

<https://doi.org/10.1038/s44260-026-00081-3>

Team careers in science: formation, composition and success of persistent collaborations

Sandeep Chowdhary¹, Federico Musciotto², Luca Gallo^{1,3,4,6} & Federico Battiston^{1,5,6} ✉

Teams are the fundamental units propelling innovation and advancing modern science. A rich literature links the characteristics of teams, such as their size and diversity, to academic success. However, such analyses fail to capture temporal patterns, treating each group of co-authors as a distinct unit and neglecting the existence of persistent collaborations. By contrast, teams are dynamic entities, made of core members who consistently work together, surrounded by transient members who sporadically participate. Leveraging a large dataset of over 205 million scientific papers published since 1900, we extract 511,550 core teams of statistically significant persistent collaborations of pairs and larger groups of scientists. We look into the career of cores, investigating their trajectories in time, characterizing their formation, productivity, and eventual dissolution. We characterize core composition along multiple dimensions, including age, academic affiliation, and scientific disciplines. Finally, we investigate the academic impact of persistent collaborations, hallmarking the key compositional features underlying their success. Our work sheds light on the nature of persistent teams, informing researchers, institutions, and funding agencies about the dynamics of their formation, evolution, and success.

While in the past centuries lone geniuses had played a decisive role in science, in modern times scientific revolutions are more likely to emerge from the work of teams^{1,2}. In the 1930s, the *Via Panisperna boys* led by Enrico Fermi discovered slow neutrons, paving the way to the development of nuclear reactors³. Between 1947 and 1948 at Bell's Labs, John Bardeen, Walter Brattain, and William Shockley invented the first working transistors, the “nerve cells” of the Information Age⁴. In 2012, a team led by Jennifer Doudna and Emmanuelle Charpentier developed the CRISPR-Cas9 gene-editing tools, pioneering a revolution in molecular biology⁵. At a larger scale, long-lasting international collaborations such as LIGO and VIRGO⁶, CERN^{7,8}, and the Human Genome project⁹ expanded the frontiers of human knowledge. In general, research is progressively dominated by teams of scientists, as they are not only more productive than individual authors but also more impactful^{10,11} and innovative¹². Research has linked the increasing predominance of teams to the need to gather researchers with specialized expertise to solve modern-day scientific problems^{13,14}. Inevitably, understanding the determinants of teams' success in academia has become a crucial area of investigation. For instance, recent works

showed that team size plays an important role¹¹, with smaller teams often producing more innovative work¹⁵. Similarly, team diversity, in terms of geography and affiliation^{14,16–18}, but also ethnicity¹⁹, and gender^{20,21}, also significantly affect scientific impact.

Despite the vast literature on team science, understanding the temporal patterns of scientific collaborations has so far been largely overlooked. Akin to individual scientists^{22,23}, some scientific teams follow consistent career trajectories, the features of which remain largely unknown. These persistent collaborations among “science buddies”²⁴ are considered pivotal to scientific research²⁵, positively contributing to productivity and impact^{26,27}. Team formation and evolution have garnered some attention. Previous research modeled scientific collaborations as a bottom-up team formation process¹⁰, which can be driven by preferential attachments²⁸, to show how system-level features of teams, like the distribution of team sizes, can emerge from the micro-dynamics of team assembly and evolution. Crucially, the model in²⁸ distinguishes between “core” teams, which persist over time, and “extended” teams, which evolve dynamically. Yet empirical research on how the career of core teams develops throughout the years is limited. The few works

¹Department of Network and Data Science, Central European University, Vienna, Austria. ²Department of Physics and Chemistry, University of Palermo, Palermo, Italy. ³ANETI Lab, Corvinus Institute for Advanced Studies (CIAS), Corvinus University of Budapest, Budapest, Hungary. ⁴Center for Social Data Science (SODAS), University of Copenhagen, Copenhagen, Denmark. ⁵Department of AI, Data and Decision Sciences, Luiss University of Rome, Viale Romania, Rome, Italy. ⁶These authors jointly supervised this work: Luca Gallo and Federico Battiston. ✉ e-mail: battistonf@ceu.edu

accounting for the temporal characteristics of collaborations are limited to pairs of scientists²⁹, and mostly focus on their impact on the progression of individual careers^{26,30–38}. Very recently, it was revealed that teams of researchers who have not previously worked together tend to publish more original and multidisciplinary successful papers³⁹. However, at large, unraveling the temporal dimensions of teams in science represents an unexplored territory.

Here, we shift the focus and reveal the determinants of career trajectories of persistent scientific collaborations. Exploiting a large dataset of over 205 million research articles published since 1900, we build a vast hypergraph of collaborations connecting more than 90 million scientists across various disciplines, where hyperedges encode groups of co-authors who published together. Leveraging higher-order network theory⁴⁰, we use a statistically validated approach^{41,42} to reveal and extract 511,550 cores of persistent collaborators. We explore the temporal patterns behind their formation, activity, and eventual dissolution. We characterize the composition of persistent cores along multiple dimensions, including age composition, geographical diversity, and disciplinary expertise, as well as the role of transient team members. We conclude by investigating the temporal and compositional key features underlying the academic success of persistent teams.

Results

Formation, productivity and dissolution of persistent scientific collaborations

Not all members of a scientific team collaborate in the same way. In science, a team is often composed of core members, who work persistently and recurrently together over the years, surrounded by transient members, who come and go in the collaboration. Since the exact set of co-authors can change from paper to paper, this makes following the trajectories of teams much more complicated than the careers of single individuals. To study team careers, our first step is thus to identify persistent scientific collaborations from empirical data. We analyze a large dataset collected from OpenAlex^{43,44}, consisting of 205 million journal papers and conference proceedings published since 1900, covering 90 million scientists across various scientific disciplines. From the publication records, we build a hypergraph⁴⁰ of scientific collaborations, where each hyperedge encodes the set of co-authors of a paper, and use a statistically validated approach^{41,42} to extract those groups of scientists that have persistently published together over their careers (Fig. 1, see the Methods for a detailed description of the data and the methodology). Through this procedure, we identify 511,550 persistent scientific collaborations. Throughout the paper, we refer to these groups of scientists either as “cores”, “core teams”, or “persistent teams/collaborations”, while we call a “team” any group of co-author, consistently with previous literature.


