

Deep Feelings: A Massive Cross-Lingual Study on the Relation between Emotions and Virality

Marco Guerini
Trento Rise
Via Sommarive 18
Trento, Italy
marco.guerini@trentorise.eu

Jacopo Staiano
LIP6, UPMC - Sorbonne Universités
and FBK MobS Lab
4 place Jussieu, Paris, France
jacopo.staiano@lip6.fr

ABSTRACT

This article provides a comprehensive investigation on the relations between virality of news articles and the emotions they are found to evoke. Virality, in our view, is a phenomenon with many facets, i.e. under this generic term several different effects of persuasive communication are comprised. By exploiting a high-coverage and bilingual corpus of documents containing metrics of their spread on social networks as well as a massive affective annotation provided by readers, we present a thorough analysis of the interplay between evoked emotions and viral facets.

We highlight and discuss our findings in light of a cross-lingual approach: while we discover differences in evoked emotions and corresponding viral effects, we provide preliminary evidence of a generalized explanatory model rooted in the deep structure of emotions: the Valence-Arousal-Dominance (VAD) circumplex. We find that viral facets appear to be consistently affected by particular VAD configurations, and these configurations indicate a clear connection with distinct phenomena underlying persuasive communication.

Categories and Subject Descriptors

H.5.m [Information Systems Applications]: Miscellaneous

General Terms

Human Factors

Keywords

Virality, Emotions, Social Media, Crowdsourcing

1. INTRODUCTION

The mass-adoption of social networking sites and the wide integration of sharing widgets on popular websites have paved the way for quantitative research efforts tackling the relations between content and virality. Very recently, a different kind of widgets is working its way through the online space: designed to take the *affective* pulse of the websites visitors, such small interfaces allow

people to explicitly tag their emotional states as they browse the web.

Although the adoption of the latter is not yet prominent, we found two very popular news outlets, one in English and one in Italian, embedding similar interfaces for affective feedback in each article page. We thus became interested in investigating the relation of this newly available *affective* data with virality of news articles in a cross-lingual experimental setting. The two online websites which made this research possible are:

1. Rappler (rappler.com), an English-written “social news” portal, which makes extensive use of its distinctive *Mood Meter* feature;
2. the online version of the Italian newspaper Corriere della Sera (corriere.it), one of the most popular daily newspapers in Italy.

In our previous work [30], we have shown the impressive value of such reader-provided affective data: we used Rappler data to automatically build an emotion lexicon and a system for automatic affective analysis of texts. In the evaluation results we outperformed existing systems even using a naïve approach, thanks to the quality and coverage of the human annotation of evoked emotions. In this work, we turn to analyze such data *along* with indices of virality in order to derive insights on the relations between the emotions evoked by textual content and its diffusion and engagement.

Early works on emotions and virality have already been published [4], paving the way for this novel line of research. Still, a few limitations of [4], which this paper attempts to overcome, can be identified in (i), the use of a small sample of articles (~7,000) with only a subset manually annotated by three annotators (~2,500 documents); (ii), the authors only consider one virality index; and (iii), they only hint to a possible explanatory role of deep constituents of emotions.

In this paper we leverage a large corpus of news articles (ten times bigger than [4]) with a massive crowd-sourced voluntary annotation of the emotion each article evokes in readers (more than 1.5 millions votes), so to:

- compare several viral phenomena at once and understand if they are consistently affected by emotions and if they are affected in the same way;
- compare emotion and virality in a cross-lingual experimental setting;
- investigate the effect of deep constituents of emotions in viral phenomena.

Our findings show that, while the relations between emotions and virality seem to vary across cultures, their deeper constituents

(i.e. the Valence, Arousal, and Dominance components, as defined in the VAD circumplex model of affect we adopt) show consistency among all indices of virality we account for, providing a generalized model for their interplay. More specifically, viral indices are coherently affected by particular VAD configurations, and these configurations point to a clear connection with general phenomena underlying different viral facets. These results are relevant not only for social science researchers interested in understanding the factors behind virality phenomena, but also for marketing and industry people as they can be very valuable in contexts such as content marketing and native advertising [13].

