

¹Magzhan Amangeldi, ^{1,2}Oleksandr Kuchanskyi*

¹ Astana IT University, Astana, Republic of Kazakhstan

² National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute", Kyiv, Ukraine

FROM CITATION COUNTS TO COMPUTATIONAL SCIENCE: METHODS, DATA, AND APPLICATIONS IN SCIENTOMETRICS

Abstract. In the last 60 years, the science of scientometrics, the quantitative analysis of science as a communication system, has developed from an esoteric interest in citation indexing to a key tool for research policy and evaluation. The aim of this review is to trace and summarise that trajectory under the single argument that scientometric indicators are only as meaningful as the data infrastructures on which they are built, and the evaluative contexts in which they are applied. We present the literature in five methodological layers, linked by a common thread on research evaluation and policy, and spanning bibliographic data sources, citation indicators and normalisation, science mapping and visualisation, altmetrics and the widening of impact, as well as the recent computational turn based on machine learning and natural language processing. Along with this narrative review, we present a companion bibliometric analysis, mapping the field by topic within the OpenAlex database, which reveals an analytic corpus of 37,166 works, whose growth we illustrate, whose core journals we identify, and whose thematic structure we discuss as a characterization of the field from a methodological perspective. The review finds that the best lasting gains in the field have been in the quality of data coverage, data normalisation and in the responsible interpretation of indicators, and the current transition from counting to understanding indicators should be subject to the same standards.

Keywords: scientometrics, bibliometrics, citation analysis, research evaluation, data infrastructures, openalex, science mapping, altmetrics, machine learning, natural language processing, citation normalization, open science, knowledge graph, scholarly communication, computational scientometrics.

Introduction

The quantification of research evaluation has become a phenomenon that has taken root over the last 20 years. The number of publications, citations and their variants are becoming more and more important in the decision-making process for hiring and promotion, in the distribution of grants to research institutions, in the ranking of universities and in the comparison of national research systems (Faghri & Bergman, 2024; Hicks et al., 2015).

This change has led to the emergence of scientometrics as an information science discipline with tangible impact on the funding, organization, and reward of science. Scientometrics should be considered as an interdisciplinary subject which focuses on the quantitative aspects of science as a communication system, and it is closely but not exclusively related to the analysis of citations (Leydesdorff et al., 2016). It is related to, but not the same as the neighbouring areas of bibliometrics and informetrics (Ellegaard, 2026; Hood & Wilson, 2001).

In practice, though, lines are blurred, and in recent years the discourse has grown more about a common science of science that incorporates macro-level aspects of scholarly production, along with models from network science, economics, and machine learning (Fortunato et al., 2018). The scholarly trajectory of the discipline can be understood as a series of extensions. It starts with Garfield's proposal for a citation index for the sciences, making the citation a tractable machine-countable measure of scholarly influence (Garfield, 1955) and its most significant individual-level instrument is Hirsch's h-index (Hirsch, 2005).

The field has then in turn evolved from raw counts to normalised indicators (Ioannidis et al., 2016; Lu et al., 2026; Waltman & van Eck, 2019); indicators to maps of the structure of science (Klavans & Boyack, 2017; Van Eck & Waltman, 2017); citations to altmetrics (González et al., 2025; Sugimoto et al., 2017); and finally, counting citations to understanding citations (Jurgens et al., 2018; Saeidnia et al., 2025). For this review that path is the organising spine. Under these expansions, there's a continuous feeling of tension. Each indicator is calculated in some database, and databases vary systematically in coverage (Culbert et al., 2025; Mongeon

& Paul-Hus, 2016); each indicator is understood in some evaluative regime, and the regimes influence the behaviour they claim to measure (Baccini et al., 2019). The Leiden Manifesto and the San Francisco Declaration on Research Assessment were both born out of the misuse and overuse of metrics (Gagliardi et al., 2023; Hicks et al., 2015) and there is significant literature warning that citations cannot be taken as a simple indicator and must be understood within their context (Aksnes et al., 2019; Ioannidis et al., 2016).

This paper provides a structured synthesis that is organized around four questions: (i) What are the data infrastructures behind the indicators of scientometrics, and how do the indicators of different data infrastructures differ in terms of their coverage and bias; (ii) How have the indicators of scientometrics and their normalisation changed; (iii) How do science-mapping methods reconstruct the structure of research; and (iv) How is the computational turn reshaping what scientometrics can measure. We are not tied to a national system and have a general scope of work. While reviews focus on a certain aspect of the field such as indicators (Waltman & van Eck, 2019), software tools (Moral-Muñoz et al., 2020), or AI methods (Safder et al., 2023), we combine across the 5 layers and we link the five layers together through the evaluation thread, and we supplement the review with a companion analysis based on an open reproducible data source (Donthu et al., 2021).

Materials and Methods

Review design. This paper follows a hybrid design: the structured narrative review of the methods, transparent and explicit in a selection protocol, and a companion bibliometric analysis of a larger harvested corpus. The narrative layer provides interpretation, the bibliometric layer provides quantitative evidence of growth, structure and trends (Donthu et al., 2021). To make the points above clearer, Table 1 compares the approaches to reviewing found here. The narrative component was question-led rather than designed as a systematic effectiveness review or meta-analysis. Sources were retained when they directly addressed one of the five methodological layers, provided methodological or empirical evidence relevant to research evaluation, or represented a landmark statement, dataset, or infrastructure. Recent review articles were used to identify the principal methods within each layer, after which backward citation tracing was used to recover foundational studies. The search and selection process continued through June 2026, prioritised peer-reviewed international literature, and retained policy statements when they directly shaped responsible-metrics practice. These procedures provide transparent boundaries for the narrative synthesis while the companion bibliometric analysis supplies reproducible descriptive evidence at corpus level.

