

The Birth of Collective Memories: Analyzing Emerging Entities in Text Streams

David Graus

Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands and FD Mediagroep, Amsterdam, The Netherlands. E-mail: david.graus@fdmediagroep.nl

Daan Odiijk

Blendle, Utrecht, The Netherlands. E-mail: daan@blendle.com

Maarten de Rijke

Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands. E-mail: derijke@uva.nl

We study how collective memories are formed online. We do so by tracking entities that emerge in public discourse, that is, in online text streams such as social media and news streams, before they are incorporated into Wikipedia, which, we argue, can be viewed as an online place for collective memory. By tracking how entities emerge in public discourse, that is, the temporal patterns between their first mention in online text streams and subsequent incorporation into collective memory, we gain insights into how the collective remembrance process happens online. Specifically, we analyze nearly 80,000 entities as they emerge in online text streams before they are incorporated into Wikipedia. The online text streams we use for our analysis comprise of social media and news streams, and span over 579 million documents in a time span of 18 months. We discover two main emergence patterns: entities that emerge in a “bursty” fashion, that is, that appear in public discourse without a precedent, blast into activity and transition into collective memory. Other entities display a “delayed” pattern, where they appear in public discourse, experience a period of inactivity, and then resurface before transitioning into our cultural collective memory.

Introduction

Remembering is a social process (Halbwachs, 1950). Collective remembrance is the process in which information moves from public discourse into a shared collective

memory. This process has been compared to the remembrance process of an individual, whose memories transfer from short-term into long-term memory (Assmann & Czaplicka, 1995). This comparison has been formalized by mapping the collective’s equivalent of long-term and short-term memory to the *cultural* and *communicative* memory, respectively.

Cultural collective memory (CM), the collective’s equivalent of an individual’s long-term memory, is characterized by being organized, specialized, formal, structured, and distanced from the immediate (Assmann & Czaplicka, 1995). Wikipedia is known to “democratize information,” through its collaborative nature: Its content is produced by volunteer editors and authors from around the world (Wallace & Van Fleet, 2005). Wikipedia has been called an online place for cultural CM (Luyt, 2016; Pentzold, 2009). We support this view, and argue that the aforementioned characteristics fit Wikipedia’s nature. First, Wikipedia is *organized*, through its hierarchy of contributors, where authors are distinguished from admins. Wikipedia is *specialized*, because appropriately citing relevant and expert sources to support and back up newly added information is a requirement. These conventions, requirements, and policies around contributing new information to Wikipedia impose a level of *formality* and enable its coherent and consistent *structure*. Finally, the requirement for new articles to be collectively deemed “important enough,” ensures Wikipedia’s *distance* from the immediate.

Communicative CM is in many aspects the opposite of cultural CM. Analogously to an individual’s short-term memory, communicative CM is mainly orally negotiated, close to the everyday, disorganized, informal, and non-specialized (Assmann & Czaplicka, 1995). Online text

Received April 2, 2017; revised September 26, 2017; accepted December 7, 2017

© 2018 The Authors. Journal of the Association for Information Science and Technology published by Wiley Periodicals, Inc. on behalf of Association for Information Science and Technology • Published online Month 0, 2017 in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/asi.24004

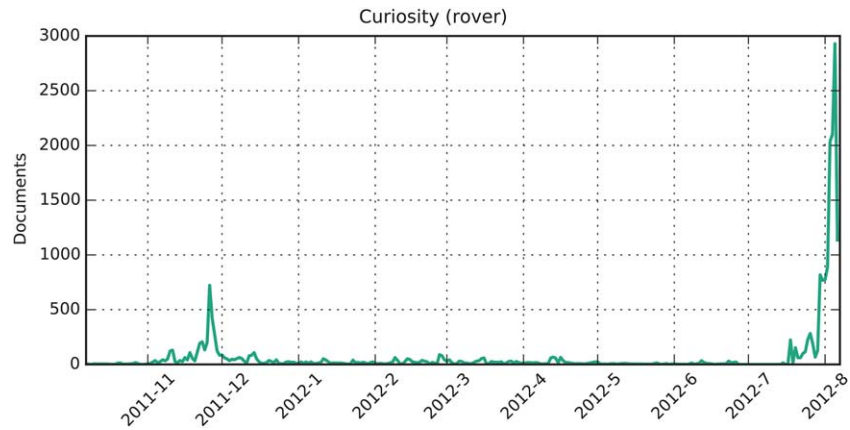


FIG. 1. Emergence pattern of the entity Curiosity (Rover), first mentioned in our text stream in October 2011. The Wikipedia page for Curiosity was created 9 months later, on August 6, 2012. There are two distinct *bursts*, one late November 2011, the second shortly before the entity is added to Wikipedia. The two bursts correspond to the Mars Rover’s launch date (November 26, 2011) and its subsequent landing (August 6, 2012). [Color figure can be viewed at wileyonlinelibrary.com]

streams fit this notion of orally negotiated memory: the rapid pace and high volume at which content is published by news websites and social media platforms means that—as opposed to the carefully curated and edited nature of Wikipedia—online text streams are close to the everyday: They not only record and reflect the actions of everyday life but also have a role in producing everyday life for a media-enabled public (Tierney, 2013, p. 33). With the advent of Web 2.0, and the ability for anyone to publish content on the web, online text streams have naturally become disorganized, informal, and non-specialized.

We study the evolution of collective memory by tracking additions to our online cultural CM, Wikipedia. Specifically, we study real world *entities*¹ as they emerge in online text streams, and are subsequently added to Wikipedia. Every day, new content is being added to Wikipedia, with the knowledge base receiving over 6 million monthly edits at its peak (Suh, Convertino, Chi, & Pirolli, 2009). Domain experts may find information missing on Wikipedia and take up the task of contributing this new information. Alternatively, new, previously unheard-of entities may emerge in news articles or social media postings that describe or comment on events, for example, the Olympics may introduce new athletes onto the world stage, or the opening of a new restaurant may be reported in local news media and appear in social media. Studying entities that transition from public discourse into Wikipedia gives us insights into how collective memory evolves—for the first time, the online world allows us to make such observations at scale.

To study emerging entities, we analyze entities in a sample of online social media and news text streams spanning over 18 months. We focus on the entities’ *emergence patterns*, that is, how an entity’s exposure evolves between its

first mention in online text streams, and the moment it is added to Wikipedia. We define an entity’s *emergence pattern* to be its “document mention time series,” that is, the time series that represent the number of documents that mention the entity per day,² starting at its first mention in the stream, until it is incorporated into Wikipedia. An example time series is shown in Figure 1, with the number of documents that mention Curiosity on the y-axis (the *emergence volume*) and the time span between the entity’s first mention in online text streams and the day it is added to Wikipedia on the x-axis (the *emergence duration*).

The main findings of this article are as follows. By clustering entity’s emergence patterns, we find two kinds of regularity: entities that show a strong *early burst* around the time of their introduction into public discourse, and *late bursting* entities that exhibit a more gradual emergence. Furthermore, we find meaningful differences between how entities emerge in social media and news streams: entities that emerge in social media tend to transition more slowly from communicative CM to cultural CM than those that emerge in news streams. Finally, we show how different entity *types* exhibit different emergence patterns; the fastest emerging entities are types that know shorter life-cycles such as devices (e.g., smartphones), and “cultural artifacts” (e.g., movies and music albums).

Related Work

The concept “collective memory” was analyzed and advanced in the 1920s by (Halbwachs (1925, 1950). Since its introduction, the concept has been studied in a variety of interdisciplinary fields, most notably in literature, history, and media (Erl, Nunning, & Young, 2008; Olick, Vinitzky-Seroussi, & Levy, 2011) but also in (experimental) psychology (Hirst & Manier, 2008; Reese & Fivush, 2008; Wessel

¹NotesBy an “entity” we follow standard practice and mean a *thing with distinct and independent existence*, for example, a person or device (Oxford Dictionary, <http://www.oxforddictionaries.com/definition/english/entity>).

²Because we are interested in broad and long-term patterns, our time series are at a granularity of days, not hours.

& Moulds, 2008), for example, by empirically studying the performance of remembering events of different members of a single social group (Brown, 1990).