We begin by investigating the typical number of members in persistent scientific teams. We compute the distribution of core sizes (Fig. 2a), finding that the greatest percentage of cores (42%) have 3 members, followed by cores of size 2 (31%) and 4 (20%). Larger persistent collaborations are rarer, with the fraction of cores of size 6 or higher being less than 1 in 100.

Typically, the formation of a core team requires a significant amount of time. Indeed, a persistent team may start as a small number of scientists working together, and gather further members around it later on. We thus ask: How fast do cores assemble? To examine this, we evaluate the time elapsing from the first publication authored by any subgroup of core members to the first publication authored by all members of the core. We refer to this quantity as the formation time. Note that, by our definition, core teams of two members are established with a formation time of zero. The average formation time for cores of different sizes is shown in Fig. 2b. We observe that smaller cores gather faster than bigger ones. In particular, cores of three members typically take 4.6 years to form, while larger cores take more and more time to gather, i.e., 7.3 years for cores of size 4, 9.2 years for cores of size 5, and 10.4 years for cores of size six. Furthermore, we note that a certain percentage of cores are formed instantaneously, i.e., they have a formation time of zero, meaning that the first publication by the core includes all of its members. Out of all cores of size 3, nearly 16.8% assembled instantaneously. For larger cores, this number drops to 5% for size 4, 2.3% for size 5, and only 1.4% for size 6.

After their formation, core teams start collaborating and publishing about their research. While the efficiency of persistent collaborations is hard to quantify, as we have no information on how long they worked on a publication, we can analyze their career in terms of the number of joint publications as a function of time. To this end, we consider all cores of a given size that have published a certain number of scientific papers and measure the average time taken to produce them (Fig. 2c). Moreover, we compute the yearly production rate, i.e., the productivity, for different core sizes (inset). Our analysis reveals that bigger cores outproduce smaller cores, taking less time on average to publish the same number of articles. Such a result calls for an in-depth analysis of how persistent scientific teams communicate, coordinate, and organize as a function of the number of their members⁴⁵. Our observation may also partly explain the increasing dominance of teams in the production of science and arts¹¹.

Though they may persist for a long time, all scientific collaborations eventually end. Therefore, we examine the typical lifespans of persistent scientific teams. We calculate the survival probability of cores as a function of the career length, i.e., the time elapsed from their first to their last publication. We find that smaller cores are typically more persistent (Fig. 2d). For instance, nearly 45% of scientific duos work together for at least 5 years, while some of them can keep working together for as long as 30 years after their first joint publication. Larger cores, instead, have shorter careers: For instance, it is highly atypical for core teams of 4 or more scientists to continue publishing together after 10 years since their first publication. Multiple factors can explain these findings. Even assuming that the members of persistent collaborations have the same chance of dropping out, because it is enough that one member leaves a core to terminate it, larger cores are more fragile⁴⁶. Additionally, as larger, persistent collaborations take more time to form (e.g., more than 10 years for cores of size six), there is a higher chance that any of their members will end their careers, thereby terminating the core. Furthermore, larger teams may require higher organizational and coordination costs⁴⁵, making them more complex to keep

Fig. 1 | Identification of cores of persistent collaborators. We build a temporal hypergraph encoding scientific collaborations, where each set of authors of a paper is represented as a hyperedge (gray-shaded areas; see Methods for details). Typically persistent core members (red), who consistently collaborate, do not work in isolation, but often publish together as part of larger teams which also include transient members (other colors). We identify sets of statistically significant collaborators and investigate the temporal careers of such persistent cores.






Fig. 2 | Formation, productivity and dissolution of persistent team cores.

a Distribution of the number of scientists per team core. **b** Formation time of cores as a function of core size. **c** Production time for a given number of papers for cores of

different sizes. Inset: average number of papers published per year as a function of the core size. **d** Survival probability of a core as a function of the career length, i.e., the time since its formation.




Fig. 3 | Composition of persistent scientific cores. **a** Percentage of each possible age composition of the core as a function of core size. **b** Percentages of cores that are spread across 1 (mono), 2 (bi), 3 (tri), and 4 or more universities as a function of core size. **c** Interdisciplinary diversity. The joint distribution of knowledge diversity

(quantifying cross-member conceptual distances) and knowledge breadth of the team (entropy of the sum of individual concept vectors of members). For mono-disciplinary cores where all members belong to the same field both knowledge breadth and core diversity are zero.

together, hence less stable. Further studies should aim at disentangling the various mechanisms underlying this pattern.

Collaboration patterns, including team dynamics, vary across scientific disciplines. We thus analyze the influence of specific scientific domains on the formation, productivity, and dissolution of persistent scientific collaborations. Our results indicate that cores in biomedical sciences, physical sciences, and engineering differ significantly in terms of size, formation time, production time, and career length (see Supplementary Fig. S1).

Composition of persistent collaborations

The members of a persistent collaboration can be identified by various characteristics, including age, academic affiliations, and scientific expertise. Understanding the composition of persistent teams in terms of the individual characteristics of their members, i.e., whether they overlap, match, or integrate, can shed light on the mechanisms that facilitate long-lasting collaborations. For instance, a persistent team may show age-homophily among its members or, instead, it may comprise young researchers and older, experienced scientists, e.g., a long-standing collaboration between

mentor and mentee. To understand the role played by age in persistent collaborations, we compute the career age of each scientist in the core (i.e., the time passed since the scientist's first publication), which we then use to characterize the age composition of persistent teams at the time of core creation (i.e., the first joint publication).

