient (MIC), considering plusoners ≤ 1200 , replies ≤ 400 e resharers ≤ 400 . MIC is a measure of dependence introduced in [32] and it is part of the Maximal Information-based Nonparametric Exploration (MINE) family of statistics. MIC is able to capture variable relationships of different nature, penalizing similar levels of noise in the same way. In this study we use the Python package *minepy* [33].

In particular, from Table II we see that: plusones and replies always have a high correlation while replies and resharers always correlate low. Plusoners and reshares, that have a mild correlation in most cases, correlate highly when it comes to funny pictures, i.e. animated ones. This can be explained by a specific “procedural” effect: the follower expresses his/her

appreciation for the funny picture and, after that, he/she reshapes the content. Since resharing implies also writing a comment in the new post, the reply is likely not to be added to the original VIP’s post.

In Table III we sum up the main findings of the paper, comparing the various CCDFs: animated images and infographics have much higher probability of being reshared, while colored images or images containing faces have higher probability of being appreciated or commented. Finally, black and white pictures (grayscale) turn out to be the least “viral” on Google+.

TABLE III. SUMMARY OF MAIN FINDINGS OF THE ANALYSIS.

	$\hat{F}_{plus}(75)$	$\hat{F}_{repl}(50)$	$\hat{F}_{resh}(10)$
very bright	0.18	0.12	0.33
grayscale	0.21	0.11	0.28
color	0.31	0.24	0.27
animated	0.17	0.08	0.48
one-face > 10% area	0.35	0.30	0.23

V. USER ANALYSIS

Finally, we investigate if there is any relevant interaction between images characteristics and VIP’s typology. In Table IV we report demographic details⁶ on the Google+ dataset, as provided by the users in their profile pages.

TABLE IV. USER DEMOGRAPHICS IN THE GOOGLE+ DATASET.

User-category	Female (%)	Male (%)	Neutral (%)	Total (%)
Technology	35 (19%)	110 (61%)	36 (20%)	181 (19%)
Photography	41 (24%)	130 (76%)	1 (1%)	172 (18%)
Music	96 (59%)	48 (29%)	19 (12%)	163 (17%)
Writing	26 (21%)	76 (63%)	19 (16%)	121 (13%)
Actor	21 (36%)	34 (59%)	3 (5%)	58 (6%)
Entrepreneur	12 (29%)	29 (71%)	-	41 (4%)
Sport	-	22 (55%)	18 (45%)	40 (4%)
Artist	11 (31%)	21 (60%)	3 (9%)	35 (4%)
TV	8 (24%)	11 (33%)	14 (42%)	33 (3%)
Company	-	-	28 (100%)	28 (3%)
Website	-	-	23 (100%)	23 (2%)
Politician	-	19 (86%)	3 (14%)	22 (2%)
No Category	6 (43%)	8 (57%)	-	14 (1%)
Organization	-	-	9 (100%)	9 (1%)
Not Available	-	-	7 (100%)	7 (1%)
Other	1 (33%)	2 (67%)	-	3 (0%)
Total	257 (27%)	510 (54%)	183 (19%)	950 (100%)

In order to investigate possible user category effects in our dataset — that is, if our analyses are also influenced by the type of user posting images rather than by the actual content solely, we evaluated the entropy for each image category over the 16 user categories (as defined in Table IV). In Table V we report the contingency table of image-category entropy distributions over user-categories. Looking at the Kullback-Leibler (KL) divergence of specific image categories with respect to the reference distribution (i.e., taken as the total number of images posted by each user-category), we observe very few but interesting effects due to specific user-categories.

In particular, while all the KL divergences are very small, two of them (for Grayscale and High Brightness, reported in

⁶No Category denotes users that do not provide any personal information and for which it was not possible to trace back their category; Not Available denotes seven accounts that were no more publicly accessible when we gathered demographic info; Other denotes very rare and unusual category definitions. The Neutral gender refers to pages afferent to “non-humans” like products, brands, websites, firms, etc.

Bold) are an order of magnitude greater than other classes. Interestingly the divergence is explained mainly by the distribution gap in only two User’s categories. For High Brightness the gap is mainly given by Technology user category that doubles its probability distribution (from 22% to 40%) and Music and Photography that reduce their probability distribution to one third. This divergence from the reference distribution is consistent with the analysis of the content we made in section IV-F: these images where mainly infographics and screenshots of software programs and social networks (so mainly connected to technology). For Grayscale the gap is mainly given by Photography users category that rises by 50% its probability distribution and Music, that reduces it to one third. This gap is consistent with the idea, expressed in Section IV-E, that black-and-white photography is a particular form of art expressivity mainly used by professionals.

