Scientometrics (2026) 131:1119–1146
https://doi.org/10.1007/s11192-026-05560-x

Does Diversity of Expertise Drive Citation Impact? Evidence from Computer Science

Angelo Salatino[1] · **Francesco Osborne**[1,2] · **Diego Reforgiato Recupero**[3] · **Simone Angioni**[4] · **Enrico Motta**[1]

Received: 28 February 2025 / Accepted: 22 January 2026 / Published online: 14 February 2026

Abstract

High-quality scientific research plays a pivotal role in advancing society, stimulating economic growth, protecting the environment, and driving technological innovation. Understanding the key factors that lead to impactful research is thus crucial as it can steer the development of more effective policies to enhance the research enterprise. Extensive literature emphasises that the composition of a research team is vital for generating innovative and impactful scientific work. Many studies have focused on how team diversity, including aspects like ethnicity, gender, and international background, affects research outcomes. These types of diversities often correlate positively with the impact of research. In this paper, we investigate a less-explored dimension of diversity: the diversity of authors' areas of expertise. This aspect has received limited attention, primarily due to the challenges involved in defining and measuring it. We present new AI-driven methods to quantify this diversity of expertise and conduct an extensive analysis of over 944,000 Computer Science papers. Specifically, this study investigates the relationship between the authors' diversity of expertise and the number of citations their paper receives within the first five years. For each paper, we modelled the expertise of each individual author and then quantified the overall diversity within the author team. We then performed a statistical analysis that revealed a significant positive correlation between two diversity metrics and the number of citations received. This suggests that, in the field of Computer Science, diversity of expertise is a key driver of high-impact research.

Keywords Team science · Science of science · Scientometrics · Research impact · Diversity of expertise · Metascience

Introduction

High-quality scientific research plays a pivotal role in advancing society, stimulating economic growth, protecting the environment, and driving technological innovation. Understanding the key factors that led to impactful research is thus crucial as it can steer the development of more effective policies to enhance the research enterprise.

Extended author information available on the last page of the article

Extensive literature emphasises that the composition of a research team is one of the crucial factors in generating innovative and impactful scientific work (Wang & Barabási, 2021). Several relevant studies have been conducted in team science, the discipline that investigates how scientific teams form, communicate, collaborate, and produce knowledge (Stokols et al., 2008). Many of these studies focused, in particular, on the different dimensions of diversity, such as nationality (Smith et al., 2014), ethnicity (AlShebli et al., 2018; Freeman & Huang, 2015), institutions (Jones et al., 2008), gender (Nielsen et al., 2017), academic age (Jones & Weinberg, 2011), disciplinary backgrounds (Uzzi et al., 2013), and team size (Wu et al., 2019), highlighting their influence on scientific outcomes. Nathan et al. (Nathan & Lee, 2013) conclude that diversity fosters the inclusion of different perspectives and approaches, leading to novel and creative solutions for complex and multifaceted problems.

One aspect that received less attention is the diversity of expertise within research teams, i.e., the variety of skills and scientific backgrounds among team members. Indeed, only a limited number of studies have examined the relationship between expertise diversity and research impact (Lungeanu et al., 2023; Mishra et al., 2025; Zheng et al., 2022). This scarcity is likely due to the challenges involved in accurately measuring individual expertise and the limitations of existing academic data in capturing nuanced specialisations. Moreover, the available studies often focus on specific domains such as biomedicine (Mishra et al., 2025), or rely on overly simplistic methods to characterise authors' expertise (Lungeanu et al., 2023; Zheng et al., 2022).

Nevertheless, in recent years, there has been an increasing focus in the academic community, including funding agencies and governmental bodies, towards recognising the importance of interdisciplinary approaches and collaborations across various scientific domains (Cooke et al., 2020; Olechnicka et al., 2019). This emphasis is particularly vital for tackling multifaceted societal challenges that require the integration of insights from diverse fields, such as climate change, health issues, poverty, and social inequality (for Economic Co-operation & Development, 2020). Consequently, it has become imperative to comprehend and measure the impact of diversity in expertise within research teams.

In this paper, we present a scientometric analysis that investigates the relationship between the diversity of expertise within authoring teams and their scientific impact. To this purpose, we employ AI-driven techniques to measure the team's diversity of expertise and examine its correlation with the citations their publications receive over the subsequent five years.

We characterised each researcher's expertise as the distribution of the 10 most prominent research topics in their publications from the five years preceding their collaboration. We then calculated the pairwise cosine similarity between all pairs of authors within each paper. For example, authors with identical research focuses would have a similarity score of 1, while those with completely divergent expertise would score 0. Finally, we assessed the diversity of expertise within each paper using two metrics: (i) the maximum cosine distance between any pair of authors, and (ii) the number of connected components formed by linking authors with a similarity above a specified threshold.

The research topics for each paper were identified using the CSO Classifier, a tool we previously developed to associate textual content with topics from the Computer Science Ontology (CSO) (Salatino et al., 2019). CSO is a large-scale ontology which covers 14K research topics. It offers a more fine-grained representation compared to the generic ones provided by well-known scholarly corpora including OpenAlex. As a result, it has been effectively utilised for exploring and analysing scholarly data (Löffler et al., 2020; Wahle et al., 2023; Zhang et al., 2021), as well as for modelling, identifying, and recommending domain experts (Konstantinidis et al., 2022; Rahdari et al., 2021; Vergoulis et al., 2020).

The findings of our study reveal a significant correlation between the two diversity metrics and citation counts of the relevant papers over a five-year period. Specifically, research articles authored by teams with diverse skill sets and expertise tend to achieve a greater impact than those produced by more homogeneous teams. We also conducted an in-depth analysis to identify specific collaboration patterns between research areas that produce the most impactful work. For instance, the most productive three-way collaboration during the period under analysis was between authors specialising in *artificial intelligence*, *data mining*, and the *internet*.

To assess the robustness of our findings, we also conducted a multivariate negative binomial regression analysis in which the 5-year citation count is predicted by our diversity metrics, while adjusting for publication year, team size, and authors' seniority. The results indicate a statistically significant positive association between diversity of expertise and citation impact, which remains robust after controlling for potential confounding factors.

A preliminary version of this analysis was presented at the 27th International Conference on Science, Technology and Innovation Indicators, also known as STI 2023 (Salatino et al., 2023). The present manuscript builds upon this previous work by utilising a substantially larger dataset, expanding from 114K to 944K papers, covering a broader time span, and presenting additional statistical tests.

In summary, the contributions of this paper are the following:

- the development of two novel metrics to quantify the diversity of expertise within research teams;
- a comprehensive analysis of how the diversity of expertise within research teams influences their research outcomes, based on an extensive dataset of 944K papers;
- an investigation into the combinations of topics that were most conducive to successful collaborations during 2005-2015;
- the release of the complete codebase used in the experiments, facilitating replication and further enhancement of the analysis by other researchers.

The remainder of the paper is organised as follows. Section 2 provides an overview of the relevant literature. Section 3 describes the dataset used in this study and outlines the research methodology. Section 4 presents the findings and provides a discussion of the results. Sect. 5 discusses the significance of the contributions and outlines the limitations. The paper concludes with Sect. 6, which summarises the key outcomes and suggests avenues for future research.

Literature Review

Diversity of research team has been analysed across several dimensions in the literature, including nationality (Smith et al., 2014), ethnicity (AlShebli et al., 2018; Freeman & Huang, 2015), institutions (Jones et al., 2008), gender (Nielsen et al., 2017), academic age (Jones & Weinberg, 2011), disciplinary backgrounds (Uzzi et al., 2013), and team size (Wu et al., 2019). All these studies concur that more diverse research teams often demonstrate higher productivity in terms of number of publications, or impact as number of citations received over a period of time.

With regard to nationality, Smith et al. (2014) found that international collaborations can increase the quality and visibility of scientific works, positively influencing citation performance.

Ethnic diversity was analysed by Freeman and Huang (Freeman & Huang, 2014), finding that papers authored by four or five authors of diverse ethnicities receive more citations than those authored by individuals of the same ethnicity. In a follow-up experiment (Freeman & Huang, 2015), the same authors discovered that publications with authors of different ethnic backgrounds and from various geographical locations, alongside longer reference lists, are more likely to be accepted in higher impact journals and receive more citations. Similarly, AlShebli et al. (2018) examined combined effects of diversity (e.g., ethnicity, gender, academic age, and affiliations) on research impact. They revealed that ethnicity is the most significant factor, associated with an impact gain of 10.63%.

Collaborating with someone from another institution is typically associated with higher impact and productivity (Freeman & Huang, 2015; Jones et al., 2008; Whitfield, 2008), as it allows for the exchange of diverse perspectives, know-how, and resources. Indeed, Jones et al. (2008) showed that in all fields of science, engineering, and social science, multi-university research teams obtained a higher citation impact on their papers. De Saá-Pérez et al. (2017) examined the effects of institutional diversity, alongside educational, seniority, and role. They found that teams with more institutional or seniority diversity achieved higher performance in terms of published articles. However, this effect was weaker for role and educational diversity.

Team size is another factor influencing research impact and productivity. Research by Cummings et al. (2013) revealed a positive correlation between the size of research groups and their publication output. Wuchty et al. (2007) further supported this finding by analysing a vast dataset of publications and patents, showing that team-authored work received more citations than solo-authored work. A more recent study by Wu et al. (2019) found that while smaller teams may work on less popular yet potentially disruptive ideas with delayed recognition, larger teams tend to focus on more popular ideas with quicker citation impact.

Gender diversity in research teams has been widely recognised as a key factor for scientific excellence and innovation, as it can enrich the perspectives, questions, and domains of the researchers (Biermann, 2023; Bozeman & Gaughan, 2011; Nielsen et al., 2017, 2018; Yang et al., 2022). Bozeman and Gaughan (2011) suggest that promoting gender diversity in research collaborations can lead to higher impact, as women's collaboration patterns and strategies can bring unique perspectives and approaches to research. Yang et al. (2022) investigated the effects of mixed-gender research teams across the medical science. They discovered that the publications of mixed-gender teams are substantially more novel and impactful than those of same-gender teams of equivalent size.

Disciplinary diversity has a complex relationship with citation rates and research impact, depending on the field and timeframe considered. Larivière and Gingras (2010) observed that interdisciplinary was positively correlated with citation rates in biology, clinical medicine, humanities, psychology, and health (social sciences), but negatively correlated with citation rates in physics, earth and space sciences, and professional fields. However, this analysis considers articles published in 2000, which may not be representative of current trends in interdisciplinary and scientific impact. Instead, Zheng et al. (2022) found out that teams with greater disciplinary diversity enjoy a high impact of their work after 10 years, and non significant impact after 5 years.

