Popularity Spikes Hurt the Future Implementations of Protomemes

Michele Coscia*

CID - Harvard University
79 JFK St, Cambridge MA, US 02138

August 24, 2017

Key Insights:

- We test the hypothesis that protomemes in social media are less likely to be popular when they are too similar to other protomemes on a large and variegated dataset.
- We calculate the canonicity of each specific usage of a protomeme, i.e. how typical or common this usage is.
- We show that canonicity has a non-linear relationship with protomeme popularity, increasing the similarity between protomemes and genes, in that not all mutations are beneficial for protomemes.

A meme is a concept introduced in [12] as an equivalent in culture of what a gene is in biology. A meme is a cultural unit–a joke, a musical tune, a behavior–that can replicate in people's mind, spreading from person to person. During the replication process, memes can mutate. Memes compete with each other for attention, because people's consciousness has finite capacity. Meme viral spreading causes behavioral change, for the better–one of the many examples, the "ALS Bucket Challenge" meme, caused a cascade of humanitarian donations[1]–and for the worse–as researchers have proven that obesity [7] and smoking [8] are socially transmittable diseases. A better theory of meme spreading could help preventing the outbreak of bad behaviors and favoring positive ones.

Studying memes in general is hard, because it is difficult to objectively detect and measure them. The inception of the Web made easier to focus on a subtype of meme: the one that is shared via social media. Researchers have been focusing on the effect of timing, social networks and limited user attention

*To whom correspondence should be addressed. E-mail: michele_coscia@hks.harvard.edu, Work Telephone number: +1 (857) 241 7563, Work Fax number: +1 (617) 496 8753 (include a header with the name "Michele Coscia").

[1] `http://www.alsa.org/fight-als/ice-bucket-challenge.html`, accessed 03/26/2015.

rather than of meme content [15, 23, 6]. Being timely and shared by actors in key positions of the network explains a large portion of a meme's success. While these factors can explain the meme ecosystem at large, both are exogenous to the meme itself. Even if endogenous meme characteristics have a lower predictive power over meme popularity, it is still important to understand their relationship with virality, because this understanding can be applied case by case to specific memes, rather than having only a description of the general mechanics of the overall system. For example, [11] shows that success eschews similarity: successful memes lie in the periphery of the meme similarity space. The more a meme gets imitated, the less the original meme (and all its imitations) will be successful in the future.

In this paper we focus on protomemes, i.e. on all catchphrases used frequently and regularly on social media which may or may not end up being adopted as memes. We aim to extend [11] by: expanding the scope of the dissimilarity-driven success theory; providing new evidences of the effects of the theory in a dynamic environment; and introducing canonicity to evaluate protomeme's content.

First, we test the dissimilarity-driven success theory on a larger data source and shifting the attention from memes to protomemes. The original work focused on a small site and on a very specific subtype of internet meme. Here we focus on Reddit and Hacker News, two of the main online social media, and we consider any kind of protomeme, in the form of frequent and regularly used n-grams, that can be shared on the platforms, rather than limiting to the image-macro meme subtype that were previously studied.

Second, we test the effect of popularity spikes on the future popularity of a protomeme, following [21] as spikes were not studied in [11]. When a protomeme suddenly becomes very popular–that is, when it places highly in a ranking of user appreciation–, many people will use it in their posts. The increased number of posts that include the protomeme will make it more similar to the average post of the day. As a consequence, according to [17], its expected popularity will drop. The new posts that include the popular protomeme are then poorly ranked, because they are stealing each other's chance of success.

Lastly, we introduce a measure of "canonicity", which captures the amount of change introduced by a post compared to previous posts with the same protomeme. We show that the more different a post is from the canonical usage of a popular protomeme, the higher its odds to go viral are. However, the effect of canonicity is not linear. High canonicity lowers the overall success of viral posts, but at the same time it helps non-viral posts to be appreciated. In other words, we correct [11]: there is a non-linear relationship between canonicity and success. This is why other works did not find content to be a powerful predictor of success: while it is true that success eschews similarity, it is not true that with dissimilarity comes success. One explanation might be that protomemes follow the same dynamics as genes: not all mutations are beneficial. Some are irrelevant and some harmful. In future work we will test this theory.