VI. CONCLUSIONS

We have presented a study, based on a novel dataset of Google+ posts, showing that perceptual characteristics of an image can strongly affect the virality of the post embedding it. Considering various kinds of images (e.g. cartoons, panorama or self-portraits) and related features (e.g. orientation, animations) we saw that users’ reactions are affected in different ways. We provided a series of analyses to explain the underlying phenomena, using three virality metrics (namely pluserers, replies and resharers). Results suggest that pluserers and replies “move together” while reshapes indicate a distinct users’ reaction. In particular, funny and informative images have much higher probability of being reshared but are associated to different image features (animation and high-brightness respectively), while colored images or images containing faces have higher probability of being appreciated and commented.

Future work will dig deeper into the assessment of relations between visual content and virality indexes, adopting multivariate analysis that includes user’s categories (e.g. which is the viral effect of b/w pictures taken by professional photographer as compared to those taken by non professional users). We will also extend our experimental setup in the following ways: (a) taking into account compositional features of the images, i.e. resembling concepts such as the well-known “rule of thirds”; (b) extracting and exploiting descriptors such as color histograms, oriented-edges histograms; (c) building upon the vast literature available in the context of scene/object recognition, dividing our dataset into specific categories in order to analyse relations between categories, such as natural images or sport images, and their virality.

ACKNOWLEDGMENT

The work of J. Staiano and M. Guerini has been partially supported by the FIRB project S-PATTERNS and the Trento RISE PerTe project, respectively.

REFERENCES

- [1] K. Lerman and R. Ghosh, “Information contagion: an empirical study of the spread of news on digg and twitter social networks,” in *Proceedings of ICWSM-10*, Mar 2010.
- [2] S. Jamali and H. Rangwala, “Digging digg : Comment mining, popularity prediction, and social network analysis,” in *Proceedings of International Conference on Web Information Systems and Mining*, 2009.

TABLE V. CONTINGENCY TABLE OF IMAGE-CATEGORY DISTRIBUTIONS OVER USER-CATEGORIES.

User-category	Grayscale	Colored	High Brightness	Low Brightness	Containing Face	Containing No Face	Squared	Vertical	Horizontal	Total
No Category	7%	6%	9%	6%	5%	7%	4%	5%	7%	6%
Actor	4%	6%	5%	5%	8%	5%	5%	6%	5%	5%
Artist	5%	6%	7%	6%	6%	6%	5%	7%	6%	6%
Company	0%	1%	1%	1%	1%	1%	1%	1%	1%	1%
Entrepreneur	8%	7%	6%	7%	7%	7%	8%	5%	8%	7%
Music	3%	16%	3%	16%	19%	12%	15%	29%	8%	14%
Not Available	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Organization	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Other	0%	0%	0%	0%	0%	0%	2%	0%	0%	0%
Photography	31%	19%	9%	22%	15%	23%	23%	14%	23%	20%
Politician	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Sport	0%	3%	1%	3%	4%	2%	2%	2%	3%	2%
Technology	27%	22%	40%	20%	19%	24%	16%	18%	25%	22%
TV	1%	2%	1%	2%	3%	1%	5%	1%	2%	2%
Website	1%	2%	2%	2%	2%	2%	1%	1%	2%	2%
Writing	11%	10%	17%	10%	11%	10%	11%	10%	11%	11%
KL-divergence	0.173	0.002	0.259	0.003	0.027	0.006	0.047	0.076	0.029	

[3] E. Khabiri, C.-F. Hsu, and J. Caverlee, "Analyzing and predicting community preference of socially generated metadata: A case study on comments in the digg community," in *Proceedings of ICWSM-09*, 2009.

[4] J. Z. Aaditeshwar Seth and R. Cohen, "A multi-disciplinary approach for recommending weblog messages," in *The AAAI 2008 Workshop on Enhanced Messaging*, 2008.

[5] K. Lerman and A. Galstyan, "Analysis of social voting patterns on digg," in *Proceedings of the first workshop on Online social networks*, ser. WOSP '08. New York, NY, USA: ACM, 2008, pp. 7–12. [Online]. Available: <http://doi.acm.org/10.1145/1397735.1397738>

[6] S. Aral and D. Walker, "Creating social contagion through viral product design: A randomized trial of peer influence in networks," *Management Science*, vol. 57, no. 9, pp. 1623–1639, 2011.

[7] S. Jamali, "Comment mining, popularity prediction, and social network analysis," Master's thesis, George Mason University, Fairfax, VA, 2009.

[8] J. A. Berger and K. L. Milkman, "Social Transmission, Emotion, and the Virality of Online Content," *Social Science Research Network Working Paper Series*, December 2009.

[9] C. Danescu-Niculescu-Mizil, J. Cheng, J. Kleinberg, and L. Lee, "You had me at hello: How phrasing affects memorability," in *Proceedings of the ACL*, 2012.

[10] M. Simmons, L. A. Adamic, and E. Adar, "Memes online: Extracted, subtracted, injected, and recollected," *Proceedings of ICWSM-11*, 2011.