We divide scientists into three age groups, namely young ("Y", career age less than 7 years), emerging ("Em", between 7 and 14 years), and established scientists ("Es", more than 14 years). We thus assign each core to one of 7 possible categories based on the age groups of its members. We distinguish cores where all scientists belong to the same age group (only young, only emerging, or only established), cores where two age groups are represented (young + emerging, young + established, or emerging + established), and cores with scientists from all three groups. The percentage of each category in the data is shown for various core sizes in Fig. 3a. For dyadic cores, the most predominant age composition at time of assembly is a young and an established scientist working together, a typical Student-Professor motif. This is followed in frequency by the young+emerging motif. For triadic cores, the mixed age composition, i.e., Y+Em+Es, likely




Fig. 4 | Temporal dynamics and composition-based correlates of core success. **a** Number of citations received by publications of a persistent collaboration involving the Nobel laureate Emmanuelle Charpentier (Nobel Prize in Chemistry, 2020). **b** Cumulative distribution $P(\leq N^*/N)$ for cores with total papers $(N) \geq 10$, where N^*/N denotes the order of the highest-impact paper in a core's career, varying between $1/N$ and 1. The cumulative distribution of N^*/N is almost a straight line with slope ≈ 1 , indicating that N^* has roughly the same probability to occur anywhere in the sequence of papers published by a core. **c** The average impact of cores that formed in

a given decade, where cores are separated into those co-located at the same university, and others with multiple institutions. **d** Average productivity as a function of knowledge breadth of the core. Cores are divided into bins of increasing knowledge breadth. Cores are also assigned categories according to core diversity into 3 equal-sized categories – low, medium, high. **e** Average impact as a function of knowledge breadth of the core. **f** Cumulative probability distribution of average impact of team publications authored solely by the core members or involving transient, non-core members.

corresponding to the typical PhD-PostDoc-PI core, starts to be a prominent age composition. Still, the dominant composition is of only Young + Established researchers. Beyond triads, mixed cores become increasingly common. We also analyze how the formation time and career length of a core vary as a function of its age composition, revealing shorter assembly times and longer careers for core teams featuring younger members (Supplementary Figs. S2, S3).

Next, we characterize how diverse persistent collaborations are in terms of academic affiliation. We will refer to this as the affiliation diversity of the cores. To quantify it, for each paper of the core, we consider the set of affiliations that the members have at the time of publication, and evaluate the minimum number of affiliations needed to represent all members of the core. For illustration, consider a core of 3 members, A, B, and C, having the affiliations A:[MIT, UCSD], B:[MIT, UCSD], C:[MIT]. In this case, the minimum number of affiliations needed to represent the core is 1, as all members share one affiliation (MIT). Instead, if we consider the case A:[Caltech, UCSD], B:[UCSD, Indiana], C:[Caltech], we would need at least 2 affiliations (Caltech and Indiana) to cover all members of the core. As the core members can change academic affiliation during the lifetime of the collaboration, we evaluate the minimum number of affiliations for each article published and define the affiliation diversity of the core as the most common value. An example of mono-university core is that of Murray Gell-Mann and Richard Phillips Feynman (both Nobel laureates for Physics), who worked together on particle physics at the California Institute of Technology. Conversely, Alberto Abadie (the inventor of the synthetic control method for econometrics) and Guido Wilhelmus Imbens (2021 Nobel laureates in economics) moved between different universities, and while they shared their affiliation (Harvard University) during their collaboration, they published most of their papers while being at different institutions. Figure 3b shows the fraction of cores of different sizes being covered by one, two, or three universities. In contrast with an increasing trend to collaborate remotely⁴⁷, we note that 75.8% of cores are situated at the same university, hinting at the importance that common institutions play in

sustaining long-term collaborations. It is also worth noting that cores that are not based at the same university are usually located at 2 universities (21.7%), while only very rarely persistent collaborations span 3 or more institutions (2.5%). Geographically speaking, 78.1% (84.7%) of cores are situated within the same country (continent). We also find that co-presence at the same institution is associated with a shorter formation time and a longer lifespan of the core (Supplementary Figs. S2, S3).

Finally, we aim to comprehend how diverse the core members are in terms of their scientific expertise, namely, whether the core teams are mono-disciplinary or interdisciplinary. To achieve a data-driven understanding of the extent to which topic composition affects persistent collaborations, we measure two complementary dimensions of team interdisciplinarity, namely knowledge breadth and knowledge diversity. Knowledge breadth captures the breadth of the combined expertise of the persistent core. Knowledge diversity, by contrast, quantifies how diverse team members are with regard to the disciplines they are associated with. Both measures range from 0 to 1, where 1 is achieved for maximum values of multidisciplinary of the team (breadth) and topic complementarity across members (diversity; see Methods for details). Knowledge breadth and diversity measure different but complementary characteristics of the scientific expertise of the cores: The first describes how different the topics associated with the core are. The second captures how different the member of the core are in terms of their expertise. For instance, Tage Shakti Rai, Alan Fiske (both anthropologists), and Steven Pinker (a popular psychologist) form a diverse core, although it does not encompass a broad spectrum of scientific disciplines. The opposite example is the core with Dharshan Kumaran and Demis Hassabis (2024 Nobel laureate for Chemistry), who share a similar, broad expertise in neuroscience, computer science, and cognitive science.

Figure 3c shows the joint distribution of knowledge breadth and knowledge diversity across the cores. We observe that only a small fraction of cores, almost 3%, are completely mono-disciplinary, namely all members work in one scientific field. Also, we notice that the density of cores tapers off as diversity increases, with few cores having a diversity larger than 0.5.

This suggests that the core members should have a minimum amount of disciplinary overlap for their collaborations to persist. The vast majority of cores, however, are broad in their knowledge base (over 50% have more than 0.6). In particular, a significant proportion of cores (~9.1%) work at the interface of several disciplinary fields, as they display a value broadness larger than 0.75. Yet, knowledge diversity never gets close to its maximum value, which would correspond to a core where no shared topics exist between its members. This highlights the importance of topic overlap to sustain persistent collaborations. These patterns emerge from non-trivial processes underlying the formation of cores, and cannot be explained by a random selection of core members (see Supplementary Fig. S4). Moreover, we find that cores with low knowledge diversity across members take less time to form and have longer careers, further supporting the association between team persistence and topic synergy, whereas low knowledge broadness, on the other hand, is associated with shorter lifespans and a slower formation process (Supplementary Figs. S2, S3).

In the SI we discuss core exclusivity, namely the tendency of the members to publish exclusively with the core, and provide an analysis of the demographics and contribution of the transient members surrounding core teams (see Supplementary Fig. S5).

