

Competition Dynamics in the Meme Ecosystem

Anonymous ACL submission

Abstract

The creation and sharing of memes is a common modality of online social interactions. The goal of the present work is to better understand the collective dynamics of memes in this accelerating and competitive environment. By taking an ecological perspective and tracking the meme-text from 352 popular memes over the entirety of Reddit, we are able to show that the frequency of memes has scaled almost exactly with the total amount of content created over the past decade. One consequence of limited human attention in the face of a growing number of memes is that the diversity of these memes has decreased at the community level, albeit slightly, in the same period. Another consequence is that the average lifespan of a meme has decreased dramatically, which is further evidence of an increase in competition among memes and a decreasing collective attention span.

1 Introduction

With the rise of social media platforms, the cost historically associated with producing and consuming information has decreased to unprecedented levels; users and organizations can easily share their thoughts, stories, and others' content with widespread audiences with little effort. The decrease in production costs and increase in content availability has induced a change in scarcity dynamics: from content scarcity to consumer scarcity (Teranova, 2012). This shift has birthed new opportunities to help users find relevant content – such as recommender systems – as part of the broader *attention economy* (Shapiro et al., 1998).

The attention economy seeks to explain the allocation of cognitive resources in the creation and consumption of information. Though this concept existed long before the advent of social media (Simon, 1969), recent work has focused on how this model governs the dynamics of information consumers and curators in the socio-digital



Figure 1: Meme image with text meant to be condescending to the reader. The text of this meme and others like it are frequently used without the image in humorous or sarcastic contexts.

space (Glenski et al., 2017, 2018). The main focus has been on the consumption of information (Weng et al., 2012; Wu and Huberman, 2007; Hodas and Lerman, 2012), but others focus on the production of information (Ciampaglia et al., 2015; Leskovec et al., 2009; Simmons et al., 2011; Huberman et al., 2009). The primary question at the center of online social media is: how does the limited attention of users shape the information landscape?

The production and resharing of memes online is particularly compelling. Memes can range from short repeated text phrases, to complex images and videos. Oftentimes, the text of an image-based meme is independent of the image and is written in plaintext in social media comments and tweets; in this case, the text itself may be considered a meme in its own right. In other cases, the image and text are more closely tied together, and do not exist as memes independent of one another.

The attention *economy* seeks to explain and model the dynamics and popularity of online content within the context of established economic principles of supply and demand. When the meme marketplace is overrun with content – an excess of supply – the content must necessarily compete for consumer attention. From this perspective, memes are the product and social media users are the

consumers. While there is useful intuition found through this perspective on the online information landscape, the economic perspective does not capture all of the dynamics of how these linguistic forms are created, evolve, shared, and compete.

More recently, memes have even been considered through the lens of epidemiology and disease transmission (Wang and Wood, 2011; Kubo et al., 2007). Intuitions about memes are also derived through association with genes and the process of gene evolution (Dawkins, 2016). Indeed, the etymology of ‘meme’ is a portmanteau of ‘mind’ and ‘gene’. This begs the question: rather than continuing the economic, genealogical, epidemiological analogies, is there a better situated analogy? In particular, are memes better situated in the realm of *ecology*? And if so, what understandings can be gleaned from this perspective?

In the present work, we present findings about competition and diffusion of memes from an ecological perspective.

By taking the ecological perspective, we consider a meme to be a single species existing within the same environment or habitat as many other meme-species. As there are many memes active at any one time, they must all coexist – and compete for resources – within the same habitat. This ecological perspective shifts the focus away from the human users back to the memes and the environments within which they exist – wherein memes seek both longevity and a large population. In competing for limited environmental resources, *i.e.*, human attention, memes must compete with one another.

Within the perspective of the *meme ecology*, we ask the following research questions:

RQ1: How do the ecological resources, *i.e.*, collective user attention, scale? Do more users permit a larger or smaller number of memes?

RQ2: How do memes compete for attention? How do new meme introductions impact the existing meme ecosystem?

RQ3: How has the social ecosystem, *i.e.*, user and meme dynamics, changed over time?

The results of this analysis and the behavior they illuminate are compelling and support the case for the ecological perspective. Overall, we find that memes comprise a relatively constant fraction of

Table 1: Meme dataset consists of 352 text memes, ranging in length from 1 to 8 tokens. Some memes reference current pop-culture events, while others seem unconnected to trends of the time.