[11] M. Guerini, C. Strapparava, and G. Özal, "Exploring text virality in social networks," in *Proceedings of ICWSM-11*, Barcelona, Spain, July 2011.

[12] M. Guerini, A. Pepe, and B. Lepri, "Do linguistic style and readability of scientific abstracts affect their virality," *Proceedings of ICWSM-12*, 2012.

[13] D. Schiöberg, S. Schmid, F. Schneider, S. Uhlig, H. Schiöberg, and A. Feldmann, "Tracing the birth of an osn: social graph and profile analysis in google+," in *Proceedings of the 3rd Annual ACM Web Science Conference*, ser. WebSci '12. New York, NY, USA: ACM, 2012, pp. 265–274. [Online]. Available: <http://doi.acm.org/10.1145/2380718.2380753>

[14] Y.-Y. Ahn, S. Han, H. Kwak, S. Moon, and H. Jeong, "Analysis of topological characteristics of huge online social networking services," in *Proceedings of the 16th international conference on World Wide Web*. ACM, 2007, pp. 835–844.

[15] J. Jiang, C. Wilson, X. Wang, P. Huang, W. Sha, Y. Dai, and B. Y. Zhao, "Understanding latent interactions in online social networks," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 2010, pp. 369–382.

[16] J. Leskovec, *Dynamics of large networks*. ProQuest, 2008.

[17] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 591–600.

[18] TrackSocial, "Optimizing Facebook Engagement, whitepaper," 2012.

[19] H. Li and J. L. Bukovac, "Cognitive impact of banner ad characteristics: An experimental study," *Journalism & Mass Communication Quarterly*, vol. 76, no. 2, pp. 341–353, 1999.

[20] N. Ambady and R. Rosenthal, "Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis," *Psychological bulletin*, vol. 111, no. 2, p. 256, 1992.

[21] D. Kenny, *Interpersonal perception: A social relations analysis*. The Guilford Press, 1994.

[22] B. Lepri, R. Subramanian, K. Kalimeri, J. Staiano, F. Pianesi, and N. Sebe, "Employing social gaze and speaking activity for automatic determination of the extraversion trait," in *International Conference on Multimodal Interfaces and the Workshop on Machine Learning for Multimodal Interaction*, ser. ICMI-MLMI '10. New York, NY, USA: ACM, 2010, pp. 7:1–7:8. [Online]. Available: <http://doi.acm.org/10.1145/1891903.1891913>

[23] J. R. Curhan and A. Pentland, "Thin slices of negotiation: predicting outcomes from conversational dynamics within the first 5 minutes," *Journal of Applied Psychology*, vol. 92, no. 3, p. 802, 2007.

[24] C. Dufour, "An investigation into the use of viral marketing for the companies and the key success factors of a good viral campaign," Ph.D. dissertation, Dublin Business School, 2011.

[25] P. Viola and M. J. Jones, "Robust real-time face detection," *Int. J. Comput. Vision*, vol. 57, no. 2, pp. 137–154, May 2004. [Online]. Available: <http://dx.doi.org/10.1023/B:VISI.0000013087.49260.fb>

[26] R. Arnheim, *Art and visual perception*. Stockholms Universitet, 1987.

[27] T. I. Ianeva, A. P. de Vries, and H. Rohrig, "Detecting cartoons: a case study in automatic video-genre classification," in *Proceedings of the 2003 International Conference on Multimedia and Expo - Volume 2*, ser. ICME '03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 449–452. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1170745.1171531>

[28] J. Uijlings, A. Smeulders, and R. Scha, "Real-time visual concept classification," *Multimedia, IEEE Transactions on*, vol. 12, no. 7, pp. 665–681, nov. 2010.

[29] D. Quercia, J. Ellis, L. Capra, and J. Crowcroft, "In the mood for being influential on twitter," *Proceedings of IEEE SocialCom'11*, 2011.

[30] J. R. Vittengl and C. S. Holt, "A time-series diary study of mood and social interaction," *Motivation and Emotion*, vol. 22, no. 3, pp. 255–275, 1998.

[31] M. De Choudhury, S. Counts, and M. Gamon, "Not all moods are created equal! exploring human emotional states in social media," in *Proceedings of ICWSM-12*, 2012.

[32] D. N. Reshef, Y. A. Reshef, H. K. Finucane, S. R. Grossman, G. McVean, P. J. Turnbaugh, E. S. Lander, M. Mitzenmacher, and P. C. Sabeti, "Detecting novel associations in large data sets," *science*, vol. 334, no. 6062, pp. 1518–1524, 2011.

[33] D. Albanese, M. Filosi, R. Visintainer, S. Riccadonna, G. Jurman, and C. Furlanello, "minerva and minepy: a c engine for the mine suite and its r, python and matlab wrappers," *Bioinformatics*, vol. 29, no. 3, pp. 407–408, 2013.