Related Approaches

This paper relates to the rich literature studying memes as they spread through the Web. In previous works, we first observed traditional meme dynamics such as competition and collaboration in the context of the Web [10], and then we provided evidence for the theory that similarity with existing content penalizes a meme's odds of success [11]. The origin of this line of research is due to the broad distribution of meme popularity observed multiple times in different contexts [15, 23] and consistent with the dynamics of fads [2].

The study of the dynamics of memes is a popular field. The papers attempting to model and predict meme success are many and it is impossible to mention all of them. The most successful track of research focuses on the relationship between memes and the social media through which they spread. For example, we can predict how large the cascade spread of a meme will be by observing its temporal and structural features: initial breadth, rather than depth, of a cascade is a better indicator of larger cascades [6]. The community structure of a social network, i.e. its tendency of forming densely connected groups, influences the dynamics of news and meme spread [19, 24].

The general conclusion of these works is that content is a secondary explanatory factor of meme virality. However, none of these works effectively rules out meme characteristics as partial explanation of their success. Further, they are interested in using networks to explain the shape of popularity distributions, not in what makes individual memes more or less fit to go viral, which is the focus of [11] and of this work. Different hashtags in Twitter have different degrees of persistence, showing evidences that meme spread is a complex form of contagion, rather than a simple epidemic one [20]. Outside computer science, researchers have found links between a meme's content and its virality. Emotionally positive content is more viral than negative content, although valence alone is just one of many factors [3]. All these studies have to face the challenges of complex spread dynamics: sometimes content that eventually becomes popular is overlooked when it first appears [14].

Our perspective here is focused on the role of mutations and innovations in meme usage, which have a relationship with success spikes [21]. An example showing the dynamics of meme mutation is [1]. Some results in literature support the role of novelty in content diffusion [5], while other papers question this result [16], though in neither case is "novelty" strictly defined. Our study builds on the evidence of a negative relationship between title similarity and success on Reddit [17]. The difference is that [17] considers how similar an instance of a meme is to the past submission history of the thematic community where it was shared. By introducing the measure of canonicity, we focus instead on the similarity of the meme instantiation with the meme's own submission history, regardless of the communities in which it was previously observed. Canonicity is closely related to the Newsjunkie framework [13], which is based on an information-theoretic background. However, Newsjunkie was developed with a different application, namely ranking novelty of full text news articles, which are significantly longer than Reddit and Hacker News post titles. For this reason, we believe Newsjunkie is not applicable to this research scenario.

Protomemes

We work with data collected by [25] from Reddit, a social bookmarking website where users can post interesting content. Every post can be upvoted (downvoted) if the user likes (dislikes) it. The upvote/downvote ratio is used to give more visibility to high quality content. In addition, there is a time discount: no matter how many upvotes a post gathers, it cannot be highly visible forever. The most popular (highest upvote/downvote ratio) posts appear on the "Front Page", which gives them a further boost in visibility. By default, the front page hosts 25 posts. Each entry in our data set consists of a post, its title and its number of upvotes/downvotes, which is combined in a post score by Reddit's sorting algorithm. Note that we can only observe the final score of a post, not its full upvote timeline. This might introduce bias when we want to establish if the post hit the front page or not. Our assumption is that the final post score is highly correlated with the post score in its first day of life. We base this assumption on the fact that the vast majority of upvotes come within 24 hours from the post submission.

Note that the terms "score" and "popularity" are not interchangeable: they refer to related but different concepts. Score is the one-off measurement of a single instance of a protomeme in a day. Popularity is the overall success of all instances of all memes during a wider period of time.

All 22,329,506 posts made on Reddit between April 5th, 2012 and April 26th, 2013 are part of our dataset. To cross-test our results, we also use a similar dataset from Hacker News. Hacker News works with the same dynamics as Reddit, although it is focused on a more specialized technical audience and it has a much smaller user base. The Hacker News dataset contains 1,194,436 posts from January 7th, 2010 to May 29th, 2014.

