

Quantifying patterns of research-interest evolution

Tao Jia^{1,2*}, Dashun Wang^{3,4,5*} and Boleslaw K. Szymanski^{6,7,8*}

To understand quantitatively how scientists choose and shift their research focus over time is of high importance, because it affects the ways in which scientists are trained, science is funded, knowledge is organized and discovered, and excellence is recognized and rewarded¹⁻⁹. Despite extensive investigation into various factors that influence a scientist's choice of research topics⁸⁻²¹, quantitative assessments of mechanisms that give rise to macroscopic patterns characterizing research-interest evolution of individual scientists remain limited. Here we perform a large-scale analysis of publication records, and we show that changes in research interests follow a reproducible pattern characterized by an exponential distribution. We identify three fundamental features responsible for the observed exponential distribution, which arise from a subtle interplay between exploitation and exploration in research-interest evolution^{5,22}. We developed a random-walk-based model, allowing us to accurately reproduce the empirical observations. This work uncovers and quantitatively analyses macroscopic patterns that govern changes in research interests, thereby showing that there is a high degree of regularity underlying scientific research and individual careers.

'The essential tension' hypothesis as described in ref. ⁵ has highlighted the conflicting demands of scientific careers that require both exploration and exploitation^{4,8,22}. Indeed, career advancement, from promotion to obtaining grants, demands a steady stream of publications, which is often achieved through uninterrupted, yet incremental contributions to the existing, established research agenda. By contrast, frequent changes in research topics invite risk of failure and decreased productivity. The disciplinary boundaries, arising from factors such as implicit culture, tacit and accumulated knowledge^{23,24} and peer recognition^{3,25}, together with intensifying specialization in science and engineering disciplines²⁶, make radical shifts, such as moving from chemical biology to high energy physics, unlikely, if at all possible. On the other hand, although a steady and focused research portfolio helps scientists stay productive, it potentially undermines chances for originality⁸. Indeed, innovative and novel insights often emerge from encountering new challenges and opportunities associated with venturing into new topics and/or incorporating them into the existing research agenda^{4,15,20,27,28}.

Given the broad effect on individual careers and strong implications for science and innovation policy, there is an urgent need for quantitative approaches to understand the nature of changes in research interests of individual scientists throughout their careers. This need becomes more urgent with the accelerating scale and complexity of scientific enterprise^{2,26,29,30}. A variety of microscopic factors have been identified that drive a scientist's choice of research problems, which range from age^{10,11} to gender^{12,13}, to training and

mentorship^{9,14}, from funding or collaboration opportunities¹⁵⁻¹⁷, to serendipity¹⁸, to a scientist's attitudes and abilities¹⁹, including risk aversion and creativity^{8,20,21}. Yet, little is known about the macroscopic patterns that underlie the research-interest evolution. Recent advances in complex systems have uncovered regularities in various dynamical processes, which give rise to a family of powerful, yet flexible statistical models that describe processes as diverse as human mobility³¹⁻³⁴, temporal dynamics³⁵⁻³⁷ and the evolution of complex networks³⁸⁻⁴⁰. This prompted us to ask: to what degree could the research-interest evolution be captured by a simple model?

Here we aim to systematically address this question by first identifying patterns in the working scientists' research agendas as their careers progress. Using articles published by American Physical Society (APS) journals covering over 30 years (1976-2009)^{41,42} and through a careful and extensive author-name disambiguation process^{6,43}, we collect publication records of individual authors over time (Supplementary Note 1). We further take advantage of the Physics and Astronomy Classification Scheme (PACS) codes used by the APS to classify topics in physics. Indeed, among all identifiers for research topics, PACS codes stand out in the frequency of their use^{30,44-48}. This is partly because unlike topics defined by keywords, which are often created in an ad hoc, unstructured manner, the PACS code classification relies on both the crowd wisdom of working scientists and expert opinions of journal editors, offering a systematic representation of the subject of a paper. There are 67 main topics defined by the first two digits of the PACS code, covering diverse topics that range from general relativity and gravitation to nuclear structure to superconductivity. By sorting these PACS codes, we obtain a topic tuple for each paper, representing the research subject of each paper as a combination of the topics the paper studies^{28,30} (Fig. 1). For a given set of papers published by a scientist, we generate a topic vector, the elements of which represent the weighted occurrence of each topic (see Methods). Therefore this vector captures not only the collection of topics a scientist studied, but also the level of involvements in each of these topics. Consequently, this vector represents a multi-dimensional measure of research interest of a scientist as revealed in the series of papers published by the author.

We compose two topic vectors based on the first and last m papers of the scientist (g_i and g_f , respectively), thereby capturing the research interest at the earliest and the latest stages of the career (Fig. 1). Using the complementary cosine similarity between g_i and g_f , we quantify the interest change J of a scientist throughout the career as:

$$J = 1 - \frac{g_i \cdot g_f}{\|g_i\| \|g_f\|} \quad (1)$$

¹College of Computer and Information Science, Southwest University, Chongqing 400715, China. ²Laboratory for Software and Knowledge Engineering, Southwest University, Chongqing 400715, China. ³Kellogg School of Management, Northwestern University, Evanston, Illinois 60208, USA.

⁴Northwestern Institute on Complex Systems (NICO), Northwestern University, Evanston, Illinois 60208, USA. ⁵McCormick School of Engineering and Applied Sciences, Northwestern University, Evanston, Illinois 60208, USA. ⁶Social Cognitive Networks Academic Research Center, Rensselaer Polytechnic Institute, Troy, New York 12180, USA. ⁷Department of Computer Science, Rensselaer Polytechnic Institute, Troy, New York 12180, USA.

⁸Spoleczna Akademia Nauk, 90-113 Łódź, Poland. *e-mail: tjia@swu.edu.cn; dashun.wang@northwestern.edu; szymab@rpi.edu

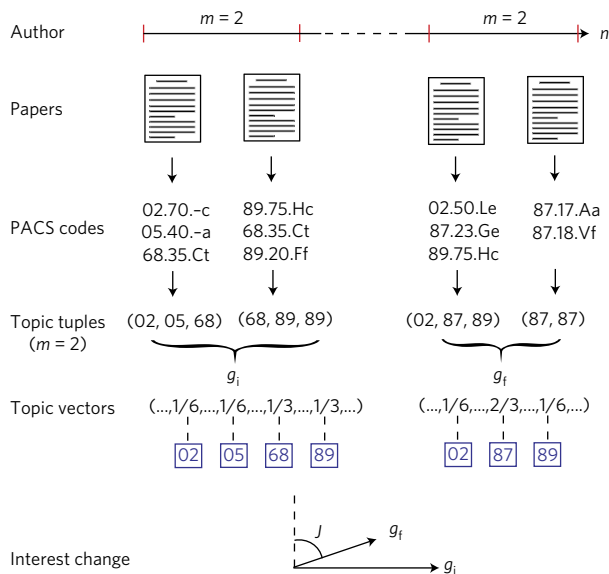


Figure 1 | An example demonstrating the procedure to compose a topic tuple and topic vector. Using two topic vectors g_i and g_f based on the first and last m papers ($m = 2$ in this case) in an author's publication sequence, the interest change J is calculated as the complementary cosine similarity between the two vectors. See Methods for more details on how to calculate topic vectors.

