

Inheritance Patterns in Citation Networks Reveal Scientific Memes

Tobias Kuhn,^{1,*} Matjaž Perc,^{2,3} and Dirk Helbing^{1,4}

¹*ETH Zurich, Clausiusstrasse 50, 8092 Zurich, Switzerland*

²*Faculty of Natural Sciences and Mathematics, University of Maribor,
Koroška cesta 160, SI-2000 Maribor, Slovenia*

³*CAMTP—Center for Applied Mathematics and Theoretical Physics, University of Maribor,
Krekova 2, SI-2000 Maribor, Slovenia*

⁴*Risk Center, ETH Zurich, Scheuchzerstrasse 7, 8092 Zurich, Switzerland*

(Received 8 July 2014; revised manuscript received 19 September 2014; published 21 November 2014)

Mememes are the cultural equivalent of genes that spread across human culture by means of imitation. What makes a meme and what distinguishes it from other forms of information, however, is still poorly understood. Our analysis of memes in the scientific literature reveals that they are governed by a surprisingly simple relationship between frequency of occurrence and the degree to which they propagate along the citation graph. We propose a simple formalization of this pattern and validate it with data from close to 50 million publication records from the Web of Science, PubMed Central, and the American Physical Society. Evaluations relying on human annotators, citation network randomizations, and comparisons with several alternative approaches confirm that our formula is accurate and effective, without a dependence on linguistic or ontological knowledge and without the application of arbitrary thresholds or filters.

DOI: [10.1103/PhysRevX.4.041036](https://doi.org/10.1103/PhysRevX.4.041036)

Subject Areas: Complex Systems,
Interdisciplinary Physics,
Statistical Physics

I. INTRODUCTION

The evaluation of scientific output and the study of patterns of scientific collaboration have received increasing attention by researchers. From citation distributions [1,2], coauthorship networks [3], and the formation of research teams [4,5], to the ranking of researchers [6–8] and the quantification and prediction of scientific success [9,10], how we do science has become a science in its own right. While the famous works of de Solla Price [11] and Merton [12] from the mid 1960s marked the beginning of a popular and long-lasting research field, the rapid progress made in recent years is largely due to the increasing availability of vast amounts of digitized data. Massive publication and citation databases, also referred to as “metaknowledge” [13], along with leaps in progress in the theory and modeling of complex systems, fuel large-scale explorations of human culture that were unimaginable even a decade ago [14]. The “science of science” is scaling up massively as well, with studies on global citation and collaboration networks [15], the “scientific food web” [16], and phylogenetic patterns in the evolution of science [17],

culminating in the visually compelling atlases of science [18] and knowledge [19].

Science is a key pillar of modern human culture, and the general concept of memes has proved to be very insightful for the study of culture. The term “meme” was coined by Dawkins in his book *The Selfish Gene* [20], where he argues that cultural entities such as words, melodies, recipes, and ideas evolve similarly as genes, involving replication and mutation but using human culture instead of the gene pool as their medium of propagation. Recent research on memes has enhanced our understanding of the dynamics of the news cycle [21], the tracking of information epidemics in the blog space [22], and the political polarization on Twitter [23]. It has been shown that the evolution of memes can be exploited effectively for inferring networks of diffusion and influence [24], and that information contained in memes is evolving as it is being processed collectively in online social media [25]. The question of how memes compete with each other for the limited and fluctuating resource of user attention has also amassed the attention of scientists, demonstrating that social network structure is crucial to understand the diversity of memes [26], which suggests that social contagion mechanisms [27] play an important role. It has also been shown that the competition among memes can bring the network to the brink of criticality [28], where even minute disturbances can lead to avalanches of events that make a certain meme go viral [29].

*tokuhn@ethz.ch

Published by the American Physical Society under the terms of the *Creative Commons Attribution 3.0 License*. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

While the study of memes in mass media and popular culture has been based primarily on their aggregated bursty occurrence patterns, the citation network of scientific literature allows for more sophisticated and fine-grained analyses. Quantum, fission, graphene, self-organized criticality, and traffic flow are examples of well-known memes from the field of physics, but what exactly makes such memes different from other words and phrases found in the scientific literature? As an answer to this question, we propose the following definition that is a modified version of Dawkins' definition of the word "gene" [20]: A scientific meme is a short unit of text in a publication that is replicated in citing publications and thereby distributed around in many copies; the more likely a certain sequence of words is to be broken apart, altered, or simply not present in citing publications, the less it qualifies to be called a meme. Publications that reproduce words or phrases from cited publications are thus the analogue to offspring organisms that inherit genes from their parents. In contrast to existing work on scientific memes, our approach is therefore grounded in the "inheritance mechanisms" of memes and not just their accumulated frequencies. The above definition covers memes made up of exact words and phrases, but the same methods apply just as well to more abstract forms of memes.

For our analyses, we rely on 47.1 million publication records from three sources. Because of its representative long-term coverage of a specific field of research, we focus mainly on the titles and abstracts from the data set of the American Physical Society (APS), consisting of almost half a million publications from the Physical Review journals

published between July 1893 and December 2009. We also present results for the over 46 million publications indexed by the comprehensive Web of Science database, and for the over 0.6 million publications from the open access subset of PubMed Central that covers research mostly from the biomedical domain and mostly from recent years.