Core team success

Peak performance in individual scientific careers is known to be randomly distributed, a result known as the random impact rule²³. In other words, the most-cited article in a scientist's career can be, with an equal probability, any paper they published, from the first to the very last publication. Analyzing the publication history of persistent scientific collaborations (Fig. 4a), we aim to understand the pattern of success in core team careers, and how their composition affects their impact. To examine whether cores' careers also exhibit the random impact rule, we adopt a similar methodology as^{23,48} and measure the relative position N^* of the highest-impact paper in a core's career, i.e., in the sequence of its N publications. We measure the impact as the number of citations after five years (c_5), normalized to account for inflation and large variations between different disciplines⁴⁹. Figure 4b shows the cumulative distribution function $P(\leq N^*/N)$, namely the probability that the most-cited work of a core appears before the N^* -th publication. We observe that the function nearly follows the cumulative probability of a uniform distribution, indicating that the highest-impact article of a core can be published at any point in the career. While the most-cited article can occur at any time, the average impact of a core throughout its career can show non-random patterns. In particular, our analysis captures a higher average impact in the first half of the career compared to the second half, hinting that freshness in an early career may bring a higher impact (Supplementary Fig. S6). Our findings are in agreement with results showing how the team freshness has a positive effect on the team impact³⁹. Another universal feature of individual creative careers is the presence of hot streaks, periods of clustered high-performance works observed in various contexts, from science^{50,51} to arts and other creative domains⁵². Our results show a tendency for cores' highest-impact publications to be clustered, i.e., to occur close to each other in the career, thus confirming the presence of hot streaks (Supplementary Fig. S7).

Having characterized the diversity of persistent collaborations in terms of age, affiliations, and scientific expertise, we are now interested in understanding how those features contribute to their academic success. First, we measure the average paper impact, i.e., the average normalized number of citations after five years, for all possible core age compositions in our data. Cores composed of members from the same age groups typically perform worse than collaborations with scientists from multiple age groups (Supplementary Fig. S8). Besides the age composition, the working location of a core team can affect the impact of the work it produces. To test this, we compare the average paper impact after five years for mono and multi-university cores. In agreement with previous research¹⁴, at the aggregated level, we observe that multi-university teams are more successful than mono-university ones (Supplementary Fig. S9). Additionally, we find that international cores are more productive and gain more citations than

collaborations within a single country (Supplementary Fig. S10), similarly to previous studies^{16,53}. However, a time-resolved analysis of the average impact, consisting in comparing mono and multi-university cores as a function of the year of formation, reveals a more nuanced picture (Fig. 4c). We find that the impact advantage of multi-university cores is a recent phenomenon. Mono-university cores formed before the 2000s have on average a higher impact compared to multi-university collaborations formed during the same years. However, while more recently formed mono-university cores do not show a significantly higher impact compared to older collaborations, multi-university cores formed in the last decade have almost three times the impact of cores started in the 60s.

Hence, multi-university cores formed after the year 2000 are more successful than mono-university cores formed in the same period. A possible explanation for this observation lies in recent technological progress, from the development of computer-based mailing systems to the advent of the internet and other communication technologies, which have enhanced the experience of remote collaboration, limiting the logistical advantage associated with working in close physical proximity.

In addition to impact, we look at the productivity of core teams as a function of knowledge broadness, grouping cores based on the level of knowledge diversity, i.e., low, medium, and high diversity, respectively (Fig. 4d). We find that the relationship between productivity and knowledge broadness approximately follows a U shape. This suggests that cores with a focused approach are more productive than those with moderate levels of diversity. However, productivity increases once again for cores exhibiting the highest degrees of broadness. For a given broadness, cores with larger knowledge diversity show higher productivity. As an additional analysis, we study how productivity depends on the core composition in terms of age and academic affiliation, finding no significant impact of these features (Supplementary Fig. S11).

Next, we investigate core team success as a function of the knowledge diversity and broadness of its core members. In Fig. 4e, we show the average c_5 of the core publications as a function of the knowledge broadness, grouping again cores based on their knowledge diversity. We observe an inverted U-shaped relationship between impact and knowledge broadness of the core, for all classes of knowledge diversity, with intermediate values of broadness supporting the highest impact. When the knowledge broadness of a core is very low, the impact of a team's work might be limited to narrow disciplinary fields. Yet, when a core's knowledge base becomes too broad, its work yields a lower impact. On one side of the spectrum lie the cores such as that formed by John Lewis Selfridge, Bryant Tuckerman, Samuel Standfield Wagstaff, Derrick Henry Lehmer (one of the founding fathers of computational number theory), and John Brillhart (Lehmer's PhD student), who worked on number factorization within the Cunningham project: Their knowledge broadness is very small (0.17), and so is their normalized c_5 score (0.23). On the other side, we find collaborations like that between Laura Diaz Anadon (a researcher on climate, energy, and innovation policy, lead author for the Intergovernmental Panel on Climate Change (IPCC) 7th Assessment Report), Valentina Bosetti (an expert in environmental and climate change economics, lead author for IPCC 5th and 6th Assessment Reports), Erin Baker (focused on technological change and energy), and Lara Aleluia Reis (an expert on energy and pollution): Their expertise spans multiple scientific disciplines (their value of knowledge broadness is 0.88), but their impact remains limited (normalized c_5 score: 0.01). Similar observations were made at the level of single works, where the highest impact is obtained in papers that display a balance between conventional and atypical combinations of prior work¹². Furthermore, we find that, for almost any value of knowledge broadness, knowledge diversity has a negative relationship with the average core impact. All in all, our findings complement previous results on the complex relation between disciplinary diversity and impact in teams¹⁹.

We conclude our analysis by assessing the role of transient members on core team performance. We classify core publications into two groups, namely those authored exclusively by core members and those including other contributors, and evaluate the average impact for these two categories. Only cores with at least one publication with non-core members and one

with only core members are kept for a fair comparison. Figure 4f shows the distribution of the average c_5 across scientific cores for the two groups of papers. The analysis reveals that publications involving transient members generally show lower impact compared to those authored by the core only, suggesting that even though transient members add diversity to the team (see Supplementary Fig. S12), they might not boost impact.

Discussion

In this paper, we introduced the notion of team careers to unravel the determinants and the temporal patterns behind persistent collaborations in science. We investigated the features of half a million persistent teams and the patterns governing their formation, lifespan, composition, production, and eventually impact.