Wikipedia was first dubbed a global memory place where collective memories are built by Pentzold (2009), with follow-up studies by Keegan (2011) and Ferron and Massa (2011). As Ferron (2012, p. 23) puts it, “Wikipedia’s processes of discussion and article construction can be seen as the discursive formation of memory, or in other terms, as the transition from *communicative memory*, which is interactive, informal, nonspecialized, reciprocal, disorganized and unstable, to *cultural memory*, which is formal, well organized and objective” (our italics).

In the context of online collective memory, studies have revolved around automated methods for analyzing texts, for example, studying temporal expressions in web documents has shown that we tend to remember the “near past” online (Au Yeung & Jatowt, 2011). Wikipedia viewership statistics have provided insights into how current events trigger remembrance patterns of past events (García-Gavilanes, Mollgaard, Tsvetkova, & Yasseri, 2017). Other sources used for online collective memory studies include search engine query logs (Campos, Dias, & Jorge, 2011) and microblog services (Jatowt, Antoine, Kawai, & Akiyama, 2015).

Our work differs from previous work on collective memory in two important ways. We are the first to empirically study the transition from communicative CM to collective CM in terms of the entities that are mentioned in news and social text streams, before being included in Wikipedia. And we are the first to empirically study this transition at scale and across text streams and entity types, signifying an important difference from case studies that involve dramatic or traumatizing events, characteristic of the study of “collective memory” (Lipsitz, 2001; Neal, 1998).

Growth and Development of Wikipedia

Previous work on studying the expansion of Wikipedia through the addition of new pages usually studies the phenomenon from the perspective of Wikipedia itself, for example, by analyzing how newly created articles fit in Wikipedia’s semantic network, studying the relation between activity on talk pages and the addition of new content to articles, or by studying controversy and disagreement on new content through “edit wars” (Kämpf, Tessenow, Kenett, & Kantelhardt, 2016; Keegan, Gergle, & Contractor, 2013; Yasseri, Sumi, & Kertész, 2012; Yasseri, Sumi, Rung, Kornai, & Kertész, 2012).

Emerging entities have emerged as object of study in the natural language processing and information retrieval communities. Different methods for identifying and linking unknown or emerging entities have been proposed (Hoffart, Altun, & Weikum, 2014; Lin, Mausam, & Etzioni, 2012; Nakashole, Tylenda, & Weikum, 2013; Voskarides, Odijk, Tsagkias, Weerkamp, & de Rijke, 2014). Graus, Tsagkias, Buitinck, and de Rijke (2014). study the problem of predicting new concepts in social streams. Färber, Rettinger, and El

Asmar (2016) study the specific challenges and aspects that come with linking emerging entities, whereas Reinanda, Meij, and de Rijke (2016) study the problem of identifying relevant documents for known and emerging entities as new information comes in, and (Graus, Tsagkias, Weerkamp, Meij, and de Rijke, 2016) present a method for updating representations based on newly identified information. Our work differs from the aforementioned studies in being observational in nature and its focus on temporal patterns.

Research Questions

In studying emergence patterns of entities, we apply different methods of grouping entities. First, we apply a burst-based unsupervised hierarchical clustering method to group entities by similarities in their emergence patterns. This allows us to answer the following question:

RQ1: Are there common patterns in how entities emerge in online text streams?

Next, we examine emerging entities in different types of text stream, viz. news and social media streams. In addition, we study the cross-pollination between the two types of streams, that is, we study whether entities appear first in either of the streams, or whether they simultaneously appear in both. We answer the following question:

RQ2: Do news and social media text streams exhibit different emergence patterns?

Finally, we characterize the emergence patterns of different types of entities. We leverage DBpedia, the structured counterpart of Wikipedia, to group emerging entities by their types, for example, companies, athletes, and video games. We answer the following question:

RQ3: Do different types of entities exhibit different emergence patterns?

Data and Methods

Our dataset spans 7.3 million time-stamped documents, with 36.2 million references to $n = 79,482$ unique emerging entities, that is, entities that did not have a Wikipedia entry at the time they were first mentioned in the corpus, but that did have one by the time the last document in the corpus was published.

We create our custom dataset by extending the TREC-KBA StreamCorpus 2014³ with an additional set of annotations to Freebase entities (FAKBA1⁴). We then enrich the FAKBA1 dataset with links to Wikipedia, including for each link (a) the creation date of the associated Wikipedia page, and (b) whether the Wikipedia page existed at the time

³<http://trec-kba.org/kba-stream-corpus-2014.shtml>

⁴<http://trec-kba.org/data/fakba1/>

TABLE 1. Descriptive statistics of our dataset acquisition. Coverage over preceding dataset in brackets. Looking at the second and third row in the table, we note that roughly two-thirds of the FAKBA1 entities can be mapped to Wikipedia. However, this portion represents 98% of the mentions. The missing one-third were Freebase entities that had no links to Wikipedia, most notably, WordNet concepts and entities from the “MusicBrainz” knowledge base (i.e., artists, albums, and artists). The last two rows show that one in ten of the entities emerge during the span of the dataset, however, they constitute a mere 1% of the mentions.

	Dataset	# Entities		# Mentions		# Documents	
1.	TREC KBA	N/A		N/A		579,838,246	
2.	FAKBA1	3,272,980		9,423,901,114		394,051,027	(68.0%)
3.	FAKBAT	2,254,177	(68.9%)	9,221,204,641	(97.8%)	394,051,027	(100%)
4.	OOKBAT	225,291	(10.0%)	94,929,292	(1.0%)	23,896,922	(6.1%)
5.	Emerging entities	79,482	(35.3%)	36,242,096	(38.2%)	7,291,700	(30.5%)

the document was created. To encourage further research in emerging entities, we publicly release the tools needed to acquire the dataset used in this paper.⁵

OOKBAT Dataset

Our custom dataset is based on the TREC KBA StreamCorpus 2014, which comprises roughly 1.2 billion time-stamped documents from global public news wires, blogs, forums, and shortened links shared on social media. It spans 572 days (October 7, 2011–May 1, 2013).

All (English) documents in the StreamCorpus have been automatically tagged for named entities with the Serif tagger (Boschee, Weischedel, & Zamanian, 2005), yielding roughly 580M tagged documents. Dalton, Frank, Gabrilovich, Ringgaard, and Subramanya (January 2015) further automatically annotated these 580M documents with Freebase entities, resulting in the *Freebase Annotations of TREC KBA 2014 Stream Corpus (FAKBA1) dataset*, which spans over 394M documents (Table 1, line 2). Because the Freebase used in FAKBA1 is dated after the StreamCorpus timespan, we can identify entities that appear in documents prior to being incorporated in Wikipedia.

We take an entity’s Wikipedia page creation date to be its time of transitioning from communicative to cultural CM. To extract Wikipedia page creation dates for the Freebase entities present in FAKBA1, we leverage the available Wikipedia-mappings in Freebase. We then append the Wikipedia page creation dates (or *entity timestamp*, denoted e_T) to each entity in the FAKBA1 dataset. In addition, we include the entity’s “age” relative to the document timestamp (doc_T): the period in days between e_T and doc_T , that is, $e_{age} = e_T - doc_T$. The resulting dataset, FAKBA1, extended with the entity age and entity timestamp, is denoted *Freebase Annotations of TREC KBA 2014 Stream Corpus with Timestamps (FAKBAT; Table 1, line 3)*.

We retain only documents that contain entities with $e_{age} < 0$, that is, emerging entities that are mentioned in documents dated before the entity’s Wikipedia creation date. We denote the resulting subset of FAKBAT documents with emerging entities *Out of Knowledge Base Annotations (with) Timestamps (OOKBAT; Table 1, line 4)*.

To study an emerging entity’s emergence patterns, we take two additional filtering steps. First, we prune entities with creation dates more recent than the last document in our stream, to ensure the entities emerged in the timespan of our document stream. Next, we prune all entities that are mentioned in fewer than 5 documents. This yields our final dataset, which comprises 79,482 emerging entities (Table 1, line 5).

Entity Types

To study entity types for RQ3, we map emerging entities to their respective classes assigned in the DBpedia ontology,⁶ for example, the entity Barack_Obama is mapped to the Person, Politician, Author, Award Winner classes. Out of the 79,482 emerging entities in our dataset, we have 39,713 class-mappings (a coverage of 50.0%).