Although the relationship between team diversity and research impact has been extensively studied across several dimensions, understanding how the *diversity of expertise* among team members influences citation impact has received less attention. This is primarily due to the inherent difficulties in accurately measuring the researcher's expertise, as well as the inability of current academic datasets to capture this kind of specialisation. Despite this challenge, we identified a few key studies that produced some initial results in this space (Li & Zheng, 2025; Lungeanu et al., 2023; Mishra et al., 2025; Nandy et al., 2024). Lungeanu et al. (2023) investigated a research question closely related to ours, although their analysis focused on innovation as evidenced by patents rather than by research publications. They conceptualised the inventor's expertise as a 300-dimensional vector, derived by averaging the Doc2Vec embeddings extracted from their previous patents. Subsequently, they used pairwise similarities to identify diverse inventor pairs. While insightful, we identify two primary limitations in their methodology, both stemming from their approach to modelling inventor expertise. First, the fields of expertise are not explicitly defined but are instead represented implicitly through embedding. Second, the practice of averaging all Doc2Vec embeddings from the inventor's prior patents could potentially lead to a loss of information, especially when compressing long documents into a single fixed-length vector.

Li and Zheng (2025) observed that teams with diverse expertise tend to produce more original work, which often results in a significantly greater long-term impact, still evident even after a decade. Their experiment has been conducted on research papers and patents from Microsoft Academic Graph (MAG). However, their methodology for modelling author expertise relies solely on 292 "level 1" MAG's Fields of Study. Nandy et al. (2024) analysed the diversity of expertise at the institutional level in order to evaluate the diversity of their research portfolios. In this context, they employed Web of Science as a dataset and considered its 254 categories to model the authors' expertise. In contrast, our approach leverages CSO, which encompasses over 14K research concepts, offering a far more granular and comprehensive representation of expertise. In the field of biomedicine, Mishra et al. (2025) conducted a study comparable in both scale and setting to ours. They modelled authors' expertise by aggregating the most frequent MeSH (Medical Subject Headings) terms from their prior publications. In contrast, our analysis focuses on the field of computer science, and crucially, our method for identifying expertise involves selecting the most relevant topics, capturing a more central and fine-grained author's specialisation, rather than merely the most frequent ones.

Materials and Methods

In this section, we describe the dataset employed for our analysis and the methodology adopted to assess the diversity of expertise and its impact.

AIDA Knowledge Graph

As main data source, we used the Academia/Industry DynAmics Knowledge Graph (AIDA KG) (Angioni et al., 2021), which is a large-scale knowledge base that describes 25 million publications in the field of Computer Science. We generated AIDA KG by applying an automatic pipeline that integrates data from several sources, such as OpenAlex,[1] DBLP,[2] CSO (Salatino et al., 2018), Research Organization Registry[3] (ROR), and DBpedia.[4] AIDA KG offers a diverse range of metadata associated with seven key entities: papers, authors, affiliations, journals, conferences, topics, and industrial sectors. This rich metadata can be leveraged to enhance various types of scientometric analyses (Angioni et al., 2022; Salatino et al., 2020). For a comprehensive overview of this resource, we direct the reader to Angioni et al. (2021).

We classified all documents in the AIDA Knowledge Graph (KG) using the CSO Classifier (Salatino et al., 2019), a tool we developed to identify topics from the CSO (Salatino et al., 2018) based on a research paper's title and abstract. CSO is the most extensive knowledge organization system in the field of Computer Science (Salatino et al., 2025), featuring over 14K research topics. It provides a highly granular representation of the domain, spanning 13 hierarchical levels. Each topic is also linked to a set of alternative labels that assist both human annotators and automated systems in classifying and organising scholarly documents. These features make CSO a powerful tool for exploring and analysing scholarly data (Löffler et al., 2020; Wahle et al., 2023; Zhang et al., 2021) as well as for modelling, identifying, and recommending domain experts (Konstantinidis et al., 2022; Rahdari et al., 2021; Vergoulis et al., 2020).

The CSO Classifier implements an unsupervised classifier that operates through three modules. The first is the *syntactic module*, which identifies topics explicitly mentioned in the text by matching word sequences (unigrams, bigrams, and trigrams) to labels in the CSO ontology using Levenshtein similarity. Next, the *semantic module* employs a pre-trained Word2Vec model (Mikolov et al., 2013) to discover semantically related topics. It first identifies candidate terms based on part-of-speech tagging and then measures their cosine similarity to CSO labels using word embeddings. Finally, the *post-processing module* enhances the resulting topics with higher-level topics from the CSO hierarchy to provide a more comprehensive topical representation of the paper. For further details on the CSO Classifier and a comprehensive evaluation of its performance, we refer the reader to Salatino et al. (2019). Since its release, the CSO Classifier has attracted increasing attention. For example, the Springer Nature editorial team routinely employs it to classify proceedings books and enhance the quality of their metadata (Salatino et al., 2019). Other applications include the classification of research software (Ciuciu-Kiss, 2022), YouTube videos (Ajwani & Arolkar, 2021), press releases (El Ghosh et al., 2022), job advertisements (Derksen & Dörpinghaus, 2023), and collections in IT museums (Djambian et al., 2024). Table 1 presents an exemplary paper (Donahue et al., 2017) and the relevant topics extracted using the CSO Classifier.

AIDA KG follows the Resource Description Framework (RDF) standard[5] and is available for download as a data dump or for querying through a triplestore at https://aida.kmi.open.ac.uk/sparql/. The ontological schema upon which AIDA KG is built is available at https://w3id.org/aida#aidaschema. It is licensed under the CC BY 4.0 license.

---

[1] OpenAlex — https://openalex.org/

[2] DBLP — https://dblp.org/

[3] Research Organization Registry — https://ror.org/

[4] DBpedia — https://www.dbpedia.org/

[5] RDF Standard (World Wide Web Consortium (W3C)) — https://www.w3.org/TR/rdf-concepts/

Table 1 Example of classified paper (Donahue et al., 2017) with the CSO Classifier. Topics are organised from the most specific to the most general within the CSO hierarchy

Title	Long-term recurrent convolutional networks for visual recognition and description
Abstract	Models based on deep convolutional networks have dominated recent image interpretation tasks; we investigate whether models which are also recurrent, or "temporally deep", are effective for tasks involving sequences, visual and otherwise. We develop a novel recurrent convolutional architecture suitable for large-scale visual learning which is end-to-end trainable, and demonstrate the value of these models on benchmark video recognition tasks, image description and retrieval problems, and video narration challenges. In contrast to current models which assume a fixed spatio-temporal receptive field or simple temporal averaging for sequential processing, recurrent convolutional models are "doubly deep"' in that they can be compositional in spatial and temporal "layers". Such models may have advantages when target concepts are complex and/or training data are limited. Learning long-term dependencies is possible when nonlinearities are incorporated into the network state updates. Long-term RNN models are appealing in that they directly can map variable-length inputs (e.g., video frames) to variable length outputs (e.g., natural language text) and can model complex temporal dynamics; yet they can be optimized with backpropagation. Our recurrent long-term models are directly connected to modern visual convnet models and can be jointly trained to simultaneously learn temporal dynamics and convolutional perceptual representations. Our results show such models have distinct advantages over state-of-the-art models for recognition or generation which are separately defined and/or optimized.
Topics	recurrent neural networks, backpropagation algorithm, stochastic processes, video streams, natural language text, natural languages, correlation analysis, network architecture, structural frames, neural networks, probability, video streaming, natural language processing, linguistics, semantics, mathematics, computer networks, computer science

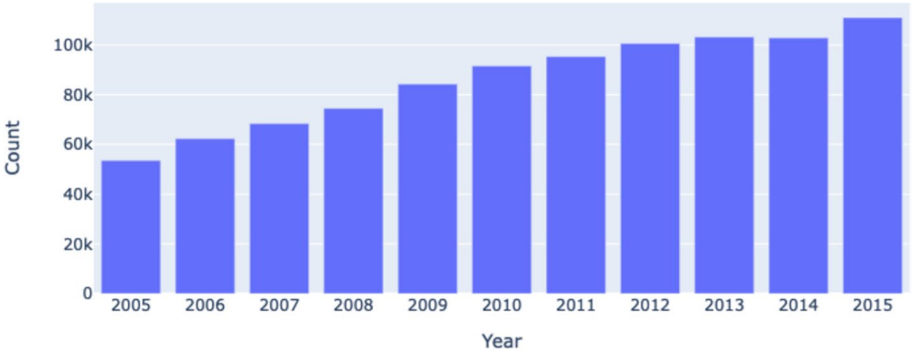

Fig. 1 Distribution of the publications under analysis across the years

Data Selection

In this study, we selected research publications from AIDA KG based on five key criteria: they were published between 2005 and 2015; they were published either in academic journals or conference proceedings; they received at least two citations within five years of publication; they were co-authored by at least two individuals; every author had at least one article in the five years preceding the publication under analysis. In practice, since it was not possible to code all of these conditions in a single query, we proceeded with the following steps. We first selected all papers (~ 3.1M) within the period 2005-15 (crit. i), and published in journals and conferences (crit. ii). We then removed the ones that did not meet the criteria iii, iv, and v. The final dataset contains 944,508 research papers.

Figure 1 reports the distribution of these papers, showing the typical upward trend of scientific publishing (Bornmann & Mutz, 2015).

The rationale for the selection criteria is as follows. The first criterion established a consistent time frame to allow for fair comparisons among papers. The second one recognised journals and conference proceedings as primary and reputable platforms for scholarly communication. The third criterion aimed to exclude editorial notes and technical documents that are occasionally published but seldom cited. The fourth one ensured the inclusion of collaborative research efforts, reflecting team-based scientific inquiries. Finally, the fifth criterion was necessary for characterising the areas of expertise of the authors at the time of the publication.

The impact of each publication was assessed by counting the number of citations it received within five years following its publication date. For instance, the impact of a paper published in 2010 was determined by the citations it had accumulated by 2015. This approach enables an equitable comparison of articles published across different years and aligns with established practices in the literature (AlShebli et al., 2018; Zheng et al., 2022), which frequently utilise a 5-year citation window to assess impact. To discretise the impact for the purpose of our analysis, we categorised the papers into 10 distinct buckets according to their accumulated citations over a five-year period. Table 2 reports the buckets of papers with their frequencies, corresponding citation ranges, and the median of citations for each group.