In this paper, we study protomemes. Our definition of protomeme is "a catchphrase that has the potential of going viral". Note that there are more possible types of memes (pictures, videos, etc), but given the nature on the data we limit ourselves to study catchphrases. Catchphrases are used as meme proxy also in Memetracker [18] and Nifty [22]. The catchphrases are extracted using information coming exclusively from the post title.

We operationalize our definition by borrowing the bag-of-word methodology from the text mining literature. A protomeme is an n-gram, with $n \geq 2$. Following the standard procedure, each word forming the n-gram is stemmed and stopwords are removed. To be classified as a protomeme, an n-gram must have been used frequently and constantly over the observation period. We use the frequent itemset mining algorithm Eclat [4] to extract the frequent n-grams and we discard an n-gram if it has not been used for a certain amount of distinct days. We also discard all n-grams that are proper subsets of another n-gram.

We perform the analysis using different thresholds to ensure independence between parameter choice and results. From the most to the least restrictive threshold choices, we obtain 2,731 to 5,585 protomemes on Reddit, and 817 to 2,538 protomemes on Hacker News. The preprocessed data we used is available for result replication[2].

Results

The Popularity Curse

Common sense tells us that popular ideas get imitated. A protomeme that is used in a very popular post today will be used in many posts tomorrow. Such intuition about the demand-supply relation in the Web is corroborated by some studies in the literature [9]. However, dissimilarity-driven success theory would predict that flooding a system with imitations of a protomeme will cause these imitations to be less popular. At first sight, this prediction is supported by two observations: the average score of the posts containing a protomeme is lower than expected the day after it experiences a popularity spike (Figure 1 (a) and (c)), and the number of posts containing that protomeme increases (Figure 1 (b) and (d)).

These observations support our theory, but they do not prove it. First, the total score awarded and the average score per post are not constant over time (see Supplementary Material). The lower score might be just a relative change: if there are fewer upvotes awarded in that particular day, a lower absolute number could still represent an increase in upvote share of the day. Second, each protomeme is characterized by its own expected popularity: scores are clustered around each protomeme's average. Third, the recent history of a protomeme might explain this variation. If a protomeme was getting declining scores, we might expect it to get even lower scores after a random popularity spike. Finally, there is a fat tail in score distributions, thus the average is not meaningful.

To give more solid evidence for the theory, we test the median popularity of a protomeme in a day using the following mixed model (which we call "MED Model"):

$$E_{w_{m,i}}(p_{m,i}) = \alpha + \beta_1 FP^l_{m,i-1} + \beta_2 E(p_{m,i-1}) + u_m + \epsilon_{m,i}.$$

An observation ($E_{w_{m,i}}(p_{m,i})$) is the median popularity ($p_{m,i}$) of the posts containing protomeme m on day i, i.e. the score of the post. Since post scores are count data, and distributed over a skewed distribution, we use a Poisson mixed model. Each observation has weight $w_{m,i}$, which is the number of posts containing m on day i, in order to weigh less heavily the (presumably noisy) information on protomemes m that were not used very much on day i.

$FP^l_{m,i-1}$ registers if protomeme m was on the front page in day $i-1$. The parameter l is the number of posts hitting the front page each day. The default front page in Reddit contains 25 posts (30 for Hacker News). So, every day at least 25 (30) posts hit the front page. However, users can increase the length

[2] http://www.michelecoscia.com/?page_id=870

of their front page. Moreover, as the day passes, front page posts get replaced by other posts. As a result, there is no way to know how many posts hit any particular user's front page in a given day (see Supplementary Material). For this reason, we run the regression multiple times for different l values, to test the robustness of our results for different front page sizes.

$E(p_{m,i-1})$ denotes the protomeme's popularity on the day before i, controlling for existing trends; u_m is a random effect of protomeme, which we use to control for the fact that different observations can refer to the same protomeme; and $\epsilon_{m,i}$ is the error term.