Equation (1) captures the research-interest change that resulted from a change in topics or from a change in engagement with the topics, thereby providing an effective quantification on the extent of change. $J = 0$ indicates that the two topic vectors g_i and g_f are identical, capturing the fact that the author not only studied the same set of topics at the two stages of the career, but also was involved in each of these topics with the same weight. $J = 1$ corresponds to a complete change in interests, in which a researcher does not engage in any initial topic of interest. We choose $m = 8$. As a result, our analyses are based on 14,715 scientists, each of which was an author on at least $2m = 16$ papers included in our dataset. We also report analyses based on other m values ($m = 6$ and $m = 10$, Supplementary Note 2) and find that our results are insensitive to the choice of m . To take other factors into account that can play a role in quantifying research-interest change, we further perform three additional measurements. First, to avoid the gaps between the two sets of papers, we take $2m$ consecutive papers starting at a randomly chosen paper and measure the interest change based on the two adjacent m paper sequences (Supplementary Note 3). Second, to eliminate the effects of different publication rates, we measure the interest change J within scientists who publish at similar rates (Supplementary Note 4). Finally, as research interest is associated with time, we measure interest change based on two sets of papers published over the same time period at the early and late stage of a career (Supplementary Note 5). We obtain similar observations in all three measurements, demonstrating that the discovered patterns in research-interest evolution do not rely on a specific metric.

The quantification of research-interest change allows us to measure its distribution within the population. We find that the fraction of scientists P decays with the extent of interest change J (Fig. 2a), which can be well fitted with an exponential function (see Supplementary Note 6 for fitting statistics). This exponential decay indicates that most scientists are characterized by little change in their research interests and the probability of making a leap decreases exponentially with its range. At the same time, the fact that the histogram $P(J)$ is positive in the full range of domain $[0, 1]$ indicates that large changes in a career, such as switching to

completely different areas, do occur, albeit very rarely. These observations raise an interesting question: What forms of interest-change distribution should we expect? To answer this question, we identify three features, that is, heterogeneity, recency and subject proximity, that characterize research-interest evolution. To illustrate how these features shape the distribution of interest change, we perform three 'experiments' in which the interest-change distribution is measured using modified publication sequences of individual scientists. These results further demonstrate the non-trivial nature of the exponential distribution observed.

The frequencies of topic tuple occurrences in the publication sequence of a scientist follow a power-law distribution (Fig. 2b), demonstrating the heterogeneity in individual's engagement in different subjects. Indeed, a scientist's research agenda contains core research subjects that are repeatedly investigated coupled with other more peripheral ones that may only be touched upon occasionally. To examine the effect of heterogeneity on the change in research interests, we modify each publication sequence by retaining only the first occurrence of each topic tuple and removing the subsequent ones (see Supplementary Fig. 1a for the illustration of this procedure and Supplementary Note 7 for more details). The interest-change distribution measured over the modified sequences reaches a peak at an intermediate value followed by a gradual decay (Fig. 2c), which is in contrast to the monotonic, exponential decrease that was observed in the real data.

By denoting the number of papers between two consecutive uses of a topic tuple as Δn (see Methods), we compare the distribution $P(\Delta n)$ to the reshuffled publication sequence $P_o(\Delta n)$. We find that the ratio $P(\Delta n)/P_o(\Delta n)$ decreases with Δn , going from above to below 1 (Fig. 2d). This implies that scientists are more likely to publish on subjects recently studied in real sequences than in reshuffled sequences. To further explore this feature, we measure the relationship between the probability that a distinct topic tuple is re-studied (II) and the order of the first occurrence of this topic tuple in a scientist's publication sequence (see Methods). As shown in Figure 2e, the relationship obtained demonstrates that scientists are more likely to publish on subjects studied recently than on those investigated a long time ago. This recency feature suggests that scientists tend to avoid going back to the original research subjects once they have moved on to other ones, which consequently drives them to explore new research subjects. This is affirmed in our analysis of the interest-change distribution measured over the reshuffled publication sequences (see Supplementary Fig. 1b for the illustration of this procedure). The obtained distribution remains exponential, but with a much steeper decrease than the original distribution, which makes the large extent of interest change significantly less frequent than what is actually observed (Fig. 2f).

Knowledge is characterized by underlying topical geometry that imposes varying inherent distances between pairs of research subjects^{26,49,50}. When a scientist changes the research subject, he or she is more likely to choose a subject that is related to the current one than to move to a totally new field, implying the subject proximity of research-interest change. To verify this insight, we modify each publication sequence by replacing distinct topic tuples in publications of a scientist with topic tuples that were randomly drawn (see Supplementary Fig. 1c for the illustration of this procedure). The resulting interest-change distribution shows that most scientists have a large extent of change (Fig. 2g), demonstrating the effect of subject proximity on research-interest evolution.