Table 1

Comparison of review types relevant to scientometric synthesis

Dimension	Narrative	Systematic (PRISMA)	Bibliometric	Meta-analysis
Primary goal	Interpretive synthesis	Answer a focused question	Map structure and trends	Pool effect sizes
Reproducibility	Low	High	High	High
Breadth	High	Narrow, deep	Very high	Narrow
Content meaning	Strong	Strong	Weak	Moderate
This paper	Core method	Protocol adopted	Companion analysis	Not used

Corpus construction. The companion analysis is based on OpenAlex, an open and fully reproducible bibliographic database, selected because its open API allows the corpus to be re-harvested, and because it is covered by a comparable number of publications to established databases (Visser et al., 2021). We segment the field into topics rather than free-text queries: we retrieve works with a major topic on OpenAlex of ‘Scientometrics and Bibliometrics Research’ (OpenAlex topic T10102) or ‘Web Visibility and Informetrics’ (OpenAlex topic T13976) and have been published in English in the time interval 2000–2025. Topic

delineation is preferred over keyword queries because it means that it reduces contamination by the search terms used to define the query and the classification itself acts as a relevance filter. Every query and the snapshot of the corpus are stored along with the date it was retrieved.

For reproducibility, the two topic identifiers were fixed in the retrieval script rather than resolved anew at each run. Works were retrieved through the OpenAlex /works endpoint on 4 June 2026, restricted to English-language articles and reviews published between 2000 and 2025. The outer-boundary query, in which either topic appeared in topics.id, returned 76,369 works. Requiring either T10102 or T13976 in primary_topic.id retained 38,319 works before venue cleaning. A venue-level audit then identified exactly 1,153 misclassified records from five OpenAlex source IDs: C&EN Global Enterprise (S4210177211, n = 56), Chemical & Engineering News (S163506364, n = 11; S4210207877, n = 814), Physics Today (S180907581, n = 155), and Physics World (S9810779, n = 117). These records were removed once, before every analysis. The resulting final cleaned analytic corpus contained 37,166 works, and this same corpus was used for annual output, Bradford analysis, open-access analysis, topic-stability assessment, and term co-occurrence. The complete API response metadata and the retrieved work-level fields were retained locally so that the filtering and cleaning sequence could be reproduced.


Corpus adequacy. Convergent criteria: saturation of the annual output, Bradford core-zone concentration, structural stability of topic composition, and a venue-level audit of anomalously dense clusters were used to assess the adequacy of the corpus size rather than merely asserting its adequacy. Because the corpus is defined by OpenAlex topic classification rather than keyword matching, its validity depends on the quality of that classification. The venue audit detected the five misclassified source IDs listed above, and the same exclusion rule was then applied globally. No independent hand-coded reference corpus was available; therefore, the study does not claim a formal recall estimate or exhaustive coverage. Instead, validation addresses observable classification noise, the stability of the principal topic structure, and the consistency of the final corpus across all reported analyses. This limitation is stated explicitly so that corpus adequacy is not confused with perfect topic-classification accuracy.

Open-access status was obtained from the OpenAlex field open_access.is_oa. Records coded true were treated as open access and records coded false as closed access, aggregating the individual OA routes for the open-versus-closed comparison. There were no missing is_oa values among works published from 2000 through 2024. Shares were calculated within consecutive five-year periods using the final cleaned corpus as the denominator for each period. The incomplete publication year 2025 was retained in the overall corpus and annual-output analysis but excluded from the five-year OA comparison. Under this procedure, the OA share increased from 20.7% in 2000–2004 to 68.0% in 2020–2024.


Results of corpus construction. The two-topic outer boundary contained 76,369 records, of which 38,319 assigned T10102 or T13976 as the primary topic. The venue audit excluded 1,153 misclassified trade and news records, leaving a final cleaned analytic corpus of 37,166 works. The distinction between the outer boundary and the analytic corpus remains important: many outer-boundary works apply a bibliometric method to a substantive domain, whereas the primary-topic corpus focuses on work classified as a methodological contribution to scientometrics or informetrics. Annual output increased from 486 works in 2000 to a maximum of 1,970 in 2017 and remained close to 1,800–1,950 during much of 2018–2022; lower counts in 2023–2025 may partly reflect indexing lag. The Bradford analysis included 27,166 works assigned to journal sources. Its first zone comprised 37 sources and 9,095 works, or 33.5% of the Bradford denominator. The most prominent sources included Scientometrics, Journal of Informetrics, Journal of the Association for Information Science and Technology, and Research Evaluation. Figure 2 summarises the identification, primary-topic restriction, venue cleaning, and final inclusion stages.

Figure 1

Annual output of the final cleaned analytic corpus ($n = 37,166$), 2000–2025. Output peaks at 1,970 works in 2017; the lower recent counts may partly reflect incomplete indexing.

**Figure 2**

Flow diagram for construction of the single-source, topic-delineated corpus. The outer-boundary set contained 76,369 records; the primary-topic restriction retained 38,319; venue cleaning removed 1,153; and the final analytic corpus contained 37,166 works.



Retrieval date: 4 June 2026. Counts refer to records after each sequential restriction.

Results and Discussion

Data sources and infrastructure. Each of the scientometric indicators is calculated with respect to some database, and no database is a neutral window on science. In the end, coverage decisions regarding journals, languages, document types, and years shape the indicators they are used to create, and the choice of a database

is a methodological decision, not a technical preliminary. For 40 years, big bibliometrics was based on just one source: the Science Citation Index (and later Web of Science, WoS); with the introduction of Scopus in 2004, the monopoly was broken and the comparison of databases became a research question.

The two are the primary reference databases for evaluating research, but they are not synonymous. Both databases demonstrate that the natural sciences, engineering, and biomedical research are overrepresented compared to the social sciences and the arts and humanities in their coverage of journals; both also favour publications in English over those in other languages, and, while the two databases cover different journals and the coverage of each journal varies significantly, the choice between the two can affect the outcome of a journal analysis (Mongeon & Paul-Hus, 2016).

A third, more widely used, source is Google Scholar (GS) which covers a much wider, and less curated, selection of content, such as books, conference papers, and work not written in English. The choice among these data sources and metrics affects the results of cross-disciplinary comparisons (Abramo et al., 2026). The lesson that runs through this literature is that bibliometric results depend on the source itself – an indicator is only as comparable as the coverage that it is based on.

Since then, the volume of data has expanded significantly. Launched in 2018, Dimensions goes beyond the publication–citation core and extends its reach to include grants, patents and clinical trials, as first-class objects that are connected within one graph and are linked with institutions and research classifications by machine learning (Hook et al., 2018). Crossref, the DOI registration agency, has become a significant open source of scholarly metadata in its own right, with over 100 million records not only of journal articles, but of datasets, preprints, peer reviews, grants, and more, and linking increasingly to citations, abstracts, and funding.