Figure 2 depicts our estimates of β_1 for different values of l. The effect of being in the front page is negative: the protomeme is expected to have a lower median score than on a business-as-usual day. This confirms the theory of dissimilarity-driven success theory. Higher l values decrease the estimated effect, because we are including lower ranked posts that might not actually have hit the front page (all p-values are significant with $p < 0.001$). Focusing on Reddit (Figure 2 (a)) β_1 values fall within the $]-0.1 : -0.02[$ interval. A value of -0.1 implies a score reduction factor equal to $e^{-0.1}$, which is close to 10%. This means that ranking in the top 25 posts in a day reduces next day's median score of a protomeme by almost 10%. The effect in Hacker News seems to be stronger.

These results have been obtained with fixed frequency thresholds. In Tables 1 and 2 we test the robustness of the results with different threshold choices. In the tables, we fix $l = 25$ for Reddit and $l = 30$ for Hacker News. For all threshold choices in Reddit and for most threshold choices in Hacker News, the results are negative and significant, consistently with our main result.

One could object to these results using the regression towards the mean argument: once the very visible front page protomeme instance is copied many times, the central score tendency will regress towards the middle. But the regression already corrects for this using the protomeme random effects. If the argument would be true, β_1 would equal to 0. In fact, the average β_1 is 0 calculated over 50 null models, where protomeme scores are generated randomly preserving average score and standard deviation of each protomeme (Figures 2 (a) and (b)), disproving the regression towards the mean argument.

To sum up: after hitting the front page, a protomeme will be used more frequently (15% more in Reddit, 8% more in Hacker News, Figures 1 (b) and (d)). β_1 values in the model suggest that that posts containing this protomeme are going to have a lower score (-10% in Reddit, -23% in Hacker News), confirming the expectations. The effect is significant and it is independent from the recent overall history of the protomeme, the changes in average post score and the front page size.

So, if we just proved that hitting the front page is bad for subsequent protomeme posts, why does common sense tell us the opposite? We propose that a protomeme appearing on the front page two days in a row is very noticeable and we just do not realize that, on average, the protomeme is doing poorly. We run the same regression, changing the target variable to the maximum score (MAX Model), instead of the median. In this model, for Reddit the sign of β_1 is the opposite of the β_1 sign for the MED Model (see Supplementary Material). If protomeme m hits the front page on day $i-1$, the top scoring post containing protomeme m on day i improves. This does not happen for Hacker News, and our hypothesis is that Hacker News is more resilient to fads, being used mostly for professional purposes, rather than humor as in Reddit.

To conclude, hitting the front page is associated with a larger number of subsequent posts containing protomeme m, which on itself is associated with lower expected popularity for those same posts. However, in some scenarios, the best ranked posts containing protomeme m can still hit the front page more easily than usual. We now turn our attention to this subset of special posts, explaining why they are able to overcome the popularity curse predicted by the theory.

The Canon Effect

Following the explanation that success is associated with dissimilarity, we hypothesize that a post including a widely used protomeme m the day after m hit the front page can still be successful if it is dissimilar from all other posts using m. In this way, the post is able to attract most of the attention that people are directing toward protomeme m.

To test this claim we first need a measure for the uniqueness of a post. In [11] a meme similarity measure is proposed, but it cannot be used here because the measure calculates meme-meme similarity, while here we are focused on post-post similarity within the same protomeme; and the [11] measure applies only to a subtype of the memes shared on Reddit. We cannot use the measure developed in [17] either, because it is a similarity of a post with the sub-community it is shared in, ignoring the meme it implements. In this section we focus on Reddit, and we expect a null result for Hacker News, as we showed it to be less prone to the fad effect.