These three features reveal insights that are offered by the observed distribution of research-interest change. Both heterogeneity and subject proximity arise from exploitation of the current field, which help stabilize the research interest. Lacking either of the two, P would have been characterized by a distribution with a larger mean. Yet, the recency feature resulting from the exploration of new areas destabilizes research interest. If it were not for the exploitation,

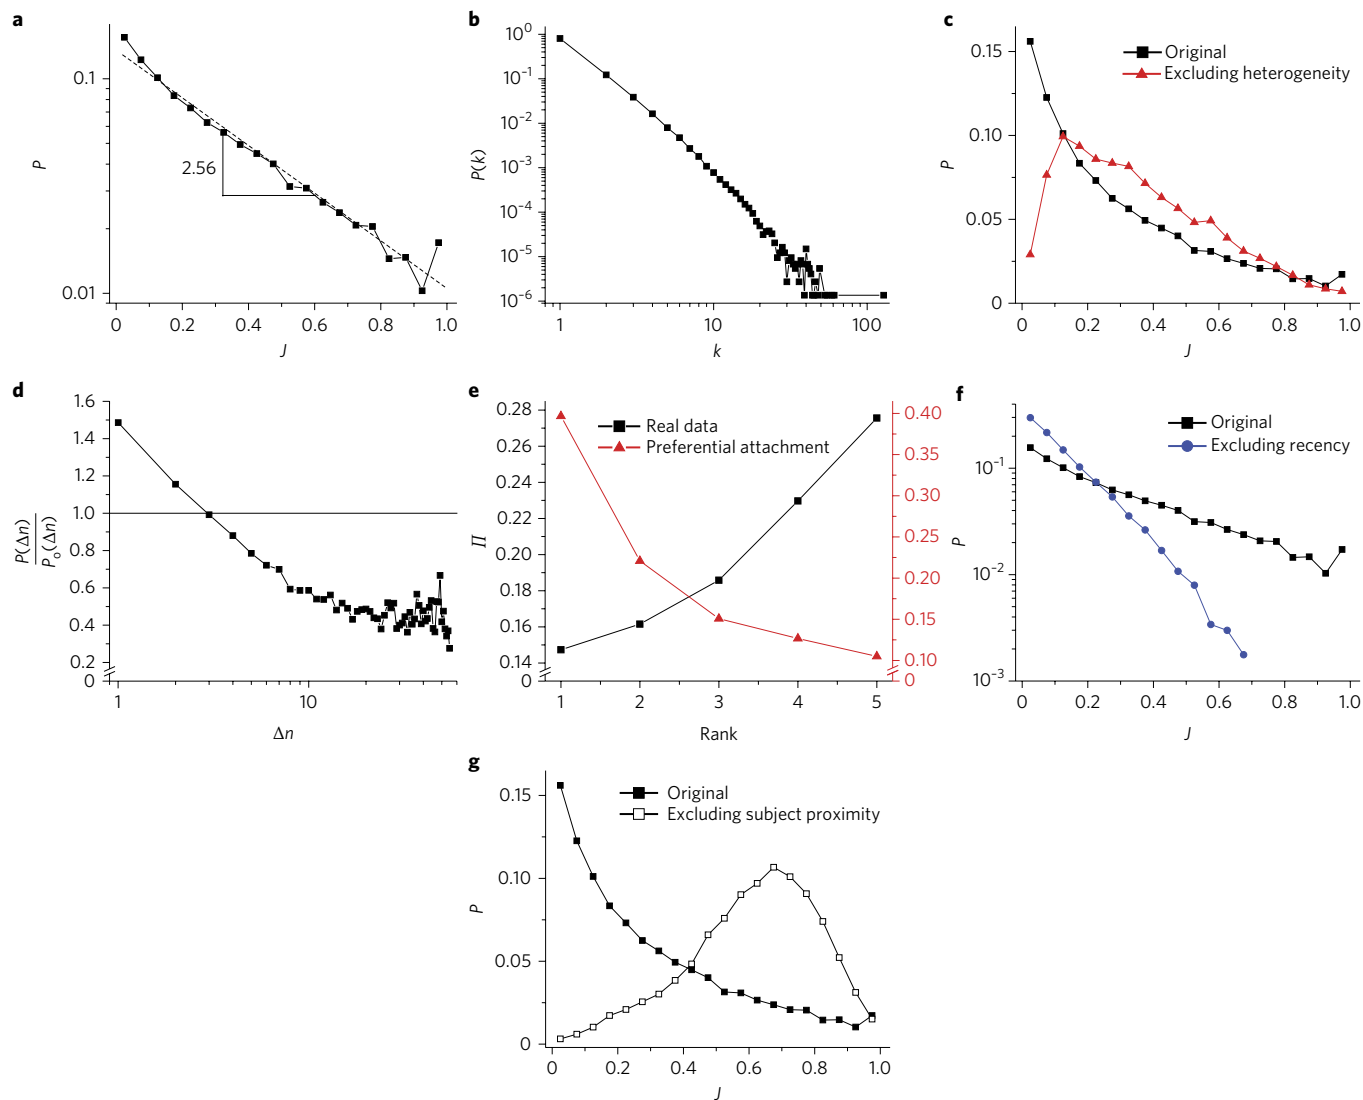


Figure 2 | Patterns in the research-interest evolution. **a**, The fraction of scientists P within a given range of interest change ($J - 0.025, J + 0.025$) decreases monotonically with J , which can be fitted well with an exponential function. At the boundary $J = 0$, the range $[0, 0.05]$ is used, and the same boundary condition applies in all the following studies. **b**, The probability $P(k)$ that a topic tuple is used for k times in an author’s publication sequence follows a power-law distribution, documenting the heterogeneity in the usage of topic tuples. **c**, $P(J)$ over sequences in which ‘heterogeneity’ is eliminated by retaining only the first occurrence of each topic tuple in the publication sequence. **d**, The separation between the usage of the same topic tuple in a scientist’s publication sequence, measured by the number of papers Δn between its two consecutive uses (see Methods). The ratio between the distribution of Δn of real data ($P(\Delta n)$) and that of the reshuffled sequence ($P_0(\Delta n)$) implies that an author is more likely to publish on subjects covered in recent papers than on those published a long time ago. **e**, The relationship between the probability of reusing a previously studied topic tuple Π (see Methods) and the rank of its first usage (rank 1 is assigned to the first distinct topic tuple used in an individual’s career, and so on). It demonstrates that a scientist is more likely to publish on subjects recently studied than on subjects investigated a long time ago. This is, however, opposite to what preferential attachment would predict (see Methods). **f**, $P(J)$ over sequences in which ‘recency’ is eliminated by removing the temporal correlation of topics by random shuffling of papers. **g**, $P(J)$ over sequences in which ‘subject proximity’ is eliminated by replacing the distinct topic tuples in each publication sequence with ones randomly drawn from all existing topic tuples. In **c**, **f** and **g**, the P value equals 0 in the two-sample Kolmogorov–Smirnov test, indicating that the observed differences are statistically significant.

the interest change would have been much more limited given the heterogeneity and subject-proximity features. Together they provide us with an empirical basis to build a statistical model for individual careers with varying research subjects over time.