Figure 1 shows visualizations of these citation graphs. The leftmost network depicts the entire giant component of the citation graph of the Web of Science, consisting of more than 33 million publications. Different scientific disciplines form relatively compact communities: The physical sciences (cyan) are close to engineering and technology (magenta) in the top right-hand corner of the network, but rather far from the social sciences and humanities (green) as well as the medical and health sciences (red), which take up the majority of the left-hand side of the network, with the natural and agricultural sciences in between (blue). Zooming in on the physical sciences and switching to the data set from the American Physical Society, we get the picture shown in the middle. The colors now encode the five most important special-focus journals of Physical Review, each covering a particular subfield of physics (general coverage and smaller journals are shown in gray). We see a complex structure with many small and large clusters. Importantly, even though the employed layout algorithm [30] does not take the scientific disciplines and the journal information explicitly into account, the different communities can be clearly inferred in the citation graphs. Following our general meme-centric perspective, the rightmost network highlights the above-mentioned memes from physics, which mostly appear in

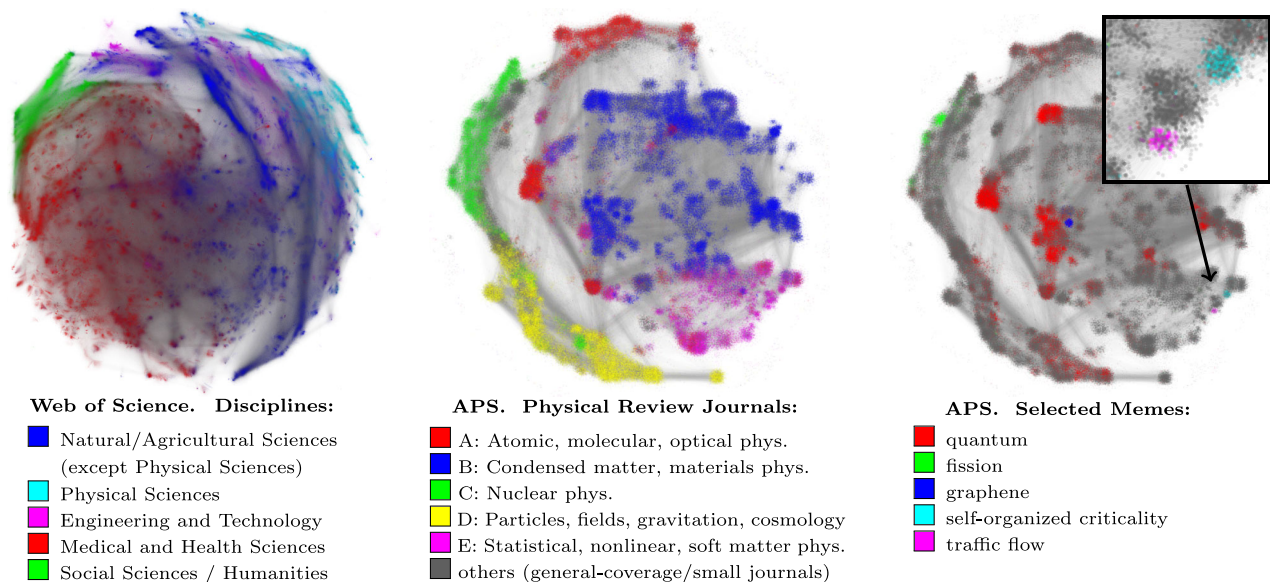


FIG. 1. Citation networks of the Web of Science and the American Physical Society (APS) data sets reveal community structures that nicely align with scientific disciplines, journals covering particular subfields, and occurrences of memes. The generation of the visualizations is based on Gephi [31] and the OpenOrd plug-in [30], which implements a force-directed layout algorithm that is able to handle very large graphs.

publications that form compact communities in the citation graph. The meme *quantum* is widely but by no means uniformly distributed, pervading several large clusters. Publications containing the meme *fission* form a few connected clusters limited to an area that makes up the journal Physical Review C covering nuclear physics. Similarly, the memes *graphene*, *self-organized criticality*, and *traffic flow* (see enlarged area) are each concentrated in their own medium-sized or small communities.

II. RESULTS

All words and phrases that occur frequently in the literature can be considered important memes, but many frequent words like “method” are not particularly interesting for any given scientific field. To quantify the degree to which a meme is interesting, we define the propagation score P_m , which determines the alignment of the occurrences of a given meme with the citation graph. P_m is high for memes that frequently appear in publications that cite meme-carrying publications (“sticking”) but rarely appear in publications that do not cite a publication that already contains the meme (“sparking”). Formally, we define the propagation score for a given meme m as its *sticking factor* divided by its *sparking factor*. The sticking factor quantifies the degree to which a meme replicates in a publication that cites a meme-carrying publication. Concretely, it is defined as $d_{m \rightarrow m}/d_{\rightarrow m}$, where $d_{m \rightarrow m}$ is the number of publications that carry the meme and cite at least one publication carrying the meme, while $d_{\rightarrow m}$ is the number of all publications (meme carrying or not) that cite at least one publication that carries the meme. Similarly, the sparking factor quantifies how often a meme appears in a publication without being present in any of the cited publications. It is thus defined as $d_{m \rightarrow \bar{m}}/d_{\rightarrow \bar{m}}$, where $d_{m \rightarrow \bar{m}}$ is the number of meme-carrying publications that do *not* cite publications that carry the meme, and $d_{\rightarrow \bar{m}}$ is the number of *all* publications (meme carrying or not) that do *not* cite meme-carrying publications. For the propagation score P_m , we thus obtain

$$P_m = \frac{d_{m \rightarrow m}}{d_{\rightarrow m}} \bigg/ \frac{d_{m \rightarrow \bar{m}}}{d_{\rightarrow \bar{m}}}. \quad (1)$$