We introduce the concept of canonicity of a post, measuring how much a post containing a protomeme m differs from the usual usage of m. A post is said to be canonic if it uses m as expected, without introducing elements not strongly associated with m itself. A post $\mathcal{M}$ is a bag-of-words. Each word μ in the bag-of-words co-occurs with protomeme m with a given probability $\pi_{m,\mu}$. If m appears in 100 posts, and in 30 of them the post title also includes μ, then $\pi_{m,\mu} = 0.3$. The canonicity of $\mathcal{M}$ is calculated as follows:

$$\Gamma(\mathcal{M}, m) = \frac{\sum\limits_{\forall \mu \in \mathcal{M}} \pi_{m,\mu}}{|\mathcal{M}|}.$$

Note that:

- Some posts contain no other word than the words composing the protomeme m itself. For this reason, we always include the m words in $\mathcal{M}$. Otherwise such posts will have $\Gamma(\mathcal{M}, m) = 0/0$, which is unacceptable.

- Posts including only a protomeme's words must have $\Gamma(\mathcal{M}, m) = 1$, because they are using m in its purest form. Since the protomeme's μs always appear in posts containing the protomeme itself, their $\pi_{m,\mu}$ always equals 1.

- A low canonicity is obtained when there are many words in the post, and they have low $\pi_{m,\mu}$. If a post includes only one unusual word, its canonicity is still high, because it is still mostly composed by the protomeme itself

How canonicity is distributed in Reddit is reported in the Supplementary Material. To test the connection between canonicity and popularity for posts using protomemes that hit the front page in the previous day, we create a rank binary variable $\phi_{\mathcal{M}}$. The variable records whether the post was among the 5% best scoring posts of the day or not. This is the target variable of the following logistic regression:

$$\phi_{\mathcal{M}_l} = \alpha + \beta\Gamma(\mathcal{M}_l, m) + u_m + \epsilon_m.$$

This model is the ϕ Model. Given a post $\mathcal{M}$ containing protomeme m, the ϕ Model estimates its probability of experiencing a popularity spike on day i, after m hit the front page on day $i - 1$. Note that the set of posts that we include in the model is still dependent on the l parameter. For different l values the set of posts included is different, because for increasing front page sizes we will have more protomemes hitting the front page, and therefore more posts in the day after will be considered in the model.

In Figure 3 we report ϕ Model's βs for increasing l. For Reddit (Figure 3(a)), β never takes values above -0.7, suggesting a noticeable effect: high canonicity halves the odds of being a highly scoring post (for a deeper discussion, see Supplementary Material). As we include more and more posts in the front page, the canonicity effect gets weaker and weaker. This is expected, as we are introducing into the regression posts for which we cannot be sure they actually hit the front page. All β values depicted in Figure 3(a) are significant ($p < 0.0001$). We expect a null result in Hacker News, given the result of the MAX Model in the previous section. Indeed, Figure 3(b) shows that the effect of canonicity in Hacker News is zero, as no p-value reported for any l is below 0.01.

We also run two Poisson mixed models with the same form of the ϕ Model, with the only differences being the dependent variable (in this case the post score) and the data included in them. In the Zero Model, we consider only the posts for which $\phi_{\mathcal{M}_l} = 0$, in the One Model we focus on the posts for which $\phi_{\mathcal{M}_l} = 1$. In practice, the ϕ Model tells us what is the effect of canonicity on the odds of experiencing two popularity spikes in a row, while the One and Zero models reveal the score effect of canonicity on the posts that did and did not experience two popularity spikes in a row.

In the One Model, β has a negative sign (Figure 4(a)) – all βs are significant with $p < 0.0001$. If the ϕ Model told us that canonicity lowers the odds of experiencing two popularity spikes in a row, the One Model tells us that if a post can nevertheless overcome those odds, it is additionally penalized with a worse score. In the Zero Model (Figure 4(b)), β is positive and significant. For the unsuccessful posts in the Zero Model, canonicity has a positive effect. For robustness, we also ran a negative binomial model, which resulted in similar estimates as the Poisson model (see Supplementary Materials).

The discordance of β signs in the Zero and One models can be interpreted as a similarity between protomeme and the gene dynamics. Most mutations are harmful or irrelevant, and they lower an organism's fitness. In protomemes, if the change is not judged "suitable" for the protomeme by the community of users, it will be selected against. This is one of the many possible speculations and it has to be properly tested to be considered as a suitable one. We leave this test as future work. However, this result could explain why meme content has not been judged a promising predictor for meme popularity [6]. Since changing the meme content can go both ways, increasing and decreasing meme fitness, the effects might cancel out.