The most systematic assessment of this expanded field is a comparison of the coverage, completeness, and accuracy of citation links across the five sources Scopus, WoS, Dimensions, Crossref and Microsoft Academic on a large scale, which suggests that the practical implication is to integrate comprehensive coverage with flexible filtering, not to rely on any single source as a primary authority (Visser et al., 2021). Independent analysis of the literature covered in Google Scholar agrees: Google Scholar is able to find a very high proportion of the open-access literature, but it is not as consistent or convenient as it would be to measure open-access literature systematically (Martín-Martín et al., 2018; Weimer et al., 2023).

A separate strand of this literature is to do with control, not coverage. The analyses using WoS and Scopus are difficult to reproduce, as these are proprietary, selectively indexed, and biased against the humanities and social sciences, non-English languages, and research from non-Western countries. The open-infrastructure response, in the form of Crossref's community-owned metadata (Hendricks et al., 2020) and the OpenCitations initiative, and the latest OpenAlex, are all designed to offer a nonproprietary, fully queryable substrate for scientometric research. It's not only political – because having an open, versioned source means that someone can re-execute it in principle, which is the requirement that the evaluation literature imposes on the metrics that it's been critiquing.

This review has a high regard for that standard in its design. The companion analysis is developed specifically using the open API of OpenAlex, which also allows the entire corpus to be re-harvested, and for the figures to be re-generated, while also providing citations and coverage that compare favourably with the established sources (Haunschild & Bornmann, 2024; Visser et al., 2021). But the compromise is real: open sources can be noisier, as the experience we had with several misclassified high-volume trade and news publications revealed (Zhang et al., 2024); they also need to be clearly cleaned and differentiated. That cost, however, gives transparency: All selection decisions are transparent and reversible, a proprietary black box cannot provide that.

The main take-home message for any scientometric study is that the data source must be selected carefully and documented completely; the source is not merely a matter of convenience, it is a key determinant of the result. Citation indicators, normalization. The indicator is not the data source; it sets the meaning of the count that is provided by the data source. One of the fundamental problems of scientometrics is that raw citation counts are not comparable: a paper in chemistry and a paper in mathematics with the same number of citations may have very different impact scores.

Much of the literature on citation impact indicators is the history of the efforts to make the counts comparable and of warnings about how the comparisons can mislead. The most basic single indicator is the h-index (Hirsch, 2005), which is the maximum value such that a researcher has *h* papers that have been cited *h*

or more times. It is an intuitive number that condenses productivity and impact in one, and this sparked its quick adoption across hiring, promotion, and funding decisions. This also its strength: the h-index is not affected by outliers, meaning that its value grows with the length of the career, and like any raw count — it is dependent on the field, which means that there are many variants and field normalised versions (Lu et al., 2026).

Indicators surveys emphasize that there is no single indicator, and point out that there needs to be a mix of quantitative indicators and qualitative judgment, as well as choosing indicators suitable for the evaluative question under consideration (Faghri & Bergman, 2024). A wide range of reviews of other citation impact indicators groups the field around a handful of common challenges: what publications and citations to count, how to distribute credit for co-authorship, how to build journal-level indicators, and, most importantly, how to adjust for field differences (Waltman & van Eck, 2019).

The technical core of modern bibliometrics is normalization. The typical approach involves comparing a paper's citations to that of the expected number of citations for the field and publication year, but there is debate over how to make that expectation. These scores could be normalised by journal subject categories or, more defensibly, by classification systems that categorise each publication into a main research area prior to calculating percentile-based normalised scores, with percentile approaches being robust to the skew in citation data, and avoiding the distortions of mean-based ratios (Lu et al., 2026).

However, the issue of normalisation is not resolved. every normalisation embeds additional assumptions: what constitutes a field boundary, what constitutes a document type, how much the databases cover, and so on, and further complications can be introduced in cases of aggregation of indicators to evaluate scientists, institutions, journals, or countries, and in cases where credit has to be given in the context of multiauthored work (Aksnes et al., 2019; Ioannidis et al., 2016). The guidance is not prescriptive; different indicators provide a different perspective and the assumptions behind each indicator need to be explained and not hidden within the composite score.

The journal impact factor is the most often-used and criticized factor, and it is the use of the impact factor as a proxy for the quality of individual articles that is one of the main criticisms of the research-assessment reform movement (Gagliardi et al., 2023; Hicks et al., 2015). At the level of countries and institutions, larger dimensions are given a higher status and it is recommended to use the size-dependent indicators in combination with the size-independent ones; for example, an empirical study of the Nature Index suggests that the size-dependent version systematically favours large countries and institutions, providing less information about the performance of small producers, while simpler versions reproduce a lot of the information provided by the Nature Index (Bornmann & Haunschild, 2017).

The general rule is similar to that found in the data-source literature: an aggregate indicator is an interpretive instrument that, to serve its purpose, must be constructed to reflect the question being asked, with the limitations of the instrument reported along with its values. Science mapping and visualization. The methods discussed above count papers, authors and journals. But science mapping raises a new question: What is the structure of a research field, and how does it change? Mapping techniques, which depict the literature as a network (of co-cited references, of bibliographically coupled documents or of co-occurring terms) and divide that network into clusters, depict the intellectual organisation of a field, which no other indicator can (Klavans & Boyack, 2017; Van Eck & Waltman, 2017). There are three relational signals that are well represented in the literature, and the selection of one or another is significant.

Co-citation analysis is performed if two earlier documents are both cited in a later one; often documents that are co-cited are interpreted as having a shared intellectual origin. Bibliographic coupling connects two documents that share references, instead of similarity of reception, it represents similarity of intellectual foundations. Direct citation connects a citing document to the cited document. Direct citation has been shown to yield the most accurate topic-level taxonomy of science in the most systematic comparison to date, where the direct citation-based taxonomy is also closer to focusing references than the other two approaches (bibliographic coupling, co-citation), and where use of discipline-level maps constructed from journal schemes is cautioned as being significantly less accurate than document-level maps (Klavans & Boyack, 2017).