This paper is structured as follows: the following section provides the reader with a brief review of recent research efforts on virality, social media and emotion studies; in Section 3 we describe the data collection procedure followed; Sections 4 and 5 report our analyses and findings, while in Section 6 we provide a comparison with [4]. Finally, we take stock of our work in Section 7.

2. RELATED WORK

In this section we provide a short review of research efforts focused on (i) understanding content virality on Social Media and (ii) the role of emotions in Social Media.

Content Virality. Several researchers have studied information flow, community building and similar processes using Social Networking sites as a reference [1, 18, 19, 23]. However, the great majority focused on network-related features without taking into account the actual content spreading within the network [22]. A hybrid approach focusing on both product characteristics and network related features is presented in [3]: the authors study the effect of passive-broadcast and active-personalized notifications embedded in an application aimed at fostering word of mouth.

Recently, the relation between content characteristics and virality has begun to be investigated, especially with regard to textual content. In [17], for instance, features derived from sentiment analysis of comments are used to predict the popularity of stories. The relevant work in [7] 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 [29], where the authors investigate which modification dynamics make a meme spread from one person to another (as compared with movie quotes which spread remaining exactly the same). Louis and Nenkova [24] focused on influential scientific articles in newspapers, considering characteristics such as readability, description vividness, use of unusual words and affective content, comparing high quality articles (NYT articles appearing in “The Best American Science Writing” anthology) against typical NYT articles.

Moreover, the work presented in [5] investigates how differences in textual description affect the spread of content-controlled videos. In [21], the authors focus on the act of resubmissions (i.e., content that is submitted multiple times with multiple titles to multiple different online communities) to understand the extent to which each factor influences the success of a content. In [32] it is investigated how content spreads in an online community by pinpointing the effect of wording in terms of content informativeness, generality and affect. Finally, Althoff et al. [2] developed a model that can predict the success of requests for a free pizza gifted from the Reddit community, using high level textual features such as politeness, reciprocity, narrative and gratitude.

More recently, some works have tried to investigate how different textual contents give rise to different reactions in the audience: the work presented in [12] correlates several viral phenomena with the wording of a post, while in [10] it is shown that specific content

features variations (like the readability level of an abstract) differentiate among virality level of downloads, bookmarking, and citations. Similarly, Shuai et al. [28] studied scientific articles in terms of downloads, Twitter mentions, and early citations in the scholarly records, trying to understand how these virality indices correlate among them. Finally, also in the realm of visual content it has been shown that different image characteristics can give rise to different viral phenomena [11]. Following this line of research, we study the effects of emotions considering several audience reactions to find out whether there are peculiar emotional characteristics of a news article that give rise to different viral reactions.

Emotions in Social Media. Kramer et al. [20] have shown, via a massive experiment on Facebook, that emotional states can be transferred to others via emotional contagion, leading people to experience the same emotions without their awareness. The experiment included reducing the amount of emotional content in the News Feed of a user: when positive expressions were reduced, people produced fewer positive posts and more negative posts; when negative expressions were reduced, the opposite pattern occurred. While in [20] the authors manipulated News Feed content, in our study we simply analyze existing and publicly available content voluntarily annotated by readers.

In [16], the authors focused on the role of emotions in the context of word-of-mouth marketing, using a dataset of Google+ posts: their analyses show consistency with [4], with increase in ANGER linked to higher likelihood of reshares, while the opposite trend holds for SADNESS. Fan et al. [9] looked at diffusion patterns on the very popular chinese platform Weibo, which shares many features with Twitter, and again found a similar result: angry posts appear to spread at a significantly faster rate, while sad posts do so at a significant lower rate.

Furthermore, the work presented in [14] comes closer to the focus of the present paper by hypothesizing that negative news content is more likely to be retweeted, while for non-news tweets positive sentiments support virality. To test this hypothesis the authors analyze three corpora and give evidence that negative sentiment enhances virality in the news segment, but not in the non-news segment. Their conclusion is that the relation between affect and virality is more complex than expected based on the findings of [4].

In [26], popular and influential users on twitter are linguistically analyzed according to their tweets. The initial hypothesis is that a Twitter account cannot be simply traced back to the graph properties of the network within which it is embedded, but also depends on the personality and emotions of the human being behind it. The reported findings suggest that popular users tend to use positive emotions while influential users lean towards negative ones.