Conclusion

In this paper we tested some of the predictions of the theory that meme success eschews similarity, because similar memes interfere with each other and get less attention [11]. We tested the theory on Reddit and Hacker News, two social bookmarking websites. Successful posts can hit the highly visible front page, and they are then copied in large quantities by people who want to use them to get to the front page again. As a consequence, the expected popularity of these posts should go down. We showed that this is the case, although on Reddit some posts are able to experience subsequent popularity spikes – Hacker News appears to be resilient to this phenomenon. We explain this apparent contradiction by showing that these posts which have persistent popularity spikes on Reddit have a low canonicity, i.e. they are usually dissimilar from the average post containing their protomeme. We showed that canonicity has a non-linear effect.

These results open the way to future works. Firstly, we can now put the theory closer to practice. We can perform a controlled experiment where we select front page memes from Reddit and semi-automatically generate imitating posts with varying degrees of canonicity. By releasing these posts on Reddit, we should observe in which cases low canonicity posts tend to gather more upvotes and in which cases high canonicity is helpful. Secondly, the theory makes claims that are not in line with another theory of meme popularity. In [15, 23, 6] meme content and structure is considered a weaker explanatory factor. Better predictors are meme timing and the social network position of the meme creators. We suggest to reunite the two theories in a unified meme analysis framework. Finally, we could extend the investigation by studying the effect of negative votes, which we expect to show non-trivial dynamics: a vote, even if negative, still comes from a person paying attention to the concept, although its effect is to prevent other people to see it. This information was not available for this study, but Reddit has recently started to provide it reliably.

References

[1] L Adamic, T Lento, E Adar, and P Ng. Information evolution in social networks. `http://www.ladamic.com/papers/infoevolution/MemeEvolutionFacebook.pdf`, 2014, Date of access: 07/02/2014.

[2] Christian Bauckhage, Kristian Kersting, and Fabian Hadiji. Mathematical models of fads explain the temporal dynamics of internet memes. In *International Conference on Weblogs and Social Media*, 2013.

[3] Jonah Berger and Katherine L Milkman. What makes online content viral? *Journal of Marketing Research*, 49(2):192–205, 2012.

[4] Christian Borgelt. Efficient implementations of apriori and eclat. In *FIMI: International Conference on Data Mining workshop*, 2003.

[5] Meeyoung Cha, Hamed Haddadi, Fabricio Benevenuto, and P Krishna Gummadi. Measuring user influence in twitter: The million follower fallacy. In *International Conference on Weblogs and Social Media*, volume 10, pages 10–17, 2010.

[6] Justin Cheng, Lada A. Adamic, P. Alex Dow, Jon M. Kleinberg, and Jure Leskovec. Can cascades be predicted? In *World Wide Web Conference*, pages 925–936, 2014.

[7] Nicholas A. Christakis and James H. Fowler. The spread of obesity in a large social network over 32 years. *New England Journal of Medicine*, 357(4):370–379, July 2007.

[8] Nicholas A. Christakis and James H. Fowler. The collective dynamics of smoking in a large social network. *The New England journal of medicine*, 358(21):2249–2258, May 2008.

[9] Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. The production of information in the attention economy. *Scientific reports*, 5, 2015.

[10] Michele Coscia. Competition and success in the meme pool: A case study on quickmeme.com. In *International Conference on Weblogs and Social Media*, 2013.

[11] Michele Coscia. Average is boring: How similarity kills a meme's success. In *Scientific Reports*, 2014.

[12] Richard Dawkins. *The Selfish Gene*. Oxford University Press, September 1976.

[13] Evgeniy Gabrilovich, Susan Dumais, and Eric Horvitz. Newsjunkie: providing personalized newsfeeds via analysis of information novelty. In *Proceedings of the 13th international conference on World Wide Web*, pages 482–490. ACM, 2004.

[14] Eric Gilbert. Widespread underprovision on reddit. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 803–808. ACM, 2013.