The fourth and more recently popular strategy is to skip the citations entirely: keyword or term co-occurrence, in which nodes are terms, and edges are formed between the nodes that co-occur in documents, capturing the conceptual structure of a field instead of the citation structure (Radhakrishnan et al., 2017). The

advantages of term co-occurrence are that only titles, abstracts or author keywords are needed to create a concept, but it also suffers from the ambiguities of natural language: synonymy, polysemy and the presence of generic terms and is therefore very sensitive to term cleaning. A surge of mapping studies is not the result of new methods, but more of tools that are accessible.

Comparative reviews include a stable division of labour: the visualization of large networks using distance-based matrices (VOSviewer); the identification of emerging trends in terms of term bursts (CiteSpace); and the broadest set of performance and mapping analyses (R package bibliometrix); and the identification of emerging trends in terms of terms and bursts (SciMAT). A recent meta-bibliometric evaluation of over ten thousand studies, carried out using VOSviewer, revealed that the tool is used almost everywhere and a growing methodological tradition focused on citation- and keyword co-occurrence analysis, as well as some frequent methodological weaknesses, such as relying on one database, keyword metadata that varies from study to study and the sensitivity to threshold values (Genc & Kocak, 2024; Omer et al., 2025). The main weakness lies not accidentally in the science map, it is a main weakness; it is the main motivation for the methodological transparency used below.

We created a term co-occurrence map on the analytic corpus to illustrate the method and describe the conceptual network of the field. Instead of author keywords, the terms were mined from the title and abstract of the article, since the keywords in open data sources are not always present and are often inconsistent; the same text-mining approach as used natively by VOSviewer. The 400 common terms used by the four most common sources were not excluded, while an explicit, reported stopword list excluded generic academic vocabulary or source-specific boilerplate. Text was converted to lowercase, non-alphabetic characters were removed, and whitespace was normalised. Binary document-term occurrences were generated with scikit-learn CountVectorizer using one- to three-word n-grams, a minimum document frequency of 40, the standard 318-entry English stopword list, and 170 additional domain-specific entries; after overlap, the combined list contained 469 unique stopwords. Of 6,293 candidate terms meeting the frequency threshold, the 400 most frequent were retained.

Documents containing the same pair of retained terms were linked, and edge strength was defined as the number of documents in which both terms occurred. Links with strength below 5 were excluded, leaving 79,742 links among 400 retained terms derived from all 37,166 documents in the final cleaned corpus. The resulting map and network files were imported into VOSviewer 1.6.20. Association-strength normalisation was used; layout attraction was 2 and repulsion was 0; clustering resolution was 1.00, the minimum cluster size was 1, and random seed 42 was used. Node size represents document occurrence, link strength represents document-level co-occurrence, and colour represents the VOSviewer cluster assignment. We explicitly report the stopword policy because, as the software literature cautions (Omer et al., 2025), term cleaning is a major source of investigator discretion. The initial map exposed a dense peripheral cluster of metadata fragments associated with the five misclassified source IDs. Removing all 1,153 records before every analysis eliminated that contamination and illustrates the direct propagation of data-source quality into mapping results.

The map in Figure 3 divided into four main clusters. The methodological nucleus of citation analysis consists of three major groups: citation, bibliometric, impact factor, the h-index, productivity, and the big databases, along with the vocabulary of mapping and ranking. A second large cluster is the open access, scholarly communication and research policy: open access to publishing to funding to institutional to policy to visibility and to research community. The third compact dimension that the map recovers as a region rather than a subset of citation impact is that of social, online, attention, media and “alternative” – this is a grouping that the map allows for visually, and that fits with the empirical observation that altmetrics provide a different signal from citations (González et al., 2025; Sugimoto et al., 2017). A fourth, peripheral cluster relates to peer-review and publishing process: peer review, editors, reviewers, conferences, proceedings, acceptance, and predatory publishing.

The spatial relations between the clusters are as informative as the clusters themselves: between the citation core and the open-access cluster lies the altmetrics cluster, which are mediated by bridge terms like “social”, “attention” – the online attention is near the citation impact and the openness of dissemination. Two cautions apply. First, the clusters represent co-occurrence of language, rather than citation-based intellectual lineage: a co-citation map or a direct-citation map would emphasize foundational work and schools of thought, and be more useful for a history-of-ideas reading (Klavans & Boyack, 2017). Second, the map cannot be

considered as an absolute result or as independent of the thresholds and of the stop word selection given above, but as a transparent, reproducible characterisation of the conceptual structure of the field.

Figure 3

Term co-occurrence map of the final cleaned analytic corpus (n = 37,166 documents; 400 retained terms; minimum document frequency = 40; minimum link strength = 5; 79,742 retained links). VOSviewer 1.6.20 settings: association-strength normalisation, attraction = 2, repulsion = 0, clustering resolution = 1.00, minimum cluster size = 1, and random seed = 42. Four clusters represent citation methods and indicators, open access and research policy, altmetrics and online attention, and peer review and publishing.





Figure 4

Open-access versus closed-access works in the final cleaned analytic corpus, by five-year period. Open access is defined from OpenAlex open_access.is_oa; no values were missing in 2000–2024. The open-access share rises from 20.7% in 2000–2004 to 68.0% in 2020–2024.



Expansion of impact and altmetrics. All of the indicators mentioned above quantify impact in the scholarly literature. Beginning in the early 2010s, a parallel current has emerged that sought to quantify impact outside of it, using the term 'altmetrics' for social media, reference managers, news outlets, policy documents and online platforms. This was for two reasons: firstly to collect influences beyond citations and secondly to help to record impact more quickly than the long timeframe of citations accumulating. This domain was recovered as a distinct cluster; it is a clearly different area from citation analysis within the area of altmetrics. The fundamental review is based on two interrelated phenomena: the use of social media platforms as a

scholarly tool and the indicators based on their interactions.

The distinction is important because it addresses two issues: an issue of scholarly practice, and an issue of measurement validity, and it is important not to mix the two in the literature. Whether or not there is a correlation between the number of mentions on social media and the number of citations has been the central empirical question. The conclusion from a meta-analysis is that they don't really, the correlation between citations and microblogging is negligible, with blog mentions it is small, and with bookmark counts from online reference managers it is medium to large (González et al., 2025). The interpretation of the data has been well-established after some major shifts and the "attention" is now widely agreed to measure attention, not the impact of the scholarly work; impact is linked, but not identical, and the link's strength will vary greatly by platform. Unlike citation data, altmetric data can have measurement difficulties.