Assessing Author Expertise

In this study, we characterise the expertise of an author at the time of a publication as the set of the most relevant research topics present in their articles over the five previous years. Figure 2 exemplifies this process. One major challenge in this space is author name disambiguation (Abramo & D'Angelo, 2023; Cappelli et al., 2024), which is the ability of bibliographic databases to reconcile variations of the same author's name (e.g., "Angelo Salatino" vs. "Angelo A. Salatino" or "A. Salatino") and to distinguish between different authors with identical names. AIDA KG directly imports authors and their paper

Table 2 Groups of papers according to the citation ranges

Group identifier	Citation ranges (c)	Citation median	# Papers
A	$2 \leq c < 5$	3	314,563
B	$5 \leq c < 10$	6	247,540
C	$10 \leq c < 15$	12	119,701
D	$15 \leq c < 20$	17	69,291
E	$20 \leq c < 30$	24	74,551
F	$30 \leq c < 40$	34	38,424
G	$40 \leq c < 50$	44	22,220
H	$50 \leq c < 100$	64	38,994
I	$100 \leq c < 150$	118	9,782
J	$c \geq 150$	226	9,442

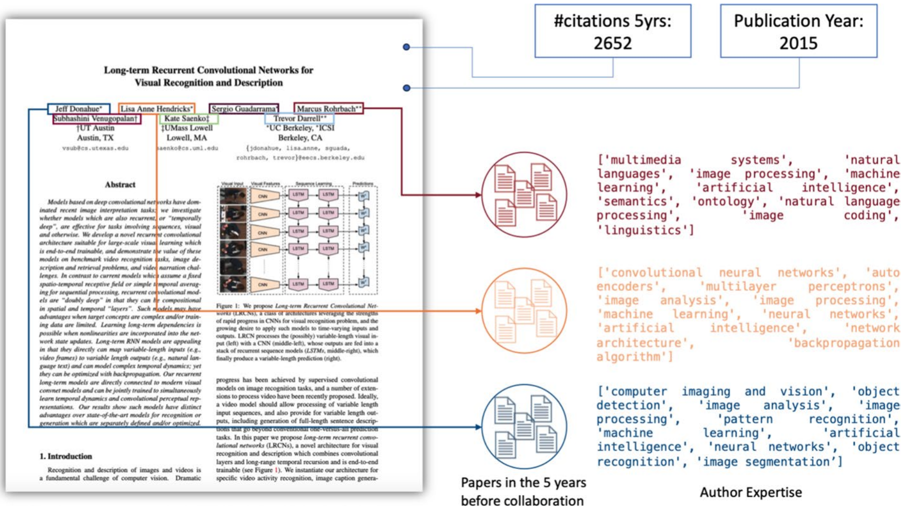

Fig. 2 Workflow for the extraction of authors' expertise starting from a given paper

associations from OpenAlex, hence we rely on OpenAlex's established capabilities for author name disambiguation.[6]

To determine the expertise of all authors in the selected sample of publications, we undertook the following steps. First, we extracted data on 4.4 million authors from the 944K papers. For each author within a publication, we extracted their research publications in the 5 years prior to the time of collaboration. We then computed the normalised distribution of their CSO topics by counting the number of times each topic appeared in these articles and divided them by the total number of articles. Next, we calculated the difference between the normalised topic distribution for the specific author and the overall distribution across the entire Computer Science domain. This approach, as recommended in Angioni et al. (2021), allows us to filter out the common topics that constitute the background noise within the field of Computer Science, thereby highlighting the topics that are more relevant to the author in question. For example, if the topic "Machine Learning" appears in 70% of an author's publications, but only in 30% of general Computer Science papers, this topic would have a 40% relevance score for that author. Conversely, if 50% of the author's papers focus on "Semantic Web", a topic present in only 0.3% of all Computer Science papers, the relevance score for this topic would be 49.7%.

In this analysis, any CSO topic is utilised to represent the author's expertise. Although CSO is a hierarchical and fine-grained knowledge organization system (Salatino et al., 2025), with research topics distributed over 13 levels, we found that the three most common depths are levels 3, 4, and 5, including topics like "wireless communication", "interactive learning environments", "peer to peer networks". Additionally, CSO groups semantically similar topics into clusters of related terms, such as "peer to peer networks" and "peer to peer". To avoid redundancy in the representation of an author's expertise, we consider only one representative topic from each cluster (i.e., the preferred label). This approach

---

[6] Author disambiguation in OpenAlex - https://help.openalex.org/hc/en-us/articles/24347048891543-Author-disambiguation

ensures that the topics used to model authors are all semantically distinct, preventing the inclusion of multiple topics with identical or closely related meanings.

Finally, we represent the author's expertise as a list of the 10 topics with the highest scores, which we will refer to as the top-10 topics. We performed a sensitivity analysis by testing other values between 5 and 20, but the overall results were very similar.

Since a researcher may contribute to multiple papers over time, each instance of their authorship is treated independently. Consequently, our analysis takes also into account that their associated expertise may change over time.

Assessing Expertise Diversity within a Team

We evaluated the diversity of expertise within a research team using two metrics: 1) the maximum pairwise cosine distance between the authors' expertise, and 2) the number of sub-teams with distinct areas of expertise.

At this stage, it is crucial to differentiate between group-level dynamics and team-level dynamics. Group-level dynamics refer to the broader research group, which often includes multiple teams, is hierarchically structured, and simultaneously pursues various research agendas (Wagner, 2009). In contrast, team-level dynamics describe a co-authoring team focused on a specific research agenda, which may be part of or overlap with the larger group (Wagner et al., 2019). For this manuscript, our focus is on team-level dynamics because our data source identifies individual authors of research papers. Here follows a detailed description of the two metrics.

Maximum Cosine Distance Given a research paper, we measure the cosine distance between all couples of authors, considering their top 10 topics. The cosine distance is the complement of the cosine similarity and ranges from 0 to 1. A high cosine distance between authors signals they have fewer topics in common. Contrariwise, the lower the cosine distance (tending to zero), the more similar the set of topics associated with the two compared authors.

To apply cosine distance to the top-10 topic arrays of two authors, we transform them using one-hot encoding. This involves identifying all unique topics from both authors, and creating new arrays where each position corresponds to a specific topic and is marked with a 1 if the author has that expertise, otherwise 0.

This process produces a set of cosine distances for the (N×(N-1))/2 possible pairs of authors. Our first diversity metric is defined as the maximum value within this set. Specifically, a maximum cosine distance of 0 indicates that all authors within the research team possess identical expertise, while a maximum cosine distance of 1 suggests the presence of at least one author whose areas of expertise are entirely distinct from the others. Our approach aligns with the work of Lungeanu et al. (2023), who also analysed the maximum pairwise distance, reflecting the extent to which a team comprises at least two individuals with diverse expertise.

In formal terms, given the N authors of a research paper, we perform the pairwise author comparison (a_i and a_j). Let $U_{i,j} = T_i \cup T_j$ be the union of topics for authors a_i and a_j. We then create one-hot encoded vectors, v_i and v_j, for authors a_i and a_j respectively, based on the topics in $U_{i,j}$. The dimension of these vectors is $|U_{i,j}|$. For each topic $t_k \in U_{i,j}$:

- $v_i[k] = 1$ if $t_k \in T_i$, else $v_i[k] = 0$.
- $v_j[k] = 1$ if $t_k \in T_j$, else $v_j[k] = 0$.

We compute the cosine distance between two authors a_i and a_j as follows:

$$\text{distance}(v_i, v_j) = 1 - \frac{v_i \cdot v_j}{\|v_i\|\|v_j\|} \tag{1}$$

whereas, the maximum cosine distance ($D_{\max}$), is defined as the maximum value within the set $\mathcal{D}$:

$$D_{\max} = \max\{distance(v_i, v_j) \mid 1 \leq i < j \leq N\} \tag{2}$$

We also explored alternative indicators based on this set of cosine distances, such as the mean, skewness, and kurtosis. However, these metrics proved fragile, often yielding very different values for research teams that we would consider very similar in terms of diversity. For example, consider a research team of two authors: one in Artificial Intelligence and the other in Software Engineering. This team would likely exhibit a high average cosine distance (e.g., ~1), indicating significant differences in their areas of expertise. If the team were expanded to include three authors from Artificial Intelligence and three from Software Engineering, the average cosine distance would decrease significantly (e.g., ~ 0.6), reflecting greater overall similarity among the team members' fields. Consequently, two teams with the same breadth of expertise, one component focusing on Artificial Intelligence and the other on Software Engineering, would receive substantially different scores when using the average cosine distance. Using the maximum cosine distance resolves this issue by offering a more robust metric that accurately captures the degree of difference between the most distinct components.

Number of sub-teams As a second diversity metric, we measure the number of sub-teams exhibiting similar expertise. This is achieved using a graph-based approach that clusters authors based on their research expertise. In this context, a higher number of sub-teams indicates greater diversity.

Specifically, for each scientific article, we created an authorship graph $G = (V, E)$, with V being the nodes representing authors and E being the edges connecting pairs of authors with similar expertise. We connected pairs of authors only if their cosine similarity is above or equal to a threshold θ.

More formally:

$$\{a_i, a_j\} \in E \iff \text{similarity}(v_i, v_j) \geq \theta \tag{3}$$

We empirically set this threshold θ to 0.7 as it typically indicates a high degree of similarity between two vectors.

We analyse the resulting graph, and we extract the number of connected components, i.e., the set of authors that are connected to each other by a path because of similar expertise.

The number of connected components can span from 1 (i.e., high cohesion between authors as they share similar expertise) to the total number of authors (i.e., the network is fragmented, implying a highly diverse pool of authors). For instance, the two examples within the previous section featuring a teams of authors from Artificial Intelligence and Software Engineering will both form 2 sub-teams.

Finally, we evaluate the diversity of a research paper based on the number of components within the authoring team, categorising it as follows: i) *low* diversity for teams with 1 or 2 components, ii) *moderate* diversity for teams with 3 or 4 components, iii) *high* diversity for teams with 5 or 6 components, and iv) *very high* diversity for teams with 7 to 10 components, with 10 being the maximum number of components identified in our dataset. Figure 3 presents an example of an authorship network with seven authors and three sub-teams, demonstrating a moderate level of diversity.

Investigating the Relationship Between Diversity of Expertise and Citations

To explore the relationship between the diversity of expertise within a team and the number of citations accumulated over five years, we analysed how the values of two expertise diversity metrics were distributed across the ten citation buckets defined in Table 2.

Specifically, we employed the chi-square test to assess the differences between the two variables. We also employed the Pearson linear correlation coefficient (*r*) to measure the relatedness of the two continuous variables. In both tests, we considered a p-value (*p*) of less than 0.05 for statistical significance. Additionally, to validate our findings at the individual paper level and to mitigate potential biases arising from group-level aggregation, we conducted a multivariate analysis using a Negative Binomial Regression model (McCullagh, 2019). This specific type of generalised linear model estimates the expected output (i.e., number of citations) as a function of multiple independent predictors. This approach allows us to control for key confounding variables, namely team size, publication year, and author experience, measured as years of research activity. By controlling for differences in team seniority and size, the analysis isolates the role played by diversity of expertise, thereby strengthening the conclusions regarding its association with scientific impact.

Results

We dedicate this section to presenting and analysing our results.

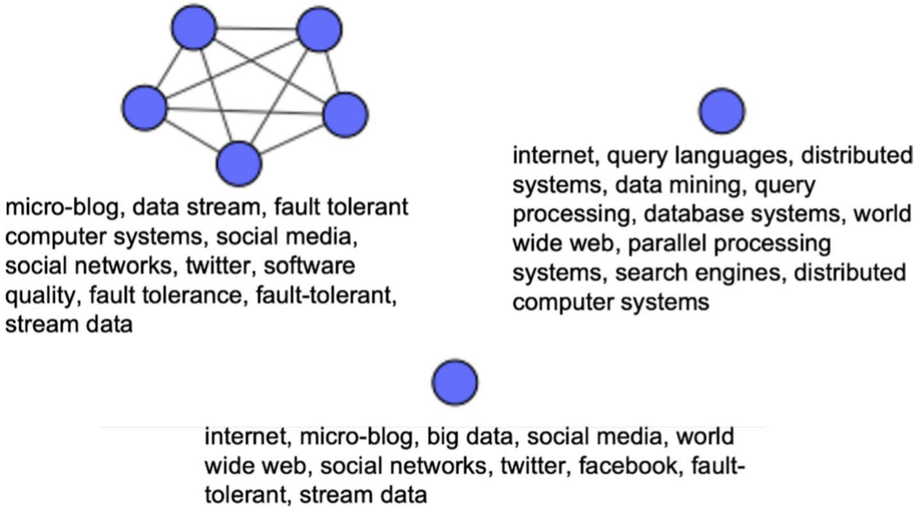

Fig. 3 An example of authorship network organised in sub-team using cosine distance on their expertise. We identified 3 sub-teams classifying it as moderately diverse

Max Cosine Distance Between Authors

Figure 4 illustrates the distributions of maximum cosine distances for all groups (A-J), revealing a consistent pattern: a dominant mode at 1 and a secondary mode at 0. For example, Group A contains approximately 120K papers with a maximum cosine distance of 1 and roughly 10K with a maximum cosine distance of 0. Notably, the frequency distributions for values greater than 0 and less than 1 exhibit remarkable similarity across the groups.