Here we consider scientific research as a random walk, following Isaac Newton’s retrospection that during his scientific career he was like “... a boy playing on the seashore... finding... a prettier shell than ordinary”⁵¹. Such a ‘seashore’ is represented by a one-dimensional lattice in our model and a ‘shell’ corresponds to a scientific finding that yields a paper (Fig. 3a). The locations at which shells can be found

are located with a certain probability p at the sites of lattice, where each location contains q shells of one type. q follows a distribution $P(q) \sim q^{-\alpha}$, which is motivated by the heterogeneity in the potential that a research yields papers: although some are fruitful, many of research topics are not. Therefore, the seashore lattice contains sites with piles of shells, separated by a random number of empty sites in between. A scientist starts at a random initial position and performs an unbiased random walk on its own lattice (50% chance to move one step left or right). Upon reaching a site that contains shells, the walker picks one shell, corresponding to publishing one paper.

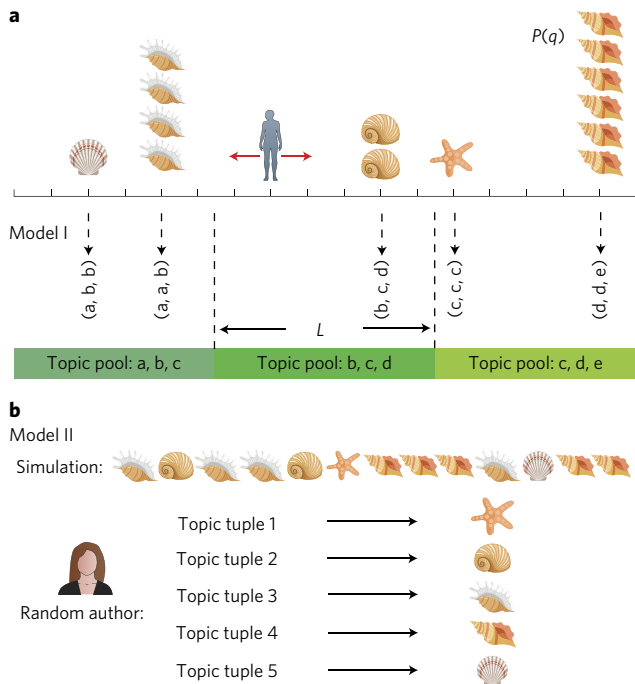


Figure 3 | An illustration of the ‘seashore walk’. **a**, The random walker traverses a one-dimensional lattice with piles of ‘shells’ located at the sites of the lattice. The probability that a site contains any shells is p . The number of shells at a non-empty site is characterized by the distribution $P(q)$. The walker picks a shell upon reaching a site that contains shells, corresponding to publishing a paper. A sequence of shells is generated until the walker stops after exhausting the total number of steps S assigned to its career span, which is characterized by a log-normal distribution $P(S)$. In Model I, we assign each shell an artificial topic tuple on the basis of the location at which the shell is picked. In particular, L sites on the lattice share a topic pool with three topics. The two neighboring topic pools vary from one to the other by exactly one topic. For example, if one topic pool is ‘a, b, c’ then the next pool could be ‘b, c, d’, and so on, where the codes ‘a, b, ...’ represent any arbitrary characterizations of different topics. The shells at a site are characterized by an artificial topic tuple as a set of three topics, each randomly drawn from the topic pool below. **b**, In Model II, distinct shells picked by a walker are assigned to real topic tuples used by a random scientist. Seashell images reproduced with permission from S-S-S/iStock/Getty Images Plus/Getty.

The walker stops when it reaches the end of the career, defined by the total number of steps S , following a truncated log-normal distribution $P(S)$, the choice of which is motivated by the distribution of the observed career lifetime span of the real scientist^{52,53} (Supplementary Fig. 2).

Despite its simplicity, the seashore walk predicts the presence of both heterogeneity and recency features. If we assume that one type of shells corresponds to one research subject, the power-law distribution $P(q)$ provides varying limits to the number of papers that can be published on different research subjects. The unbiased random walker is likely to return to a site repeatedly, which consequently enables it to collect all the shells from the site. Such repeated visits and $P(q)$ give rise to the heterogeneity and recency features observed empirically (Supplementary Fig. 3a, b).

To further capture the evolution of research interest, we need to assign topics to each type of shell. Recent advances in knowledge expansion^{4,8,29} have provided two specific features governing the evolution of a scientist’s research agenda: first, existing topics are connected to form a research subject; and second, new topics are occasionally added²⁹. By absorbing these two features into

the seashore walk, we define how the research subject of a type of shell is associated with its location on the lattice and how topics evolve on the underlying lattice, leading to the introduction of Model I (Fig. 3a).

For Model I we assume that there is a topic pool containing 3 topics for each site on the lattice. The shells on the site, if there are any, are characterized by a random combination of the 3 topics (with repetitions), representing an artificial topic tuple for the research subject of the site. The value 3 is based on the observation that each paper is characterized on average by 2.89 PACS codes⁴² (Supplementary Fig. 4a, b). Furthermore, we assume that one topic pool covers L sites on the lattice and two neighboring topic pools differ by one topic. Therefore, new topics are encountered as the walker moves away from the starting point, which is in line with the empirical observation that new topics emerge when the number of distinct topic tuples used by a scientist increases (Supplementary Fig. 4c).

Model I generates a sequence of shells for each walker traversing its own seashore, and each shell is characterized by an artificial topic tuple. By measuring interest change for an ensemble of walkers (see Supplementary Note 8 for the implementation of the simulation), we obtain a $P(J)$ similar to that of real data (Fig. 4a and see Supplementary Note 9 for statistical analyses). Despite its simplicity, Model I captures the process of research-interest evolution. It also leads to an interesting question: to what extent could the framework of seashore walk be improved? Note that the generation of topic tuples in Model I is based on simple processes, which can limit its capability to accurately reproduce the interest-change distribution. If this is the case, we would expect a better result if more complexities are added to ensure that a walker’s topic tuples as well as the correlations among these topic tuples are statistically similar to the data. To test this hypothesis, we constructed Model II (Fig. 3b).

For Model II we generated a sequence of shells picked by a walker and identify the number of distinct shells in the sequence. For a walker who has picked x distinct shells, we randomly find a scientist from the data who used x different topic tuples in the publication sequence and randomly map them to the x distinct shells (see Supplementary Note 10 for the implementation of simulation).