Based on the propagation score P_m and the frequency of occurrence f_m (which is simply the ratio of publications carrying the meme) of a particular meme m , we define the *meme score* M_m as

$$M_m = f_m P_m. \quad (2)$$

The propagation score, as defined in Eq. (1), can be improved by adding a small amount of controlled noise δ , thus obtaining

$$P_m = \frac{d_{m \rightarrow m}}{d_{\rightarrow m} + \delta} \bigg/ \frac{d_{m \rightarrow \bar{m}} + \delta}{d_{\rightarrow \bar{m}} + \delta}. \quad (3)$$

This corrects for the fact that any of the four basic terms can be zero, and it also prevents phrases with a very low frequency from getting a high score by chance. The controlled noise corresponds to δ fictitious publications that carry all memes and cite none, plus another δ publications that carry no memes and cite all. This decreases the sticking factors and increases the sparking factors of all memes, thereby reducing all meme scores—very slightly so for frequent memes but heavily for rare memes. Our tests show that a small value of δ (e.g., $\delta = 3$ as used throughout this work unless stated otherwise) is sufficient. Another matter that deserves attention is the potential “free-riding” of shorter memes on longer ones. For example, the multitoken meme “the littlest Higgs model” contains the specific token “littlest” that rarely occurs otherwise. The meme “littlest” therefore gets about the same propagation score as the long meme, yet the larger meme is clearly more interesting. This can be addressed by discounting for free-riding by redefining the term $d_{m \rightarrow m}$ in Eq. (1) to exclude publications where the given meme appears in the publication and its cited publications only within the same larger meme. If “littlest,” for example, is always followed by “Higgs” in a given publication and all its cited publications, then this publication shall not contribute to the $d_{m \rightarrow m}$ term for $m = \text{littlest}$.

The meme score considers whether a meme is important (f_m) and whether it is interesting (P_m), and it has additionally a number of desirable properties: (i) it can be calculated exactly without the introduction of arbitrary thresholds, such as a minimum number of occurrences, without limiting the length of n -grams to consider, and without filtering out words containing special characters; (ii) it does not depend on external resources, such as dictionaries or other linguistic data; (iii) it does not depend on filters, like stop-word lists, to remove the most common words and phrases; and (iv) it is very simple with only one parameter (δ).

Calculating the meme score for all n -grams in the three data sets considered gives us the results presented in Fig. 2. Their relative frequency and their propagation score are plotted against each other in the form of heat maps with logarithmic scales. There is no upper limit to the length of n -grams, and the presented maps cover without exception all n -grams with a nonzero meme score (N being the number of such n -grams). Meme scores are increasing towards the top right and decreasing towards the bottom left corner. Figures 2(a), 2(c), and 2(d) feature a broad band with a downward slope, indicating that, in general, more frequent memes tend to propagate less via the citation graph. In the lower half of each figure, we see a wedge of very high densities that follows the larger band on the bottom left edge, but getting narrower towards the middle