Core teams of three scientists were prevalent, highlighting the need to study team careers beyond pairwise collaborations²⁹. Larger cores formed more slowly and had shorter career lengths. Smaller cores often featured a mix of young and established researchers, while larger cores included members from all age groups. Additionally, members of persistent collaborations were affiliated with the same university for the majority of cases. A wide range of diverse degrees of disciplinary composition was observed, including core teams with large knowledge broadness, likely linked to the rising popularity of interdisciplinary research topics^{54,55}. Cores with a higher knowledge diversity were found to have shorter lifespans and longer formation times, highlighting the importance of topic synergy for persistence. A limitation of our analysis is that we focused on a specific classification of science, namely the 19 major scientific disciplines used by OpenAlex. However, classification schemes are not unique⁵⁶, and they could, in principle, lead to different results. Additionally, alternative metrics of knowledge diversity and broadness can impact our findings. Future work should thus aim at replicating our analysis, considering different classification schemes and diversity metrics.

In our examination of temporal patterns in the success of persistent teams, we found that the highest-impact paper can occur at any time, with the same probability, throughout the team's career. Over extended periods, however, we found that the average magnitude of impact is higher for publications in the first half of the career, in agreement with earlier observations that freshness is associated with high multidisciplinary impact³⁹. Besides, hot streaks were observed in core team careers, with the highest-impact papers showing a tendency to be clustered in time, mirroring findings on individual scientific^{50,51} as well as artistic careers⁵².

Previous research shows that a time-aggregated analysis captures multi-university teams to be more successful¹⁴, yet, we find that, for persistent collaborations, this is the case only since the 2000s, coinciding with the advent of ICT technologies and the internet. Recent studies indicate that remote work results in fewer bridges and encourages asynchronous communication among employees⁵⁷, while also being linked to a lower likelihood of disruptive scientific ideas as compared to onsite collaborations⁵³. Impactful core teams seemed to strike a balance in terms of knowledge broadness, while diversity among members was associated with low citation accumulation, in accordance with previous analysis at the level of single manuscripts¹², where the highest impact is obtained for papers that display a balance between the conventional and atypical combinations of prior work. Contrasting with citation impact, productivity was highest for cores that were topic-focused or ones that maximized knowledge breadth. Moreover, publications authored solely by core members yielded higher impact compared to those involving transient non-core members, hinting that expanding teams may not always enhance outcomes for persistent collaborations.

While our work reveals features of successful persistent collaborations, whether the observed patterns reflect a causal relationship between core team composition and success, or whether success is simply a prerequisite for collaborations to persist and survive, is an open question that might be clarified in future investigations. Besides, future works investigating core team composition might benefit from integrating additional information about the members, including gender and ethnicity. Indeed, we know at the

level of teams (defined as a single set of coauthors) that gender diversity among members plays a role in the impact achieved by the publication²¹, paving the way for a similar investigation on the role of gender diversity for persistent collaborations. We note that bibliometric databases do not include self-declared gender or race by authors, and inferring such metadata from names often introduces algorithmic biases⁵⁸ which are distributed unevenly among other demographic traits⁵⁹, highlighting the difficulties in properly carrying out such analysis.

In the future, our results can be combined with prior theoretical work^{10,28,60} to build a more complete theory of team assembly in science. Besides extracting persistent collaborations active over time, we can study their dynamics, namely how teams evolve and transform themselves and collectively adapt to a scientific ecosystem that constantly changes in time⁶¹. Furthermore, a cross-disciplinary analysis comparing teams in rapidly evolving fields like AI with those in more stable areas like mathematics can guide tailored team formation strategies.

Examining the role of funding on team careers can highlight how external support mechanisms can promote meaningful persistent collaborations. Indeed, the scientific ecosystem is largely shaped in response to funding. Over time, disciplines differ in terms of the funding support they receive, with interdisciplinary research so far achieving lower funding success⁶², likely impacting the longevity of such interdisciplinary teams. Moreover, grant success is strongly dependent on the collaboration network^{37,38}, particularly for young researchers, and this may incentivize participation in teams on strategic rather than scientific grounds. From the perspective of individual careers, how persistent collaborations affect the way in which researchers navigate the knowledge landscape is also an intriguing question to address⁶³. Long-term career outcomes of individuals who participate in high-impact, persistent collaborative experiences can be compared to those who work in more transient or less successful teams. Notably, early-career collaborations with elite scientists predict subsequent career success³⁵. Investigating this aspect further could provide insights into how early-career researchers can strategically navigate collaborations and switch between teams to enhance their professional growth and impact.

Overall, our work identifies persistent teams in science and captures temporal as well as compositional patterns of success in shared careers. Our research informs both scientists in building their collaborations and funders in choosing which research teams to support and promote.

Methods

Extracting persistent cores

In this work, we analyze publication data from OpenAlex^{43,44}. This database provides publication metadata, topic classification and citation records for around 205 million journal papers and conference proceedings since 1900, covering 90 million scientists disambiguated using machine learning algorithms and integration with ORCID ids of scientists to identify authors⁶⁴. We curate all scientists with substantial publication records, keeping those with at least 20 papers, resulting in 4,000,926 scientists. In this context, identifying the members of long-lasting scientific teams corresponds to detecting the maximal sets of significantly co-publishing authors⁴². In fact, these sets represent groups of authors that consistently work together, pruned from those who only occasionally have published with them. To extract them, we start by constructing the underlying hypergraph of scientific collaborations, i.e., a generalized network that naturally encodes group relationships, usually called hyperedges⁴⁰. Specifically, we construct a hypergraph whose nodes are authors and whose hyperedges represent joint publications among them.