Finally, the work presented in [4] 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. We consider this work as a starting point for our research and for comparison of results: we will discuss it throughout the paper.

Human, crowd-sourced, voluntary, large scale annotations.

These distinguishing characteristics of the affective data employed in our analyses mark the significant difference between this article and previous works. While previously mentioned research efforts have resorted to computational linguistics classifiers to automatically annotate textual content, the data we crawled (from publicly available websites) allows us to leverage affective annotations voluntarily provided by the news article readers.

3. DATASET COLLECTION

The “social-news” website `rappler.com` embeds a small interface, called *Mood Meter*, in every article it publishes. Such interface allows the readers to express with a simple click their emotional reaction to the story they are reading. The percentages of votes obtained by each emotion are also visualized by the interface, which is depicted in Figure 1. Similarly, the online version of a very popular Italian newspaper (*Corriere della Sera*) has recently adopted a similar approach, based on *emoticons*, to sense the emotional states of its readers, as shown in Figure 2¹.



Figure 1: Rappler’s Mood Meter.

DOPO AVER LETTO QUESTO ARTICOLO MI SENTO...



Figure 2: The emoticon-based interface on `corriere.it` articles. The sentence translates to “After reading this article, I feel..”

Following previous work presented in [30], we harvested a total of 53,226 news articles from `rappler.com`, and 12,437 articles from `corriere.it`. Our crawler was programmed to retrieve articles at least a month old, in order to allow the indicators to settle and thus not to penalize the most recent articles. The articles span over roughly one year for Rappler, and nine months for Corriere della Sera. For the scope of this paper, the main difference between the two mechanisms lies in the fact that, while it is possible

¹As we are forced to adopt a crawling approach, we cannot have any control on the layout of the widgets shown above. Thus, we cannot rule out the possibility of biases arising from the (fixed) order used to present readers with affective labels and emoticons. Still, the results reported in [30] – obtained thanks to data crawled with the same strategy, indicate that such bias, if present, is negligible.

to extract absolute votes from `corriere.it`, along with percentage values, `rappler.com` only exposes the latter. In particular, for the bulk of articles crawled from `corriere.it`, our data includes a total of 320,697 votes for the five emotional dimensions available.

Although we have no means to verify the actual absolute number of votes collected by the Mood Meter, we can provide a very conservative estimate for it: by computing the lower common denominator over the percentages of affective votes obtained by a Rappler article, we can derive the minimum number of votes needed to obtain them, compounding to a total of 1,145,543 votes over the entire Rappler dataset. For comparison, consider that the same conservative estimate for the `corriere.it` data amounts to 210,113 – less than two thirds of the actual value depicted above.

Thus, the datasets used in this work comprise more than 65,000 news articles and more than 1.5 million annotations.

Since the sets of emotions accounted for in the two websites differ both in size (`rappler.com` allows to tag eight affective dimensions, whereas five are available on `corriere.it`) and, although slightly, in semantics (being in two different languages), we need to proceed to map the subset available from the latter to the former, as reported in Table 1.

<code>corriere.it</code> label	maps to
TRISTE	SAD
DIVERTITO	AMUSED
SODDISFATTO	HAPPY*
PREOCCUPATO	AFRAID*
INDIGNATO	ANNOYED*

Table 1: Mapping of original emotion labels from `corriere.it` to those present in `rappler.com`. * denotes appropriate mapping, although the English label is not the primary translation.

In Table 2 we report the mean percentage of votes for each emotional dimension on the two corpora: it can be seen that HAPPINESS has the highest percentage of votes by a large margin in comparison to the other dimensions in `rappler.com`, while it is very closely followed by INDIGNATO/ANNOYED in the `corriere.it` data.

	Rappler _μ	Corriere _μ
AFRAID	.05	.09
DONT_CARE	.05	-
AMUSED	.11	.11
HAPPY	.31	.29
ANGRY	.11	-
INSPIRED	.11	-
ANNOYED	.06	.25
SAD	.12	.10

Table 2: Mean vote percentages obtained by emotions in Rappler and Corriere.

The statistics on `rappler.com` emotional data confirm the trend noted in [30] about HAPPINESS predominance, for which several explanations may be hypothesized: from cultural characteristics, to a bias in the dataset itself – as it might contain mainly ‘positive’ news, through psychological phenomena leading people to express more positive moods on social networks [8, 26, 33].