Tks	Cnt	Examples
1	69	thicc; yeet; wat; mfw; impossibru; stahp
2	80	moms spaghetti; zerg rush; y tho; nailed it
3	69	winter is coming; u wot m8; download more ram
4	48	do you even lift; kill it with fire; its free real estate
5	25	hello darkness my old friend; you must construct additional pylons
6	25	shrek is love shrek is life; shut up and take my money
7	18	still a better love story than twilight; one does not simply walk into mordor
≥8	18	this is why we cant have nice things; i dont want to live on this planet

all activity on the Reddit platform, even as it has increased in popularity. We also find that as more memes have been created, their average lifecycles have become shorter, which further suggests that the collective human attention span on social media is decreasing.

Although the current work focuses on short, frequently repeated texts, *i.e.*, memes, we hypothesize that our findings are likely to apply to a number of other communication modalities like image-memes and hashtags.

2 Methodology

Text Memes

We collected 352 popular text-based memes tracked by KnowYourMeme.com for use in our analysis across Reddit’s history. KnowYourMeme is a website that hosts a mixture of user generated and owner curated information and analysis on memes that have become popular online. For our analysis, we selected the most-popular 352 text-based memes from the "Confirmed" meme category on KnowYourMeme. In terms of popularity, these text-based memes range from 250 thousand to 13 million page views each at the time of dataset creation.

Extended meme text (*e.g.*, *copy pasta*) is truncated to include only the 8-token prefix. The final set contains meme-phrases that range in length from 1 to 8 word tokens as shown in Table 1.

Reddit Dataset

Given its size, popularity, and abundance of communities, we selected Reddit to study. Reddit is a

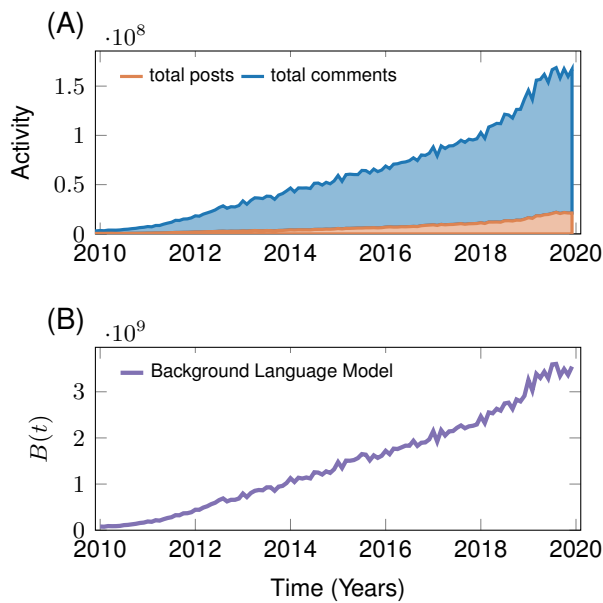


Figure 2: (A) Stacked line plot representation of Reddit contributions between 2010 and 2020. The lower (orange) region shows the number of posts per month; the upper (blue) region shows the number of comments per month. (B) The unigram background model is used to compute normalized meme-frequencies. This background behavior closely mirrors the growth of Reddit, but is one order of magnitude larger.

150 primarily text-based communication platform that
 151 allows users to create communities called subred-
 152 dits and, within those communities, create and par-
 153 ticipate in discussion threads. We collected all
 154 posts and comments on Reddit from Jan. 2010
 155 to Jan. 2020; the number of monthly posts and
 156 comments is plotted in Figure 2. The data was
 157 collected from Reddit using the PushShift.io data
 158 collection toolset (Baumgartner et al., 2020). The
 159 dataset contains messages and all data necessary to
 160 reconstruct discussion threads and the communities
 161 those discussions existed within.

162 By tracking the use of all 352 memes across all
 163 Reddit activity, we can assess not only individual
 164 meme popularity, but also more general patterns of
 165 meme growth and behavior.

166 Note that tracking rapidly-evolving image tem-
 167 plates is outside the scope of the present work.
 168 More plainly, image-memes are not included in
 169 this analysis for a myriad of reasons. Most impor-
 170 tantly, tracking evolving meme images and their
 171 included text is a task that has yet to have a satis-
 172 factory technical solution developed.

173 Despite the advantages of a text-based analysis,

174 there are limitations to the meme set that we col-
 175 lected. A large segment of online meme culture
 176 is dedicated to more visual forms of meme: im-
 177 ages, gifs, and videos. We therefore assume that
 178 textual meme representations follow similar pat-
 179 terns and behaviors as image memes, but cannot
 180 guarantee that results will generalize to these other
 181 content types. We also restrict our analysis to iden-
 182 tical text-matching, thereby omitting any evolution
 183 of text memes (Leskovec et al., 2009). The set of
 184 memes was curated with this in mind, and therefore
 185 focuses on clear text ‘units’ that exist without vari-
 186 ation, omitting template-style text memes in favor
 187 of simpler repeated phrases. Still, we acknowledge
 188 the implications of this decision as it pertains to the
 189 current work. Despite the limitations, we believe
 190 the dataset is varied and large enough to support
 191 our analysis and the conclusions we draw from it.