The counts are skewed and zero-inflated: most papers are cited on very few platforms, so that mean-based indicators would be misleading, while they require methods with sparse data, such as indicator families that could differentiate between the quality levels which are assessed by citing them on platforms that have been validated by peer assessment, and on those that have not. In addition to sparsity, the sources of altmetrics are unstable and unevenly covered: platforms come and go, coverage is uneven across disciplines, and the sources of counts are often not clearly stated. For example, an analysis of Publons metrics revealed that there are important coverage biases by discipline and publisher, and weak, non-significant relationships between bibliometric and altmetric metrics, meaning that these metrics should not be used as evaluation tools (Ortega, 2019). A closely related literature is open access (OA) which was placed on the map adjacent to the altmetrics and policy clusters.

A large-scale analysis determined that a significant proportion of the literature is freely accessible, and that the open access citation advantage is about 18%, taking into account discipline and age, which was mainly due to specific OA mechanisms (Piwowar et al., 2018). However, more recent science-wide analysis muddies the picture: the advantage of open access is neither consistent across disciplines nor for citations of more recent papers, nor is it more pronounced for the metrics that measure online attention (such as news, blogs, social media) (Taylor, 2026). The number of open documents in our own corpus, which reflects the underlying growth of openness, increases from 20.7% to 68.0% between the two periods (2000-2004 and 2020-2024, see figure 4). The message of the two "impact" literatures - altmetrics and OA - converges and it is this: "impact" is not a single thing. It's a family of different signals, each with distinct behaviors across fields, and no single indicator will do: a pluralism that the responsible-metrics movement begins with.

The computational turn: AI, machine learning and NLP. The breakthroughs examined thus far are based on bibliographic data, that is, who cited whom and which words co-occur. The most significant change in recent times has been the transition from metadata to full text, supported by the new technologies of natural language processing (NLP) and machine learning (ML). This computational shift will alter the fundamental level of analysis: citations can no longer simply be counted; they can be read, interrogated, and asked "why, how, and in what spirit" are they being cited? There are two broad areas of growth that are well documented in the systematic review of the area: semantic analysis of citations and representation and discovery of scientific knowledge (Saeidnia et al., 2025; Safder et al., 2023).

The premise of citation-based evaluation has always been known to be false; citations serve many rhetorical and social functions, and a citation that contradicts a work counts the same as one that builds on it. This long-standing critique is made operational in the computational turn. Decades of small manual research have been synthesized in a single classification that distinguishes between a scientific motivation (rhetoric) and a tactical motivation (social/benefit oriented) (Lyu et al., 2021). NLP allows such distinctions to be measured on scale. Previous supervised studies demonstrated that features like the number of citations and citation location in the text can be automatically used to identify influential citations (Valenzuela et al., 2015).

Later large-scale research categorized citations at the whole field level, and found that framing a paper's citations predicts the number of citations it will receive, and that such changes mirror the development of a field towards agreement (Jurgens et al., 2018). This results in a number of subfields that are then mapped as a coherent program in a narrative review: citation context and content analysis, citation classification, sentiment analysis, summarisation and recommendation, and all based on full-text processing (Saeidnia et al., 2025). These techniques have also been extended from individual research projects to large-scale infrastructures that classify citation contexts as supporting, contrasting, or mentioning cited work (Saeidnia et al., 2025). Altmetric

data pipelines for sentiment analysis also try to extract the valence (or positivity) of online mentions, not just the number of mentions (Hassan et al., 2021).

The second front is dealing with the representation of scientific documents computationally. There are several of pre-trained language models that have been adapted for scientific text, such as SciBERT, trained on a large multidomain collection of scientific articles, which have significantly outperformed general-purpose pre-trained models in downstream scientific NLP tasks. To generate document-level representations instead of token-level representations, SPECTER utilizes the citation graph as a training signal to generate embeddings which can be used for classification and recommendation without fine-tuning for a specific task (Cohan et al., 2020).

Neural topic-modelling methods such as BERTopic, rather than named-entity-recognition methods, use the transformer embeddings and introduce a class-based weighting scheme to generate coherent topics with clear interpretation, which is an alternative to surface co-occurrence mapping by applying semantic analysis (Grootendorst, 2022). Maybe the most impressive example of latent knowledge in the literature is in materials science: unsupervised word embeddings trained on past publications encoded chemical and structural relationships, and that knowledge was sufficiently explicit to be used successfully for predicting functional materials years before any were discovered, indicating that the published corpus contains recoverable knowledge not yet explicit to its authors (Tshitoyan et al., 2019). All of this rests on the fact that there are large, structured scholarly datasets available, such as large collections of computer science metadata that can be used to analyse trends in the field at a scale of thousands of papers or more (Wahle et al., 2022); it also relies on ML approaches that rank and prioritise research-assessment parameters empirically, not by convention (Ahmed et al., 2023). The computational turn isn't just transformative because it allows an analysis of the content of communications to be performed on something other than metadata; it's transformative because it allows the semantics of scholarly communication to be measured. It brings in new dangers that field is just starting to face, though.

The outputs of classifiers are often dependent on training data and model selection that is not easily understood, and a citation-intent label or a sentiment score comes with some level of uncertainty that a raw count does not; adding up the outputs of the classifiers to make evaluation indicators can hide model error as apparent precision. Reproducibility is a special issue, because the field studies (Kapoor & Narayanan, 2023; Wang et al., 2026) are subjected to the same publish-or-perish pressures as scientific analyses based on ML, and also because of the possibility of very subtle data leakage. The methodological norms that it has established for citations (transparency of construction, explicit reporting of assumptions, validation against external criteria) are equally applicable and, indeed, more so with the indicators derived from AI, whose inner workings are less accessible to inspection.

The computational turn, briefly, is one that seeks to expand the scope of what scientometrics can measure, but not at the cost of withholding the interpretive reservations that the remainder of this review has amplified. Research evaluation, policy and the constitutive effects of metrics. The previous subsections have dealt with indicators as tools of measurement. This one treats them as the means of governance, because in contemporary scientometrics, what is at stake is the allocation of jobs, grants and prestige, and its application alters the behaviour of the scientists under measurement.