Studies on other English datasets, e.g. on LiveJournal posts [31], have in the past noted predominance of the happy mood. Con-

Rappler Comments				Rappler Tweets				Rappler G+			
	Estimate	Std. Err	Sigf.		Estimate	Std. Err	Sigf.		Estimate	Std. Err	Sigf.
ANNOYED	.61	.03	***	INSPIRED	.52	.02	***	INSPIRED	.55	.02	***
ANGRY	.54	.02	***	ANGRY	.32	.02	***	ANGRY	.25	.02	***
INSPIRED	.23	.02	***	ANNOYED	.32	.03	***	HAPPY	.24	.02	***
DONT_CARE	.18	.04	***	HAPPY	.31	.02	***	AMUSED	.21	.02	***
AMUSED	.13	.02	***	DONT_CARE	.30	.04	***	ANNOYED	.20	.03	***
HAPPY	.10	.02	***	SAD	.24	.02	***	SAD	.16	.02	***
SAD	.07	.02	***	AMUSED	.21	.02	***	AFRAID	.14	.03	***
AFRAID	-.03	.03	†	AFRAID	.19	.03	***	DONT_CARE	.13	.04	***

Table 3: Emotions impact on rappler.com articles viral indices. ***: $p < .001$; **: $p < .01$; *: $p < .05$; †: not significant – this legend applies to all tables in this paper.

Corriere Comments				Corriere Tweets				Corriere G+			
	Estimate	Std. Err	Sigf.		Estimate	Std. Err	Sigf.		Estimate	Std. Err	Sigf.
ANNOYED	1.11	.04	***	ANNOYED	.33	.03	***	SAD	.57	.05	***
HAPPY	.40	.04	***	SAD	.20	.05	***	ANNOYED	.39	.03	***
AFRAID	.19	.06	**	HAPPY	.14	.03	***	AMUSED	.34	.05	***
AMUSED	.13	.06	*	AFRAID	.06	.05	†	AFRAID	.33	.05	***
SAD	.12	.05	*	AMUSED	.02	.05	†	HAPPY	.27	.03	***

Table 4: Emotions impact on corriere.it articles viral indices.

versely, no previous studies have dealt at this scale with Italian language, and it will be worth investigating in future works what factors may influence the trend shown for emotional dimensions in the *corriere.it* data.

Turning to the statistics of the various viral indices available for the two datasets, we provide a summary in Table 5.

Viral Index	Rappler $_{\mu}$	Corriere $_{\mu}$
COMMENTS	4.11	83.22
THREADS	2.81	-
G+ SHARES	.91	3.79
TWITTER SHARES	32.33	40.65
FACEBOOK SHARES	-	502.93

Table 5: Mean figures for virality indices.

From now on, we will only consider the intersection of the two sets of indices.

3.1 Narrow- and Broad-casting

In the literature, *broadcasting* refers to the act of communicating or transmitting a content to numerous recipients simultaneously, over a communication network; on the contrary, *narrowcasting* has traditionally been understood as the dissemination of information to a narrow audience, rather than to the broader public at large. With the advent of new media, this definition has been updated (see, for instance, [15]). In fact, from the perspective of readers of online content, they can decide whether and how to share such content by either broadcasting it to their audience or to address only a small portion of it (narrowcasting). In [4], for example, the act of forwarding a news article by email is treated as a form of narrowcasting, since in such case subjects are addressing a selected section of their audience/contacts.

Following the above definition, we will consider article comments as a form of *narrowcasting*, as the readers who upload a comment on the article page are contributing to a discussion which happens at most between the readers of that article (and most probably, in fact, to a small subset of it). On the other hand, we consider the act of sharing an article to social networking sites as a form of *broadcasting*.

The rationale behind this distinction of viral facets in narrowcasting and broadcasting lies in previous research efforts showing

that virality can be assessed through different metrics which represent distinct phenomena: it is not only the magnitude of spreading (virality) that depends on content but also the users reactions (comments, shares, tweets, etc.) [10, 11, 28].