192 The questions and data considered in the present
 193 work are considered human subjects research, and
 194 relevant ethical considerations are present. We
 195 sought and received research approval from the
 196 REDACTED Institution Review Board.

3 Collective Attention to Memes is Stationary

197 Previous work has shown that innovation and tech-
 198 nological development is accelerating. Moore’s
 199 Law is one example of this phenomenon where
 200 a compounding increase in circuit density has
 201 led to remarkable increases in computational
 202 power (Schaller, 1997); similar effects have been
 203 shown in genome sequencing (Mardis, 2011) and
 204 telecommunications bandwidth (Eldering et al.,
 205 1999). In online social systems, the early empirical
 206 evidence suggests that a similar pattern exists: that
 207 social innovations are accelerating (Lorenz-Spreen
 208 et al., 2019; Rosa, 2003; Hutchins, 2011; Rosa,
 209 2013). This pattern is the basis for **RQ1**: How does
 210 collective user attention of memes scale? Does the
 211 presence of larger groups result in super-scaling ef-
 212 fects like those found in population densities (Pan
 213 et al., 2013) and software development (Thomas
 214 et al., 2019), where collections of individuals pro-
 215 duce more than the sum of their parts?
 216

217 At first glance, Figure 2 appears to show that our
 218 data supports these claims: more posts, comments,
 219 and memes are being made at an accelerating pace
 220 year over year. But how much attention is paid to
 221 individual memes? To answer this question, we
 222 first need to measure collective attention.
 223

3.1 Measuring Collective Attention

Because we cannot collect the number of users who viewed or thought about a particular meme, we instead estimate the collective attention for each meme by the number of times it appears, *i.e.*, its frequency. So our first task is to extract the daily frequency $F_m(t)$ of each meme, such that $F_m(t)$ is the frequency of meme m on day t . This frequency encompasses all occurrences of meme m on Reddit, including both post and comment text written by any user.

Counting tokens would be insufficient to fully determine whether memes are truly increasing in prevalence or if their frequency increases mirror the increase in content that Reddit has experienced overall. In order to produce an accurate assessment of the meme ecology, we need to normalize the frequency of meme occurrences.

To do this, we constructed a set of 5000 randomly selected words from Reddit to serve as a background language model. Then, for each day, we count the number of occurrences of each token in the background set, such that $B(t)$ gives the background frequency sum of all 5000 words on day t . As the number of posts and comments on Reddit grows over time, so too do the occurrences of the background set. As seen in Figure 2, the background growth closely mirrors the growth of Reddit.

Given the similar growth patterns of overall Reddit activity and the background language model, we expect the number of meme occurrences to increase at a similar rate. Therefore, we compute the *normalized meme-frequency* for each meme $\hat{F}_m(t) = F_m(t)/B(t)$ as a control.

This normalization step removes the Reddit growth effects from the meme frequencies, allowing popularity to be directly compared across various time frames. Because the background model mirrors Reddit’s growth over time, the normalized frequency $\hat{F}_m(t)$ can be viewed as the attention given to meme m relative to the activity level at time t .

To assess the aggregate attention paid to all memes in the dataset, we compute the mean average normalized meme-frequency $\frac{1}{|M|} \sum_{m \in M} \hat{F}_m(t)$ over the entire collection of memes M . Instead of measuring the popularity of an individual meme, this summarizes the average popularity of memes in total, relative to Reddit’s growth.




Figure 3: Mean normalized meme-frequency from 2010–2020 and 95% confidence interval (shaded region). Light grey lines show the individual normalized meme-frequency for a random 10% sample of individual memes. Overall, meme occurrence has remained consistent over the past decade (Pearson $R = +0.03$, p -value < 0.01).

Figure 3 illustrates the mean normalized meme-frequency along with its 95% confidence interval from 2010 to 2020. A random selection of individual memes are also plotted in light grey. We find that the occurrence of memes remains remarkably consistent when controlled for Reddit’s overall activity, even as the occurrence of individual memes varies widely. Correlation analysis finds almost no association between time and the normalized meme-frequency (Pearson $R = +0.03$, p -value < 0.01).