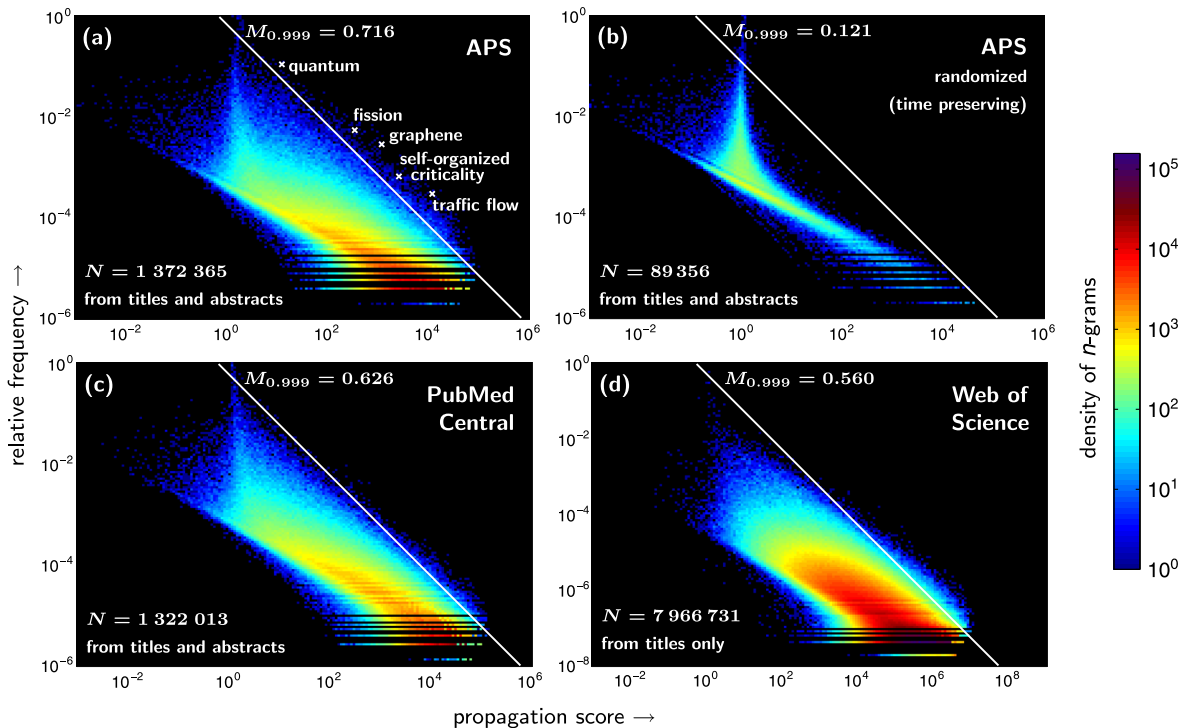


FIG. 2. Universality in the distribution patterns of scientific memes across data sets. Heat maps encode the density of all n -grams with $M \neq 0$ (N being the number of such n -grams) with respect to their propagation score and frequency. Panels (a), (c), and (d) each show a broad band with a downward slope for the data sets from the APS, the open access subset of PubMed Central, and the Web of Science, respectively. The 99.9% quantile with respect to the meme score distribution ($M_{0.999}$) is depicted as a white line. Memes are located mostly around the very edge of the top right-hand side of the band (in the vicinity of the 99.9% quantile line). Panel (b) shows the results obtained with a time-preserving randomization of the APS citation graph (see Sec. IV for details).

where it ends. Though this wedge has a somewhat rounder and broader shape for the Web of Science database, overall these patterns look remarkably similar across all data sets despite their differences with respect to topic, coverage, and size. This is an indication of universality in the distribution patterns of scientific memes. The 99.9% quantile line ($M_{0.999}$) is also surprisingly stable, considering that the underlying values range over 5 orders of magnitude or more. Localizing the previously mentioned physics memes in the APS data set Fig. 2(a), we see that they are located on the very edge of the top right side of the band, where the density of n -grams is very low. (Very frequent words like “of” or “the” are found in the faint spike at the top of the plot where $P \approx 1$ and the frequency is close to 100%.)

The heat map in Fig. 2(b) illustrates a typical case of what happens when the APS citation graph is randomized but the time ordering of publications is preserved. The number of terms with a nonzero meme score decreases dramatically [from approximately 1.4 million in Fig. 2(a) to just 89 356 in Fig. 2(b)], the universal distribution pattern of scientific memes vanishes, and the top right part, where the top-ranked memes should be located, disappears completely. Naturally, if the APS citation graph is randomized without preserving the time ordering, the overlap with the original results presented in Fig. 2(a) is even smaller (see Supplemental Material [32]). Statistical

analysis reveals that median values of the meme score obtained with the randomized networks differ by more than 1 order of magnitude from those obtained with the original citation graph, with very little variation between different randomization runs. These results show that topology and time structure alone fail to account for the reported universality in the distribution patterns, and thus that the top memes get their high meme scores based on intricate processes and conventions that underlie the dynamics of scientific progress and the way credit is given to previous work.

Table I shows the 50 top-ranked memes from the APS data set, also indicating their agreement with human annotation and whether they can be found under a subcategory of physics in Wikipedia. Most of these memes are noun phrases denoting real and reasonable physics concepts, which is remarkable given that the computation of the simple meme-score formula uses no linguistic or ontological knowledge whatsoever. The dominance of noun phrases is consistent with the finding that (scientific) concepts are typically captured by noun phrases when represented as key words in terminologies [33,34]. The extracted memes consist of one, two, or three words, which indicates that the meme score does not favor short or long phrases, again without applying explicit measures to balance n -gram lengths. A further observation is that

TABLE I. Top 50 memes with respect to their meme score from the APS data set. The symbol + indicates memes where the human annotators agree that this is an interesting and important physics concept, while the symbol * indicates memes that are also found on the list of memes extracted from Wikipedia (see Sec. IV for details).

1. Loop quantum cosmology ^{++*}	14. Strange nonchaotic	27. Na _x CoO ₂ ⁺	38. Inspirals*
2. Unparticle ^{++*}	15. In NbSe ₃	28. The unparticle ⁺	39. Spin Hall effect ^{++*}
3. Sonoluminescence ^{++*}	16. Spin Hall ⁺	29. Black	40. PAMELA
4. MgB ₂ ⁺	17. Elliptic flow ^{++*}	30. Electromagnetically induced transparency ^{++*}	41. BaFe ₂ As ₂ ⁺
5. Stochastic resonance ^{++*}	18. Quantum Hall ^{++*}	31. Light-induced drift ⁺	42. Quantum dots ^{++*}
6. Carbon nanotubes ^{++*}	19. CeCoIn ₅ ⁺	32. Proton-proton bremsstrahlung ⁺	43. Bose-Einstein condensates ⁺
7. NbSe ₃ ⁺	20. Inflation ⁺	33. Antisymmetrized molecular dynamics ⁺	44. X(3872)*
8. Black hole ^{++*}	21. Exchange bias ^{++*}	34. Radiative muon capture ⁺	45. Relaxor ⁺
9. Nanotubes ⁺	22. Sr ₂ RuO ₄ ⁺	35. Bose-Einstein ⁺	46. Blue phases ⁺
10. Lattice Boltzmann ^{++*}	23. Traffic flow ^{++*}	36. C ₆₀ ⁺	47. Black holes ^{++*}
11. Dark energy ^{++*}	24. TiOCl	37. Entanglement ⁺	48. PrOs ₄ Sb ₁₂ ⁺
12. Rashba	25. Key distribution ⁺		49. The Schwinger multichannel method ⁺
13. CuGeO ₃ ⁺	26. Graphene ^{++*}		50. Higgsless ⁺

chemical formulas such as MgB₂ and CuGeO₃ are relatively frequent, which we investigate in more detail below.