Given these characteristics, we focused our analysis on the frequency of 0 s and 1 s within each group. These values represent the cases where the authors either share the same expertise (a score of 0) or where at least one author specialises in entirely different areas (a score of 1).

Table 3 presents the frequency of 0 and 1 scores, along with their corresponding ratio (#1/#0). A higher ratio indicates that a specific bucket is associated with a greater number of 1 s relative to 0 s, which, in turn, reflects a higher diversity of expertise.

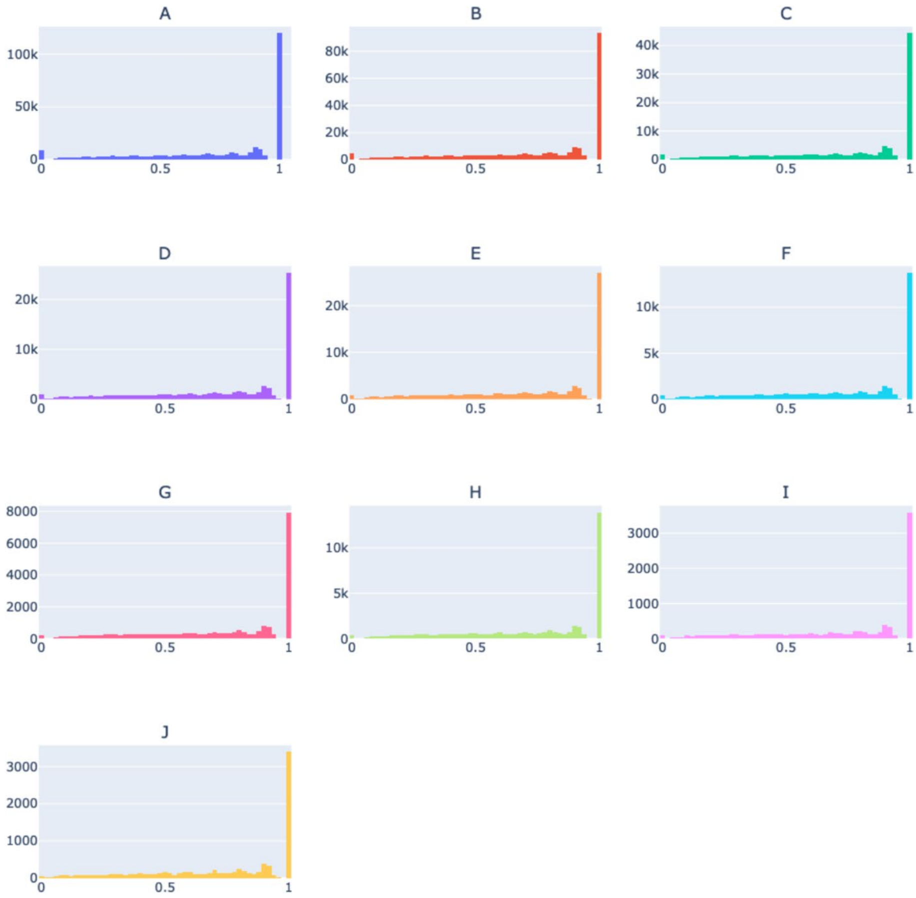

Fig. 4 Distributions of maximum values found within the various groups. The x-axis represents the maximum cosine value, and the y-axis indicates the number of papers (frequency) at each value

Table 3 Frequencies of research papers with zeros and ones spread across the ranges of citations

Gr.	#0	#1	#1/#0
A	7,610	119,707	15.73
B	3,933	93,219	23.70
C	1,392	44,262	31.79
D	735	25,364	34.50
E	632	27,018	42.75
F	279	13,741	49.25
G	132	7,935	60.11
H	204	13,907	68.17
I	56	3,593	64.16
J	37	3,394	91.73

Figure 5 illustrates the relationship between the #1/#0 ratio for each group and the group's median citation count (refer to Table 2). We can notice a strong positive linear correlation between these two variables, with a Pearson correlation coefficient r=0.883 and p-value p=0.0007. This supports the hypothesis that greater diversity in expertise is associated with increased citations and higher research impact.

Number of Components in the Author Graph

Table 4 presents the percentages of papers with low, moderate, high, and very high diversity across the 10 groups. The median number of citations is negatively correlated with the percentage of papers with low diversity (r=−0.817, p=0.003) and positively correlated with the percentage of papers with high diversity (r=0.918, p<0.0001).

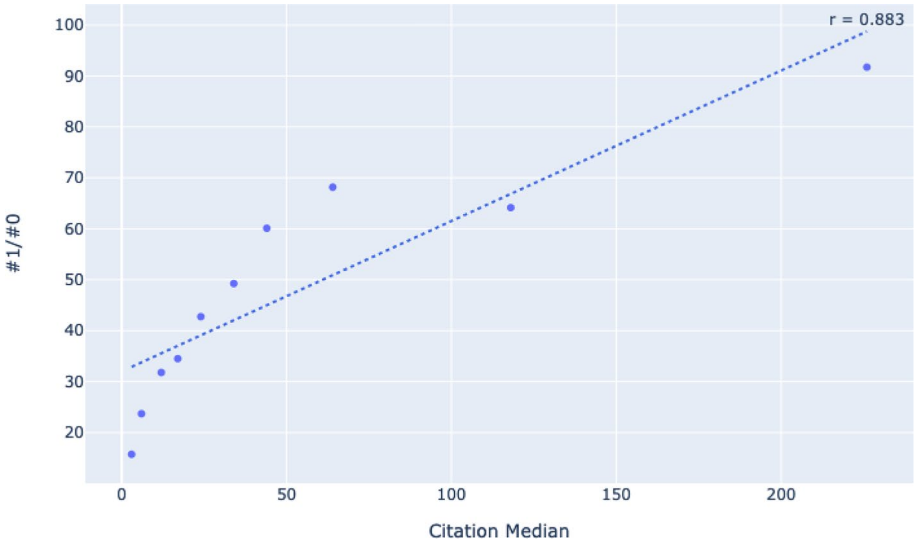

Fig. 5 The #1/#0 ratios scattered against the citations medians

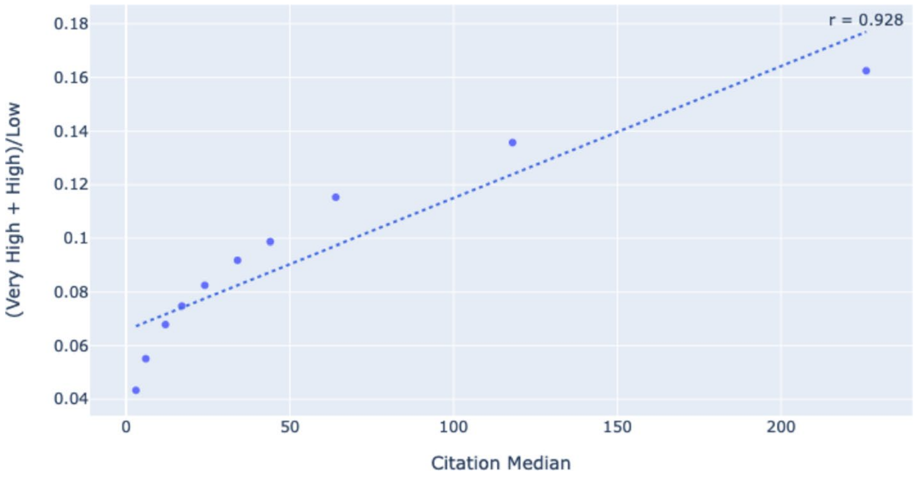

Fig. 6 The ratios ($(very\,high + high)/low$) of components scattered against the citations medians

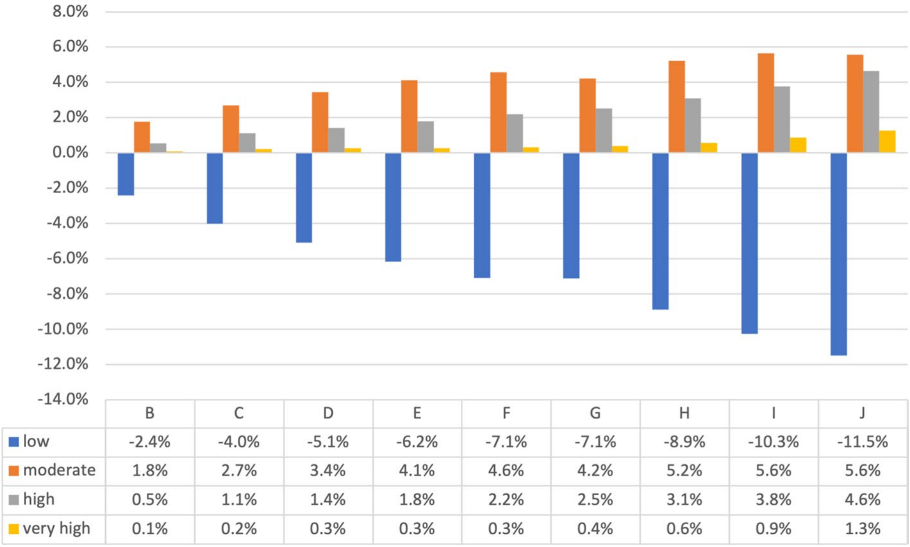

	B	C	D	E	F	G	H	I	J
■ low	-2.4%	-4.0%	-5.1%	-6.2%	-7.1%	-7.1%	-8.9%	-10.3%	-11.5%
■ moderate	1.8%	2.7%	3.4%	4.1%	4.6%	4.2%	5.2%	5.6%	5.6%
■ high	0.5%	1.1%	1.4%	1.8%	2.2%	2.5%	3.1%	3.8%	4.6%
■ very high	0.1%	0.2%	0.3%	0.3%	0.3%	0.4%	0.6%	0.9%	1.3%

Fig. 7 Difference in the ratio of the four categories of diversity between B-J and A

Table 4 also reports the ratio of papers with high or very high diversity to those with low diversity ($(very\,high + high)/low$). Similar to the #1/#0 ratio, this metric can be used to assess the overall diversity of expertise. Figure 6 illustrates the relationship between this ratio and the median number of citations. Once again, these variables exhibit a strong positive linear correlation ($r = 0.928$, $p < 0.0001$). Indeed, only 2.8% of papers with fewer than five citations (Group A) demonstrate high or very high diversity (see Table 4), while nearly 9% of papers with more than 150 citations (Group J) exhibit high or very high diversity.