Entity Popularity

As a proxy for an entity’s popularity, we extract Wikipedia pageview statistics. We extract the total number of pageviews each entity received during 2015. We choose to use the pageview counts of a year that falls outside of the timespan of our dataset so as to minimize the effects of timeliness.

Time Series Clustering

The core unit in our analysis are so-called *emergence patterns*, that is, time series that represent the number of documents that mention an entity over time. To answer our first research question, Are there common patterns in how entities emerge in online text streams? (RQ1), we apply clustering to group entities with similar emergence patterns. Clustering time series consists of three steps: First, we normalize the time series, as they might span very different periods of time. Next, we measure the similarities between time series. Third, we apply hierarchical agglomerative clustering.

⁵<https://github.com/graus/emerging-entities-timeseries>

⁶<http://mappings.dbpedia.org/server/ontology/classes/>

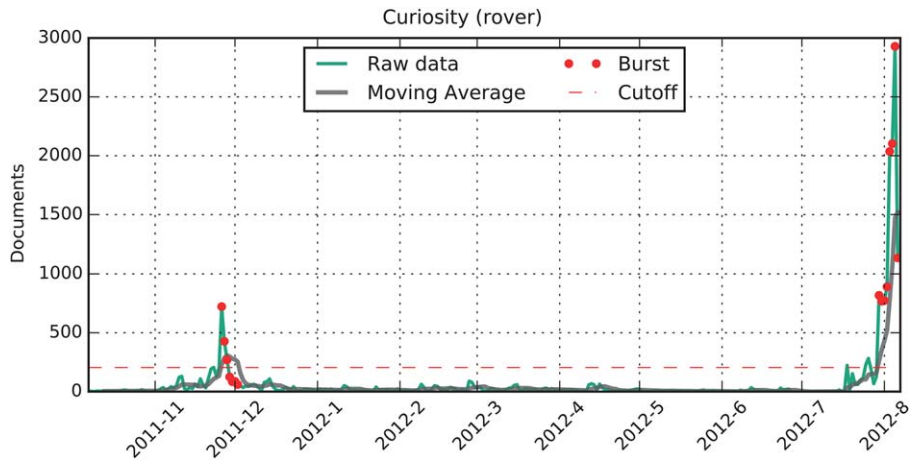


FIG. 2. Detected bursts of Curiosity (rover)’s time series. The bursts correspond to the earlier described launch and landing of the Mars Rover (see also Figure 1). [Color figure can be viewed at wileyonlinelibrary.com]

Normalization

The entities’ time series we study here are characterized by several properties. First, they are of variable length: Some entities may take days to transition, others take months. Second, the time series in our dataset are temporally unaligned: each time series starts at the timestamp of the article that first mentions the entity, and ends at the entity’s Wikipedia page creation date. To be able to visualize the time series of variable lengths, we linearly interpolate the time series to have equal length (Rani & Sikka, 2012). Furthermore, because we are not interested in the absolute differences in document volumes or entity popularity when visualizing the time series clusters, we standardize our time series by subtracting the mean and dividing by the standard deviation (Vlachos, Meek, Vagena, & Gunopulos, 2004).

Similarity

Typically, time series similarity metrics rely on fixed-length time series, and may leverage seasonal or repetitive patterns (Liao, 2005). Because our time series are of variable length, and not temporally aligned, common time series similarity metrics such as Dynamic Timewarping (DTW) are not suitable (Berndt & Clifford, 1994). Furthermore, we are interested in periods in which the exposure of an entity in public discourse increases or changes. These “bursts” may be correlated to real-world activity and events around the entity. To address the nature of the time series, and our focus on bursts we employ BSim (Vlachos et al., 2004; Burst Similarity) as our similarity metric. BSim relies on measuring the overlap between bursts of different time series. To detect these bursts we compute a moving average of each emerging entity time series (T_e), denoted T_e^{MA} . We set parameter w (the size of the rolling window) to 7 days. Bursts are the points in T_e^{MA} that surpass a cutoff value (c). We set $c = 1.5 \cdot \sigma_{MA}$, where σ_{MA} is the standard deviation of MA . The parameter choices for w and c are in line with previous work (Vlachos et al., 2004). Figure 2 shows an example time series (T_e), with the bursts

detected for the previously shown Curiosity (rover). The detected bursts correspond to the earlier mentioned launch and landing of the Mars rover.

Hierarchical Agglomerative Clustering

To cluster time series, we compute pair-wise similarities between all time series, and yield Similarity Matrix SM . We then apply L_2 normalization to SM , and convert it to a distance matrix DM ($DM = 1 - SM$). Finally, we apply hierarchical agglomerative clustering (HAC) on DM using the fastcluster package (Müllner, 2013). As our linkage criterion, we employ *Ward’s method* (Ward Jr., 1963).

Analysis

Given our grouping methods (clustering, by entity type, by stream type), we apply two methods to analyze the resulting groups of emerging entities: (a) visualization of group signatures, and (b) descriptive statistics that reflect properties of the underlying time series.

Visualization

To compare groups of emerging entities, we visualize their so-called *group signatures*, that is, the average of all time series that belong to a group. See Figure 4 for an example group signature of all emerging entities in our dataset ($n = 79,482$). As described above, the time series (may) differ in length, and are not temporally aligned. To visualize the time series, we linearly interpolate each to the (overall) longest emergence duration, effectively “stretching” them to have equal length. Next, we align them in relative duration, that is, we overlay each entity’s first and last mention at the start and end of the x -axis, respectively.

Descriptive Statistics

Visualizing group signatures does not paint the full picture. More fine-grained aspects of emergence, for example,

the average *emergence duration* (the time between an entity’s first mention in the text stream, and its subsequent incorporation into Wikipedia), or *emergence volume* (the number of documents that mention the entity before it is incorporated into Wikipedia) disappear through our visualization method. To study these aspects, we describe the time series groups using different features that reflect the emergence behavior of the group. For an overview of the descriptive statistics that we consider, see Table 2.

Results

In this section we present the analyses that answer our research questions.

RQ1: Emergence Patterns

Figure 3 shows a cluster tree that results from clustering the time series distance matrix. At its highest level, the tree shows two distinct clusters, each of which is broken down into multiple smaller sub-clusters. In the following section, we study the *global emergence patterns*, by taking all time series at the root of the tree (*Top level* in Figure 3), and next, the two main clusters (*Level 1* in Figure 3).

Global Emergence Pattern

Figure 4 shows how both the emerging entities’ introduction into public discourse (the first mention at the left-most

TABLE 2. Descriptive statistics used for analyzing and comparing different groupings of emerging entities. We distinguish between *time series statistics* (top) and *burst statistics* (bottom). All statistics are computed for the period ranging from the emerging entity’s first mention in the corpus to the creation date of the Wikipedia page devoted to it.

Emergence volume	Number of documents that mention the entity
Emergence duration	Number of days from first mention to incorporation
Emergence velocity	$\frac{\text{Volume}}{\text{Duration}}$ (average number of documents per day)
Bursts number	Total number of bursts
Bursts duration	Normalized average durations of bursts (i.e., bursts widths)
Bursts value	Normalized average heights of burst (i.e., bursts heights)

side of the plot) and subsequent incorporation into cultural CM (the right-most side of the plot) occur in bursts of documents, that is, overall, the largest number of documents that mention a newly emerging entity are either at the start or at the end of their time series. This can be explained as follows. The entrance into public discourse represents the first emergence of an entity, whereas being added to cultural CM is likewise likely to happen in a period of increased attention, for example, a real world event that puts the entity in public discourse. Between these two bursts, the number of documents that mention the entity seems to increase gradually as time progresses, suggesting that on average, the number of documents that mention a new entity, and thus the attention the entity receives in public discourse increases over time before it reaches “critical mass.”

Turning to the descriptive statistics in Table 3, it takes 245 days on average for an entity to emerge, but with large variations between entities, motivating our clustering approach. On average, an entity is associated with multiple bursts (3.8), indicating that entities are likely to resurface multiple times in public discourse before being deemed important enough to transition into cultural CM.