In Table II, we present results of a manual annotation of terms identified by meme score as compared to randomly selected terms (both nonweighted and weighted by frequency; see also Supplemental Material [32]). Each of the annotators considers around 86% of the meme-score terms to be important physics concepts, agreeing on this in 81% of the cases. With respect to their linguistic categories, each annotator considers 86% of the meme-score terms to be noun phrases, and the two annotators agree on that for 83% of the terms. The respective values are much lower for the randomly extracted terms. Only 25% (nonweighted) and 19% (weighted) of terms are, in agreement, found to be important physics concepts, and only 33% (nonweighted) and 25% (weighted) are found to be noun phrases. The

reported differences between the meme score and the two random selection methods are highly significant ($p < 10^{-15}$ using Fisher's exact test on the number of agreed classifications). These results confirm that the meme score strongly favors noun phrases and important concepts.

Next, we compare the meme score to a number of possible alternative metrics, as defined in the Sec. IV, and align the identified words and phrases with a ground-truth list of terms extracted from physics-related Wikipedia titles. Figure 3 summarizes the results, showing on the top right that $\approx 70\%$ of the top 10 memes identified by meme score correspond to terms extracted from Wikipedia, as well as $\approx 55\%$ of the top 20, $\approx 40\%$ of the top 50, and $\approx 26\%$ of the top 100. The largest area under the curve A is obtained for a controlled noise level $\delta = 4$, which is highlighted by the thick blue line. The box plot compares the outcomes of

TABLE II. The two annotators (A1 and A2) classified more than 80% of the memes with the highest meme scores as relevant physics concepts and noun phrases. The differences involving the meme score are highly significant (*). See Sec. IV and Supplemental Material [32] for details.

Method	Main class	Annotator	(Annotator agreement)	Classified as main class:		p value for difference to	
				Individually	In agreement	Random	Weighted random
Meme score	Physics concept	A1	(90.0%)	85.3%	81.3%	$< 10^{-15}$ *	$< 10^{-15}$ *
		A2		87.3%			
	Noun phrase	A1	(93.3%)	86.0%	82.7%	$< 10^{-15}$ *	$< 10^{-15}$ *
		A2		86.0%			
Random	Physics concept	A1	(85.3%)	32.7%	25.3%		0.123
		A2		32.7%			
	Noun phrase	A1	(86.0%)	39.3%	33.3%		0.163
		A2		36.0%			
Weighted random	Physics concept	A1	(90.7%)	20.7%	19.3%	0.123	
		A2		27.3%			
	Noun phrase	A1	(86.0%)	28.7%	25.3%	0.163	
		A2		28.7%			

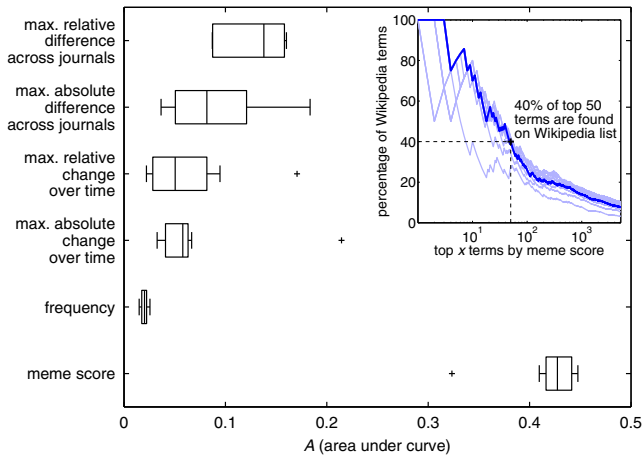


FIG. 3. The meme score outperforms alternative metrics. The box plot shows the agreement with the ground-truth list of physics terms extracted from Wikipedia as achieved by the different metrics. Agreement is measured as the area under the curve A , as shown for the meme score in the embedded graph on the top right. The curves are defined as the percentage from the x top-ranked terms that also appear on the ground-truth list for the different parameter settings ($1 \leq \delta \leq 10$ for the meme score; the thick line stands for $\delta = 4$, which has the largest area A ; see Sec. IV for details).

different metrics with respect to A , as described in the Sec. IV. The isolated outlier at 32.3% on the lower end of the distribution for the meme score originates from the parameter setting $\delta = 1$. All other parameter settings $2 \leq \delta \leq 10$ lead to results in a narrow band between 40.9% and 44.8% (meaning that the performance of the metric is not sensitive to movements within this range of the parameter space). In contrast, all alternative metrics score considerably worse, consistently below 22% (including outliers).