To further validate the significant impact of diversity on citation counts, we conducted a chi-square test to compare the distribution of four diversity levels (low, moderate, high,

Table 4 Percentages of papers with low, moderate, high, and very high diversity

Group	Low (%)	Moderate (%)	High (%)	Very high (%)	$(very\ high + high)/low$ (%)
A	65.18	32.00	2.60	0.23	4.34
B	62.77	33.77	3.14	0.32	5.52
C	61.17	34.68	3.72	0.44	6.79
D	60.07	35.44	4.02	0.48	7.48
E	59.01	36.12	4.37	0.50	8.26
F	58.09	36.57	4.80	0.54	9.19
G	58.06	36.20	5.11	0.62	9.87
H	56.29	37.22	5.69	0.80	11.53
I	54.92	37.63	6.37	1.08	13.57
J	53.71	37.57	7.24	1.48	16.25
ALL	62.13	33.99	3.49	0.39	6.23

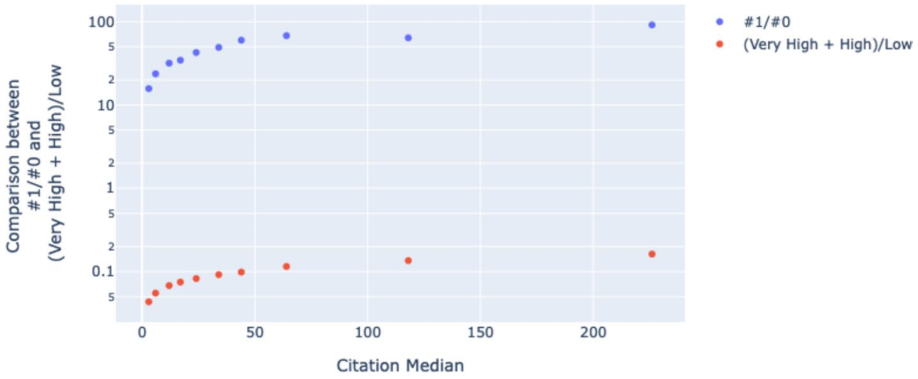

Fig. 8 Scatter plot displaying the similarities between #1/#0 (blue dots) and $((very\ high + high)/low)$ in red dots, based on the median values of citation ranges. The y axis is in logarithmic scale

and very high) between each part of adjacent citation groups *i* and *i+1*. For instance, we compared group A with group B, group B with group C, and so forth. In all comparisons, the distributions were found to be significantly different (p<0.0001). We also observed that the distribution of group A is significantly different from the remaining distributions (B-J) combined (p<0.0001), and A-B combined are significantly different from C-J (p<0.0001).

Figure 7 illustrates the variation in the distribution of diversity categories across groups B to J, compared to group A. For each group, we calculated the percentage of papers within each diversity category (low, moderate, high, and very high) and analysed how these proportions differ relative to group A. For example, Group J, which includes the highest-impact papers, shows a higher proportion of moderate (5.6%), high (4.6%), and very high (1.3%) diversity papers, alongside a notably lower proportion (−11.5%) of low-diversity papers compared to Group A, which represents the least-cited papers. This trend is consistent across other groups (B to I), further emphasising the differences in diversity category distributions relative to Group A.

Relation Between Max Values and Numbers of Components

We analysed the correlation between the two metrics used for measuring diversity in the previous section, i.e. the ratio of papers with ones and zeros (#1/#0, last column of Table 3) and the ratio of papers with high diversity against low diversity ($(very\,high + high)/low$, last column of Table 4). Figure 8 illustrates the distribution of these two scores, depicted in blue and red, respectively, as a function of the median values of citation ranges. Despite differences in scale, both distributions demonstrate a high degree of similarity, with a Pearson correlation coefficient of 0.975.

In conclusion, the results derived from employing the four categories of diversity align closely with those obtained using the maximum cosine similarity approach. In both cases, the diversity metric demonstrates a strong and significant correlation with the number of citations received over a five-year period.

Multivariate Analysis

To validate the robustness of our findings at the individual paper level, we conducted a multivariate negative binomial regression analysis (McCullagh, 2019) (see Table 5). This model predicts the 5-year citation count based on our diversity metrics while simultaneously controlling for publication year, team size, and authors' seniority. The latter is operationalised as academic age, capturing experience through three dimensions: the total and average activity of all authors in the team, and the maximum academic age accounting for the experience level of the most senior member.

The regression results confirm a statistically significant positive relationship between diversity of expertise and scientific impact, even when adjusting for these confounding factors. Specifically, the maximum cosine distance (#1 in Table 5) displayed a strong positive coefficient ($\beta = 0.1547, p < 0.001$), suggesting that teams with highly divergent expertise profiles achieve notably higher citation rates. Similarly, the number of distinct sub-teams (#2) was positively associated with impact ($\beta = 0.0198, p < 0.001$). As expected, the number of authors (#4) also yielded a positive and significant coefficient ($\beta = 0.0207, p < 0.001$). The latter finding is consistent with extensive bibliometric literature indicating that larger teams generally produce higher-impact work due to increased labour capacity and broader dissemination networks (AlShebli et al., 2018; Cummings et al., 2013; Wuchty et al., 2007; Wu et al., 2019). Crucially, these effects persist even when

Table 5 Results of the multivariate negative binomial regression analysis. β *coef* represents the estimated regression coefficient; *Std err* indicates the standard error of the estimate; z denotes the Wald z-statistic; $P > |z|$ represents the two-tailed p-value testing for statistical significance

#	Variable	β coef	Std err	z	P>|z|
1	Max diversity	0.1547	0.011	14.189	< 0.001
2	Number of components	0.0198	0.003	6.089	< 0.001
3	Publication Year	-0.0045	0.000	-10.95	< 0.001
4	Team Size	0.0207	0.002	10.494	< 0.001
5	Total academic age of the team	0.0006	0.000	9.837	< 0.001
6	Average academic age of the team	0.0019	0.000	9.555	< 0.001
7	Ac. age of the most senior member	-0.0008	0.000	-8.99	< 0.001

controlling for the team's experience, measured as aggregate (#5), average (#6), and maximum (#7) years of activity. These indicators showed a statistically significant yet practically negligible effect ($-0.0008 < \beta < 0.0019, p < 0.001$). A similar pattern was observed for publication year (#3), which also showed a minimal effect size. This suggests that the diversity of expertise, as formalised in this paper, constitutes a distinct phenomenon that provides value beyond publication timing, authors' seniority or cumulative activity.

Analysis on the Most Successful Collaborations

We expanded our investigation to examine the impact of specific combinations of expertise. The goal was to identify which combinations were most successful during the period under analysis. Our study focused on collaborations between authors, either in pairs (dyads) or groups of three (triads), from different subfields within Computer Science.

For this analysis, we focused on 14 main disciplines within *Computer Science*, namely *data mining*, *computer networks*, *human computer interaction*, *computer aided design*, *software engineering*, *information technology*, *computer security*, *robotics*, *computer imaging and vision*, *internet*, *information retrieval*, *bioinformatics*, *artificial intelligence*, *computer hardware*, and for all combinations, both as dyad and triad, we computed the normalised cumulative number of citations.

As a first step, for each discipline d, we identified all papers whose authors have expertise in such a discipline. For this, we leveraged CSO, which maps the diverse expertise to their primary discipline. We then calculated the cumulative number of citations as $\sum_{i=1}^{n} cit_i^d$; where cit_i^d is the number of citations of the *i-th* paper.

In the same fashion, we counted the cumulative number of citations for all possible dyads (91 in this case) and triads (364 in total) of disciplines. In particular, given a couple of disciplines (d_1, d_2), we identified all papers where authorship included experts from both disciplines, and we computed their cumulative number of citations as ($\sum_{i=1}^{n} cit_i^{d_1,d_2}$), where $cit_i^{d_1,d_2}$ are the citations of the *i-th* paper. For triplets, we computed the cumulative citations as $\sum_{i=1}^{n} cit_i^{d_1,d_2,d_3}$, where $cit_i^{d_1,d_2,d_3}$ is the number of citations of the *i-th* paper that has at least one different author in each discipline.

Finally, we quantify the collaborative impact achieved by a dyad or triad, as the harmonic mean of their normalised cumulative number of citations. This calculation is formalised in Eq. 4 and Eq. 5, respectively.

$$impact^{d_1,d_2} = h\_mean\left(\frac{\sum_{i=1}^{n} cit_i^{d_1,d_2}}{\sum_{i=1}^{o} cit_i^{d_1}}, \frac{\sum_{i=1}^{n} cit_i^{d_1,d_2}}{\sum_{i=1}^{p} cit_i^{d_2}} \right) \tag{4}$$

$$impact^{d_1,d_2,d_3} = h\_mean\left(\frac{\sum_{i=1}^{n} cit_i^{d_1,d_2,d_3}}{\sum_{i=1}^{o} cit_i^{d_1}}, \frac{\sum_{i=1}^{n} cit_i^{d_1,d_2,d_3}}{\sum_{i=1}^{p} cit_i^{d_2}}, \frac{\sum_{i=1}^{n} cit_i^{d_1,d_2,d_3}}{\sum_{i=1}^{q} cit_i^{d_3}} \right) \tag{5}$$

Figure 9 reports a two-dimensional matrix chart displaying the combined impact for each pair of disciplines. The entries on the main diagonal are all set to zero because they represent the combined impact of a discipline with itself, which is not a meaningful comparison in this context. Moreover, the chart is symmetric as $x_{d_1,d_2} = x_{d_2,d_1}$. From this chart, we can notice that the majority of the disciplines, e.g., *computer networks*, *human*

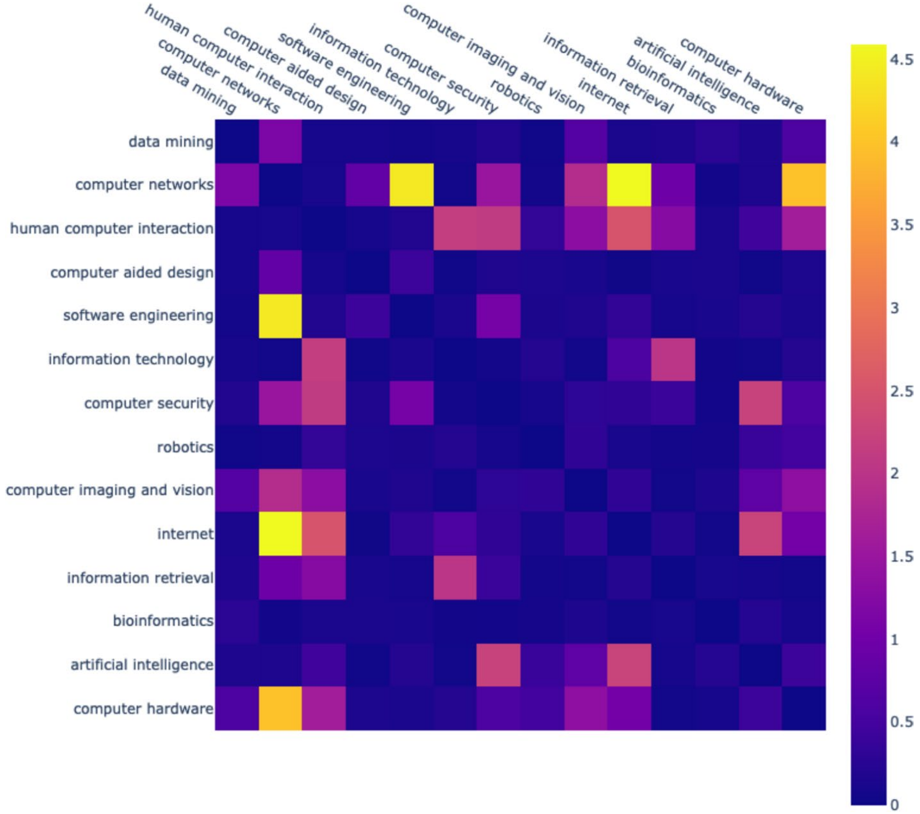

Fig. 9 Two-dimensional matrix chart displaying the impact that each pair of disciplines obtained in the analysed period. Brighter squares indicate high values of impact

computer interaction, *software engineering*, *information technology*, *computer security*, *computer imaging and vision*, *internet*, *information retrieval*, *artificial intelligence*, and *computer hardware* have been involved in collaboration efforts that led to an averaged impact (as in Eq. 4) higher than 2.