We apply Model II to the same shell sequences generated in Model I and find that the resulting research-interest change distribution matches closer to the empirical observations (Fig. 4b and Supplementary Note 9 for statistical analyses). The improved result confirms the validity and potential of the seashore walk to capture the research-interest evolution. It is noteworthy that the random-mapping procedure in Model II is used only to avoid introducing any sophisticated means to generate topic tuples. It is inherently different from duplicating the actual sequence (Supplementary Discussion 1). Moreover, despite some assumptions defining the seashore walk, such as the assumption that the number of shells at a site follows a power-law distribution, the systematic reproduction of empirical observations is obtained from the interplay of multiple mechanisms. Removing either of these assumptions would invalidate the model (Supplementary Discussion 2).

Finally, the seashore walk makes two additional predictions regarding an individual’s career. First, the individual’s publication process is bursty, as the inter-publication time follows a power-law distribution^{35,54} (Supplementary Fig. 5a). In the model, a random walker’s first passage steps follow a power-law distribution asymptotically with exponential cutoff $\sim S^{-3/2}$ (ref. 55), giving rise to the ‘burstiness’ of publication time (Fig. 4c). Second, the number of papers authored by a scientist follows a power-law distribution with an exponential cutoff^{56,57} (Supplementary Fig. 5b). We obtain the same form of distribution from the model (Fig. 4c). This is owing to a combination of factors, including the uniform probability of encountering sites with shells, the property that the mean number of sites visited by a random walker scales as $\sim S^{-1/2}$,

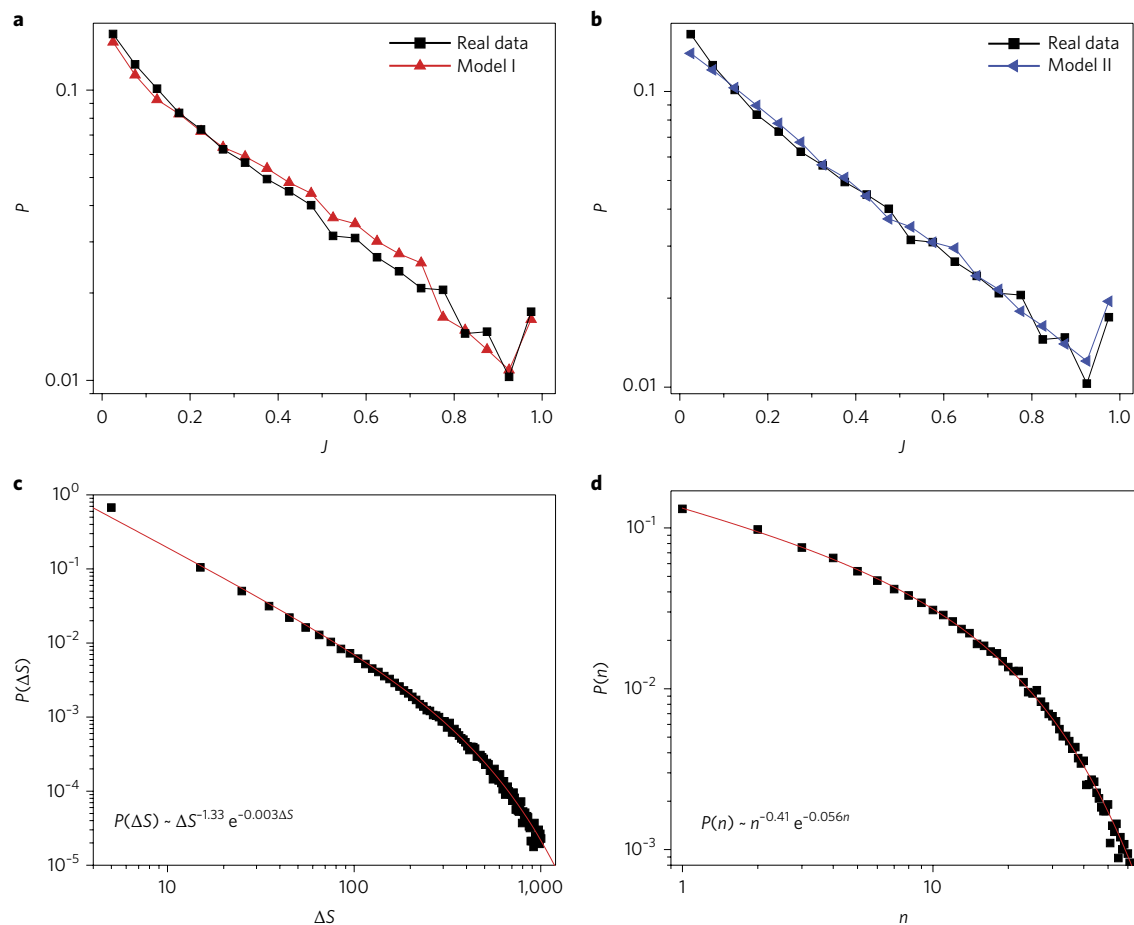


Figure 4 | Results of the 'seashore walk'. **a**, Model I generates an interest-change distribution qualitatively similar to that of real data. The variables are $p = 0.2$, $L = 35$, $P(q) \sim q^{-2}$, and the log-normal distribution $P(S)$ with mean $\mu = 6$, standard deviation $\sigma = 3$ and cutoff $S_{\max} = 2,000$. **b**, For the same sequence of shells in **a**, Model II produces an interest-change distribution that matches closely with that of real data. **c**, The distribution of intervals between a scientist's successive publications $P(\Delta S)$ is characterized by a power-law distribution with exponential cutoff in our model. **d**, The number of papers authored by a scientist $P(n)$ follows a power-law distribution with exponential cutoff in our model. The survival functions of **c** and **d** are plotted in Supplementary Fig. 6.

and the existence of a fat-tail in the log-normal distribution $P(S)$ (Supplementary Discussion 3).

The success of our simple model in capturing patterns observed in individual careers raises another question: can other related approaches be adopted and applied to the modelling process investigated here? To this end, we identified two classes of models that might be suitable on the basis of existing works in science of science and in network science. The first class pertains to models for the mobility patterns of an individual^{32,37} by treating topic tuples as locations. In these models, a scientist's sequence of research subjects becomes the sequence of locations visited in an individual mobility trajectory (see Methods). Although this approach could reproduce the heterogeneity feature based on preferential attachment³⁸, it could not capture the recency feature. Indeed, under the preferential attachment mechanism, a positive feedback would arise by which the more frequently a research subject was studied, the more likely it would be studied again. As a result, an old subject would receive more attention than the recent subject. Therefore, the probability of reusing a topic tuple (II) would decrease with the rank of first usage of this topic tuple³². This, however, directly contradicts the recency feature that was observed in the data (Fig. 2e), demonstrating the inherent inability of models that are based on preferential attachment to capture research-interest change. The second class of models treats the individual interest change as a Markov process on the knowledge network^{4,8,49}. This approach provides a comprehensive picture of the geometry of the

knowledge network that gives rise to subject proximity. However, the heterogeneity feature leading to the power-law distributed topic tuple usage can not be generated by a Markov process with fixed transition probabilities. Moreover, the knowledge network characterized by the movements of an individual between research subjects is not static but dynamic³⁹, which can not be accounted for without introducing a much more complicated model. Taken together, both approaches exhibit clear limitations in reproducing important characteristics of interest evolution that were studied in this paper. Our model, on the other hand, overcomes these limitations and preserves the patterns observed in the research-interest evolution.