3.2 Discussion

Competition for limited attention is not a new concept in socio-digital media studies. However, previous analyses generally focused on network-level simulation or mathematical modeling as an attempt to predict the behavior of large social systems (Weng et al., 2012; Gleeson et al., 2014). User-level attention analysis seeks to understand the impact of competition and content volume on individual actors (Hodas and Lerman, 2012). Other work attempts to determine which features will make a meme successful within a competitive environment (Coscia, 2013; Lakkaraju et al., 2013). These works are valuable extensions of the classic attention economy, yet fail to demonstrate the system-wide effect of competition.

By showing that meme activity accounts for a stable fraction of all Reddit content, we have demon-

strated the ability for competition to act as a kind-of global bandwidth cap. Even with a rapidly expanding user base, and a seemingly insatiable hunger for fresh content, Reddit appears to have a limited amount of attention that can be paid to memes as a whole.

Given that aggregate meme activity levels are stable, this suggests that the online environment can only support a limited meme population. This is consistent with real-world ecology dynamics: a given habitat has limited resources, which must be shared between all resident species, which in turn limits the total population of the habitat. In summary, because memes must compete with one another for scarce resources, the Reddit environment can only sustain a limited population of memes.

4 Competition Among Memes

The prevalence of individual memes rises and falls over time. Popularity is fleeting, and most memes do not stay relevant forever. But as one meme dies out, another rises to take its place. The ebb and flow of popularity has long been studied as the *diffusion of innovations* (Gleeson et al., 2014).

The dynamic behavior of memes is especially interesting: although individual memes may die, memes in general continue to thrive. This contradictory dynamic reminds us that these memes exist in constant competition with one another. The competition among memes is not unlike competition observed in markets, where innovations eventually replace outdated products, or ecological systems, where competition among individuals in a group exerts a selective pressure that rewards certain genetic innovations.

4.1 Community Diversity

In the ecological perspective, one way to assess the overall health of an ecological system is to examine the *diversity* of the species living within the habitat. In this case, the ecology of a diverse, thriving social media environment ought to have several different memes simultaneously appearing in abundance. This ‘meme diversity’ would indicate healthy user interaction with a wide variety of content topics, and leads us to **RQ2**: As Reddit grows and new memes are introduced, does the overall diversity of memes increase or decrease?

We measure the diversity of an ecological system using Simpson’s Diversity Index (D) (Gregorius and Gillet, 2008). This index takes into account

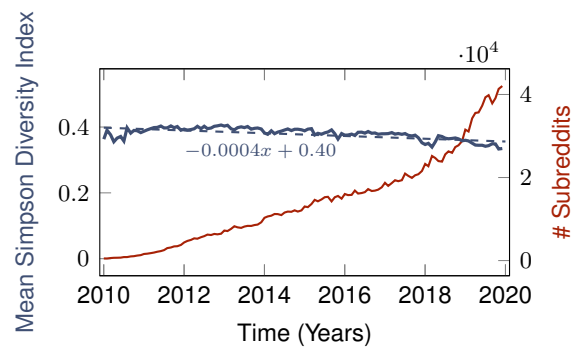


Figure 4: Average Simpson’s Diversity Index across subreddits, and number of subreddits, over time. Meme diversity on Reddit is decreasing (Pearson $R = -0.63$, p -value < 0.01) with a small slope (0.04% per month) despite an increase in the number of subreddits containing memes.

both the number of unique species present, as well as the population of each species.

With this we can measure the ‘meme diversity’ of a community, and assess that community’s suitability as a meme habitat. Does the community provide enough resources to support a wide variety of memes, or can only a few species survive?

We compute Simpson’s Diversity Index for each subreddit in each month. By computing the diversity for each individual month, we hope to capture the changeover in meme popularity. From these subreddit diversity measures, we then compute the mean diversity index over all subreddits, resulting in a monthly mean diversity of Reddit at the community level. This serves as an assessment of Reddit overall – how diverse is the average community, and how is this changing over time?

Figure 4 shows the average diversity and 95% confidence intervals over time, along with the number of active subreddits in each month. The confidence intervals are present in Fig. 4, but are very small: typically about 0.1%. Overall, the diversity of Reddit communities appears to be decreasing at a small (0.48% per year) but steady rate (Pearson $R = -0.63$, p -value < 0.001). Even as the number of subreddits grows, the overall diversity of these communities is decreasing.