Clusters at Level 1 in Figure 3: Early Versus Late Bursts

In our first attempt at uncovering distinct patterns in which collective remembrance happens, we study the two

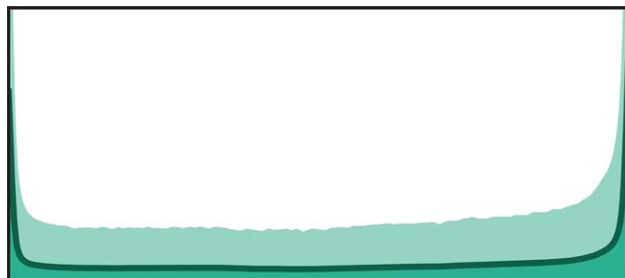


FIG. 4. Global cluster signature (of all emerging entities) where $n = 79,482$. That is, the top level in the dendrogram in Figure 3. The axes are not labeled because all time series values (i.e., document counts) are standardized, and the series are linearly interpolated to have equal length. The solid line represents the cluster signature (i.e., the average time series), the lighter band represents standard deviation. [Color figure can be viewed at wileyonlinelibrary.com]

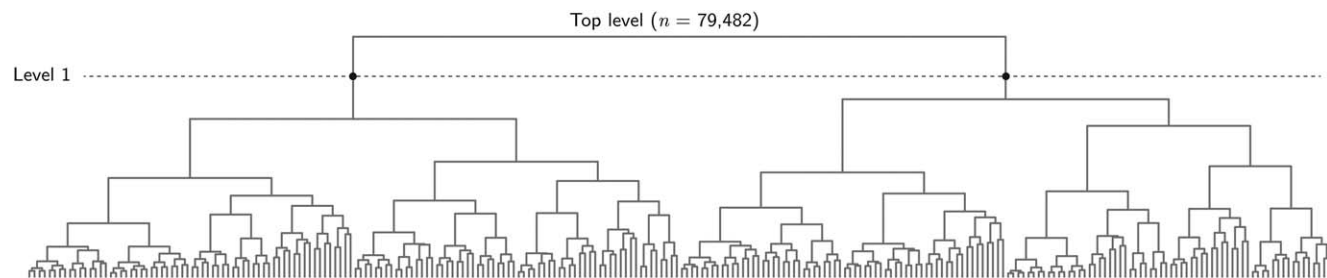


FIG. 3. Dendrogram resulting from applying hierarchical agglomerative clustering using BSim similarity (Vlachos et al., 2004), on our corpus of emerging entity time series ($n = 79,482$). The cutoff-points at which we analyze the clusters are denoted Top level, and Level 1 (2 clusters). For clarity, the tree is truncated by showing no more than 7 levels of the hierarchy.

TABLE 3. Global time series and burst descriptive statistics

Duration (# days)			Volume (# docs)			Velocity (docs/day)		
Mean \pm	Std	Med.	Mean \pm	Std	Med.	Mean \pm	Std	Med.
245 \pm	153	221	183 \pm	1,180	32	0.87 \pm	5.6	0.19
n_bursts			Bursts durations			Bursts values		
3.8 \pm	2.62	3	0.03 \pm	0.03	0.02	0.03 \pm	0.08	0.02

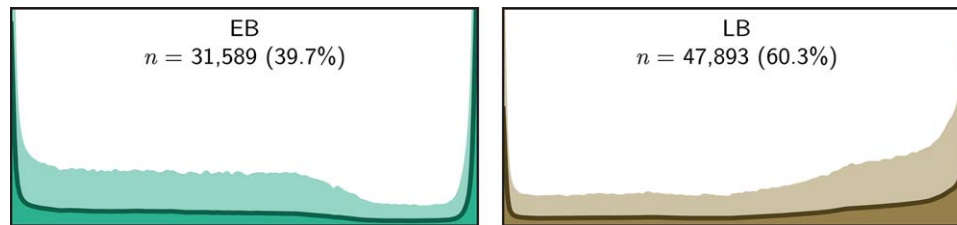


FIG. 5. Cluster signatures of the early bursting entities (left plot) and late bursting entities (right plot) clusters, denoted Level 1 in Figure 3. [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 4. Comparison of early bursting and late bursting entities clusters statistics

	Proportion	Duration (#days)			Volume (#docs)			Velocity (docs/day)		
		Mean \pm	Std	Med.	Mean \pm	Std	Med.	Mean \pm	Std	Med.
EB	0.40	224 \pm	146	195	118 \pm	804	22	0.70 \pm	6.45	0.15
LB	0.60	259 \pm	156	238	225 \pm	1,371	42	0.99 \pm	4.96	0.23
		n_bursts			Burst durations			Burst values		
EB	0.40	3.32 \pm	2.20	3	0.03 \pm	0.03	0.02	0.05 \pm	0.11	0.02
LB	0.60	4.12 \pm	2.82	4	0.03 \pm	0.03	0.02	0.03 \pm	0.05	0.01

main clusters at Level 1 of the cluster tree (Figure 3). The resulting cluster signatures are shown in Figure 5.

Much like the global cluster signatures in the previous section, the Level 1 clusters show two main bursts: the *initial burst* around the first mention, and the *final burst* around the time an entity is added to Wikipedia. However, the left cluster, which we call *early bursting (EB) entities*, is characterized by a stronger initial burst, with the majority of the documents that mention the entity concentrated at the time when the entity surfaces in communicative CM. This suggests that the cluster contains new entities that suddenly emerge and experience a (brief) period of lessened attention, before transitioning into the collective’s CM. The right cluster, which we denote as *late bursting (LB) entities*, shows a more gradual pattern in activity toward the point at which the entity is incorporated into cultural CM, much like we saw in the global signature.

We note two main differences between the group signatures of the EB and LB entities in Figure 5. First, the distribution of documents between the initial and final burst. The EB entities show a more “abrupt” final burst: the majority of the documents are in the wake of the initial burst, that is, at the left-hand side of the plot, then, the document volume gradually winds down, before it finally seems to abruptly transition into the final burst at the right hand-side of the

plot. In contrast, the LB entities cluster shows a relatively subtle initial burst, which likewise quiets down, followed by a gradual increase of document volume that leads up to the final burst.

A second difference is the height difference between the initial and final bursts. The EB cluster shows roughly equally high initial and final bursts; the LB cluster shows a substantially smaller initial burst, which suggests the introduction into public discourse is more subtle than its addition to Wikipedia.

We turn to the clusters’ descriptive statistics in Table 4. We first test for statistical significance in the differences between the cluster statistics. We perform a Kruskal-Wallis one-way analysis of variance test, and follow this omnibus test with a post-hoc test using Dunn’s multiple comparison test (with *p*-values corrected for family-wise errors using Holm-Bonferroni correction). We find that all differences are statistically significant at the $\alpha = 0.05$ level.

Table 4 shows LB entities emerge more slowly (259 days) than EB entities (224 days). LB entities also receive more exposure during emergence (225 vs. 118 documents for EB entities). The shorter emergence duration and lower volumes seen with the EB entities suggest they represent more popular, timely, or “urgent” entities, that will be incorporated quickly after emerging in public discourse, for