In summary, by taking advantage of the PACS codes that classify general areas of physics into multiple clearly defined sub-areas, we quantitatively measure the extent of interest change for over 14,000 scientists, and show an exponential distribution of interest change within the population. We identify three key features in interest evolution that are essential for the presence of the observed distribution. We further develop a simple statistical model that describes scientific research as a random walk and that successfully captures empirical observations. Together, our results fill a critical gap in our quantitative understanding of science at a large scale by identifying a set of macroscopic patterns, which govern research-interest change throughout individual careers. Despite the well-known fact that scientists' choices of research subjects are driven by a myriad of factors, our results indicate that research-interest evolution can be

captured well by a simple statistical model, uncovering a new degree of regularity underlying individual careers.

The methodology introduced here implies that there are some limitations and potential for future work. When composing a topic vector, we assume papers on which an author's name appears are equally representative of his or her research interests. This assumption is justified by the difference between the interest in the problem addressed by this paper and the contribution to the paper or recognition of each author^{3,58}: every author has to be interested in the problem to engage in co-authoring the paper. In the future work, it would be of interest to systematically quantify the difference of each co-author's interest in a single paper on which they collaborate. The macroscopic patterns emphasized in this study are not significantly affected by potential errors in name disambiguation given the large number of scientists analysed (Supplementary Note 1). Yet, the accuracy of author name disambiguation needs to be constantly challenged and scrutinized whenever publication data is applied. The systematic nature of classification codes and their rich, hierarchical structures make them good approximations of topics in research ranging from scientific discoveries^{30,44–48} to inventions^{28,59}. But we need to understand the degree to which classification codes are good proxies for research topics better. Our ability to identify author names and research topics of papers may improve markedly, however, thanks to rapid advances in artificial intelligence and natural language processing (NLP) that may offer more comprehensive publication datasets in the near future.

Promising future directions include extending the simple model proposed here to a multidimensional random walk in the knowledge space, which may lead to a model capturing a richer set of phenomena that characterize individual careers. Other directions include extending this work to other scientific domains to address the universality and robustness of our results, investigating how the observed patterns depend on contextual information such as nations^{41,60}, institutions^{9,43}, scientific disciplines¹⁷, the size of the research community^{16,23,61}, the status of a scientist^{3,52,53} and publication habits. It would also be important to understand the short-term benefit and long-term scientific impact^{4,6,7,62–65} of research-interest change by focusing on citations instead of publications. Answering these questions could not only offer a better understanding of the fundamental mechanisms that underpin a scientific career, but might also substantially improve our ability to trace, assess, predict and nurture high-impact scientists.

Methods

Calculating topic vectors. The value of each element in the topic vector represents a topic's normalized frequency of occurrence in the set of papers analysed. Given a topic tuple, we can use a vector $X = (a_1, a_2, \dots, a_{67})$ to express the occurrence of each of the 67 topics in the topic tuple, whereby $a_i = 0$ indicates that the i th topic is not included in the topic tuple, $a_i = 1$ that the i th topic appears once, and so on. $N_X = \sum_{i=1}^{67} a_i$ is the size of the topic tuple. By normalizing X , we obtain a vector $Y = (b_1, b_2, \dots, b_{67})$ in which $b_i = a_i/N_X$ is the normalized frequency of occurrence of the i th topic. The topic vector g is calculated by averaging m different Y vectors drawn from m papers, with $g = \sum_{j=1}^m Y_j/m$. Take the calculation of g_i in Fig. 1 as an example. The two topic tuples for g_i are (68, 89, 89) and (02, 05, 68). The element value of topic 68 is calculated as $\frac{1/3+1/3}{2} = \frac{1}{3}$ as it appears once in each of the topic tuples. The element value of topic 89 is calculated as $\frac{2/3+0}{2} = \frac{1}{3}$ as it appears twice in one topic tuple and is not included in the other. The elemental values of topic 02 and 05 are calculated as $\frac{1/3+0}{2} = \frac{1}{6}$.

Measuring Δn . Δn is measured as the separation between the two subsequent appearances of the same topic tuple in a scientist's publication sequence. For example, representing distinct topic tuples as different capital letters and assuming a publication sequence is 'A A A B C C B D E', we obtain the following series of measurements: $\Delta n_A = 1$ (distance between first and second appearance of A), $\Delta n_A = 1$ (distance between the second and third appearance of A), $\Delta n_B = 3$ and $\Delta n_C = 1$. These values are then used to calculate the distribution $P(\Delta n)$.

Measuring Π . Each publication sequence is characterized by two parameters. One is the number of papers n (that is, the length of the sequence) and the other is the number of distinct topic tuples in the sequence (x). As both parameters

vary among individuals, we first fix a set of distinct topic tuples to measure their re-usage frequency Π . Here we focus on the first 5 distinct topic tuples in the sequence (that is, the maximum rank is 5). Therefore, only those sequences with $x \geq 5$ are considered. We also analyse other cases by filtering $x \geq 4$ and $x \geq 6$ and similar patterns are observed. For each qualified sequence ($x \geq 5$), we go through it from the beginning until the fifth distinct topic tuple is firstly used. We then start to count the instance where one of the five topic tuples is reused. For each individual, we obtain a fraction of time each of the five topic tuples is reused. This fraction is then averaged over all qualified sequences to generate Π .

Generating the sequence using a preferential-attachment-based model.