[15] James P Gleeson, Kevin P O'Sullivan, Raquel A Baños, and Yamir Moreno. Determinants of meme popularity. *arXiv preprint arXiv:1501.05956*, 2015.

[16] Nicholas Harrigan, Palakorn Achananuparp, and Ee-Peng Lim. Influentials, novelty, and social contagion: The viral power of average friends, close communities, and old news. *Social Networks*, 34(4):470–480, 2012.

[17] Himabindu Lakkaraju, Julian J McAuley, and Jure Leskovec. What's in a name? understanding the interplay between titles, content, and communities in social media. In *International Conference on Weblogs and Social Media*, 2013.

[18] Jure Leskovec, Lars Backstrom, and Jon Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506. ACM, 2009.

[19] Azadeh Nematzadeh, Emilio Ferrara, Alessandro Flammini, and Yong-Yeol Ahn. Optimal network modularity for information diffusion. *Physical Review Letters*, 113(8):088701, 2014.

[20] Daniel M Romero, Brendan Meeder, and Jon Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 695–704. ACM, 2011.

[21] Ceyda Sanlı and Renaud Lambiotte. Local variation of hashtag spike trains and popularity in twitter. *arXiv preprint arXiv:1503.03349*, 2015.

[22] Caroline Suen, Sandy Huang, Chantat Eksombatchai, Rok Sosic, and Jure Leskovec. Nifty: a system for large scale information flow tracking and clustering. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1237–1248. ACM, 2013.

[23] L. Weng, A. Flammini, A. Vespignani, and F. Menczer. Competition among memes in a world with limited attention. *Scientific Reports*, 2, March 2012.

[24] Lilian Weng, Filippo Menczer, and Yong-Yeol Ahn. Predicting successful memes using network and community structure. In *The 8th International AAAI Conference on Weblogs and Social Media (International Conference on Weblogs and Social Media'14)*. AAAI, 2014.

[25] Tim Weninger, Xihao Avi Zhu, and Jiawei Han. An exploration of discussion threads in social news sites: a case study of the reddit community. In *International Conference on Advances in Social Network Analysis and Mining*, pages 579–583, 2013.

Acknowledgments

The author wants to thank Clara Vandeweerdt, Frank Neffke and Andres Gomez for useful discussions on the statistical methodology. The author has been partly supported by FNRS grant #24927961.

Bio

Michele Coscia (michele_coscia@hks.harvard.edu) is a research fellow in the Center for International Development at Harvard University, Cambridge, MA.

Figures

Figure 1: Distributions of score and number of posts per day of the observed memes, overall and focusing only on those posts created the day after the meme was among the 25 top scoring posts. We report the average and the 95% confidence intervals. (a) Average score (Reddit); (b) Number of posts per day (Reddit); (c) Average score (HN); (d) Number of posts per day (HN).

Figure 2: Evolution of β_1 coefficients for increasing l in the model and in its associated null model. Thin lines represent the 95% confidence intervals. (a) Reddit; (b) Hacker News.

Figure 3: Distribution of the ϕ Model's β for varying l. Thin lines represent the 95% confidence intervals. (a) Reddit; (b) Hacker News.

Figure 4: Distribution of the One and Zero Models' β for varying l. Thin lines represent the 95% confidence intervals. (a) One Model; (b) Zero Model.

Tables

Table 1: Effect on β_1 of different threshold choices for the Reddit dataset. Each row is a different frequency threshold: the minimum share of posts that must contain the protomeme ($0.004 = 0.4\%$ of posts). Each column is the minimum share of days in which at least one post containing the protomeme appeared ($0.91 = 91\%$ of days). Significant values with $p < 0.01$ are highlighted in bold. The threshold values used for Figure 2 (a) are highlighted in italic.

Table 2: Effect on β_1 of different threshold choices for the Hacker News dataset. Rows and columns are interpreted as in Table 1, and significant values with $p < 0.01$ are highlighted in bold here too. The threshold values used for Figure 2 (b) are highlighted in italic. The different threshold values were chosen to accommodate for the significant difference in size between the two datasets.