In the following sections, all virality indices have been standardized in order to make them comparable ($\mu = 0, \sigma = 1$) across the two datasets. Emotion scores range between 0 and 1.

4. EMOTION ANALYSIS

To analyze the relationship between an article emotional characteristics and the corresponding impact on virality indices we use simple linear models. In this section we will discuss the importance of each emotion for the various virality indices, while the overall explanatory power of the models will be discussed in Section 6.

Results on the *rappler.com* dataset, shown in Table 3, are in accordance with the findings of Berger et al. [4]: high influence of INSPIRING (awe) and of negative-valence and high-arousal emotions such as ANGER, jointly with a low influence of SADNESS. This similarity stands out for the specific case of *narrowcasting*, and it is maintained for *broadcasting*, albeit to a lower extent.

As mentioned in Section 2, other previous research works [9, 16], focusing on diverse scenarios and using heterogeneous datasets, have reported results consistent with these findings.

Nonetheless, turning to the Italian resource used in our analyses, we notice that results obtained from the *corriere.it* data, summarized in Table 4, are partially in line with [4] for what concerns *narrowcasting*, while drastically diverge on *broadcasting*: in the latter case, SADNESS (a low-arousal emotional dimension) is found to be most relevant for virality, disproving the hypothesis that arousal alone can be used to explain virality phenomena. It should be noted that the lack of an INSPIRED dimension for *corriere.it* cannot account for such discrepancy on *broadcasting*, since it seems very unlikely that its potential votes would have been collected by SADNESS.

It can be hypothesized, as a consequence, that strong cultural differences emerge when it comes to emotions (more precisely, to explicitly tagging one’s own feeling on a website). What factors might underlie this phenomenon (e.g. historical period, editorial choices, deeper cultural sensibilities, etc.) represents a very fascinating research question, which is out of the scope of this work.

Rather than focusing on specific emotions, in the next section we attempt to provide a more general explanation.

corriere.it comments				rappler.com comments			
	Estimate	Std. Err	Sigf.		Estimate	Std. Err	Sigf.
AROUSAL	.3741	.0372	***	AROUSAL	.1205	.0101	***
VALENCE	-.3155	.0486	***	VALENCE	-.1636	.0165	***
DOMINANCE	.1000	.0761	†	DOMINANCE	.0728	.0221	**
corriere.it tweets				rappler.com tweets			
	Estimate	Std. Err	Sigf.		Estimate	Std. Err	Sigf.
DOMINANCE	.2158	.0709	**	DOMINANCE	.2244	.0221	***
VALENCE	-.2176	.0453	***	VALENCE	-.1495	.0165	***
AROUSAL	.0400	.0348	†	AROUSAL	.0065	.0101	†
corriere.it g+ shares				rappler.com g+ shares			
	Estimate	Std. Err	Sigf.		Estimate	Std. Err	Sigf.
DOMINANCE	.4817	.0704	***	DOMINANCE	.2619	.0221	***
VALENCE	-.3781	.0450	***	VALENCE	-.1530	.0165	***
AROUSAL	-.0242	.0345	†	AROUSAL	-.0371	.0101	***

Table 7: Valence, Arousal and Dominance significant effects from simple Linear Models.

EMOTION	Valence	Arousal	Dominance
AFRAID	2.25	5.12	2.71
AMUSED	7.05	4.27	5.93
ANGRY	2.53	6.02	4.11
ANNOYED	2.80	5.29	4.08
DONT_CARE	3.53	4.27	3.62
HAPPY	8.47	6.05	7.21
INSPIRED	6.89	5.56	7.30
SAD	2.10	3.49	3.84

Table 6: Valence, Arousal and Dominance scores for emotion labels, as provided by [34].

5. VAD ANALYSIS

We now turn to investigate how basic constituents of emotions, such as Valence, Arousal, and Dominance (VAD), connect to virality. The widely adopted VAD circumplex model of affect [6, 27] maps emotions on a three-dimensional space, namely: Valence, denoting the degree of positive/negative affectivity (e.g. FEAR has high negative valence, while JOY has high positive valence); Arousal, ranging from calming to exciting (e.g. ANGER is denoted by high arousal while SADNESS by low arousal); and Dominance, going from “controlled” to “in control” (e.g. INSPIRED, highly in control, vs FEAR, overwhelming).