Notably, *computer networks* researchers achieved an outstanding impact (~4) through collaborations with researchers from *internet*, *software engineering*, and *computer hardware*. For instance, they co-authored influential work with *internet* researchers on novel approaches for network intrusion detection (García-Teodoro et al., 2009), new encryption schemes (Canetti et al., 2007), and the implementation of secure multi-party computation (Pinkas et al., 2009). Collaborations with *software engineering* researchers led to studies on the transformative potential of cloud computing in the IT industry (Armbrust et al., 2009), as well as investigations into the architectural elements and future directions of the Internet of Things (Gubbi et al., 2013). Finally, joint efforts with *computer hardware* researchers resulted in the development of new techniques for interference alignment in wireless networks (Cadambe & Jafar, 2008) and the application of blockchain to enhance data protection and decentralize privacy (Zyskind et al., 2015).

Table 6 highlights the 30 most impactful dyads. Beyond the just discussed top three interdisciplinary collaborations, several other intriguing patterns emerge. In parallel,

Table 6 Top-30 dyads having the most impactful collaboration. Performance index is computed using Eq. 4

Dyads	Impact (Eq. 4)
Computer networks, internet	4.59
Computer networks, software engineering	4.414
Computer hardware, computer networks	3.972
Human computer interaction, internet	2.503
Artificial intelligence, internet	2.243
Artificial intelligence, computer security	2.228
Human computer interaction, information technology	2.148
Computer security, human computer interaction	2.104
Information retrieval, information technology	2.025
Computer imaging and vision, computer networks	1.885
Computer hardware, human computer interaction	1.62
Computer networks, computer security	1.488
Computer hardware, computer imaging and vision	1.355
Computer imaging and vision, human computer interaction	1.324
Human computer interaction, information retrieval	1.257
Computer networks, data mining	1.139
Computer security, software engineering	1.068
Computer hardware, internet	1.028
Computer networks, information retrieval	0.979
Computer aided design, computer networks	0.837
Artificial intelligence, computer imaging and vision	0.803
Computer imaging and vision, data mining	0.677
Information technology, internet	0.585
Computer hardware, data mining	0.582
Computer hardware, computer security	0.579
Computer hardware, robotics	0.483
Artificial intelligence, human computer interaction	0.444
Computer aided design, software engineering	0.425
Artificial intelligence, computer hardware	0.425
Computer security, information retrieval	0.398

separate research initiatives within *artificial intelligence* and *human-computer interaction* fostered dynamic collaborations, producing highly influential research that has significantly shaped the modern field of *Computer Science*. On this regard, notable examples include a novel approach for human detection in pictures (Dalal & Triggs, 2005) and a presentation of the concept and technical principles of Linked Data (Bizer et al., 2009), which laid the groundwork for the development of Knowledge Graphs (Peng et al., 2023).

Table 7 presents the top-30 most impactful triads. The highest-impact collaboration involves *artificial intelligence*, *data mining*, and *internet*, which also reflects the rise of a new job role: data scientist (Davenport & Patil, 2012).

Other significantly impactful triads see various forms of collaborations among *artificial intelligence*, *computer hardware*, *human computer interaction*, *computer networks*, and *computer imaging and vision*. For instance, *computer imaging and vision*, and *internet* collaborating with *artificial intelligence*, in the years 2005 to 2015 led to the development of

Table 7 Top-30 triads having the most impactful collaboration. Performance index is computed using Eq. 5

Triad	Impact (Eq. 5)
Artificial intelligence, data mining, internet	1.173
Artificial intelligence, computer imaging and vision, internet	1.028
Computer imaging and vision, information retrieval, internet	0.747
Computer networks, computer security, internet	0.664
Data mining, human computer interaction, information retrieval	0.461
Artificial intelligence, computer security, data mining	0.409
Computer hardware, computer networks, internet	0.376
Computer hardware, human computer interaction, information retrieval	0.373
Computer networks, human computer interaction, information technology	0.324
Artificial intelligence, computer hardware, computer imaging and vision	0.315
Computer hardware, computer networks, human computer interaction	0.309
Artificial intelligence, computer imaging and vision, human computer interaction	0.287
Bioinformatics, computer hardware, data mining	0.28
Artificial intelligence, human computer interaction, internet	0.249
Artificial intelligence, bioinformatics, information retrieval	0.242
Computer hardware, information technology, software engineering	0.219
Computer hardware, computer imaging and vision, computer security	0.219
Artificial intelligence, computer hardware, robotics	0.208
Artificial intelligence, computer imaging and vision, robotics	0.201
Artificial intelligence, computer hardware, internet	0.196
Computer imaging and vision, human computer interaction, robotics	0.192
Computer hardware, human computer interaction, robotics	0.192
Artificial intelligence, computer hardware, human computer interaction	0.187
Artificial intelligence, human computer interaction, robotics	0.174
Computer aided design, computer hardware, software engineering	0.17
Artificial intelligence, computer hardware, computer security	0.169
Computer networks, information technology, internet	0.169
Artificial intelligence, computer networks, human computer interaction	0.166
Computer hardware, computer security, data mining	0.164
Artificial intelligence, computer imaging and vision, computer security	0.162

image recognition and classification tools (Simonyan & Zisserman, 2014), and improved tools for face recognition (Parkhi et al., 2015). Additionally, the same two areas collaborating instead with *information retrieval* let to the development of content-based image retrieval (Nielsen et al., 2018; Wan et al., 2014), visual search engines (Sivic & Zisserman, 2008), augmented reality applications (Carmigniani et al., 2011) and others. Finally, another noteworthy collaboration consists of *computer networks*, *computer security*, and *internet*, which brought to the development and consolidation of a number of protocols (such as https, ssl/tls, wpa2, ssh, kerberos) to ensure security and privacy on online communications.

Discussion and Limitations

This study investigates how the diversity of expertise within a research team influences its future impact within the first five years. Our theoretical contribution includes the number of sub-teams as a new graph-based diversity metric. As an additional metric, we also incorporated the maximum cosine distance, which has also been concurrently employed by Lungeanu et al. (2023). A key innovation lies in our explicit modelling of author's expertise. We leverage a large, granular ontology of over 14,000 research topics in Computer Science, providing enhanced explainability and comprehension. Furthermore, our methodology identifies the most relevant topics by contextualising them against the entire corpus, a significant difference from less explicit embedding-based (Lungeanu et al., 2023) or frequency-focused approaches common in the existing literature (Mishra et al., 2025).

Our findings offer several practical implications. Organisations can use these metrics to measure and tune team diversity, perhaps by fostering collaborations or making strategic hiring decisions. For example, a group leader could assess their team's current diversity and then recruit a new member to fill a knowledge gap. Similarly, funding agencies might use these insights to prioritise highly diverse, potentially more disruptive teams, anticipating greater impact from their work.

Despite these contributions, our work presents a few limitations that should be addressed in future research. First, this analysis focuses exclusively on the field of Computer Science. While Mishra et al. (2025) found similar results in biomedicine, more studies are needed across other scientific disciplines to fully understand the relationship between expertise diversity and impact. Second, our study analyses diversity at the authoring team level. We acknowledge that understanding dynamics at a broader research group level, which extends beyond the immediate authoring team, could offer deeper insights into the connection between research group composition and success. This limitation, however, stems from the significant challenge of obtaining such meticulously curated data.

A third limitation is that we excluded papers with fewer than two citations within five years of publication. This decision aligns with common practices in Computer Science, where the incremental nature of scholarly dissemination makes it uncommon for relevant papers to remain uncited. It is also supported by our analysis of the selected dataset, which showed that such papers were predominantly artefacts, non-English publications, or low-quality articles. Nevertheless, we acknowledge that in other datasets and research fields, where different scientific practices are observed, excluding these papers may not be appropriate.

Another limitation of our analysis is the exclusion of research papers with more than 10 authors. This decision is justified by the relative rarity of such publications in the field of Computer Science. In AIDA KG, only 0.9% (~8K papers) of the indexed papers fall into this category. These papers follow a power-law distribution in terms of citations, with an average of 15.13 and a standard deviation of 24.03. In the literature, there is evidence suggesting that extreme diversity may not always be beneficial. Indeed, while it may foster creativity, it can also lead to conflict and misunderstanding (Cummings et al., 2013; Uzzi et al., 2012). In this regard, the literature seems to agree with diversity and impact related through an inverted U-shaped relationship: extreme homogeneity and extreme diversity lead to low impact (Østergaard et al., 2011). Our results are consistent with the literature, reflecting observations prior to the hypothesised inflection point. However, our 10-author cutoff prevents us from empirically confirming the existence of this decline in impact at extremely high levels of diversity.

A further limitation of our methodology stems from *categorising* the number of subteams into just four groups, which could potentially introduce bias. For instance, a two-member team's expertise diversity is inherently restricted to low diversity. While designing our approach, we considered normalising the number of subgroups by the total number of authors, opting for a ratio rather than an absolute value. However, this alternative also presented its own biases. To illustrate, consider two hypothetical teams, one with four authors and another with eight, both featuring two subgroups. Normalisation would unfairly disadvantage the larger team (0.25) compared to the smaller one (0.5). Consequently, we chose to use the absolute number of subgroups to account for team size as well. Future work should also focus on studying different categorisation strategies.

A final limitation concerns our choice of diversity metrics. The bibliometric literature (Cassi et al., 2017; Leydesdorff & Rafols, 2011; Mugabushaka et al., 2016) commonly employs Simpson's diversity, True Diversity, and the Rao-Stirling index (Leydesdorff et al., 2019). In our setting, Simpson's diversity proved unsuitable because, in preliminary tests, it produced a narrow range of values, typically between approximately 0.8 and 1. This limited variability reduced its discriminative power for our analysis. In particular, teams exhibiting high cohesion according to our measure would still be classified as highly diverse when evaluated using Simpson's diversity. Nevertheless, we plan to incorporate True Diversity and the Rao-Stirling index in future work. The latter is particularly relevant because it accounts for the semantic distance between topics, extending beyond the syntactic matching of standard metrics. This perspective may enable a more accurate assessment of the true breadth of an author's expertise, as authors who combine cognitively distant topics contribute more substantially to the overall diversity of the team. While existing implementations of the Rao-Stirling index[7] rely on the Web of Science classification scheme, a potential avenue for our future work involves adapting this metric to leverage the Computer Science Ontology (Salatino et al., 2018) structure.