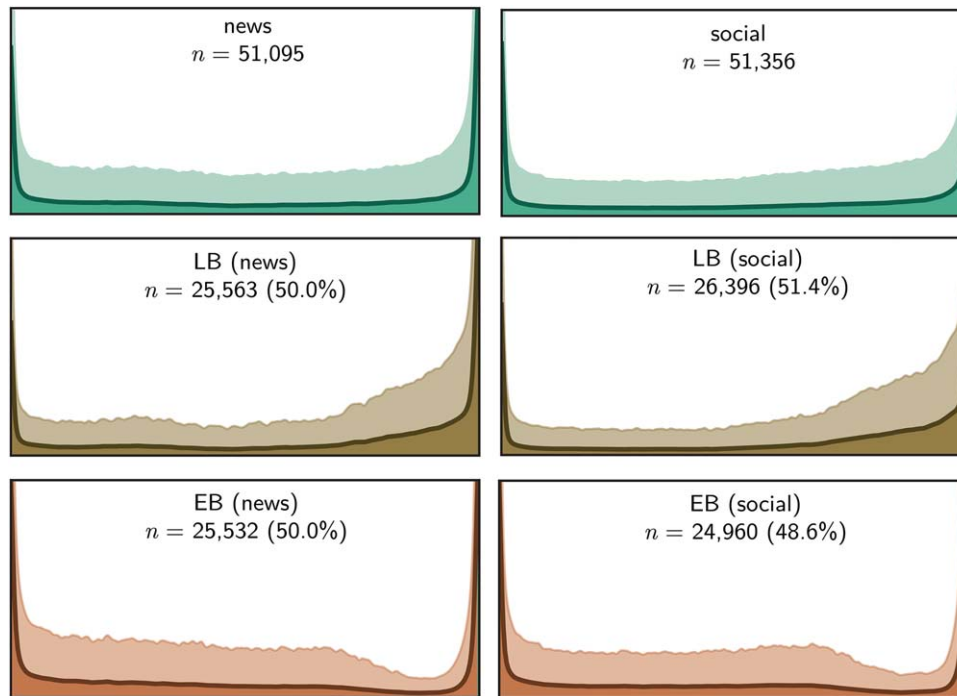


FIG. 6. News vs. Social stream cluster signatures. The top row shows the global cluster signature of the news (left) and social (right) streams. The bottom two rows show the signatures of the late bursting and early bursting entity clusters for each stream (news left, social right). [Color figure can be viewed at wileyonlinelibrary.com]

example, large-scale events and popular entities. The descriptive statistics of the LB entities on the other hand suggests less timely or urgent entities. The burst statistics confirm this view of slower, less timely LB entities, and more urgent, faster EB entities: the average burst heights of EB entities are higher, suggesting LB entities see a more evenly spread volume of documents that mention them. Furthermore, EB entities show fewer bursts (3.22 vs. 4.12, on average).

And indeed, the EB entities that occur most frequently in our dataset include many “central” entities related to popular culture, for example, products such as Xbox One (121,813 mentions), movies, for example, *The Twilight Saga: Breaking Dawn - Part 2* (124,222 mentions), and news events, for example, Disappearance of Lisa Irwin (15,917 mentions). The most frequent LB entities on the other hand include more obscure, long-tail, or niche entities: most notably people, for example, Jeffrey Chiesa (31,560 mentions), Sergio Ermotti (22,274 mentions), and James Rolfe (filmmaker) (15,797 mentions).

Summary

In this section, we have answered our first research question: “Are there common patterns in how entities emerge in online text streams?” We performed hierarchical clustering using a burst similarity-metric of the emerging entity time series and discovered two distinct emergence patterns: early bursting entities and late bursting entities. Our visual inspection of the cluster signatures suggest LB entities emerge

more slowly, that is, build up attention more slowly before transitioning from communicative into cultural CM, whereas EB entities are associated with more sudden and higher bursts of activity, prior to transitioning into cultural CM. We find that the two clusters differ substantially and significantly in their cluster signature and descriptive statistics.

RQ2: Emerging Entities in Social Media and News

In this section we answer our second research question: “Do news and social media text streams exhibit different emergence patterns?” In the previous section we have shown that 79,482 entities emerge in the combined news and social media streams. By splitting out these entities by stream, we find 51,095 of these entities emerge in the news stream (i.e., are mentioned in the news stream), similar to the number of entities that emerge in the social media stream, at 51,356. Finally, 30,148 of the emerging entities are mentioned in both streams before being incorporated into Wikipedia.

Global: News Versus Social

First, we compare the emergence patterns of entities in news and social streams. We apply the same hierarchical clustering method from the previous section on the two subsets of entities that emerge in news and social media streams (where $n_{news} = 51,095$ and $n_{social} = 51,356$).

Figure 6 first shows the global emergence patterns (top row, in green), which are largely the same in the two streams and highly similar to the global patterns studied in the previous section. The bottom two rows of Figure 6 show that

TABLE 5. Emergence features for our five groups of entities: entities that emerge in both streams, but first in the news stream (news-first), entities that emerge in both streams, but first in the social media stream (social-first), entities that emerge in both streams on the same day (same-time), entities that emerge only in the news stream (news-only), and finally, entities that emerge only in the social media stream (social-only)

Stream	Duration (#days)			Volume (#docs)			Velocity (docs/day)		
	Mean \pm	Std	Med.	Mean \pm	Std	Med.	Mean \pm	Std	Med.
news first	298 \pm	139	305	123 \pm	291	53	0.58 \pm	1.59	0.21
social first	281 \pm	157	276	182 \pm	445	74	0.95 \pm	3.22	0.32
same time	197 \pm	147	163	192 \pm	662	67	2.87 \pm	23.59	0.51
only news	250 \pm	152	216	415 \pm	2,215	65	1.45 \pm	6.45	0.35
only social	214 \pm	148	190	33 \pm	134	12	0.41 \pm	2.60	0.08

both streams exhibit groups that are similar to the early bursting and late bursting entities shown in Figure 5. Looking at the top row, however, shows that entities that emerge in news have slightly more of their emergence volume mass after the initial burst, compared to the corresponding pattern of the social media stream, which exhibits a more gradual increase in emergence volume toward the final burst at the right-hand side of the plot. This may be attributed to the slightly higher proportion of early bursting entities in the news stream, which has 50.0% of its entities falling in this cluster, whereas the social media stream has 48.6%.

Who Is First?

Of the 79,482 entities that emerge in the 18 month period our dataset spans, 45,678 appear in both the news and social media stream before they transition to cultural CM; 9,681 entities are mentioned exclusively in the news stream, and never appear in social media (news-only) before transitioning into cultural CM. Finally, 23,096 appear only in the social media stream (social-only). In Table 5, we compare entities as they emerge in different streams.

Of the 45,678 entities that emerge in both streams, the majority appears in the social media stream before they appear in the news stream. This may be explained by the different nature of the publishing cycles of the two streams; whereas news stories need to be checked and edited before being published, social media follows a more unedited and direct publishing cycle.

The entities that appear in social media first (social-first), cover 64.9% ($n = 29,665$) of the entities that emerge in both streams. Interestingly, entities that emerge in news first, subsequently appear in social media streams slower than vice versa: on average 66 days for the former, and 49 days for the latter. A relatively small number of entities is mentioned in both streams on the same day (sametime): 8.7% ($n = 3,967$). Such entities are expected to be more urgent and central, as they appear more widely in public discourse. This group's shortest emergence durations and highest velocities, support this view of entities that play a more central role in public discourse. And indeed, looking at the entities that appear in this set, we see a large number of news events-related entities, for example, 12-12-12: The Concert for Sandy Relief, 2013 Alabama bunker hostage crisis, and Suicide of Jacintha Saldanha.

Summary

News and social media streams show broadly similar emergence patterns for entities but the population and the behavior of entities emerging in news and social differ significantly. Entities are slower on average in emerging in social media streams, and entities that appear in both streams on the same day are the fastest to transition to cultural CM.

RQ3: Emergence Patterns of Different Entity Types

In this section we answer our third research question: "Do different types of entities exhibit different emergence patterns?" We compare the descriptive statistics of different entity types in our dataset, to assert whether different types exhibit different emergence patterns.

Entity Types: Temporal Patterns

First, we study the descriptive statistics per entity type. Table 6 provides an overview of the most frequent entity types in our dataset (i.e., all entity types with a frequency of ≥ 400). We find that the entity type signatures are very similar to the global pattern of Figure 4, which suggests the time series are highly variable within an entity type. See Figure 7 for an example of two common entity types (top row) and two less frequently emerging types (bottom row): whereas the signature becomes smoother as the number of mentions increases, the overall pattern is highly similar across the four types.

Turning to Table 6, we note the null class, that is, entities that are not assigned an entity type in DBpedia exhibit very low emergence volumes (98 documents on average). This may be explained by their nature: long-tail, or unpopular entities are more likely to not have a class assigned in the DBpedia ontology.