We thus proceed to map the emotions an article is found to evoke to the VAD circumplex model. To do so, we exploit the work of Warriner et al. [34], who provided a resource mapping roughly 14 thousands words to the VAD circumplex model, including words representing the emotional dimensions we consider in this work – see Table 6, which for us serve as a gold standard. Such scores are in a Likert scale, ranging from 1 (low/negative) to 9 (high/positive).

Then, in order to compute VAD scores for a given article *doc* we simply multiply the percentage of votes each emotion *e* is found to evoke in *doc* by the corresponding VAD score provided in Table 6, and then take the sum over the *n* emotional dimensions considered. The formula for Valence is provided below; the equations used for Arousal and Dominance are akin.

$$doc_v = \sum_{i=1}^n votes\%(e_i) \times valence(e_i) \quad (1)$$

As with virality indices, each VAD dimension has been standardized. Subsequently, we build linear models in order to assess whether and how VAD dimensions are connected to virality. The results reported in Table 7 show very interesting trends: VAD dimension estimates are found to be consistent among the two datasets (see estimate order and sign in Table 7), in both *broadcasting* (tweets/g+ shares) and *narrowcasting* (comments) scenarios. Comparing these results with those provided in the previous sections, we see that the cross-cultural divergences in the relations between emotional dimensions and virality indices disappear when accounting for their more profound VAD components.

This finding hints at a generalized, culturally independent phenomenon: readers of the articles in our datasets tend to choose communication forms of *narrowcasting* when the content is arousing but on which they feel less in control².

Conversely, they turn to *broadcasting* (e.g. share on social networks) when they feel more in control. This result is further supported by the fact that the less important dimensions (i.e. Dominance for *narrowcasting*, and Arousal for *broadcasting*) are most of the time not significant, or with a slight negative impact on the model, indicating that these two dimensions *switch roles* when transitioning from narrowcasting to broadcasting (and viceversa). This finding is important as it appears to be valid, in our analyses, among the two different languages (and cultural factors) our two datasets provide.

Finally, the role of Valence appears to be consistent among all indices of virality in both datasets, with negative valence contributing to higher virality. This result is in line with [14], who found that news-related content spreads more when imbued with negative Valence.

6. R² ANALYSIS

In this section we examine the explanatory power of our models and compare them with the results obtained by the models presented in Berger et al. [4] – in particular, with *model 2* (Positivity and Emotionality) and *model 3* (Positivity and Emotionality plus the emotions Awe, Anger, Anxiety, Sadness).

²Again, this result is in line with the intuition in [4], that arousal has a strong effect on *narrowcasting*. Nonetheless, it appears of paramount importance to take the role of Dominance (not considered in [4]) into account in order to understand *broadcasting* phenomena.

We compare our VAD models with *model 2*, which was automatically computed starting from a lexicon of relevant affective words (the LIWC lexicon described in [25]), and our emotion-based models with *model 3* in [4]. It should be noted that the extensive work presented in [4] includes other models accounting for variables (such as publication time, homepage position, author, among others) not available in our datasets. Nonetheless, it has been shown in [4] that, while augmenting the explanatory power, these variables do not influence weight and role of the emotions in the models.

In [4], the dependent variable is represented as binary (i.e., 1 for those articles that make it on the *most emailed list*, 0 for all the others). On the contrary, our viral indices are originally represented as continuous variables. In order to provide a fair comparison we thus proceeded to binarize the virality indices V of article n , according to the following scheme:

$$V_n = \begin{cases} 1, & \text{if } V_n > \text{mean}(V) + \text{sd}(V) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

In this way, we obtain a small sample of the *most viral* articles in our dataset, specular to the *most emailed list* in [4]. We also used logistic regression in accordance with [4]. For comparison, we report in table 8 the McFadden’s R^2 used by Berger et al. [4], along with the standard R^2 .