This indicates that subreddit communities, on average, are using slightly less diverse meme subsets. This finding, coupled with the growth in the number of subreddits using memes, suggests a slow Balkanization of Reddit communities. That is to say, as more subreddit communities come into existence,

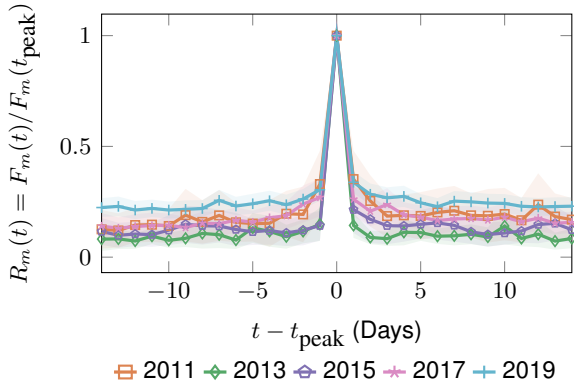


Figure 5: Average relative meme-frequency, time-shifted so that the maximum frequency occurs on day 0. Shaded areas indicate 95% confidence intervals. The width of the primary peak has not changed, suggesting that memes dynamics have not accelerated.

human attention – and the memes that attention supports – moves towards more niche subreddits, which aligns with findings from other subreddit analyses (Marchal, 2020).

5 Changing Dynamics of the Meme Ecosystem

Here we shift focus to the dynamics of collective attention. Recent work suggests that these dynamics are accelerating (Lorenz-Spreen et al., 2019). New concepts are becoming viral faster and stay viral for a shorter duration. But does this acceleration hold true for memes? **RQ3:** How have the dynamics of collective attention on memes changed over time?

5.1 Meme Dynamics

Lorenz-Spreen et al. (Lorenz-Spreen et al., 2019) used a variety of methods to analyze collective dynamics in the online social sphere. Here, we apply their methodology to our meme dataset.

First, we focus on the peaks of memes on Reddit to assess the pace of collective attention. For each meme, we compute its frequency across all of Reddit each day, such that $F_m(t)$ gives the frequency of meme m on day t . We also identify the *peak frequency* for each meme $F_i(t_{\text{peak}})$ and when this peak occurred. To ensure all memes are on the same scale, the frequencies of each meme are then normalized by that meme’s peak to get a *relative meme-frequency*: $R_m(t) = F_m(t)/F_m(t_{\text{peak}})$. In Figure 5, we illustrate the average relative peak frequency $R_m(t)$ for all memes in our dataset, grouped by the year of each meme’s peak. Overall,

there appears to be no change in peak dynamics over time. The difference between the peak and the baseline frequency (*i.e.*, frequency before and after the peak) remains relatively stable, nor does it change a statistically significant amount from year to year. Furthermore, the changes seen do not show a trend over time. This suggests that memes have not exhibited a significant acceleration over the past decade, at least with respect to their peak popularity.

What does this mean for our ecology of memes? Even as the wider ecosystem of Reddit grows, the impact on an individual meme’s trajectory has not measurably accelerated.

5.2 Meme Lifespans are Shrinking

The previous analysis raises interesting questions about the collective dynamics of memes. More memes are being created, yet their peak popularity dynamics, on average, do not appear to be accelerating; how do we account for the stability of their use? Answer: Meme lifespans are shrinking.

To understand this empirically we define the lifespan of a meme as follows. For a meme with a peak frequency of $F_m(t_{\text{peak}})$, the lifespan begins on the first day u where $\hat{F}_m(u) \geq \alpha \hat{F}_m(t_{\text{peak}})$. Recall that $\hat{F}_m(t)$ is the normalized meme-frequency and is computed as $\hat{F}_m(t) = F_m(t)/B(t)$. The lifespan ends on the last day v where the meme experiences a normalized frequency $\hat{F}_m(v) \geq \alpha \hat{F}_m(t_{\text{peak}})$, such that all days between the beginning, peak, and the end are continuous and above $\alpha \hat{F}_m(t_{\text{peak}})$. The threshold value α is very small; here we present results for $\alpha = 0.005, 0.01$, and 0.02 . Other values produced similar results.

By defining the lifespan this way, each meme’s lifespan captures the majority of its occurrences, but does not include very early, late, or anomalous uses. The lifespan is also determined using normalized meme-frequencies to control for Reddit growth. We compute the lifespan length in number of days for each meme. For a lifespan threshold of $\alpha = 0.005$, these lifespans range from a high of 4140 days to a low of 1 day. (For completeness, we include the full history of Reddit beginning in 2005 when defining lifespans.)

Figure 6(A) illustrates the number of memes active within each month. As expected, the number of active memes has increased over time, mirroring Reddit’s overall growth.