We apply a preferential-attachment-based model to generated topic tuple sequences with power-law distributed usage of each topic tuple^{32,37,38}. In the model, an individual's activity is randomly chosen from the two actions. One is to explore a new subject and publish a paper with a new topic tuple. The other is to return to a previously studied subject and publish a paper with a topic tuple that has already been used. The probability to explore is defined as $\rho n^{-\gamma}$ in which the term $n^{-\gamma}$ captures the decreasing trend to explore a new subject as the number of papers increases. Consequently the probability of return, that is, to reuse an old topic tuple, is $1 - \rho n^{-\gamma}$. If one returns, the choice of existing topic tuples is governed by preferential attachment: the probability p_i to use a specific topic tuple i is proportional to the tuple i 's current usage, therefore, $p_i = n_i / \sum_j n_j$ where n_i is the number of times that the tuple i is used. The parameters applied are $\rho = 0.4$ and $\gamma = 0.1$. Each individual's time step is controlled by the number of papers published, following the distribution $P(n) \sim n^{-1.5}$ with a cutoff $n_{\max} = 150$. These variables make the sequence generated similar to those in real data. We generate a total of 20,000 independent sequences, a comparable number to the size of the real data. See Supplementary Discussion 4 for more information about the preferential-attachment-based model.

Data availability. The Physical Review dataset is available upon request from the APS at <http://journals.aps.org/datasets>. The name disambiguation procedure and the associated data are described in refs^{6,43}.

Code availability. Computational codes for data processing, analysis, and model simulation are available upon request.

Received 25 July 2016; accepted 21 February 2017; published 20 March 2017

References

- Kuhn, T. S. *The Structure of Scientific Revolutions* (Univ. Chicago Press, 2012).
- de Solla Price, D. J. *Little Science, Big Science... and Beyond* (Columbia Univ. Press, 1986).
- Merton, R. K. *The Sociology of Science: Theoretical and Empirical Investigations* (Univ. Chicago Press, 1973).
- Rzhetsky, A., Foster, J. G., Foster, I. T. & Evans, J. A. Choosing experiments to accelerate collective discovery. *Proc. Natl. Acad. Sci. USA* **112**, 14569–14574 (2015).
- Kuhn, T. S. *The Essential Tension: Selected Studies in Scientific Tradition and Change* (Univ. Chicago Press, 1979).
- Sinatra, R., Wang, D., Deville, P., Song, C. & Barabási, A.-L. Quantifying the evolution of individual scientific impact. *Science* **354**, aaf5239 (2016).
- Wang, D., Song, C. & Barabási, A.-L. Quantifying long-term scientific impact. *Science* **342**, 127–132 (2013).
- Foster, J. G., Rzhetsky, A. & Evans, J. A. Tradition and innovation in scientists' research strategies. *Am. Sociol. Rev.* **80**, 875–908 (2015).
- Clauset, A., Arbesman, S. & Larremore, D. B. Systematic inequality and hierarchy in faculty hiring networks. *Sci. Adv.* **1**, e1400005 (2015).
- Jones, B. F. & Weinberg, B. A. Age dynamics in scientific creativity. *Proc. Natl. Acad. Sci. USA* **108**, 18910–18914 (2011).
- Packalen, K. Multiple successful models: how demographic features of founding teams differ between regions and over time. *Entrep. Reg. Dev.* **27**, 357–385 (2015).
- Duch, J. *et al.* The possible role of resource requirements and academic career-choice risk on gender differences in publication rate and impact. *PLoS ONE* **7**, e51332 (2012).
- West, J. D., Jacquet, J., King, M. M., Correll, S. J. & Bergstrom, C. T. The role of gender in scholarly authorship. *PLoS ONE* **8**, e66212 (2013).
- Malmgren, R. D., Ottino, J. M. & Amaral, L. A. N. The role of mentorship in protégé performance. *Nature* **465**, 622–626 (2010).
- Guimera, R., Uzzi, B., Spiro, J. & Amaral, L. A. N. Team assembly mechanisms determine collaboration network structure and team performance. *Science* **308**, 697–702 (2005).
- Jones, B. F., Wuchty, S. & Uzzi, B. Multi-university research teams: shifting impact, geography, and stratification in science. *Science* **322**, 1259–1262 (2008).
- Hoonlor, A., Szymanski, B. K. & Zaki, M. J. Trends in computer science research. *Commun. ACM* **56**, 74–83 (2013).