Conclusion and Future Work

This study investigates how the diversity of expertise within research teams influences the impact of their publications, assessed through citation counts within five years.

To achieve this, we characterised researchers' expertise as a set of topics extracted from their publications using the CSO Classifier, an automated tool that identifies relevant topics based on the Computer Science Ontology. We then assessed the diversity of expertise within research teams using two complementary metrics: i) the maximum cosine distance between the expertise vectors of the authors, and ii) the number of connected components generated by linking authors based on a similarity threshold. The first metric evaluates whether at least one team member possesses significantly divergent expertise compared to the others. The second metric estimates the number of distinct sub-teams with varying expertise within the research team. Finally, we analysed which interdisciplinary collaborations yield the greatest impact in terms of citation performance in the period under analysis.

The analysis presented in this paper reveals a significant and positive correlation between both diversity metrics and the number of citations received over a five-year period

---

[7] pySciSci — https://github.com/SciSciCollective/pyscisci

($p<0.0007$). Consequently, we conclude that, in the field of Computer Science, diversity of expertise is a key driver of high-impact research.

As a future work, we aim to pursue multiple directions. First, we plan to explore additional dimensions of diversity in scientific research and assess how these, in combination, correlate with future impact. Secondly, we aim to broaden our analysis to other fields of science, such as Engineering, Materials Science, and Medicine. We have already constructed AIDA KG in Engineering, and are in the process of building the same knowledge graph in Medicine. Finally, we aim to analyse the impact of diversity in expertise alongside other relevant dimensions. Specifically, we plan to investigate the behaviour of teams from industry, academia, and collaborative efforts that bridge these two domains. Furthermore, we intend to examine how team engagement with open science practices, such as releasing open resources (Kelley & Garijo, 2021), contributes to achieving a greater impact.

Acknowledgements We would like to acknowledge SpringerNature for funding our research and supporting our data ingestion process.

Author contributions Conceptualization: Angelo Salatino; Methodology: Angelo Salatino, Francesco Osborne; Formal analysis and investigation: Angelo Salatino, Francesco Osborne, Simone Angioni, Diego Reforgiato Recupero; Writing - original draft preparation: Angelo Salatino; Writing - review and editing: Francesco Osborne, Simone Angioni, Diego Reforgiato Recupero, Enrico Motta; Funding acquisition: Francesco Osborne, Enrico Motta; Resources: Angelo Salatino, Simone Angioni;

References

Abramo, G., & D'Angelo, C. A. (2023). How reliable are unsupervised author disambiguation algorithms in the assessment of research organization performance? *Quantitative Science Studies, 4*(1), 144–166.

Ajwani, P., & Arolkar, H.A. (2021). Classification of domains in computer science using random forest algorithm for youtube dataset. In: *International Conference on Image Processing and Capsule Networks*, pp. 662–670. Springer.

AlShebli, B. K., Rahwan, T., & Woon, W. L. (2018). The preeminence of ethnic diversity in scientific collaboration. *Nature Communications, 9*(1), 5163. https://doi.org/10.1038/s41467-018-07634-8. Number: 1 Publisher: Nature Publishing Group. Accessed 2023-03-30.

Angioni, S., Salatino, A., Osborne, F., Birukou, A., Recupero, D.R., & Motta, E. (2022). Leveraging knowledge graph technologies to assess journals and conferences at springer nature. In: *International Semantic Web Conference*, pp. 735–752. Springer

Angioni, S., Salatino, A., Osborne, F., Recupero, D. R., & Motta, E. (2021). Aida: A knowledge graph about research dynamics in academia and industry. *Quantitative Science Studies, 2*(4), 1356–1398.

Armbrust, M., Fox, A., Griffith, R., Joseph, A.D., Katz, R.H., Konwinski, A., Lee, G., Patterson, D.A., Rabkin, A., Stoica, I.,& Zaharia, M. (Feb 2009). Above the clouds: A berkeley view of cloud computing. Technical Report UCB/EECS-2009-28. http://www2.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html.

Biermann, M. (2023). On the gender diversity of research teams in economics seminars. *AEA Papers and Proceedings, 113*, 473–476. https://doi.org/10.1257/pandp.20231123

Bizer, C., Heath, T., & Berners-Lee, T. (2009). Linked data-the story so far. *International Journal on Semantic Web and Information Systems, 5*(3), 1–23.

Bornmann, L., & Mutz, R. (2015). Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references. *Journal of the Association for Information Science and Technology, 66*(11), 2215–2222.

Bozeman, B., & Gaughan, M. (2011). How do men and women differ in research collaborations? An analysis of the collaborative motives and strategies of academic researchers. *Research Policy, 40*(10), 1393–1402. https://doi.org/10.1016/j.respol.2011.07.002

Cadambe, V. R., & Jafar, S. A. (2008). Interference alignment and degrees of freedom of the *k*-user interference channel. *IEEE Transactions on Information Theory, 54*(8), 3425–3441. https://doi.org/10.1109/TIT.2008.926344

Canetti, R., Halevi, S., & Katz, J. (2007). A forward-secure public-key encryption scheme. *Journal of Cryptology, 20*, 265–294.

Cappelli, F., Colavizza, G., & Peroni, S. (2024). Recent Developments in Deep Learning-based Author Name Disambiguation. https://arxiv.org/abs/2503.13448.

Carmigniani, J., Furht, B., Anisetti, M., Ceravolo, P., Damiani, E., & Ivkovic, M. (2011). Augmented reality technologies, systems and applications. *Multimedia tools and applications, 51*, 341–377.

Cassi, L., Champeimont, R., Mescheba, W., & De Turckheim, E. (2017). Analysing institutions interdisciplinarity by extensive use of rao-stirling diversity index. *PLoS ONE, 12*(1), 0170296.

Ciuciu-Kiss, J.T. (2022). A methodology for research software classification. Unpublished. https://oa.upm.es/71376/.

Cooke, S. J., Nguyen, V. M., Anastakis, D., Scott, S. D., Turetsky, M. R., Amirfazli, A., Hearn, A., Milton, C. E., Loewen, L., Smith, E. E., et al. (2020). Diverse perspectives on interdisciplinarity from members of the college of the royal society of canada. *Facets, 5*(1), 138–165.

Cummings, J. N., Kiesler, S., Zadeh, R. B., & Balakrishnan, A. D. (2013). Group heterogeneity increases the risks of large group size: A longitudinal study of productivity in research groups. *Psychological Science, 24*(6), 880–890. https://doi.org/10.1177/0956797612463082. PMID: 23575599.

Dalal, N., & Triggs, B. (2005). Histograms of oriented gradients for human detection. In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05),* vol. 1, pp. 886–8931. https://doi.org/10.1109/CVPR.2005.177.

Davenport, T. H., & Patil, D. (2012). Data scientist: The sexiest job of the 21st century. *Harvard Business Review, 90*(5), 70–76.

De Saá-Pérez, P., Díaz-Díaz, N. L., Aguiar-Díaz, I., & Ballesteros-Rodríguez, J. L. (2017). How diversity contributes to academic research teams performance. *R&D Management, 47*(2), 165–179. https://doi.org/10.1111/radm.12139

Derksen, F., & Dörpinghaus, J. (2023). Digitalization and sustainability in german continuing education.

Djambian, C., Rossi, M., d'Ippolito, G., Poncet, E., & Maret, P. (2024). New terminological approaches for new heritages and corpora: The itinheritage project. In: *MULTILINGUAL DIGITAL TERMINOLOGY TODAY (MDTT) 2024.*

Donahue, J., Hendricks, L. A., Rohrbach, M., Venugopalan, S., Guadarrama, S., Saenko, K., & Darrell, T. (2017). Long-term recurrent convolutional networks for visual recognition and description. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 39*(4), 677–691. https://doi.org/10.1109/TPAMI.2016.2599174

Economic Co-operation, O., Development: (2020). Addressing Societal Challenges Using Transdisciplinary Research. OECD Publishing.

El Ghosh, M., Delestre, N., Kotowicz, J.-P., Zanni-Merk, C., & Abdulrab, H. (2022). Rel topic: A graph-based semantic relatedness measure in topic ontologies and its applicability for topic labeling of old press articles. *Semantic Web, 14*(2), 293–321.

Freeman, R. B., & Huang, W. (2014). Collaboration: Strength in diversity. *Nature, 513*(7518), 305–305. https://doi.org/10.1038/513305a. Number: 7518 Publisher: Nature Publishing Group. Accessed 2023-12-24.

Freeman, R. B., & Huang, W. (2015). Collaborating with people like me: Ethnic coauthorship within the united states. *Journal of Labor Economics, 33*(S1), 289–318. https://doi.org/10.1086/678973

Freeman, R. B., & Huang, W. (2015). Collaborating with People Like Me: Ethnic Coauthorship within the United States. *Journal of Labor Economics, 33*(S1), 289–318. https://doi.org/10.1086/678973. Publisher: The University of Chicago Press. Accessed 2023-03-30.

García-Teodoro, P., Díaz-Verdejo, J., Maciá-Fernández, G., & Vázquez, E. (2009). Anomaly-based network intrusion detection: Techniques, systems and challenges. *Computers & Security, 28*(1), 18–28. https://doi.org/10.1016/j.cose.2008.08.003

Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of things (iot): A vision, architectural elements, and future directions. *Future Generation Computer Systems, 29*(7), 1645–1660. https://doi.org/10.1016/j.future.2013.01.010

Jones, B. F., & Weinberg, B. A. (2011). *Age dynamics in scientific creativity*. https://doi.org/10.1073/pnas.1102895108. Accessed 2023-03-30.

Jones, B. F., Wuchty, S., & Uzzi, B. (2008). Multi-university research teams: Shifting impact, geography, and stratification in science. *Science, 322*(5905), 1259–1262. https://doi.org/10.1126/science.1158357. Publisher: American Association for the Advancement of Science. Accessed 2023-03-30.

Kelley, A., & Garijo, D. (2021). A framework for creating knowledge graphs of scientific software metadata. *Quantitative Science Studies, 2*(4), 1423–1446. https://doi.org/10.1162/qss_a_00167

Konstantinidis, I., Maragoudakis, M., Magnisalis, I., Berberidis, C., & Peristeras, V. (2022). Knowledge-driven unsupervised skills extraction for graph-based talent matching. In: *Proceedings of the 12th Hellenic Conference on Artificial Intelligence.*

Larivière, V., & Gingras, Y. (2010). On the relationship between interdisciplinarity and scientific impact. *Journal of the American Society for Information Science and Technology, 61*(1), 126–131.

Leydesdorff, L., & Rafols, I. (2011). Indicators of the interdisciplinarity of journals: Diversity, centrality, and citations. *Journal of Informetrics, 5*(1), 87–100.

Leydesdorff, L., Wagner, C. S., & Bornmann, L. (2019). Interdisciplinarity as diversity in citation patterns among journals: Rao-stirling diversity, relative variety, and the gini coefficient. *Journal of Informetrics, 13*(1), 255–269. https://doi.org/10.1016/j.joi.2018.12.006

Li, W., & Zheng, H. (2025). Expertise diversity of teams predicts originality and long-term impact in science and technology. https://arxiv.org/abs/2210.04422.