The two plots of Fig. 4 show the same data points as Fig. 2(a) but with colors standing for the relative density of Wikipedia terms [Fig. 4(a)] and of terms containing

chemical formulas [Fig. 4(b)]. Figure 4(a) confirms that phrases in the area of a high meme score (towards the top right) tend to show up as titles of Wikipedia articles on physics. Additionally, the plot shows that this is the *only* such area. There are a few scattered outliers, but the only significant area with a high density of Wikipedia terms is found around the 99.9% quantile. Figure 4(b) shows that phrases containing chemical formulas (such as BaFe_2As_2) tend to have a relatively low frequency (individually) but high propagation score. The area with the highest density can again be found along the 99.9% quantile, which is consistent with the expectation of chemical compounds to be important and interesting entities for physics research. Moreover, the fact that they are standardized and compressed representations reduces their “vulnerability” to synonyms or spelling variants, making them stronger memes on the level of pure character sequences. (We come back to the issue of memes on different levels of abstraction below.)

Figure 5 shows the top memes over time, revealing bursty dynamics, akin to the one reported previously in human dynamics [35] and the temporal distribution of words [36]. These bursts might be a reflection of the fast rise and fall of many scientific memes in terms of their popularity. As new scientific paradigms emerge, the old ones seem to quickly lose their appeal, and only a few memes manage to top the rankings over extended periods of time. The bursty dynamics also support the idea that both the rise and fall of scientific paradigms is driven by robust principles of self-organization [37].

III. DISCUSSION

Going back to the original analogy with genes put forward by Dawkins [20], we investigate the relation between the occurrence frequency of scientific memes and the degree to which they propagate along the citation graph. We find that scientific memes are indeed governed by a surprisingly

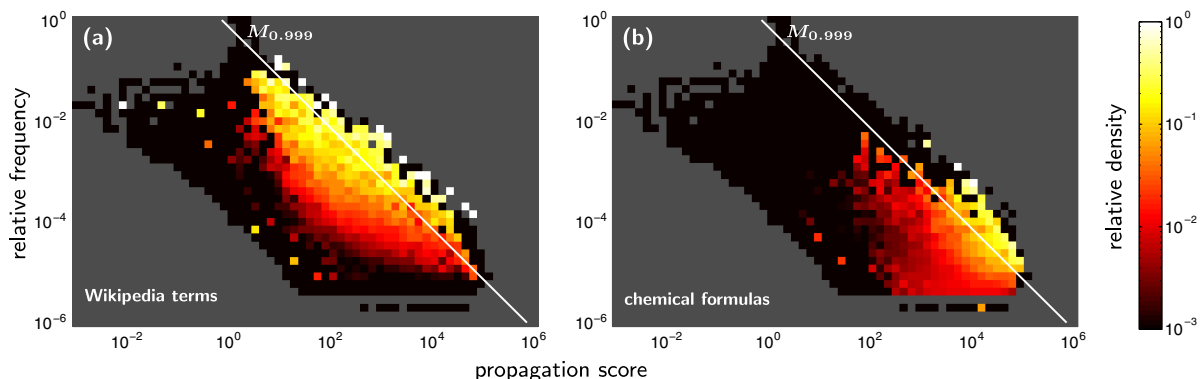


FIG. 4. Panel (a) shows that phrases with a high meme score (i.e., around the 99.9% quantile $M_{0.999}$) tend to be found as titles of Wikipedia articles on physics, while other phrases tend not to. Panel (b) shows that phrases of relatively low frequency but high meme score have the highest density of terms containing chemical formulas (e.g., MgB_2). The data points of the two plots are the same as in Fig. 2(a), except for the colors showing the relative density of the specific type of terms. Note that the colors are log scaled: white stands for 100% terms of the given type, yellow for $\approx 30\%$, orange for $\approx 10\%$, light red for $\approx 2\%$, dark red for $< 1\%$, and black for $< 0.1\%$.




FIG. 5. Time history of top physics memes based on their meme scores obtained from the American Physical Society data set. The time axis is scaled by publication count. Bars and labels are shown for all memes that top the rankings for at least 10 out of the displayed 911 points in time. The gray area represents the second-ranked meme at a given time.

simple relationship of these two factors. This is formalized by the meme score—a metric to characterize and identify scientific memes—defined as the product of the frequency of occurrence and the propagation score.

We show that the meme score can be calculated exactly and exhaustively without the introduction of arbitrary thresholds or filters and without relying on any kind of linguistic or ontological knowledge. The method is fast and reliable, and it can be applied to massive databases. We demonstrate the effectiveness of the meme score on more than 47 million publication records from the Web of Science, PubMed Central, and the American Physical Society. Moreover, we evaluate the accuracy of the proposed meme score by means of full and time-preserving randomizations of the citation graphs, by means of manual annotation of publications, as well as by means of several alternative metrics. We provide statistical evidence for the agreement between human annotators and the meme-score results, and we show that it is superior to alternative metrics. We also confirm that the observed patterns cannot be explained by topological or temporal features alone, but are grounded in more intricate processes that determine the dynamics of the scientific progress and the way credit is given to preceding publications. Furthermore, the top-ranking scientific memes reveal a bursty time dynamics, which might be a reflection of the fierce competition among memes for the limited and fluctuating resource of scientists' attention.