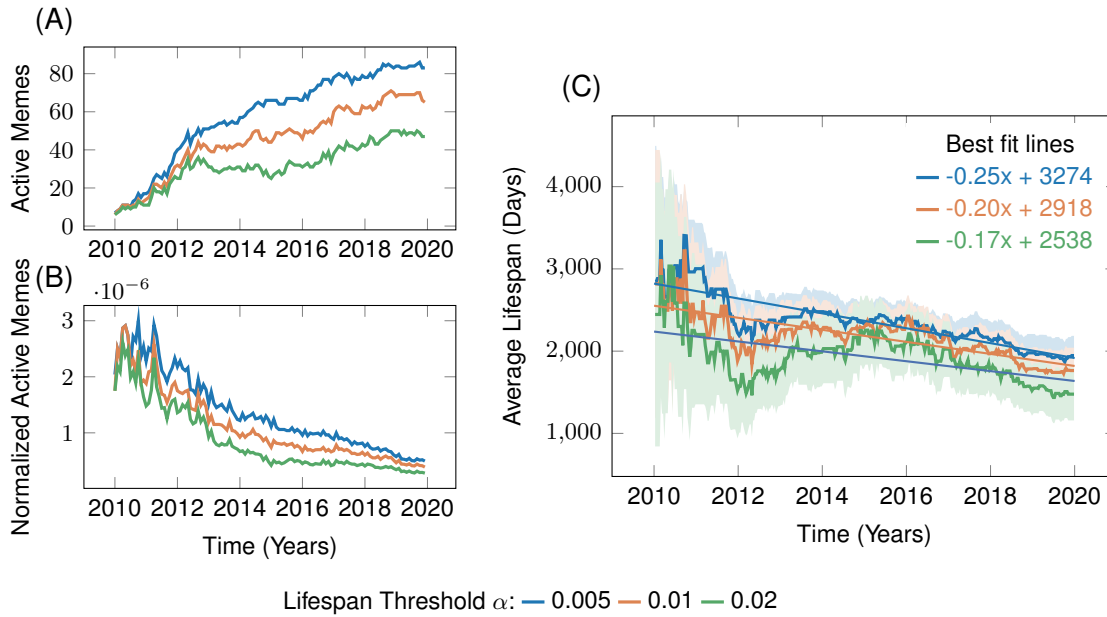


Figure 6: Results of lifespan analysis on meme set for $\alpha = 0.005, 0.01,$ and 0.02 . Lifespan starts on the first day with a frequency $\geq \alpha \hat{F}_m(t_{\text{peak}})$, and ends on the last day with frequency $\geq \alpha \hat{F}_m(t_{\text{peak}})$, such that all days in the lifespan have a frequency $\geq \alpha$ times the maximum normalized meme-frequency. (A) Number of active memes per month, where a meme is active only during its defined lifespan. (B) Number of active memes per month, normalized by total Reddit contributions. Reddit growth is outpacing the number of active text memes. (C) Average meme lifespan (in days) over time and corresponding best fit lines. Shaded area indicates 95% confidence interval. Over the course of 10 years, the average lifespan has decreased ($\alpha = 0.005$: Pearson $R = -0.85$, p -value < 0.01 ; $\alpha = 0.01$: Pearson $R = -0.77$, p -value < 0.01 ; $\alpha = 0.02$: Pearson $R = -0.59$, p -value < 0.01), indicating that memes are cycling faster.

To accurately assess the number of active memes we control for Reddit’s growth. Figure 6(B) shows the normalized count of active memes, *i.e.*, the number of active memes per month divided by the background model $B(t)$. Applying this normalization it is apparent that, over time, the *relative* number of active memes is decreasing. This decrease occurs simultaneously with our finding that normalized meme-frequency has remained stable (Figure 3). The most likely explanation for these seemingly contradictory findings is that the lifespan of memes is decreasing.

We can test this hypothesis directly by computing how lifespans have changed over time. Figure 6(C) plots the average meme lifespan and 95% confidence interval in days for the active memes within each month. For example, if there were 40 memes active in June 2012 and the mean-average lifespan of these 40 memes was 2,400 days, then we would plot 2,400 (and the corresponding confidence interval) for June 2012. To avoid recency bias wherein memes that were active when the data collection ended would skew results downward, we removed any memes that we active at the end of the data

collection from the lifespan results. The lifespan is heavily dependent on the α threshold. So, we repeat these calculations for $\alpha = 0.005, 0.01,$ and 0.02 . In all cases, we see that the average lifespan decreases over time. Overall, memes are rising and falling more quickly now than previously.

6 Discussion

Our findings show that the relative fraction of Reddit focused on memes has not changed over time (Figure 3), the meme diversity of subreddits has decreased (Figure 4), and the relative number of active memes has declined (Figure 6(B)). These results suggest that the production and consumption of new memes are accelerating. As fast as new content is created, old content is falling into disuse. Given the growing volume of content on Reddit every day, a necessary consequence is that the collective attention span of Reddit is decreasing.