The key theme of the evaluation literature is that measurements do not simply reflect scientific processes but are rather actively involved in them. It was the reform movement that crystallised in two significant statements. The San Francisco Declaration on Research Assessment (DORA) challenged specifically the use of journal-level measures, especially impact factors, as a proxy for the quality of individual articles, or individual researchers (Gagliardi et al., 2023).

The Leiden Manifesto extended the argument to a number of principles for the responsible use of quantitative indicators: metrics should supplement, never replace, expert judgment; differences in the field should be taken into account during evaluation; indicators should be transparent and verifiable; and the indicators evaluated should have the possibility to check the data. Both statements have a common diagnosis that runs through this review: An indicator without a question to which it responds, and used regardless of its construction and limitations, has no objectivity, but distortion. This position is strengthened by the increasing acceptance of the idea that citations are imperfect, context-dependent proxies with meaning that relies on the motivating factors and discipline (Aksnes et al., 2019; Ioannidis et al., 2016).

Strategic behaviour studies are the best available evidence that metrics are constitutive. Once the indicator is becoming a goal, the persons measured for it change their behavior to maximize this indicator (the phenomenon is known as “gaming”). Italy's “inwardness” (national share of citations) increased dramatically following the national introduction of bibliometric requirements for career advancement in 2011; this resulted in Italy becoming an outlier among similar countries, likely because of the strategic use of self-citations and citation clubs in response to the evaluation system (Baccini et al., 2019).

The case shows that an evaluation system not only measures a research community but, it changes the citation behaviour of a research community. This is not an unintended consequence of having metrics to govern that needs to be fixed, but rather an inherent characteristic of metric-based governance, and as such, a hard limit on the amount of any given indicator that should be relied on without it being compromised by incentives it creates. A more radical approach to the literature on science and technology studies claims that the process of evaluation do not merely represent or distort existing science but is also a process that helps construct fields of science. A study of how US ‘translational science’ came into being reveals that evaluation practices do not only measure the impact of translational science, but shape its emergence, develop representations of it for researchers, funders and policymakers, and contribute to its legitimation as a form of science that is worthy of public investment (Molldrem et al., 2025) - a way of co-production between science and metrics that its authors describe as ‘the evaluative way’.

The strongest version of the claim, whether accepted or not, serves to focus the central lesson: that scientometric indicators are instruments that act upon the science they measure, and thus require an act of science policy to design; and that the responsibility for producing the impacts of scientometric indicators cannot be divorced from the responsibility for the indicators' accuracy. Synthesis and Open Challenge. Three cross-cutting tensions structure the field's agenda. Coverage vs. comparability: each indicator is contingent on a source that can be biased, and open infrastructure can be about curation, but it cannot be about noise; methods need to be explicit about the need for source limits, and need to be robust to them. The terms quantification versus validity: impact is plural, so the challenge is to not reduce it to single scores but to correlate indicators with questions. The computable turn in the discipline: the way to making semantics measurable is making them opaque in the process of automation: the transparent and validated rules of the discipline should apply to the indicators created with AI. In all three, the constitutive knowledge is that the methodological decisions in scientometrics are decisions about the governance of science.

Conclusion

Scientometrics has evolved from a narrow focus on citation indexing to a discipline and approach that influence the way science is funded, evaluated and organised. This review was a synthesis of five methodological layers, bound together by a consistent commitment to evaluation and policy. Two findings stand out. First, the most enduring advances in the field have been not in the introduction of new indicators, but in the improvements of the underlying data infrastructure and greater transparency about the context in which these data are evaluated. New indicators are not the most important advances in the field, however, but improvements in the underlying data infrastructure, and greater candor about the evaluation context surrounding it. Second, the new methodological paradigm, from counting to understanding citations, is the most radical development in a generation but it should be subject to the same standards of transparency, validity, and responsible use that were developed for the simpler methods; indeed, the more powerful the methods, the more important those standards. The accompanying analysis is based on an open and reproducible data source and is presented as evidence of the structure of the field, as well as an example of the transparency that the field increasingly demands of itself.

Conflict of Interest Statement

The authors declare no potential conflicts of interest with respect to the research, authorship, or publication of this article.

Author Contributions

Magzhan Amangeldi and Oleksandr Kuchanskyi contributed equally to the conceptualization, methodology, data collection and processing, bibliometric analysis, interpretation of the findings, preparation of the original manuscript, review, editing, and finalization of the study. Both authors approved the final

version of the manuscript.

Declaration of Use of Generative AI and AI-Assisted Technologies

The authors declare that no artificial intelligence (AI) tools were used in the preparation of this manuscript.