18. McNally, A., Prier, C. K. & MacMillan, D. W. Discovery of an α -amino c-h arylation reaction using the strategy of accelerated serendipity. *Science* **334**, 1114–1117 (2011).
19. Ericsson, K. A. in *The Cambridge Handbook of Expertise and Expert Performance* Ch. 38, 683–703 (Cambridge Univ. Press, 2006).
20. Azoulay, P., Graff Zivin, J. S. & Manso, G. Incentives and creativity: evidence from the academic life sciences. *RAND J. Econ.* **42**, 527–554 (2011).
21. Bergstrom, C. T., Foster, J. G. & Song, Y. Why scientists chase big problems: individual strategy and social optimality. Preprint at <https://arxiv.org/abs/1605.05822> (2016).
22. March, J. G. Exploration and exploitation in organizational learning. *Organ. Sci.* **2**, 71–87 (1991).
23. Vilhena, D. A. *et al.* Finding cultural holes: how structure and culture diverge in networks of scholarly communication. *Sociological Science* **1**, 221–238 (2014).
24. Hidalgo, C. *Why Information Grows: The Evolution of Order, from Atoms to Economies* (Basic Books, 2015).
25. de Solla Price, D. J. & Beaver, D. Collaboration in an invisible college. *Am. Psychol.* **21**, 1011 (1966).
26. Jones, B. F. The burden of knowledge and the death of the renaissance man: is innovation getting harder? *Rev. Econ. Stud.* **76**, 283–317 (2009).
27. Uzzi, B., Mukherjee, S., Stringer, M. & Jones, B. Atypical combinations and scientific impact. *Science* **342**, 468–472 (2013).
28. Youn, H., Strumsky, D., Bettencourt, L. M. & Lobo, J. Invention as a combinatorial process: evidence from US patents. *J. R. Soc. Interface* **12**, 20150272 (2015).
29. Cokol, M., Iossifov, I., Weinreb, C. & Rzhetsky, A. Emergent behavior of growing knowledge about molecular interactions. *Nat. Biotechnol.* **23**, 1243–1247 (2005).
30. Sinatra, R., Deville, P., Szell, M., Wang, D. & Barabási, A.-L. A century of physics. *Nat. Phys.* **11**, 791–796 (2015).
31. Gonzalez, M. C., Hidalgo, C. A. & Barabasi, A.-L. Understanding individual human mobility patterns. *Nature* **453**, 779–782 (2008).
32. Song, C., Koren, T., Wang, P. & Barabási, A.-L. Modelling the scaling properties of human mobility. *Nat. Phys.* **6**, 818–823 (2010).
33. Simini, F., González, M. C., Maritan, A. & Barabási, A.-L. A universal model for mobility and migration patterns. *Nature* **484**, 96–100 (2012).
34. Yan, X.-Y., Zhao, C., Fan, Y., Di, Z. & Wang, W.-X. Universal predictability of mobility patterns in cities. *J. R. Soc. Interface* **11**, 20140834 (2014).
35. Barabasi, A.-L. The origin of bursts and heavy tails in human dynamics. *Nature* **435**, 207–211 (2005).
36. Malmgren, R. D., Stouffer, D. B., Motter, A. E. & Amaral, L. A. A Poissonian explanation for heavy tails in e-mail communication. *Proc. Natl. Acad. Sci. USA* **105**, 18153–18158 (2008).
37. Zhao, Z.-D. *et al.* Emergence of scaling in human-interest dynamics. *Sci. Rep.* **3**, 3472 (2013).
38. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
39. Scholtes, I. *et al.* Causality-driven slow-down and speed-up of diffusion in non-markovian temporal networks. *Nat. Commun.* **5**, 5024 (2014).
40. Holme, P. & Saramäki, J. Temporal networks. *Phys. Rep.* **519**, 97–125 (2012).
41. Zhang, Q., Perra, N., Gonçalves, B., Ciulla, F. & Vespignani, A. Characterizing scientific production and consumption in physics. *Sci. Rep.* **3**, 1640 (2013).
42. Radicchi, F., Fortunato, S., Markines, B. & Vespignani, A. Diffusion of scientific credits and the ranking of scientists. *Phys. Rev. E* **80**, 056103 (2009).
43. Deville, P. *et al.* Career on the move: geography, stratification, and scientific impact. *Sci. Rep.* **4**, 4770 (2014).
44. Herrera, M., Roberts, D. C. & Gulbahce, N. Mapping the evolution of scientific fields. *PLoS ONE* **5**, e10355 (2010).
45. Radicchi, F. & Castellano, C. Rescaling citations of publications in physics. *Phys. Rev. E* **83**, 046116 (2011).
46. Pan, R. K., Sinha, S., Kaski, K. & Saramäki, J. The evolution of interdisciplinarity in physics research. *Sci. Rep.* **2**, 551 (2012).
47. Wei, T. *et al.* Do scientists trace hot topics? *Sci. Rep.* **3**, 2207 (2013).
48. Shen, Z. *et al.* Interrelations among scientific fields and their relative influences revealed by an input–output analysis. *J. Informetr.* **10**, 82–97 (2016).
49. Shi, F., Foster, J. G. & Evans, J. A. Weaving the fabric of science: dynamic network models of science’s unfolding structure. *Social Networks* **43**, 73–85 (2015).
50. Boyack, K. W., Klavans, R. & Börner, K. Mapping the backbone of science. *Scientometrics* **64**, 351–374 (2005).
51. Mandelbrote, S. *Footprints of the Lion* (Cambridge Univ. Library, 2001).
52. Petersen, A. M., Riccaboni, M., Stanley, H. E. & Pammolli, F. Persistence and uncertainty in the academic career. *Proc. Natl. Acad. Sci. USA* **109**, 5213–5218 (2012).
53. Petersen, A. M. *et al.* Reputation and impact in academic careers. *Proc. Natl. Acad. Sci. USA* **111**, 15316–15321 (2014).
54. Rybski, D., Buldyrev, S. V., Havlin, S., Liljeros, F. & Makse, H. A. Scaling laws of human interaction activity. *Proc. Natl. Acad. Sci. USA* **106**, 12640–12645 (2009).
55. Redner, S. *A Guide to First-Passage Processes* (Cambridge Univ. Press, 2001).
56. Newman, M. E. The structure of scientific collaboration networks. *Proc. Natl. Acad. Sci. USA* **98**, 404–409 (2001).
57. Araújo, E. B., Moreira, A. A., Furtado, V., Pequeno, T. H. & Andrade, J. S. Jr Collaboration networks from a large CV database: dynamics, topology and bonus impact. *PLoS ONE* **9**, e90537 (2014).
58. Shen, H.-W. & Barabási, A.-L. Collective credit allocation in science. *Proc. Natl. Acad. Sci. USA* **111**, 12325–12330 (2014).
59. Strumsky, D., Lobo, J. & Van der Leeuw, S. Using patent technology codes to study technological change. *Econ. Innov. N. Technol.* **21**, 267–286 (2012).
60. King, D. A. The scientific impact of nations. *Nature* **430**, 311–316 (2004).
61. Milojević, S. Principles of scientific research team formation and evolution. *Proc. Natl. Acad. Sci. USA* **111**, 3984–3989 (2014).
62. Bergstrom, C. T., West, J. D. & Wiseman, M. A. The eigenfactor metrics. *J. Neurosci.* **28**, 11433–11434 (2008).
63. Radicchi, F., Fortunato, S. & Castellano, C. Universality of citation distributions: toward an objective measure of scientific impact. *Proc. Natl. Acad. Sci. USA* **105**, 17268–17272 (2008).
64. Yao, L., Wei, T., Zeng, A., Fan, Y. & Di, Z. Ranking scientific publications: the effect of nonlinearity. *Sci. Rep.* **4**, 6663 (2014).
65. Ke, Q., Ferrara, E., Radicchi, F. & Flammini, A. Defining and identifying sleeping beauties in science. *Proc. Natl. Acad. Sci. USA* **112**, 7426–7431 (2015).

Acknowledgements

We thank A.-L. Barabasi for providing the initial dataset, A.-L. Barabási and G. Korniss for discussions. This work was supported by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053. T.J. is supported by the Natural Science Foundation of China (61603309) and CCF-Tencent RAGR (20160107). D.W. is supported by the Air Force Office of Scientific Research under award number FA9550-15-1-0162 and FA9550-17-1-0089. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions

T.J., D.W. and B.K.S. designed the research. T.J. performed numerical simulations and analysed the empirical data. T.J., D.W. and B.K.S. prepared the paper.

Additional information

Supplementary information is available for this paper.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence and requests for materials should be addressed to T.J., D.W. or B.K.S.

How to cite this article: Jia, T., Wang, D. & Szymanski, B. K. Quantifying patterns of research-interest evolution. *Nat. Hum. Behav.* **1**, 0078 (2017)

Publisher’s note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Competing interests

The authors declare no competing interests.