Viral Index	R^2	McFadden’s R^2
corriere.it emotion models		
Comments	.0813	.0922
G+	.0207	.0527
Tweets	.0084	.0604
corriere.it VAD models		
Comments	.0530	.0840
G+	.0195	.0446
Tweets	.0072	.0591
rappler.com emotion models		
Comments	.0191	.0801
G+	.0115	.0314
Tweets	.0112	.0300
rappler.com VAD models		
Comments	.0101	.0569
G+	.0088	.0288
Tweets	.0100	.0266
Berger et al. models		
Model 2	-	.0400
Model 3	-	.0700

Table 8: R^2 scores for the various linear models described and comparison with models presented in Berger et al. [4]

Hence, when projecting emotions to their basic VAD dimensions, the models experience only a slight drop in terms of explanatory power, while gaining the advantage of language-invariance.

Moreover, our regression models for *narrowcasting* have greater explanatory power than the best performing model in [4]. This fact can be partially explained by the quality of our data: our datasets are much bigger and the affective annotations of the news articles were massively crowdsourced to the readers.

Finally, emotions have significant effects on both *narrowcasting* and *broadcasting*, with a stronger impact on the former. Again, this finding is cross-cultural consistent.

7. CONCLUSIONS

In this article we have provided a comprehensive investigation on the relations between virality of news articles and the emotions they are found to evoke. By exploiting a high-coverage and bilingual corpus of 65k documents containing metrics of their spread on social networks as well as a massive affective annotation provided by readers (more than 1.5 millions votes), we presented a thorough analysis of the interplay between evoked emotions and viral facets.

We highlighted and discussed our findings in light of a cross-lingual approach.

Our results show differences in evoked emotions and corresponding viral effects across the bilingual data we crawled; amongst other findings, we note the remarkably discordant influence of the SADNESS affective dimension between the english- and italian- language datasets, in contrast with the hypothesis that arousal alone can be used to explain virality phenomena. While these findings seem to indicate effects of cultural differences on the relation existing between emotions and virality, such hypothesis could only be properly tested were more information on the annotators demographics available (e.g. geographic provenance).

Still, when accounting for the deeper constituents of emotions our analyses provided compelling evidence of a generalized explanatory model rooted in the Valence-Arousal-Dominance (VAD) circumplex. Viral facets seem to be coherently affected by particular VAD configurations (namely, the alternation between Dominance and Arousal), and these configurations indicate a clear connection with distinct phenomena underlying persuasive communication. In particular, high arousal is more connected to *narrowcasting* phenomena, while Dominance to *broadcasting* phenomena.

Extensions of this work will include deeper investigations of the relations between VAD dimensions and virality. For instance, while in the analyses reported herein we have used a resource [34] to map affective tags to their VAD constituents, we plan to increase the resolution by associating each article with a VAD representation derived from its textual content.

The results presented in this article can be very valuable in contexts such as content marketing and native advertising, and thus be relevant not only for social science researchers interested in understanding the factors behind virality phenomena, but also for marketing and industry people.

Acknowledgements

The work of M. Guerini has been supported by the Trento RISE PerTe project. The work of J. Staiano has been partially funded by the French government through the National Agency for Research (ANR) under the Sorbonne Universités Excellence Initiative program (Idex, reference ANR-11-IDEX-0004-02) in the framework of the *Investments for the future* program, and by the German Federal Ministry of Research and Education, within the EBITA project.

8. REFERENCES

- [1] J. Z. Aaditshwar Seth and R. Cohen. A multi-disciplinary approach for recommending weblog messages. In *The AAAI 2008 Workshop on Enhanced Messaging*, 2008.
- [2] T. Althoff, C. Danescu-Niculescu-Mizil, and D. Jurafsky. How to ask for a favor: A case study on the success of altruistic requests. *Proceedings of ICWSM*, 2014.
- [3] S. Aral and D. Walker. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science*, 57(9):1623–1639, 2011.
- [4] J. Berger and K. L. Milkman. What makes online content viral? *Journal of Marketing Research*, 49(2):192–205, 2012.