In the ecological context, what does this tell us about meme species? The resources available to memes have not increased (Figure 3), the diversity of individual habitats is decreasing (Figure 4), and the proportion of active memes within the ecosys-

tem has declined (Figure 6(B)). The emergence, and eventual extinction, of new meme species is accelerating – lifespans are growing shorter.

7 Conclusions

The three research questions in the present work coalesce into an emerging theory of meme dynamics in online social bulletin boards like Reddit. Taking an ecological perspective, we have shown that although Reddit as a whole continues to grow, the fraction of Reddit devoted to text memes is consistent. This reinforces the notion of a meme ecology, which, like the economic and epidemiological perspectives of information diffusion, has consequences that derive from limited user attention – the scarce resource of the online ecosystem.

We have also shown that the diversity of memes across communities is decreasing slightly, even as the number of communities continues to grow. This represents yet another consequence of the ecological perspective: as social media communities continue to Balkanize, *i.e.*, split into narrow and self-referential communities, the number of shared references to memes appears to be decreasing. With so many habitats available, each one is able to specialize in a particular type of meme species, leading to reduced meme diversity of that environment.

Finally, we have shown that meme lifespans have decreased significantly. Yet, unlike similar work on general n-grams, hashtags, citations, etc., we do not find that all aspects of meme dynamics are accelerating, and even the accelerations are bounded by the effect of increasing competition for limited user attention. The patterns of relative meme growth, even surrounding a meme’s peak, remain relatively unchanged by the expansion of Reddit. And although the absolute number of memes active on Reddit has increased over time, this growth has been outpaced by the growth of Reddit itself. The population change of new meme species has not changed, but the overall lifespan of the species is reduced. This suggests that memes are spending less time at their peak popularity – the status as the dominant species is more transient than before.

Limitations

It is important to recognize the limitations of this study. Our evaluation is principally limited by the KnowYourMeme dataset. Because the popularity metric is accumulative on KnowYourMeme, it is likely that newer memes, which have not yet accumulated their full/final popularity, are less likely to

appear in our dataset. Our focus on the textual artifacts of memes may also limit the generalizability of our findings. Much of the online meme culture, especially recently, is conveyed through visual depictions of memes (Theisen et al., 2020). We necessarily assume that textual representations of visual memes follow proportionately, but we are unable to validate this claim in the current work. Future work should endeavor to confirm these findings on visual imagery. Finally, the present work used a strict text-matching algorithm to identify memes. Additional follow-up work is needed to better understand and trace how these memes evolve and how the shape of a meme’s popularity and evolution affects its lifespan and reach.

Future Work

We believe that the ecological perspective is a valuable tool for the analysis of online social media. By applying existing tools in this area to current computational problems, we can reveal unexpected facets of content dynamics. In the future, we hope to apply our methodology to larger, more varied content datasets. For example, image memes are even better-suited to an ecological analysis, with their rapidly evolving templates and formats. We also aim to expand the ecological approach to examine other components of content dynamics and the attention economy. Our current work lays the foundation for a more comprehensive analysis of the meme ecology.

References

- Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. 2020. The pushshift reddit dataset. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 830–839.
- Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. 2015. The production of information in the attention economy. *Scientific reports*, 5:9452.
- Michele Coscia. 2013. Competition and success in the meme pool: a case study on quickmeme.com. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7.
- Richard Dawkins. 2016. *The selfish gene*. Oxford university press.
- Charles A Eldering, Mouhamadou Lamine Sylla, and Jeffrey A Eisenach. 1999. Is there a moore’s law for bandwidth? *IEEE Communications magazine*, 37(10):117–121.