In order to assess the statistical significance of collective interactions among sets of nodes in a hypergraph, one needs to take into account the heterogeneity of node activity. Indeed, if on one side the easiest way would be setting a fixed threshold on the minimum number of repeated interactions needed to consider a group of nodes a statistically validated set, this approach is sub-optimal, due to the multiscale nature of collaboration networks and the varying activity levels among scientists (the same threshold can be too restrictive for an author with limited publication

records and very permissive for a more prolific author). To overcome this limitation, our method is based on a null hypothesis approach, that naturally tunes the threshold on the activity of the involved nodes. For simplicity, we start from the case of 3 nodes i, j, k that co-interact N_{ijk} times in hyperedges of size $n \geq 3$. The three nodes appear respectively in N_i, N_j, N_k hyperedges. Under the null hypothesis that each node selects randomly the hyperedges to which it participates - and thus its $n - 1$ counterparts in a hyperedge of size n - the probability of observing i, j, k interacting N_{ijk} times is

$$p(N_{ijk}) = \sum_X H(X|N, N_i, N_j) \times H(N_{ijk}|N, X, N_k) \\ = \frac{1}{\binom{N}{N_i} \binom{N}{N_j} \binom{N}{N_k}} \sum_X \binom{N_i}{X} \binom{N - N_i}{N_j - X} \binom{X}{N_{ijk}} \binom{N - X}{N_k - N_{ijk}}, \quad (1)$$

where $H(N_{AB}|N, N_A, N_B)$ is the hypergeometric distribution that computes the probability of having an intersection of size N_{AB} between two sets A and B of size N_A and N_B given N total elements. The probability $p(N_{ijk})$ in Eq. (1) represents the probability of having a random intersection of size N_{ijk} between the three sets of hyperedges of nodes i, j, k out of N total hyperedges⁶⁵, and is obtained through the convolution of two instances of the hypergeometric distribution. Starting from Eq. (1), we then compute a p -value for the triplet that contains i, j and k through the survival function,

$$p(x \geq N_{ijk}) = 1 - \sum_{x=0}^{N_{ijk}-1} p(x). \quad (2)$$

The p -value represents the probability of observing N_{ijk} or more hyperedges that contain - but are not limited to - the nodes i, j, k . The smaller the p -value, the higher the possibility that i, j, k constitute a significant set of size 3.

In the approach of refs. 41,42, N corresponds to the total number of papers in the corpus. In our case, this choice for the value of N implies that all hyperedges are equally accessible by all nodes. It is, however, unrealistic that a scientist could have participated in all papers, due to limited time and resources. Thus, we decided to bound the number of papers a scientist or a group of scientists could have worked on. Specifically, we approximate N by summing the number of unique publications by the scientists and all their coauthors, for each of the members of the group we are testing for significance.

For a generic hyperedge of n nodes, Eq. (1) becomes

$$p(N_{1...n}) = \sum_{X_{12}} H(X_{12}|N, N_1, N_2) \times \\ \times \sum_{X_{123}} H(X_{123}|N, X_{12}, N_3) \times \dots \\ \dots \times \sum_{X_{12...n-1}} H(X_{12...n-1}|N, X_{12...n-2}, N_{n-1}) \times \\ \times H(N_{12...n}|N, X_{12...n-1}, N_n). \quad (3)$$

How do we set a rigorous criterion to assess whether a group of n scientists is statistically significant, once we have calculated the associated p -value? We test all p -values against a threshold of statistical significance α , corrected by a control for the False Discovery Rate (FDR). This means that we include a multiple hypothesis test correction, needed because of the high number of tests - one per each group. Specifically, we apply the Benjamini-Hochberg procedure⁶⁶ starting from a Bonferroni threshold equal to $\alpha_t = \alpha/N_t$, where $\alpha = 0.01$ and N_t is the total number of tests run for hyperedges of size n . It is worth noting that the procedure is robust in the presence of correlation between the different tests. This is meaningful in our case because the formation of teams, assembled from a fixed set of authors, constitutes a sequence of events that are not independent from each other. Following⁴², we start testing from the

largest set and we then proceed towards the smallest. If a set of size n passes our statistical test and is thus selected as significant because it rejects the null hypothesis, we do not test any of its smaller subsets. In this way, the obtained statistically significant sets can be considered maximal, i.e. for each of them there is no superset that includes all its members that is also statistically significant.

The extraction procedure resulted in over half-a-million persistent collaborations, with size ranging from 2 to 10. Yet, we limit our analysis to cores of size 2 to 6, as bigger cores are rare and statistics are insufficient for an in-depth analysis.

Knowledge broadness and knowledge diversity

We start from the OpenAlex dataset, where researchers are associated with different scientific concepts with a score varying between 0 and 100. The concepts are organized in a hierarchy with 19 root-level concepts representing major disciplines such as physics, chemistry, computer science and so on, and 5 layers of descendants branching out from them, for a total of 65,000 concepts⁶⁷. For our analysis, we consider the topmost layer of the concept hierarchy, namely the scientific disciplines. For each member of a core team, we evaluate a knowledge vector where each component represents how strongly that scientist is associated with a scientific discipline. Each entry of the vector consists of the topic score normalized by the sum of the scores.

Knowledge broadness captures the breadth of the combined expertise of the core. To calculate it, we first consider the sum of the knowledge vectors of the core members, normalized so that the components of the combined knowledge vector sum to 1. We then evaluate the entropy of this combined knowledge vector, calculated as $-\sum_x p_x \log_2 p_x$, where x is a scientific discipline and p_x represent how strongly the core is connected to it. Finally, to obtain the knowledge broadness, the entropy of the vector is normalized by its maximum theoretical value of $\log_2(N)$ which corresponds to the case where the core is uniformly spread across all N disciplines. In particular, as Openalex has 19 disciplines in the top layer of its concept hierarchy. For illustration, consider a core of three members, A, B, and C, whose knowledge vectors are A: [physics: 0.5, chemistry: 0.5], B: [physics: 0.7, chemistry: 0.1, biology: 0.2], C: [physics: 1]. The combined knowledge vector for it would be: [physics: 0.85, chemistry: 0.077, biology: 0.077]. The knowledge broadness of the core is thus ~ 0.18 . The value of knowledge broadness ranges from 0 to 1. A value of 0 indicates a core where all members are associated to a single topic, i.e., a monodisciplinary core, while a value of 1 corresponds to a core for which the joint knowledge vector is uniformly distributed across all possible topics in our data.

Knowledge diversity, on the other hand, quantifies how much core members are different with regard to the scientific disciplines they are associated with. Given the knowledge vectors of the core members, we quantify the knowledge diversity as 1 minus the average cosine similarity between all pairs of knowledge vectors. For instance, in the example above, the cosine similarity between members A and B (~ 0.77) is calculated, then B and C (~ 0.95), then C and A (~ 0.71). The values are then averaged and subtracted from 1 to obtain the knowledge diversity of the core (~ 0.19). Knowledge diversity ranges between 0 and 1, where a value of 0 indicates all members have the identical knowledge vectors, i.e., all members are associated with the same strength to the same disciplines, whereas a value of 1 indicates that knowledge vectors of all members are non-overlapping, i.e., they are associated to completely different scientific disciplines.