We consider only fixed character sequences as potential memes, but it is clear that memes do not exist only on this low level, and it is reasonable to expect that the inclusion of additional layers of processing using linguistic and ontological resources will lead to even better results and let us

capture memes on a more abstract level. Such memes might consist of sets of morphological variants, co-occurrences of words, compositions of multiple memes, grammatical constructions, or even argumentation schemes and rhetorical styles. We deliberately keep the meme score as simple as possible to emphasize that it is surprisingly precise on its own. At the same time, there are many ways to improve the metric in the future with more sophisticated processing to capture memes on a higher level. In general, we believe that our approach, allowing for the study of memes in a comprehensive manner, opens up the field for a wide range of future research on topics such as information diffusion, complex systems, innovation, scientific progress, social dynamics, ecosystems, cultural evolution, and, of course, the study of memes themselves.

IV. METHODS

A. Graph randomization

The analyzed randomized networks have exactly the same topology as the original ones, but the article texts (i.e., titles and abstracts with their memes) are randomly assigned to the nodes. Each node, therefore, owes its position in the network to one particular publication but has text attached that comes from a different one. For the time-preserving randomizations, we shuffle only publications that were published within narrow consecutive time windows. Concretely, we use time windows of 1000 publications, meaning that—after shuffling—no publication has moved more than 1000 positions forward or backward from the original chronological order.

B. Human annotation

For the first part of the manual annotation, we use the following two categories: (i) the phrase is not a meaningful term or not an important concept of physics, and (ii) the phrase is an important concept or entity of physics—it could appear as the title of an entry of a comprehensive encyclopedia of physics. For the second part, we define the following linguistic phrase types: (i) noun phrase, (ii) verb, (iii) adjective or adverb, and (iv) other.

The set of phrases used for this evaluation consists of the top 150 memes with respect to their meme score, extracted from the American Physical Society data set, plus another two sets for comparison of 150 randomly drawn phrases each. For the two comparison sets, we consider all phrases that appear in at least 100 publications. From these, 150 terms are drawn randomly without taking into account their frequency; i.e., frequent terms have the same chance of being selected as infrequent ones, whereas the 150 terms of the second set are drawn with a weight that corresponds to their frequency. Moreover, to rule out effects of different n -gram lengths, we make sure that the two batches of random terms follow exactly the same length distribution as the main sample extracted based on the meme score. The resulting 450 terms are shuffled and given to two human annotators, both Ph.D. students with a degree in physics, who independently annotate the terms.

C. Metrics for comparison

We use the following metrics with different parameter settings as alternatives to the meme score (see Supplemental Material for details [32]): (i) frequency—the most frequent terms, optionally skipping the first x terms; (ii) maximum absolute change over time—the highest-scoring terms with respect to maximum *absolute* change in frequency; (iii) maximum relative change over time—the same as (ii) but based on *relative* changes; (iv) maximum absolute difference across journals—the highest-scoring terms with respect to maximum *absolute* difference in frequency between journals; (v) maximum relative difference across journals—the same as (iv) but based on *relative* differences. Metric (i) is based on the assumption that important memes are frequent but not as frequent as the small class of general words that can be found in all types of texts. Metrics (ii) and (iii) are based on an idea proposed in Ref. [37], that interesting memes exhibit trends over time. Metrics (iv) and (v) are based on the intuition that phrases occurring mostly in specific journals but not in others must be specific concepts of the particular field of research.

As a ground-truth list of memes, we automatically extract 5178 terms from Wikipedia. We collect the titles of all articles—and terms redirecting to them—from the categories “physics,” “applied and interdisciplinary physics,” “theoretical physics,” “emerging technologies,” and their direct subcategories, but filtering out terms that appear

in less than 10 publications of the American Physical Society data set. To quantify the agreement between the top memes identified by a particular metric and the Wikipedia list, we use the normalized area A under the curve as shown in Fig. 3. The step-shaped curve has a log-scaled x axis running up to the number of terms s on the ground-truth meme list ($s = 5178$ in our case) and a y axis running from 0 (no overlap) to 1 (perfect overlap). Limiting cases are $A = 1$, representing perfect agreement, and $A = 0$, representing no agreement at all between the two compared lists.

ACKNOWLEDGMENTS

This research was supported by the European Commission through the ERC Advanced Investigator Grant “Momentum” (Grant No. 324247) and by the Slovenian Research Agency through the Program No. P5-0027. In addition, we would like to thank Karsten Donnay, Matthias Leiss, Christian Schulz, and Olivia Woolley-Meza for their useful feedback and help with the realization of the evaluations.

-
- [1] S. Redner, *How Popular is Your Paper? An Empirical Study of the Citation Distribution*, *Eur. Phys. J. B* **4**, 131 (1998).
 - [2] F. Radicchi, S. Fortunato, and C. Castellano, *Universality of Citation Distributions: Toward an Objective Measure of Scientific Impact*, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 17268 (2008).
 - [3] M. E. J. Newman, *Coauthorship Networks and Patterns of Scientific Collaboration*, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 5200 (2004).
 - [4] R. Guimerà, B. Uzzi, J. Spiro, and L. A. N. Amaral, *Team Assembly Mechanisms Determine Collaboration Network Structure and Team Performance*, *Science* **308**, 697 (2005).
 - [5] S. Milojević, *Principles of Scientific Research Team Formation and Evolution*, *Proc. Natl. Acad. Sci. U.S.A.* **111**, 3984 (2014).
 - [6] J. E. Hirsch, *An Index to Quantify an Individual’s Scientific Research Output*, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 16569 (2005).
 - [7] F. Radicchi, S. Fortunato, B. Markines, and A. Vespignani, *Diffusion of Scientific Credits and the Ranking of Scientists*, *Phys. Rev. E* **80**, 056103 (2009).
 - [8] A. M. Petersen, F. Wang, and H. E. Stanley, *Methods for Measuring the Citations and Productivity of Scientists across Time and Discipline*, *Phys. Rev. E* **81**, 036114 (2010).
 - [9] D. Wang, C. Song, and A.-L. Barabási, *Quantifying Long-Term Scientific Impact*, *Science* **342**, 127 (2013).
 - [10] O. Penner, R. K. Pan, A. M. Petersen, K. Kaski, and S. Fortunato, *On the Predictability of Future Impact in Science*, *Sci. Rep.* **3**, 3052 (2013).
 - [11] D. J. de Solla Price, *Networks of Scientific Papers*, *Science* **149**, 510 (1965).
 - [12] R. K. Merton, *The Matthew Effect in Science*, *Science* **159**, 56 (1968).