References

- Abramo, G., Cicero, T., & D'Angelo, C. A. (2026). Data source effects in research performance assessment of individuals and institutions: Comparing OpenAlex with established bibliographic databases. *Scientometrics*, 131(5), 3411–3441. <https://doi.org/10.1007/s11192-026-05638-6>
- Ahmed, B., Wang, L., Al-Shamayleh, A. S., Afzal, M. T., Mustafa, G., Alrawagfeh, W., & Akhunzada, A. (2023). Machine learning approach for effective ranking of researcher assessment parameters. *IEEE Access*, 11, 133294–133312. <https://doi.org/10.1109/ACCESS.2023.3336950>
- Aksnes, D. W., Langfeldt, L., & Wouters, P. (2019). Citations, citation indicators, and research quality: An overview of basic concepts and theories. *SAGE Open*, 9(1), 2158244019829575. <https://doi.org/10.1177/2158244019829575>
- Baccini, A., De Nicolao, G., & Petrovich, E. (2019). Citation gaming induced by bibliometric evaluation: A country-level comparative analysis. *PLOS ONE*, 14(9), e0221212. <https://doi.org/10.1371/journal.pone.0221212>
- Bornmann, L., & Haunschild, R. (2017). An empirical look at the Nature Index. *Journal of the Association for Information Science and Technology*, 68(3), 653–659. <https://doi.org/10.1002/asi.23682>
- Cohan, A., Feldman, S., Beltagy, I., Downey, D., & Weld, D. S. (2020). SPECTER: Document-level representation learning using citation-informed transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 2270–2282). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.207>
- Culbert, J. H., Hobert, A., Jahn, N., Haupka, N., Schmidt, M., Donner, P., & Mayr, P. (2025). Reference coverage analysis of OpenAlex compared to Web of Science and Scopus. *Scientometrics*, 130(4), 2475–2492. <https://doi.org/10.1007/s11192-025-05293-3>
- Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285–296. <https://doi.org/10.1016/j.jbusres.2021.04.070>
- Ellegaard, O. (2026). The method of bibliometric analysis: A critical evaluation. *Scientometrics*, 131, 4175–4196. <https://doi.org/10.1007/s11192-026-05670-6>
- Faghri, A., & Bergman, T. L. (2024). Highly ranked scholars and the influence of countries and regions in research fields, disciplines, and specialties. *Quantitative Science Studies*, 5(2), 464–483. https://doi.org/10.1162/qss_a_00291
- Fortunato, S., Bergstrom, C. T., Börner, K., Evans, J. A., Helbing, D., Milojević, S., Petersen, A. M., Radicchi, F., Sinatra, R., Uzzi, B., Vespignani, A., Waltman, L., Wang, D., & Barabási, A.-L. (2018). Science of science. *Science*, 359(6379), eaao0185. <https://doi.org/10.1126/science.aao0185>
- Gagliardi, A. R., Chen, R. H., Boury, H., Albert, M., Chow, J., DaCosta, R. S., Deyell, R., Finelli, A., Keating, A., Légaré, F., McGowan, J., McLeod, R., Moore, A., Parulekar, W., Saginur, R., Straus, S. E., & Protze, S. (2023). DORA-compliant measures of research quality and impact to assess the performance of researchers in biomedical institutions: Review of published research, international best practice and Delphi survey. *PLOS ONE*, 18(5), e0270616. <https://doi.org/10.1371/journal.pone.0270616>
- Garfield, E. (1955). Citation indexes for science: A new dimension in documentation through association of ideas. *Science*, 122(3159), 108–111. <https://doi.org/10.1126/science.122.3159.108>
- Genc, H. N., & Kocak, N. (2024). Bibliometric analysis of studies on artificial intelligence in science education with VOSviewer. *Journal of Education in Science, Environment and Health*, 10(4), 183–195. <https://doi.org/10.55549/jeseh.756>
- González, P., Fors, M., & Torres, A. (2025). Altmetrics in the evaluation of scholarly impact: A systematic and critical literature review. *Frontiers in Research Metrics and Analytics*, 10, 1693304. <https://doi.org/10.3389/frma.2025.1693304>

- Grootendorst, M. (2022). BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv*. <https://doi.org/10.48550/arXiv.2203.05794>
- Hassan, S. U., Saleem, A., Soroya, S. H., Safder, I., Iqbal, S., Jamil, S., Bukhari, F., Aljohani, N. R., & Nawaz, R. (2021). Sentiment analysis of tweets through Altmetrics: A machine learning approach. *Journal of Information Science*, 47(6), 712–726. <https://doi.org/10.1177/0165551520930917>
- Haunschild, R., & Bornmann, L. (2024). The use of OpenAlex to produce meaningful bibliometric global overlay maps of science on the individual, institutional, and national levels. *PLOS ONE*, 19(12), e0308041. <https://doi.org/10.1371/journal.pone.0308041>
- Hendricks, G., Tkaczyk, D., Lin, J., & Feeney, P. (2020). Crossref: The sustainable source of community-owned scholarly metadata. *Quantitative Science Studies*, 1(1), 414–427. https://doi.org/10.1162/qss_a_00022
- Hicks, D., Wouters, P., Waltman, L., De Rijcke, S., & Rafols, I. (2015). Bibliometrics: The Leiden Manifesto for research metrics. *Nature*, 520(7548), 429–431. <https://doi.org/10.1038/520429a>
- Hirsch, J. E. (2005). An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences of the United States of America*, 102(46), 16569–16572. <https://doi.org/10.1073/pnas.0507655102>
- Hood, W. W., & Wilson, C. S. (2001). The literature of bibliometrics, scientometrics, and informetrics. *Scientometrics*, 52(2), 291–314. <https://doi.org/10.1023/A:1017919924342>
- Hook, D. W., Porter, S. J., & Herzog, C. (2018). Dimensions: Building context for search and evaluation. *Frontiers in Research Metrics and Analytics*, 3, 23. <https://doi.org/10.3389/frma.2018.00023>
- Ioannidis, J. P. A., Boyack, K. W., & Wouters, P. F. (2016). Citation metrics: A primer on how (not) to normalize. *PLOS Biology*, 14(9), e1002542. <https://doi.org/10.1371/journal.pbio.1002542>
- Jurgens, D., Kumar, S., Hoover, R., McFarland, D., & Jurafsky, D. (2018). Measuring the evolution of a scientific field through citation frames. *Transactions of the Association for Computational Linguistics*, 6, 391–406. https://doi.org/10.1162/tacl_a_00028
- Kapoor, S., & Narayanan, A. (2023). Leakage and the reproducibility crisis in machine-learning-based science. *Patterns*, 4(9), 100804. <https://doi.org/10.1016/j.patter.2023.100804>
- Klavans, R., & Boyack, K. W. (2017). Which type of citation analysis generates the most accurate taxonomy of scientific and technical knowledge? *Journal of the Association for Information Science and Technology*, 68(4), 984–998. <https://doi.org/10.1002/asi.23734>
- Leydesdorff, L., Wouters, P., & Bornmann, L. (2016). Professional and citizen bibliometrics: Complementarities and ambivalences in the development and use of indicators—A state-of-the-art report. *Scientometrics*, 109(3), 2129–2150. <https://doi.org/10.1007/s11192-016-2150-8>
- Lu, X., Li, L., & Shen, Z. (2026). Revisiting the field normalization approaches and practices. *Scientometrics*. <https://doi.org/10.1007/s11192-026-05676-0>
- Lyu, D., Ruan, X., Xie, J., & Cheng, Y. (2021). The classification of citing motivations: A meta-synthesis. *Scientometrics*, 126(4), 3243–3264. <https://doi.org/10.1007/s11192-021-03908-z>
- Martín-Martín, A., Costas, R., van Leeuwen, T., & Delgado López-Cózar, E. (2018). Evidence of open access of scientific publications in Google Scholar: A large-scale analysis. *Journal of Informetrics*, 12(3), 819–841. <https://doi.org/10.1016/j.joi.2018.06.012>
- Molldrem, S., Moses, J., Smith, E., Tumilty, E., & Farroni, J. (2025). Metrics producing science: The evaluative way in the emergence of US translational science. *Science & Technology Studies*. <https://doi.org/10.23987/sts.145702>
- Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of Web of Science and Scopus: A comparative analysis. *Scientometrics*, 106(1), 213–228. <https://doi.org/10.1007/s11192-015-1765-5>
- Moral-Muñoz, J. A., Herrera-Viedma, E., Santisteban-Espejo, A., & Cobo, M. J. (2020). Software tools for conducting bibliometric analysis in science: An up-to-date review. *El Profesional de la Información*, 29(1), e290103. <https://doi.org/10.3145/epi.2020.ene.03>
- Omer, A. A. A., Issa, I. I. M., Abuker, Y. Y. A., & Asad, S. (2025). Evaluating the use of VOSviewer in bibliometric and science-mapping studies: Trends, current state and future directions. *International Journal of Current Research in Science, Engineering & Technology*, 8(4), 483–497. <https://doi.org/10.30967/IJCRSET/Altyeb-Ali-Abaker-Omer/208>