Knowledge broadness and knowledge diversity cover complementary dimensions of team topic composition, as the first describes how different the topics associated to the team are, while the second captures how the members of the team are diverse among each other. For instance, a collaboration among scientists working in the same discipline will have low knowledge diversity and low knowledge broadness, while a team where members are associated with the same set of multiple disciplines will have low knowledge diversity but high knowledge broadness.

Data availability

The data from openalex.org used in this work is openly accessible for download using the API <https://api.openalex.org/works>.

Code availability

The code used in this study is available at <https://github.com/chowdhary-sandeep/sciscicareers>.

Received: 3 April 2025; Accepted: 19 March 2026;

Published online: 15 June 2026

References

- Simonton, D. K. Scientific genius is extinct. *Nature* **493**, 602–602 (2013).
- Bozeman, B. & Boardman, C. *Research Collaboration and Team Science: A State-of-the-art Review and Agenda*, Vol. 17 (Springer, 2014).
- Segrè, G. & Hoerlin, B. *The pope of Physics: Enrico Fermi and the birth of the atomic age* (Henry Holt and Company, 2016).
- Riordan, M., Hoddeson, L. & Herring, C. The invention of the transistor. *Rev. Mod. Phys.* **71**, S336 (1999).
- Ledford, H. & Callaway, E. Pioneers of crispr gene editing win chemistry nobel. *Nature* **586**, 346–347 (2020).
- Collaboration, L. S. & Collaboration, V. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.* **116**, 061102 (2016).
- Collaboration, A. Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc. *Phys. Lett. B* **716**, 1–29 (2012).
- Collaboration, C. Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc. *Phys. Lett. B* **716**, 30–61 (2012).
- Consortium, I. H. G. S. Initial sequencing and analysis of the human genome. *nature* **409**, 860–921 (2001).
- 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).
- Wuchty, S., Jones, B. F. & Uzzi, B. The increasing dominance of teams in production of knowledge. *Science* **316**, 1036–1039 (2007).
- Uzzi, B., Mukherjee, S., Stringer, M. & Jones, B. Atypical combinations and scientific impact. *Science* **342**, 468–472 (2013).
- Hunter, L. & Leahey, E. Collaborative research in sociology: trends and contributing factors. *Am. Sociol.* **39**, 290–306 (2008).
- Jones, B. F., Wuchty, S. & Uzzi, B. Multi-university research teams: shifting impact, geography, and stratification in science. *Science* **322**, 1259–1262 (2008).
- Wu, L., Wang, D. & Evans, J. A. Large teams develop and small teams disrupt science and technology. *Nature* **566**, 378–382 (2019).
- Hsiehchen, D., Espinoza, M. & Hsieh, A. Multinational teams and diseconomies of scale in collaborative research. *Sci. Adv.* **1**, e1500211 (2015).
- Coccia, M. & Wang, L. Evolution and convergence of the patterns of international scientific collaboration. *Proc. Natl. Acad. Sci. USA* **113**, 2057–2061 (2016).
- Gazni, A., Sugimoto, C. R. & Didegah, F. Mapping world scientific collaboration: authors, institutions, and countries. *J. Am. Soc. Inf. Sci. Technol.* **63**, 323–335 (2012).
- AlShebli, B. K., Rahwan, T. & Woon, W. L. The preeminence of ethnic diversity in scientific collaboration. *Nat. Commun.* **9**, 1–10 (2018).
- Nielsen, M. W. et al. Gender diversity leads to better science. *Proc. Natl. Acad. Sci. USA* **114**, 1740–1742 (2017).
- Yang, Y., Tian, T. Y., Woodruff, T. K., Jones, B. F. & Uzzi, B. Gender-diverse teams produce more novel and higher-impact scientific ideas. *Proc. Natl. Acad. Sci. USA* **119**, e2200841119 (2022).
- Fortunato, S. et al. Science of science. *Science* **359**, eaao0185 (2018).
- Sinatra, R., Wang, D., Deville, P., Song, C. & Barabási, A.-L. Quantifying the evolution of individual scientific impact. *Science* **354**, aaf5239 (2016).
- Yanai, I. & Lercher, M. J. It takes two to think. *Nat. Biotechnol.* **42**, 18–19 (2024).
- Hoy, A. Q. Science diplomacy leverages alliances to build global bridges. *Science* **365**, 875–876 (2019).
- Petersen, A. M. Quantifying the impact of weak, strong, and super ties in scientific careers. *Proc. Natl. Acad. Sci. USA* **112**, E4671–E4680 (2015).
- Liu, J. et al. Understanding super-partnerships in scientific collaboration: evidence from the field of economics. *J. Assoc. Inf. Sci. Technol.* **75**, 717–733 (2024).
- Milojević, S. Principles of scientific research team formation and evolution. *Proc. Natl. Acad. Sci. USA* **111**, 3984–3989 (2014).
- Son, G., Yun, J. & Jeong, H. Untangling pair synergy in the evolution of collaborative scientific impact. *EPJ Data Sci.* **12**, 62 (2023).
- Malmgren, R. D., Ottino, J. M. & Nunes Amaral, L. A. The role of mentorship in protégé performance. *Nature* **465**, 622–626 (2010).
- Azoulay, P., Graff Zivin, J. S. & Wang, J. Superstar extinction. *Q. J. Econ.* **125**, 549–589 (2010).
- Borjas, G. J. & Doran, K. B. Which peers matter? the relative impacts of collaborators, colleagues, and competitors. *Rev. Econ. Stat.* **97**, 1104–1117 (2015).
- Liénard, J. F., Achakulvisut, T., Acuna, D. E. & David, S. V. Intellectual synthesis in mentorship determines success in academic careers. *Nat. Commun.* **9**, 4840 (2018).
- Sekara, V. et al. The chaperone effect in scientific publishing. *Proc. Natl. Acad. Sci. USA* **115**, 12603–12607 (2018).
- Li, W., Aste, T., Caccioli, F. & Livan, G. Early coauthorship with top scientists predicts success in academic careers. *Nat. Commun.* **10**, 5170 (2019).
- Krauss, A., Danús, L. & Sales-Pardo, M. Early-career factors largely determine the future impact of prominent researchers: evidence across eight scientific fields. *Sci. Rep.* **13**, 18794 (2023).
- Chowdhary, S., Defenu, N., Musciotto, F. & Battiston, F. Dependency of erc-funded research on us collaborations. *Nat. Phys.* **19**, 1746–1749 (2023).
- Chowdhary, S., Defenu, N., Musciotto, F. & Battiston, F. Funding bias: nurture European researchers' independence. *Nature* **616**, 33 (2023).
- Zeng, A., Fan, Y., Di, Z., Wang, Y. & Havlin, S. Fresh teams are associated with original and multidisciplinary research. *Nat. Hum. Behav.* **5**, 1314–1322 (2021).
- Battiston, F. et al. Networks beyond pairwise interactions: structure and dynamics. *Phys. Rep.* **874**, 1–92 (2020).
- Musciotto, F., Battiston, F. & Mantegna, R. N. Detecting informative higher-order interactions in statistically validated hypergraphs. *Commun. Phys.* **4**, 218 (2021).
- Betti, L., Musciotto, F., Battiston, F. & Mantegna, R. N. Identifying maximal sets of significantly interacting nodes in higher-order networks. Preprint at <https://arxiv.org/abs/2209.12712> (2022).
- Openalex: a fully-open index of scholarly works, authors, venues, institutions, and concepts. <https://openalex.org/>.
- Chawla, D. S. Massive open index of scholarly papers launches. *Nature*. <https://doi.org/10.1038/d41586-022-00138-y> (2022).
- Hall, K. L. et al. The science of team science: a review of the empirical evidence and research gaps on collaboration in science. *Am. Psychologist* **73**, 532 (2018).
- Coleman, J. S. & James, J. The equilibrium size distribution of freely-forming groups. *Sociometry* **24**, 36–45 (1961).
- Teasley, S. D. & Wolinsky, S. Scientific collaborations at a distance. *Science* **292**, 2254–2255 (2001).
- Janosov, M., Battiston, F. & Sinatra, R. Success and luck in creative careers. *EPJ Data Sci.* **9**, 9 (2020).