614	James P Gleeson, Jonathan A Ward, Kevin P O'sullivan, and William T Lee. 2014. Competition-induced criticality in a model of meme popularity. <i>Physical review letters</i> , 112(4):048701.	669
615		670
616		671
617		672
618	Maria Glenski, Corey Pennycuff, and Tim Weninger. 2017. Consumers and curators: Browsing and voting patterns on reddit. <i>IEEE Transactions on Computational Social Systems</i> , 4(4):196–206.	673
619		674
620		675
621		
622	Maria Glenski, Greg Stoddard, Paul Resnick, and Tim Weninger. 2018. Guessthekarma: a game to assess social rating systems. <i>Proceedings of the ACM on Human-Computer Interaction</i> , 2(CSCW):1–15.	676
623		677
624		
625		
626	Hans-Rolf Gregorius and Elizabeth M Gillet. 2008. Generalized simpson-diversity. <i>Ecological Modelling</i> , 211(1-2):90–96.	680
627		681
628		682
629	Nathan Oken Hodas and Kristina Lerman. 2012. How visibility and divided attention constrain social contagion. In <i>2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Confernece on Social Computing</i> , pages 249–257. IEEE.	683
630		684
631		685
632		
633		
634		
635	Bernardo A Huberman, Daniel M Romero, and Fang Wu. 2009. Crowdsourcing, attention and productivity. <i>Journal of Information science</i> , 35(6):758–765.	686
636		687
637		688
638	Brett Hutchins. 2011. The acceleration of media sport culture: Twitter, telepresence and online messaging. <i>Information, communication & society</i> , 14(2):237–257.	689
639		690
640		
641		
642	Masao Kubo, Keitaro Naruse, Hiroshi Sato, and Takashi Matubara. 2007. The possibility of an epidemic meme analogy for web community population analysis. In <i>International Conference on Intelligent Data Engineering and Automated Learning</i> , pages 1073–1080. Springer.	691
643		692
644		693
645		694
646		695
647		
648	Himabindu Lakkaraju, Julian McAuley, and Jure Leskovec. 2013. What's in a name? understanding the interplay between titles, content, and communities in social media. In <i>Proceedings of the International AAAI Conference on Web and Social Media</i> , volume 7.	696
649		697
650		698
651		699
652		700
653		701
654	Jure Leskovec, Lars Backstrom, and Jon Kleinberg. 2009. Meme-tracking and the dynamics of the news cycle. In <i>Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining</i> , pages 497–506.	702
655		703
656		704
657		705
658		
659	Philipp Lorenz-Spreen, Bjarke Mørch Mønsted, Philipp Hövel, and Sune Lehmann. 2019. Accelerating dynamics of collective attention. <i>Nature communications</i> , 10(1):1–9.	706
660		707
661		708
662		709
663	Nahema Marchal. 2020. The polarizing potential of intergroup affect in online political discussions: Evidence from reddit r/politics. <i>Available at SSRN</i> .	710
664		711
665		712
666	Elaine R Mardis. 2011. A decade's perspective on dna sequencing technology. <i>Nature</i> , 470(7333):198–203.	
667		
668		
	Wei Pan, Gourab Ghoshal, Coco Krumme, Manuel Cebrían, and Alex Pentland. 2013. Urban characteristics attributable to density-driven tie formation. <i>Nature communications</i> , 4(1):1–7.	
	Hartmut Rosa. 2003. Social acceleration: ethical and political consequences of a desynchronized high-speed society. <i>Constellations</i> , 10(1):3–33.	
	Hartmut Rosa. 2013. <i>Social acceleration: A new theory of modernity</i> . Columbia University Press.	
	Robert R Schaller. 1997. Moore's law: past, present and future. <i>IEEE spectrum</i> , 34(6):52–59.	
	Carl Shapiro, Shapiro Carl, Hal R Varian, et al. 1998. <i>Information rules: a strategic guide to the network economy</i> . Harvard Business Press.	
	Matthew P Simmons, Lada A Adamic, and Eytan Adar. 2011. Memes online: extracted, subtracted, injected, and recollected. <i>icwsm</i> , 11:17–21.	
	Herbert A Simon. 1969. Designing organizations for an information-rich world. <i>Brookings Institute Lecture</i> .	
	Tiziana Terranova. 2012. Attention, economy and the brain. <i>Culture Machine</i> , 13.	
	William Theisen, Joel Brogan, Pamela Biló Thomas, Daniel Moreira, Pascal Phoa, Tim Weninger, and Walter Scheirer. 2020. Automatic discovery of political meme genres with diverse appearances. <i>arXiv preprint arXiv:2001.06122</i> .	
	Pamela Biló Thomas, Rachel Krohn, and Tim Weninger. 2019. Dynamics of team library adoptions: an exploration of github commit logs. In <i>Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining</i> , pages 470–473.	
	Lin Wang and Brendan C Wood. 2011. An epidemiological approach to model the viral propagation of memes. <i>Applied Mathematical Modelling</i> , 35(11):5442–5447.	
	Lilian Weng, Alessandro Flammini, Alessandro Vespignani, and Filippo Menczer. 2012. Competition among memes in a world with limited attention. <i>Scientific reports</i> , 2:335.	
	Fang Wu and Bernardo A Huberman. 2007. Novelty and collective attention. <i>Proceedings of the National Academy of Sciences</i> , 104(45):17599–17601.	