Second, we note a group of "fast" emerging entity types with short emergence durations and/or high emergence velocities, for example, DesignedArtifact, CreativeWork, MusicalWork, and VideoGame, consider, for example, the DesignedArtifacts emergence velocity, at 217 days with over 7 documents a day on average. This type includes entities such as devices and products, for example, smartphones, tablets, and laptops. The relatively fast transitioning speed may be explained by their nature: they have short "lifecycles" and may be superseded or replaced at high

TABLE 6. Descriptive statistics per entity type (for types that occur ≥ 400 times in our dataset)

Stream	n_{Samples}	Duration (#days)			Volume (#docs)			Velocity (docs/day)		
		Mean \pm	Std	Med.	Mean \pm	Std	Med.	Mean \pm	Std	Med.
Person	21,295	270 \pm	151	254	243 \pm	692	71	1.03 \pm	3.32	0.32
Athlete	8,018	260 \pm	150	235	264 \pm	674	76	1.05 \pm	2.45	0.37
InformationEntity	7,847	242 \pm	154	210	294 \pm	1,923	90	1.42 \pm	6.53	0.51
CreativeWork	7,795	243 \pm	154	211	294 \pm	1,928	90	1.42 \pm	6.54	0.52
Organization	5,606	279 \pm	153	270	335 \pm	1,812	71	1.40 \pm	14.44	0.31
Place	3,689	274 \pm	149	273	122 \pm	448	33	0.48 \pm	1.61	0.16
Company	2,536	284 \pm	156	275	462 \pm	1,964	108	1.98 \pm	20.88	0.47
MusicalWork	2,474	218 \pm	148	181	170 \pm	533	78	1.13 \pm	2.23	0.49
Movie	2,033	267 \pm	154	247	279 \pm	1,322	87	1.20 \pm	6.57	0.42
OfficeHolder	1,929	287 \pm	158	284	210 \pm	476	73	0.85 \pm	1.88	0.31
MusicGroup	1,649	293 \pm	150	289	221 \pm	393	86	0.95 \pm	1.97	0.34
Artist	1,624	299 \pm	152	302	240 \pm	564	75	0.95 \pm	2.19	0.30
ArchitecturalStructure	1,591	279 \pm	149	284	133 \pm	436	36	0.47 \pm	1.14	0.18
PopulatedPlace	1,521	262 \pm	145	244	119 \pm	481	31	0.53 \pm	2.15	0.15
Building	1,067	281 \pm	150	290	125 \pm	374	34	0.43 \pm	0.98	0.17
TelevisionShow	1,043	229 \pm	158	193	228 \pm	503	87	1.33 \pm	3.05	0.57
WrittenWork	959	267 \pm	147	245	307 \pm	864	88	1.26 \pm	2.99	0.44
EducationalInstitution	915	290 \pm	144	308	108 \pm	564	30	0.41 \pm	2.21	0.12
Software	769	250 \pm	156	221	732 \pm	5,413	192	3.04 \pm	16.33	0.92
School	554	280 \pm	142	305	56 \pm	158	25	0.29 \pm	2.24	0.11
Book	524	272 \pm	147	263	286 \pm	827	93	1.17 \pm	3.03	0.44
VideoGame	505	229 \pm	150	198	381 \pm	657	189	2.08 \pm	3.24	1.03
DesignedArtifact	409	217 \pm	149	187	1,420 \pm	7,142	214	7.04 \pm	20.42	1.39
Infrastructure	403	271 \pm	146	269	127 \pm	560	34	0.47 \pm	1.42	0.17
null	39,807	225 \pm	151	196	98 \pm	861	15	0.58 \pm	3.70	0.10

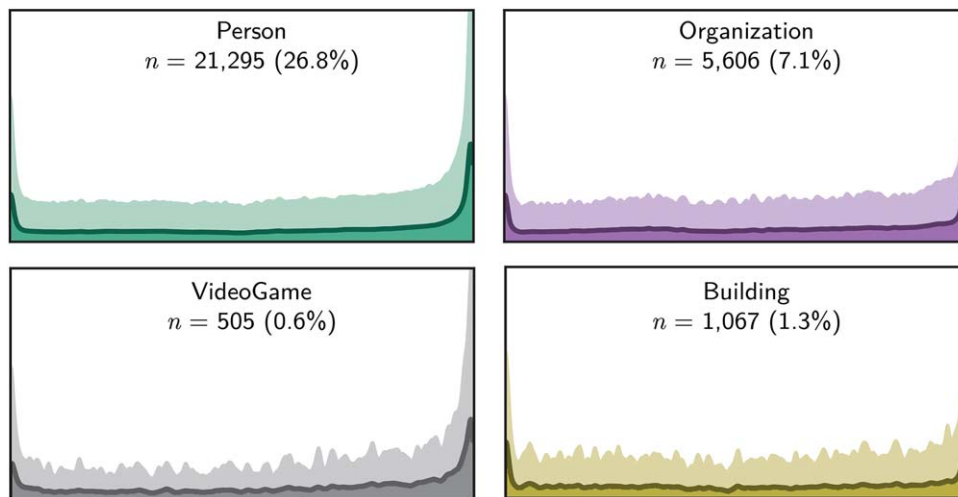


FIG. 7. Type signatures of the Person, Organization, VideoGame and Building types. Even though the number of entities per type differ substantially, the signatures show similar patterns. [Color figure can be viewed at wileyonlinelibrary.com]

frequencies. Consider, for example, the release or announcement of a new smartphone: this event typically generates a lot of attention in a short timeframe, which may result in a fast transition into cultural CM. Similar to the DesignedArtifact-type, CreativeWorks (including, for example, MusicalWork, WrittenWork, Movie) share this characteristic: they play a central but short-lived role in public discourse.

Third, the “slower” entities, that is, those with longer emergence durations and lower emergence volumes, are largely person types such as Writer, Artist, and political

figures (OfficeHolder), but also School and EducationalInstitution, and geographical entities (e.g., Building, ArchitecturalStructure, Place, and PopulatedPlace). These entities may have longer life-cycles and a more gradual “rise to fame” by their nature, and have a less central role in public discourse. Consider, for example, politicians who generally have a long and gradual career and are more likely to first emerge in local media. Similarly, an opening of a new school building may emerge in regional news, but is unlikely to be globally and widely reported.

TABLE 7. “Popularity” (number of pageviews) of each entity in our dataset, aggregated per entity type. Ranked in descending order

Rank	Type	Mean \pm std	Median
1	Movie	98,387 \pm 322,166	14,352
2	TelevisionShow	97,098 \pm 309,172	9,765
3	VideoGame	51,236 \pm 166,802	11,852
4	CreativeWork	50,024 \pm 213,490	5,716
5	InformationEntity	49,704 \pm 212,816	5,634
6	Software	43,657 \pm 149,499	9,582
7	MusicGroup	38,883 \pm 133,336	5,400
8	Artist	35,607 \pm 122,032	4,116
9	DesignedArtifact	29,830 \pm 82,081	7,191
10	Book	18,248 \pm 109,400	3,126
11	WrittenWork	14,227 \pm 86,637	1,801
12	Person	13,772 \pm 77,791	1,568
13	MusicalWork	10,443 \pm 25,009	3,523
14	Athlete	9,415 \pm 41,887	1,545
15	Organization	9,003 \pm 45,140	1,816
16	Company	7,624 \pm 21,371	2,566
17	OfficeHolder	3,763 \pm 16,167	958
18	ArchitecturalStructure	3,189 \pm 16,978	1,042
19	Building	3,180 \pm 20,106	987
20	Infrastructure	2,813 \pm 6,769	1,085
21	Place	2,339 \pm 12,649	827
22	EducationalInstitution	1,799 \pm 3,031	862
23	PopulatedPlace	1,743 \pm 9,081	694
24	School	1,137 \pm 1,426	747

To better understand the difference between “fast” and “slow” entities, we examine the popularity of entities. Table 7 lists the average number of pageviews received per entity in 2015, per type. Looking at the ranking, we note how “faster” emerging entity types remain more popular over time: types that are associated with short emergence durations and high velocities all fall in the top 10 (ranks 3, 4, and 9, for VideoGame, CreativeWork, and DesignedArtifact, respectively), whereas slower types reside in lower ranks in the table, for example, rank 19, 22, and 24 for Building, EducationalInstitution and School, respectively.