- [13] J. A. Evans and J. G. Foster, *Metaknowledge*, *Science* **331**, 721 (2011).
- [14] J. B. Michel, Y. K. Shen, A. Presser Aiden, A. Veres, M. K. Gray, J. P. Pickett, D. Hoiberg, D. Clancy, P. Norvig, J. Orwant, S. Pinker, M. A. Nowak, and E. Lieberman Aiden (The Google Books Team), *Quantitative Analysis of Culture Using Millions of Digitized Books*, *Science* **331**, 176 (2011).
- [15] R. K. Pan, K. Kaski, and S. Fortunato, *World Citation and Collaboration Networks: Uncovering the Role of Geography in Science*, *Sci. Rep.* **2**, 902 (2012).
- [16] A. Mazloumian, D. Helbing, S. Lozano, R. P. Light, and K. Börner, *Global Multi-level Analysis of the “Scientific Food Web”*, *Sci. Rep.* **3**, 1167 (2013).
- [17] D. Chavalarias and J.-P. Cointet, *Phylomemetic Patterns in Science Evolution—The Rise and Fall of Scientific Fields*, *PLoS One* **8**, e54847 (2013).
- [18] K. Börner, *Atlas of Science* (MIT Press, Cambridge, MA, 2010).
- [19] K. Börner, *Atlas of Knowledge* (MIT Press, Cambridge, MA, 2014).
- [20] R. Dawkins, *The Selfish Gene* (Oxford University Press, Oxford, England, 1989), 2nd ed.
- [21] J. Leskovec, L. Backstrom, and J. Kleinberg, in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York, NY, 2009), pp. 497–506.
- [22] E. Adar and L. A. Adamic, in *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence* (IEEE, New York, NY, 2005), pp. 207–214.
- [23] M. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, F. Menczer, and A. Flammini, in *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM)* (AAAI, Menlo Park, CA, 2011), pp. 89–96.
- [24] M. Gomez Rodriguez, J. Leskovec, and A. Krause, in *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York, NY, 2010), pp. 1019–1028.
- [25] M. P. Simmons, L. A. Adamic, and E. Adar, in *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM)* (AAAI, Menlo Park, CA, 2011), pp. 353–360.
- [26] L. Weng, A. Flammini, A. Vespignani, and F. Menczer, *Competition among Memes in a World with Limited Attention*, *Sci. Rep.* **2**, 335 (2012).
- [27] N. A. Christakis and J. H. Fowler, *Social Contagion Theory: Examining Dynamic Social Networks and Human Behavior*, *Stat. Med.* **32**, 556 (2013).
- [28] H. E. Stanley, *Introduction to Phase Transitions and Critical Phenomena* (Clarendon Press, Oxford, England, 1971).
- [29] J. P. Gleeson, J. A. Ward, K. P. O’Sullivan, and W. T. Lee, *Competition-Induced Criticality in a Model of Meme Popularity*, *Phys. Rev. Lett.* **112**, 048701 (2014).
- [30] S. Martin, W. M. Brown, R. Klavans, and K. W. Boyack, in *Visualization and Data Analysis 2011*, Proceedings of SPIE, Vol. 7868 (SPIE, Bellingham, WA, 2011).
- [31] M. Bastian, S. Heymann, and M. Jacomy, in *Proceedings of the Third International AAAI Conference on Weblogs and Social Media (ICWSM)* (AAAI, Menlo Park, CA, 2009), pp. 361–362.
- [32] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.4.041036> for details on the alternative metrics, possible meme score modifications, the network randomization, the annotators’ task description, and further illustrative figures and tables.
- [33] D. Bourigault, I. Gonzalez-Mullier, and C. Gros, in *Proceedings of the 7th EURALEX International Congress* (Göteborg University, Göteborg, Sweden, 1996), pp. 771–779.
- [34] A. Hulth, in *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, Stroudsburg, PA, 2003), pp. 216–223.
- [35] A. L. Barabási, *The Origin of Bursts and Heavy Tails in Humans Dynamics*, *Nature (London)* **435**, 207 (2005).
- [36] E. G. Altmann, J. B. Pierrehumbert, and A. E. Motter, *Beyond Word Frequency: Bursts, Lulls, and Scaling in the Temporal Distributions of Words*, *PLoS One* **4**, e7678 (2009).
- [37] M. Perc, *Self-Organization of Progress across the Century of Physics*, *Sci. Rep.* **3**, 1720 (2013).