- [5] Y. Borghol, S. Ardon, N. Carlsson, D. Eager, and A. Mahanti. The untold story of the clones: Content-agnostic factors that impact youtube video popularity. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1186–1194. ACM, 2012.
- [6] M. M. Bradley and P. J. Lang. Measuring emotion: the self-assessment manikin and the semantic differential. *Journal of behavior therapy and experimental psychiatry*, 25(1):49–59, 1994.
- [7] 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.
- [8] M. De Choudhury, S. Counts, and M. Gamon. Not all moods are created equal! exploring human emotional states in social media. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM)*, 2012.
- [9] R. Fan, J. Zhao, Y. Chen, and K. Xu. Anger is more influential than joy: sentiment correlation in weibo. *Plos One*, 2014.
- [10] M. Guerini, A. Pepe, and B. Lepri. Do linguistic style and readability of scientific abstracts affect their virality? In *ICWSM*, 2012.
- [11] M. Guerini, J. Staiano, and D. Albanese. Exploring image virality in google plus. In *Social Computing (SocialCom), 2013 International Conference on*, pages 671–678. IEEE, 2013.
- [12] M. Guerini, C. Strapparava, and G. Özbal. Exploring text virality in social networks. In *Proceedings of ICWSM-11*, Barcelona, Spain, July 2011.
- [13] C. Hackley and R. A. Hackley. *Advertising and promotion*. Sage, 2014.
- [14] L. K. Hansen, A. Arvidsson, F. Å. Nielsen, E. Colleoni, and M. Etter. Good friends, bad news-affect and virality in twitter. In *Future information technology*, pages 34–43. Springer, 2011.
- [15] M. Hirst, J. Harrison, and P. Mazepa. *Communication and new media: From broadcast to narrowcast*. Oxford University Press, 2014.
- [16] R. Hochreiter and C. Waldhauser. The role of emotions in propagating brands in social networks. *arXiv preprint arXiv:1409.4617*, 2014.
- [17] S. Jamali. Comment mining, popularity prediction, and social network analysis. Master’s thesis, George Mason University, Fairfax, VA, 2009.
- [18] 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.
- [19] 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 *ICWSM*, 2009.
- [20] A. D. Kramer, J. E. Guillory, and J. T. Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, page 201320040, 2014.
- [21] H. Lakkaraju, J. J. McAuley, and J. Leskovec. What’s in a name? understanding the interplay between titles, content, and communities in social media. In *ICWSM*, 2013.
- [22] K. Lerman and A. Galstyan. Analysis of social voting patterns on digg. In *Proceedings of the first workshop on Online social networks, WOSP ’08*, pages 7–12, New York, NY, USA, 2008. ACM.
- [23] K. Lerman and R. Ghosh. Information contagion: an empirical study of the spread of news on digg and twitter social networks. In *Proceedings of 4th International Conference on Weblogs and Social Media (ICWSM-10)*, Mar 2010.
- [24] A. Louis and A. Nenkova. What makes writing great? first experiments on article quality prediction in the science journalism domain. *TACL*, 1:341–352, 2013.
- [25] J. Pennebaker and M. Francis. Linguistic inquiry and word count: LIWC, 2001. Erlbaum Publishers.
- [26] D. Quercia, J. Ellis, L. Capra, and J. Crowcroft. In the mood for being influential on twitter. *Proceedings of IEEE SocialCom’11*, 2011.
- [27] J. A. Russell. A circumplex model of affect. *Journal of personality and social psychology*, 39(6):1161, 1980.
- [28] X. Shuai, A. Pepe, and J. Bollen. How the scientific community reacts to newly submitted preprints: Article downloads, twitter mentions, and citations. *PloS one*, 7(11):e47523, 2012.
- [29] M. Simmons, L. A. Adamic, and E. Adar. Memes online: Extracted, subtracted, injected, and recollected. *ICWSM 2011*, 2011.
- [30] J. Staiano and M. Guerini. Depeche mood: a lexicon for emotion analysis from crowd annotated news. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014*, volume 2, pages 427–433. The Association for Computer Linguistics, 2014.
- [31] C. Strapparava and R. Mihalcea. Learning to identify emotions in text. In *Proceedings of the 2008 ACM symposium on Applied computing*, pages 1556–1560. ACM, 2008.
- [32] C. Tan, L. Lee, and B. Pang. The effect of wording on message propagation: Topic-and author-controlled natural experiments on twitter. *ACL*, 2014.
- [33] J. R. Vittengl and C. S. Holt. A time-series diary study of mood and social interaction. *Motivation and Emotion*, 22(3):255–275, 1998.
- [34] A. B. Warriner, V. Kuperman, and M. Brysbaert. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, 45(4):1191–1207, 2013.