Summary

We have shown that different entity types exhibit substantially different emergence patterns, but entities that belong to a particular type show broadly similar emergence patterns. Furthermore, entities with a fast transition from communicative to cultural CM, are more likely to remain popular over time.

Conclusion

In this article we studied entities as they transition from communicative into cultural collective memory. We did so by studying a large set of time series of mentions of entities in online news streams before transitioning into cultural CM (as represented by the creation of a Wikipedia page). We studied implicit groups of similarly emerging entities by applying a burst-based agglomerative hierarchical clustering method and explicit groups by isolating entities by whether they emerge in news or social media streams.

Findings

We found that, globally, entities have a long time span between surfacing in communicative CM and transitioning into cultural CM. During this time span, an entity may emerge with multiple bursts, however both the entities’ introduction into public discourse, and subsequent transitioning into cultural CM occur in the largest document bursts. Emergence statistics show large standard deviations, indicating that they differ substantially between entities. For this reason, we turned to time series clustering to uncover distinct groups of entities. We discovered two emergence patterns: early bursting (EB) entities and late bursting (LB) entities. Analysis suggests that EB entities comprise mostly “head” or popular entities; they exhibit fewer and higher bursts, with shorter emergence durations and lower emergence volumes. The LB entities emerge more slowly, and witness a more gradual increase of exposure before transitioning into cultural CM. The emergence patterns we visualized differ substantially from the global average and from, for example, the type signatures shown in Figure 7, suggesting that the entities in each of the underlying clusters exhibit substantially different and distinct emergence patterns from entities in the other clusters.

We showed that entities emerging in news and social media streams display very similar emergence patterns, but that on average, entities that emerge in social media take longer to be incorporated into cultural CM. We hypothesize that this can be attributed to the nature of the underlying sources. News media are more mainstream and professional, with a larger audience, reach, and authority, than social media. Our findings are in line with those of Petrovic, Osborne, Mccreadie, Macdonald, and Ounis (2013), who compare breaking news on traditional media with that on social media. Their findings suggest reported events overlap largely between both media, however, social media exhibits a long tail of minor events, which may explain the longer uptake on average. Leskovec, Backstrom, and Kleinberg (2009) find that the “attention span” for news events on social media both increases and decays at a slower rate than for traditional news sources, which additionally supports our observations of the slower uptake on social media.

Finally, we showed that different entity types exhibit substantially different patterns, but entities of similar types show similar patterns. Some entity types, for example, devices or creative works, transition faster from communicative to cultural CM, than entities such as buildings, locations, and people. At the same time, the former “fast” entity types remain more popular over time.

One aspect that distinguishes between “fast” and “slow” entity types, is that the former are more likely to appear in so-called “soft news” that covers sensational or human-interest events and topics (e.g., news related to celebrities and cultural artifacts). The slower entity types on the other hand, are more likely to appear in more substantive “hard news” that encompasses more urgent events and topics (e.g., political elections) Tuchman, 1972). Granka (2010) studied

the differences in “attention span” of the public and the traditional news media for “hard” and “soft news,” and found that hard news is associated with a relatively shorter period of public attention. Soft news exhibits a slower decrease of the public’s attention, which supports our finding that faster entity types (more likely associated with soft news) tend to remain more popular over time.

As emerging entities are not “born equal.” The patterns and circumstances under which an entity transitions from communicative to cultural CM differ depending on source and type.

Implications

Our findings have implications for designing systems to detect emerging entities, and more generally for studying and understanding how collective memories are formed. We show that entities are likely to resurface multiple times in public discourse before transitioning into cultural CM. This suggests that monitoring bursts of new entities could prove effective for predicting the formation of collective memories. Furthermore, we show that the type of stream in which entities emerge shows different patterns. This suggests that taking the different nature of streams into account can be beneficial for predicting emerging entities. Finally, we show that different types of entity exhibit different emergence patterns, suggesting the underlying entity type could likewise prove valuable in predicting emerging entities.

Limitations

Part of our findings are derived from an unsupervised clustering method. Interpreting cluster signatures is a subjective matter, and clustering is a difficult task to evaluate (Von Luxburg, Williamson, & Guyon, 2012). In our defense, the clustering’s dendrogram suggests the presence of distinct and meaningful groups, as the structure of the dendrogram shows symmetry and clear separations. More importantly, the cluster signatures yielded visually discernible, and different patterns between clusters, which was not the case for the signatures of other grouping strategies in this paper (see, e.g., Figure 7).

The fragmented nature of the source of our dataset (TREC-KBA StreamCorpus 2014) means that coverage, and hence representativeness of the data cannot be guaranteed. Popular social media channels such as Tumblr, Twitter and Facebook are not part of the dataset, there may be a sampling bias in the sources that represent the streams, resulting in a similar bias in the entities. Different sources may well yield different findings. This is unavoidable.

Another limitation relates to the entity annotations used as a starting point in this paper: they cannot be assumed to be 100% accurate. So-called “cascading errors” (Finkel, Manning, & Ng, 2006) in NLP pipelines cause the accuracy of downstream tasks to suffer, in our case having imperfect tags (named entities) for imperfect tagging (entity linking). The FAKBA1 annotations are estimated (from manual inspection) to contain around 9% incorrectly linked entities,

with around 8% of SERIF mentions being wrongfully not linked. Even more so, the “difficult” entities are long-tail entities, which are more likely to be part of our filtered set. However, manually correcting the annotations was beyond the scope of this study.

Finally, there may be a cultural bias inherent in our choice of datasets: we used English language news sources and social media as well as the English version of Wikipedia. One could object that we studied the birth of collective memories for the *English-speaking* part of the world, and that different datasets may also yield different findings. It is unfortunate that the English-speaking part of the world is disproportionately represented in our field of research, as is witnessed by the biggest constraint in conducting this study: dataset availability. We invite the community to create suitable datasets in other languages or reflecting cultural practices in other parts of the planet so as to enable comparative studies.

Future Work

As a next step, we should take a closer look at the circumstances in which entities emerge, by not exclusively considering in how many documents they appear over time, but also in which contexts, for example, by looking at the content of the articles themselves. Furthermore, in this paper we have chosen to restrict ourselves to the entities that transition and remain in cultural CM. Another interesting aspect of CM, is the notion of “consensus.” For example, one could study the emergence patterns of entities that are removed from cultural CM after transitioning. Finally, the observations made in this paper could be explored in a prediction task, where, for example, given a partial entity time series, the task would be to predict the point at which the entity transitions from communicative to cultural CM.

Acknowledgments

We thank our reviewers for their helpful comments that helped to improve the paper.