- Ortega, J. L. (2019). Exploratory analysis of Publons metrics and their relationship with bibliometric and altmetric impact. *Aslib Journal of Information Management*, 71(1), 124–136. <https://doi.org/10.1108/AJIM-06-2018-0153>
- Piwowar, H., Priem, J., Larivière, V., Alperin, J. P., Matthias, L., Norlander, B., Farley, A., West, J., & Haustein, S. (2018). The state of OA: A large-scale analysis of the prevalence and impact of Open Access articles. *PeerJ*, 6, e4375. <https://doi.org/10.7717/peerj.4375>
- Radhakrishnan, S., Erbis, S., Isaacs, J. A., & Kamarthi, S. (2017). Novel keyword co-occurrence network-based methods to foster systematic reviews of scientific literature. *PLOS ONE*, 12(3), e0172778. <https://doi.org/10.1371/journal.pone.0172778>
- Saeidnia, H. R., Hosseini, E., Abdoli, S., & Ausloos, M. (2025). Unleashing the power of AI: A systematic review of cutting-edge techniques in AI-enhanced scientometrics, webometrics and bibliometrics. *Library Hi Tech*, 43(4–5), 1292–1314. <https://doi.org/10.1108/LHT-10-2023-0514>
- Safder, I., Ali, M., Aljohani, N. R., Nawaz, R., & Hassan, S. U. (2023). Neural machine translation for in-text citation classification. *Journal of the Association for Information Science and Technology*, 74(10), 1229–1240. <https://doi.org/10.1002/asi.24817>
- Sugimoto, C. R., Work, S., Larivière, V., & Haustein, S. (2017). Scholarly use of social media and altmetrics: A review of the literature. *Journal of the Association for Information Science and Technology*, 68(9), 2037–2062. <https://doi.org/10.1002/asi.23833>
- Taylor, M. (2026). Evaluating open access advantages for citations and altmetrics (2011–2021): A dynamic and evolving relationship. *Quantitative Science Studies*, 7, 643–670. <https://doi.org/10.1162/QSS.a.470>
- Tshitoyan, V., Dagdelen, J., Weston, L., Dunn, A., Rong, Z., Kononova, O., Persson, K. A., Ceder, G., & Jain, A. (2019). Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571(7763), 95–98. <https://doi.org/10.1038/s41586-019-1335-8>
- Valenzuela, M., Ha, V., & Etzioni, O. (2015). Identifying meaningful citations. In *AAAI Workshop: Scholarly Big Data* (Vol. 15, p. 13).
- van Eck, N. J., & Waltman, L. (2017). Citation-based clustering of publications using CitNetExplorer and VOSviewer. *Scientometrics*, 111(2), 1053–1070. <https://doi.org/10.1007/s11192-017-2300-7>
- Visser, M., van Eck, N. J., & Waltman, L. (2021). Large-scale comparison of bibliographic data sources: Scopus, Web of Science, Dimensions, Crossref, and Microsoft Academic. *Quantitative Science Studies*, 2(1), 20–41. https://doi.org/10.1162/qss_a_00112
- Wahle, J. P., Ruas, T., Mohammad, S., & Gipp, B. (2022). D3: A massive dataset of scholarly metadata for analyzing the state of computer science research. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference* (pp. 2642–2651). European Language Resources Association. <https://doi.org/10.63317/3hvgsmuea6o7>
- Waltman, L., & van Eck, N. J. (2019). Field normalization of scientometric indicators. In W. Glänzel, H. F. Moed, U. Schmoch, & M. Thelwall (Eds.), *Springer handbook of science and technology indicators* (pp. 281–300). Springer International Publishing. https://doi.org/10.1007/978-3-030-02511-3_11
- Wang, X., Gu, P., & Shehzadi, K. (2026). Global open science research in the digital & intelligent era: Advances, focuses and prospects. *Scientometrics*. 10.1007/s11192-026-05597-y
- Weimer, V., Heck, T., van Leeuwen, T., & Rittberger, M. (2023). The quantification of open scholarship—A mapping review. *Quantitative Science Studies*, 4(3), 650–670. https://doi.org/10.1162/qss_a_00266
- Zhang, L., Cao, Z., Shang, Y., Sivertsen, G., & Huang, Y. (2024). Missing institutions in OpenAlex: Possible reasons, implications, and solutions. *Scientometrics*, 129(10), 5869–5891. <https://doi.org/10.1007/s11192-023-04923-y>

Information about authors

Magzhan Amangeldi – Master's degree, Doctoral Student, Astana IT University, Astana, Republic of Kazakhstan; e-mail: 255242@astanait.edu.kz; ORCID iD: 0009-0002-7979-0502.

Oleksandr Kuchanskyi – Doctor of Technical Sciences, Professor, School of Artificial Intelligence and Data Science, Astana IT University, Astana, Republic of Kazakhstan; National Technical University of Ukraine “Igor Sikorsky Kyiv Polytechnic Institute”, Kyiv, Ukraine; e-mail: kuczanski@gmail.com; ORCID iD: 0000-0003-1277-8031 (*corresponding author*).

Received: 10.07.2026

Revised: 28.08.2026

Accepted: 05.09.2026

Published: 30.09.2026

Type of Article: Research Article