Löffler, F., Wesp, V., Babalou, S., Kahn, P., Lachmann, R., Sateli, B., Witte, R., & König-Ries, B. (2020). ScholarLensViz: A Visualization Framework for Transparency in Semantic User Profiles. In: Taylor, K.L., Gonçalves, R.S., Lécué, F., Yan, J. (eds.) *Proceedings of the ISWC 2020 Demos and Industry Tracks: From Novel Ideas to Industrial Practice Co-located with 19th International Semantic Web Conference (ISWC 2020), Globally Online, November 1-6, 2020 (UTC). CEUR Workshop Proceedings*, vol. 2721, pp. 20–25. CEUR-WS.org, [Online]. CEUR-WS.org. http://ceur-ws.org/Vol-2721/paper485.pdf

Lungeanu, A., Whalen, R., Wu, Y. J., DeChurch, L. A., & Contractor, N. S. (2023). Diversity, networks, and innovation: A text analytic approach to measuring expertise diversity. *Network Science, 11*(1), 36–64. https://doi.org/10.1017/nws.2022.34

McCullagh, P. (2019). Generalized Linear Models. Routledge.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. https://arxiv.org/abs/1301.3781.

Mishra, A., Lee, H., Jeoung, S., Torvik, V. I., & Diesner, J. (2025). Patterns of diversity in biomedical coauthorships: An analysis across authors' ethnicity, gender, age, and expertise. *PLoS ONE, 20*(1), 1–22. https://doi.org/10.1371/journal.pone.0316890

Mugabushaka, A.-M., Kyriakou, A., & Papazoglou, T. (2016). Bibliometric indicators of interdisciplinarity: the potential of the leinster-cobbold diversity indices to study disciplinary diversity. *Scientometrics, 107*(2), 593–607.

Nandy, A., Lathabai, H. H., & Singh, V. K. (2024). xd-index and its variants: a set of overall scholarly expertise diversity indices for the research portfolio management of institutions. *Scientometrics., 129*, 5937.

Nathan, M., & Lee, N. (2013). Cultural diversity, innovation, and entrepreneurship: Firm-level evidence from london. *Economic Geography, 89*(4), 367–394. https://doi.org/10.1111/ecge.12016

Nielsen, M. W., Alegria, S., Börjeson, L., Etzkowitz, H., Falk-Krzesinski, H. J., Joshi, A., Leahey, E., Smith-Doerr, L., Woolley, A. W., & Schiebinger, L. (2017). Gender diversity leads to better science. *Proceedings of the National Academy of Sciences, 114*(8), 1740–1742. https://doi.org/10.1073/pnas.1700616114

Nielsen, M. W., Alegria, S., Börjeson, L., Etzkowitz, H., Falk-Krzesinski, H. J., Joshi, A., Leahey, E., Smith-Doerr, L., Woolley, A. W., & Schiebinger, L. (2017). Gender diversity leads to better science. *Proceedings of the National Academy of Sciences, 114*(8), 1740–1742. https://doi.org/10.1073/pnas.1700616114. Publisher: Proceedings of the National Academy of Sciences. Accessed 2023-03-30.

Nielsen, M. W., Bloch, C. W., & Schiebinger, L. (2018). Making gender diversity work for scientific discovery and innovation. *Nature Human Behaviour, 2*(10), 726–734.

Olechnicka, A., Ploszaj, A., & Celińska-Janowicz, D. (2019). *The Geography of Scientific Collaboration.* Taylor & Francis.

Østergaard, C. R., Timmermans, B., & Kristinsson, K. (2011). Does a different view create something new? The effect of employee diversity on innovation. *Research Policy, 40*(3), 500–509. https://doi.org/10.1016/j.respol.2010.11.004

Parkhi, O., Vedaldi, A., & Zisserman, A. (2015). Deep face recognition. In: *BMVC 2015-Proceedings of the British Machine Vision Conference 2015.* British Machine Vision Association.

Peng, C., Xia, F., Naseriparsa, M., & Osborne, F. (2023). Knowledge graphs: Opportunities and challenges. *Artificial Intelligence Review, 56*(11), 13071–13102.

Pinkas, B., Schneider, T., Smart, N. P., & Williams, S. C. (2009). Secure two-party computation is practical. In M. Matsui (Ed.), *Advances in Cryptology - ASIACRYPT 2009* (pp. 250–267). Berlin, Heidelberg: Springer.

Rahdari, B., Brusilovsky, P., & Javadian Sabet, A. (2021). Connecting students with research advisors through user-controlled recommendation. In: *Proceedings of the 15th ACM Conference on Recommender Systems.*

Salatino, A., Angioni, S., Osborne, F., Recupero, D. R., & Motta, E. (2023). Diversity of expertise is key to scientific impact: a large-scale analysis in the field of computer science. arXiv preprint arXiv:2306.15344.

Salatino, A., Osborne, F., & Motta, E. (2020). Researchflow: Understanding the knowledge flow between academia and industry. In: *International Conference on Knowledge Engineering and Knowledge Management*, pp. 219–236. Springer.

Salatino, A.A., Osborne, F., Birukou, A., & Motta, E. (2019). Improving editorial workflow and metadata quality at springer nature. In: *The Semantic Web–ISWC 2019: 18th International Semantic Web Conference*, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part II 18, pp. 507–525. Springer.

Salatino, A.A., Osborne, F., Thanapalasingam, T., & Motta, E. (2019). The cso classifier: Ontology-driven detection of research topics in scholarly articles. In: Doucet, A., Isaac, A., Golub, K., Aalberg, T., Jatowt, A. (eds.) *Digital Libraries for Open Knowledge*, pp. 296–311. Springer, Cham. https://doi.org/10.1007/978-3-030-30760-8_26.

Salatino, A.A., Thanapalasingam, T., Mannocci, A., Osborne, & F., Motta, E. (2018). The computer science ontology: A large-scale taxonomy of research areas. In: Vrandečić, D., Bontcheva, K., Suárez-Figueroa, M.C., Presutti, V., Celino, I., Sabou, M., Kaffee, L.-A., Simperl, E. (eds.) The Semantic Web – ISWC 2018, pp. 187–205. Springer, Cham. https://doi.org/10.1007/978-3-030-00668-6_12.

Salatino, A., Aggarwal, T., Mannocci, A., Osborne, F., & Motta, E. (2025). A survey of knowledge organization systems of research fields: Resources and challenges. *Quantitative Science Studies, 6*, 567–610. https://doi.org/10.1162/qss_a_00363

Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.

Sivic, J., & Zisserman, A. (2008). Efficient visual search of videos cast as text retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 31*(4), 591–606.

Smith, M. J., Weinberger, C., Bruna, E. M., & Allesina, S. (2014). The Scientific Impact of Nations: Journal Placement and Citation Performance. *PLoS ONE, 9*(10), Article 109195. https://doi.org/10.1371/journal.pone.0109195.

Stokols, D., Hall, K. L., Taylor, B. K., & Moser, R. P. (2008). The science of team science: overview of the field and introduction to the supplement. *American Journal of Preventive Medicine, 35*(2), 77–89.

Uzzi, B., Wuchty, S., Spiro, J., & Jones, B. (2012). Scientific teams and networks change the face of knowledge creation. In: Networks in Social Policy Problems, pp. 47–59. Cambridge University Press, Cambridge, UK. https://doi.org/10.1017/CBO9780511842481.003.

Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical Combinations and Scientific Impact. *Science, 342*(6157), 468–472. https://doi.org/10.1126/science.1240474.

Vergoulis, T., Chatzopoulos, S., Dalamagas, T., & Tryfonopoulos, C. (2020). Veto: Expert set expansion in academia. In: *Digital Libraries for Open Knowledge: 24th International Conference on Theory and Practice of Digital Libraries, TPDL 2020, Lyon, France, August 25–27, 2020, Proceedings*, pp. 48–61. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-030-54956-5_4 . https://doi.org/10.1007/978-3-030-54956-5_4

Wagner, C.S. (2009). *The New Invisible College: Science for Development*. Rowman & Littlefield.

Wagner, C. S., Whetsell, T. A., & Mukherjee, S. (2019). International research collaboration: Novelty, conventionality, and atypicality in knowledge recombination. *Research Policy, 48*(5), 1260–1270. https://doi.org/10.1016/j.respol.2019.01.002

Wahle, J.P., Ruas, T., Abdalla, M., Gipp, B., & Mohammad, S.M. (2023). We are who we cite: Bridges of influence between natural language processing and other academic fields. arXiv preprint arXiv:2310.14870.

Wan, J., Wang, D., Hoi, S.C.H., Wu, P., Zhu, J., Zhang, Y., & Li, J. (2014). Deep learning for content-based image retrieval: A comprehensive study. In: *Proceedings of the 22nd ACM International Conference on Multimedia*, pp. 157–166.

Wang, D., & Barabási, A.-L. (2021). *The Science of Science*. Cambridge: Cambridge University Press.

Whitfield, J. (2008). Collaboration: Group theory. *Nature, 455*(7214), 720–723. https://doi.org/10.1038/455720a.

Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science, 316*(5827), 1036–1039. https://doi.org/10.1126/science.1136099

Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature, 566*(7744), 378–382. https://doi.org/10.1038/s41586-019-0941-9.

Yang, Y., Tian, T. Y., Woodruff, T. K., Jones, B. F., & Uzzi, B. (2022). Gender-diverse teams produce more novel and higher-impact scientific ideas. *Proceedings of the National Academy of Sciences, 119*(36), 2200841119. https://doi.org/10.1073/pnas.2200841119

Zhang, X., Chandrasegaran, S., & Ma, K.-L. (2021). Conceptscope: Organizing and visualizing knowledge in documents based on domain ontology. In: *Proceedings of the 2021 Chi Conference on Human Factors in Computing Systems,* pp. 1–13.

Zheng, H., Li, W., & Wang, D. (2022). Expertise diversity of teams predicts originality and long-term impact in science and technology.

Zyskind, G., Nathan, O., & Pentland, A.S. (2015). Decentralizing privacy: Using blockchain to protect personal data. In: *2015 IEEE Security and Privacy Workshops,* pp. 180–184. https://doi.org/10.1109/SPW.2015.27.

Authors and Affiliations

Angelo Salatino[1] · **Francesco Osborne**[1,2] · **Diego Reforgiato Recupero**[3] · **Simone Angioni**[4] · **Enrico Motta**[1]

✉ Angelo Salatino
angelo.salatino@open.ac.uk

Francesco Osborne
francesco.osborne@open.ac.uk

Diego Reforgiato Recupero
diego.reforgiato@unica.it

Simone Angioni
simone.angioni@isti.cnr.it

Enrico Motta
enrico.motta@open.ac.uk

[1] Knowledge Media Institute, The Open University, Milton Keynes, UK

[2] Department of Business and Law, University of Milano Bicocca, Milan, Italy

[3] Department of Mathematics and Computer Science, University of Cagliari, Cagliari, Italy

[4] Institute of Information Science and Technologies “Alessandro Faedo”, National Research Council of Italy, Pisa, Italy