This research was supported by Ahold Delhaize, Amsterdam Data Science, Blendle, the Bloomberg Research Grant program, the Dutch national program COMMIT, Elsevier, the European Community’s Seventh Framework Programme (FP7/2007–2013) under grant agreement nr 312827 (VOX-Pol), the Microsoft Research Ph.D. program, the Netherlands Institute for Sound and Vision, the Netherlands Organisation for Scientific Research (NWO) under project nrs 727.011.005, 612.001.116, HOR-11-10, CI-14–25, 652.002.001, 612.001.551, 652.001.003, 314-98-071, the Yahoo! Faculty Research and Engagement Program, and Yandex. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- Assmann, J., & Czaplicka, J. (1995). Collective memory and cultural identity. *New German Critique*, 65(Spring - Summer), 125–133.
- Au Yeung, C.-M., & Jatowt, A. (2011). Studying how the past is remembered: Towards computational history through large scale text mining. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management* (pp. 1231–1240). New York, NY, USA: ACM.
- Berndt, D. J., & Clifford, J. (1994). Using dynamic time warping to find patterns in time series. In *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining* (pp. 359–370). Seattle, WA, USA: AAAI Press.
- Boschee, E., Weischedel, R., & Zamanian, A. (2005). Automatic information extraction. In *Conference on Intelligence Analysis*. <https://cryptome.org/intel-2005.htm>.
- Brown, N. R. (1990, September). Organization of public events in long-term memory. *Journal of Experimental Psychology: General*, 119(3), 297–314.
- Campos, R., Dias, G., & Jorge, A. M. (2011). What is the temporal value of web snippets? In *Www 2011 workshop on temporal web analytics* (pp. 9–16). CEUR Workshop Proceedings 813, CEUR-WS.org.
- Dalton, J., Frank, J. R., Gabrilovich, E., Ringgaard, M., & Subramanya, A. (January 2015). Fakba1: Freebase annotation of TREC KBA Stream Corpus, Version 1 (Release date 2015-01-26, Format version 1, Correction level 0). <http://trec-kba.org/data/fakba1/>.
- Erlil, A., Nünning, A., & Young, S. (Eds.). (2008). *Cultural memory studies, an international and interdisciplinary handbook*. Berlin, Boston: De Gruyter.
- Färber, M., Rettinger, A., & El Asmar, B. (2016). On emerging entity detection. In E. Blomqvist, P. Ciancarini, F. Poggi, & F. Vitali (Eds.), *Proceedings of the 20th International Conference on Knowledge Engineering and Knowledge Management* (pp. 223–238). Cham, Switzerland: Springer International Publishing.
- Ferron, M. (2012). *Collective memories in Wikipedia* (Unpublished doctoral dissertation). Doctoral School in Cognitive and Brain Sciences, Università degli Studi di Trento.
- Ferron, M., & Massa, P. (2011). Collective memory building in Wikipedia: the case of North African uprisings. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*. Mountain View, CA. New York, NY, USA: ACM.
- Finkel, J. R., Manning, C. D., & Ng, A. Y. (2006). Solving the problem of cascading errors: Approximate bayesian inference for linguistic annotation pipelines. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing* (pp. 618–626). Stroudsburg, PA, USA: Association for Computational Linguistics.
- García-Gavilanes, R., Mollgaard, A., Tsvetkova, M., & Yasseri, T. (2017). The memory remains: Understanding collective memory in the digital age. *Science Advances*, 3(4). <http://advances.sciencemag.org/content/3/4/e1602368>.
- Granka, L. A. (2010). Measuring agenda setting with online search traffic: Influences of online and traditional media. Presented at the 2010 Annual Meeting of the American Political Science Association. Available at SSRN: <https://ssrn.com/abstract=1658172>.
- Graus, D., Tsagkias, M., Buitinck, L., & de Rijke, M. (2014, April). Generating pseudo-ground truth for predicting new concepts in social streams. In *36th European Conference on Information Retrieval*. Cham, Switzerland: Springer.
- Graus, D., Tsagkias, M., Weerkamp, W., Meij, E., & de Rijke, M. (2016, February). Learning dynamic collective entity representations for entity ranking. In *the 9th International Conference on Web Search and Data Mining* (pp. 595–604). New York: ACM.
- Halbwachs, M. (1925). *Les cadres sociaux de la mémoire*. Paris: Félix Alcan.
- Halbwachs, M. (1950). *La mémoire collective*. Paris, France: Albin Michel.
- Hirst, W., & Manier, D. (2008). Towards a psychology of collective memory. *Memory*, 16(3), 183–200.
- Hoffart, J., Altun, Y., & Weikum, G. (2014). Discovering emerging entities with ambiguous names. In *Proceedings of the 23rd International Conference on World Wide Web* (pp. 385–396). New York: ACM.
- Jatowt, A., Antoine, E., Kawai, Y., & Akiyama, T. (2015). Mapping temporal horizons: Analysis of collective future and past related attention in twitter. In *Proceedings of the 24th International Conference on World Wide Web* (pp. 484–494). Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee.
- Kampf, M., Tessenow, E., Kenett, D. Y., & Kantelhardt, J. W. (2016, December). The detection of emerging trends using Wikipedia traffic data and context networks. *PLoS ONE*, 10(12), 1–19.
- Keegan, B. (2011). Breaking news, breaking planes, and breaking hearts: Psycholinguistics and sensemaking in collaborative accounts of catastrophe. In *International Communication Association*. Boston.
- Keegan, B., Gergle, D., & Contractor, N. (2013). Hot off the wiki: Structures and dynamics of wikipedia’s coverage of breaking news events. *American Behavioral Scientist*, 57(5), 595–622.
- Leskovec, J., Backstrom, L., & Kleinberg, J. (2009). Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 497–506). New York: ACM.
- Liao, T. W. (2005). Clustering of time series data—a survey. *Pattern Recognition*, 38(11), 1857–1874.
- Lin, T., Mausam, & Etzioni, O. (2012). No noun phrase left behind: Detecting and typing unlinkable entities. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning* (pp. 893–903). Stroudsburg, PA: Association for Computational Linguistics.
- Lipsitz, G. (2001). *Time passages: Collective memory and american popular culture*. Minneapolis: University of Minnesota Press.
- Luyt, B. (2016). Wikipedia, collective memory, and the Vietnam war. *Journal of the Association for Information Science and Technology*, 67(8), 1956–1961.
- Müllner, D. (2013). fastcluster: Fast hierarchical, agglomerative clustering routines for r and python. *Journal of Statistical Software*, 53(1), 1–18.
- Nakashole, N., Tylenda, T., & Weikum, G. (2013). Fine-grained semantic typing of emerging entities. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics* (pp. 1488–1497). Sofia, Bulgaria: ACL.
- Neal, A. G. (1998). National trauma and collective memory: Major events in the American century. Armonk, N.Y.: M.E. Sharpe.
- Olick, J. K., Vinitzky-Seroussi, V., & Levy, D. (Eds.). (2011). *The collective memory reader*. Oxford University Press.
- Pentzold, C. (2009). Fixing the floating gap: The online encyclopaedia Wikipedia as a global memory place. *Memory Studies*, 2(2), 255–272.
- Petrovic, S., Osborne, M., McCreddie, R., Macdonald, C., & Ounis, I. (2013). Can twitter replace newswire for breaking news? In *Proceedings of the Seventh International Conference on Weblogs and Social Media (ICWSM 2013)*. AAAI Press.
- Rani, S., & Sikka, G. (2012, August). Recent techniques of clustering of time series data: A survey. *International Journal of Computer Applications*, 52(15), 1–9.
- Reese, E., & Fivush, R. (2008). The development of collective remembering. *Memory*, 16(3), 201–212.
- Reinanda, R., Meij, E., & de Rijke, M. (2016, October). Document filtering for long-tail entities. In *Proceedings of the 25th ACM Conference on Information and Knowledge Management*. New York: ACM.
- Suh, B., Convertino, G., Chi, E. H., & Pirolli, P. (2009). The singularity is not near: Slowing growth of Wikipedia. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration* (pp. 8:1–8:10). New York: ACM.
- Tierney, T. (2013). *The public space of social media: Connected cultures of the network society*. Abingdon, United Kingdom: Taylor & Francis.

- Tuchman, G. (1972). Objectivity as strategic ritual: An examination of newsmen's notions of objectivity. *American Journal of Sociology*, 77(4), 660–679.
- Vlachos, M., Meek, C., Vagena, Z., & Gunopoulos, D. (2004). Identifying similarities, periodicities and bursts for online search queries. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data* (pp. 131–142). New York: ACM.
- Von Luxburg, U., Williamson, R. C., & Guyon, I. (2012). Clustering: Science or art? In *Unsupervised and transfer learning workshop* (Vol. 27, pp. 65–80).
- Voskarides, N., Odijk, D., Tsagkias, M., Weerkamp, W., & de Rijke, M. (2014). Query-dependent contextualization of streaming data. In *Proceedings of the 36th European Conference on Information Retrieval*. Cham, Switzerland: Springer.
- Wallace, D. P., & Van Fleet, C. (2005). From the editors: The democratization of information? Wikipedia as a reference resource. *Reference & User Services Quarterly*, 45(2), 100–103.
- Ward Jr., J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301), 236–244.
- Wessel, I., & Moulds, M. L. (2008). Collective memory: A perspective from (experimental) clinical psychology. *Memory*, 16(3), 288–304.
- Yasseri, T., Sumi, R., & Kertész, J. (2012, January). Circadian patterns of Wikipedia editorial activity: A demographic analysis. *PLoS ONE*, 7(1), 1–8.
- Yasseri, T., Sumi, R., Rung, A., Kornai, A., & Kertész, J. (2012, June). Dynamics of conflicts in Wikipedia. *PLoS ONE*, 7(6), 1–12.