49. Radicchi, F., Fortunato, S. & Castellano, C. Universality of citation distributions: toward an objective measure of scientific impact. *Proc. Natl. Acad. Sci.* **105**, 17268–17272 (2008).
50. Liu, L. et al. Hot streaks in artistic, cultural, and scientific careers. *Nature* **559**, 396–399 (2018).
51. Liu, L., Dehmamy, N., Chown, J., Giles, C. L. & Wang, D. Understanding the onset of hot streaks across artistic, cultural, and scientific careers. *Nat. Commun.* **12**, 5392 (2021).
52. Williams, O. E., Lacasa, L. & Latora, V. Quantifying and predicting success in show business. *Nat. Commun.* **10**, 2256 (2019).
53. Lin, Y., Frey, C. B. & Wu, L. Remote collaboration fuses fewer breakthrough ideas. *Nature* **623**, 987–991 (2023).
54. Porter, A. & Rafols, I. Is science becoming more interdisciplinary? measuring and mapping six research fields over time. *Scientometrics* **81**, 719–745 (2009).
55. Van Noorden, R. et al. Interdisciplinary research by the numbers. *Nature* **525**, 306–307 (2015).
56. Milojević, S. Practical method to reclassify web of science articles into unique subject categories and broad disciplines. *Quant. Sci. Stud.* **1**, 183–206 (2020).
57. Yang, L. et al. The effects of remote work on collaboration among information workers. *Nat. Hum. Behav.* **6**, 43–54 (2022).
58. Kozłowski, D. et al. Avoiding bias when inferring race using name-based approaches. *Plos one* **17**, e0264270 (2022).
59. Lockhart, J. W., King, M. M. & Munsch, C. Name-based demographic inference and the unequal distribution of misrecognition. *Nat. Hum. Behav.* **7**, 1084–1095 (2023).
60. Milojević, S. Modes of collaboration in modern science: Beyond power laws and preferential attachment. *J. Am. Soc. Inf. Sci. Technol.* **61**, 1410–1423 (2010).
61. Galesic, M. et al. Beyond collective intelligence: Collective adaptation. *J. R. Soc. Interface* **20**, 20220736 (2023).
62. Bromham, L., Dinnage, R. & Hua, X. Interdisciplinary research has consistently lower funding success. *Nature* **534**, 684–687 (2016).
63. Singh, C. K., Tupikina, L., Lécuyer, F., Starnini, M. & Santolini, M. Charting mobility patterns in the scientific knowledge landscape. *EPJ Data Sci.* **13**, 12 (2024).
64. Openalex: Author disambiguation. <https://docs.openalex.org/api-entities/authors/author-disambiguation>.
65. Wang, M., Zhao, Y. & Zhang, B. Efficient test and visualization of multi-set intersections. *Sci. Rep.* **5**, 16923 (2015).
66. Benjamini, Y. & Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc.* **57**, 289–300 (1995).
67. Openalex: Concepts. <https://docs.openalex.org/api-entities/concepts>.

Acknowledgements

L.G. and F.B. acknowledge support from the Air Force Office of Scientific Research under award number FA8655-22-1-7025. L.G. acknowledges

support from the Villum Foundation (project no. 57396) at the University of Copenhagen. F.B. acknowledges support from the Austrian Science Fund (FWF) through projects 10.55776/PAT1052824 and 10.55776/PAT1652425.

Author contributions

S.C. led the study, processed the data and performed the analysis. F.M. provided methodological insights and curated the data. L.G. provided methodological insights, performed the analysis, and directed the study. F.B. provided methodological insights, directed and supervised the study. S.C., L.G., F.B. substantially contributed to the interpretation of the analysis, all authors revised the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at

<https://doi.org/10.1038/s44260-026-00081-3>.

Correspondence and requests for materials should be addressed to Federico Battiston.

Reprints and permissions information is available at <http://www.nature.com/reprints>

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

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026