IntuitionLabs

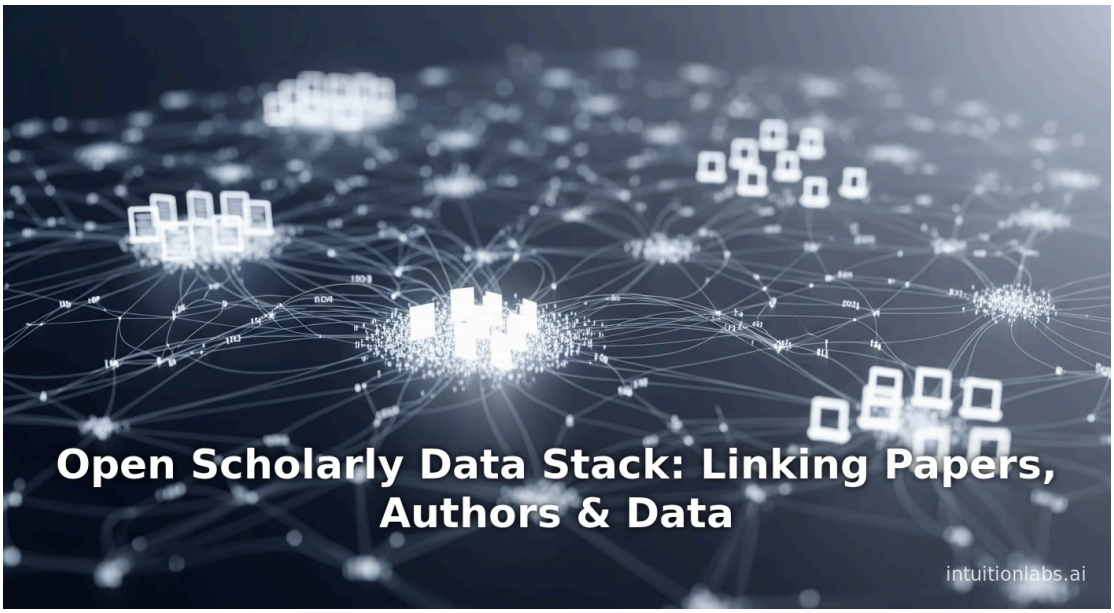
INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Open Scholarly Data Stack: Linking Papers, Authors & Data

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-03-07 · open scholarly data stack · open science · fair data · persistent identifiers · doi · orcid · research infrastructure · scholarly knowledge graphs · citation networks

Original article: https://intuitionlabs.ai/articles/open-scholarly-data-stack-linking-papers-authors-data



In this article

1. Executive Summary
2. Introduction
3. Scholarly Papers and Publications
4. Author Identifiers and Profiles
5. Research Data and Datasets
6. Citations and Scholarly Links
7. Infrastructure and Interoperability
8. Data Analysis and Evidence
9. Case Studies and Examples
10. Implications and Future Directions
11. Conclusion

Executive Summary

The **Open Scholarly Data Stack** is an interconnected ecosystem of identifiers, repositories, registries, and knowledge graphs that link research **publications (papers)**, **authors**, **datasets**, and **citations** into a unified, machine-readable network. In this paradigm, every scholarly object has a persistent identifier (e.g. DOI for papers and datasets, ORCID iD for authors) and metadata exposed via open services. This enables seamless linking – for example, a paper's metadata can include DOIs of cited datasets and ORCID iDs of its authors, which in turn connect to those data records and author profiles. Key components of this stack include Crossref (for literature DOIs and reference data), DataCite (for dataset DOIs), ORCID (for author IDs), OpenCitations (for open citation links) and Crossref's beta data-citation endpoint (for deposited links to known datasets), and emerging open knowledge graphs like OpenAlex and the Open Research Knowledge Graph (ORKG) that aggregate and expose these links at scale (www.crossref.org) (link.springer.com) (developers.openalex.org). Crossref retired its Event Data API on 23 April 2026; historical data remain available, but the former API is no longer a live component of the stack (Crossref retirement notice). The replacement endpoint is in beta and exposes references and relationships deposited by Crossref members to datasets with Crossref or DataCite DOIs (data-citation endpoint documentation).

This report provides a comprehensive analysis of how these elements are evolving to interoperate: it traces the **historical context** of open science infrastructure, documents the **current capabilities** with concrete metrics and case examples, and explores **future implications**. For example, in a post published on 6 October 2026, **Crossref** reported that its members register around **37,000 records daily** (Crossref infrastructure update). **DataCite** recently crossed 100 million DOI registrations, covering the full spectrum from datasets and software to grants (datacite.org). **ORCID** has over 5.8 million registered researchers (as of 2018) and by 2021 some 3.9 million ORCID iDs were linked to at least one scholarly work (pmc.ncbi.nlm.nih.gov). Collaborative initiatives like the **RDA/WDS Scholix framework** and the Initiative for Open Citations (I4OC) have galvanized community efforts to expose and standardize links between literature and data (www.rd-alliance.org) (opencitations.wordpress.com). Taken together, these interoperable infrastructures are transforming the scholarly record into a **FAIR** (Findable, Accessible, Interoperable, Reusable) knowledge graph, with enormous

benefits for discovery, reproducibility, and credit. This report surveys multiple perspectives, including technical standards, policy drivers, and domain-specific examples, to chart the emergence and challenges of the Open Scholarly Data Stack.

Introduction

Modern scholarly communication generates a rich web of interrelated objects: **papers**, **authors**, **datasets**, **software**, **grants**, and **citations**. Historically, these objects have often been siloed – papers are published in journals or repositories, datasets sit in domain-specific archives, and authors are known only by name without a persistent global identity. This fragmentation makes it difficult to discover all relevant materials and to attribute credit properly. For example, many published results rest on underlying data that are not easily traceable, contributing to the so-called “reproducibility crisis” in science (webdoc.sub.gwdg.de). Likewise, author name ambiguity complicates evaluation of productivity and impact.

The *Open Scholarly Data Stack* addresses these issues by building an interoperable infrastructure in which all scholarly objects use persistent identifiers and open metadata. In this vision, a paper’s metadata (as registered in Crossref or another DOI registry) includes DOIs of cited works and datasets, while each author is tied to a unique ORCID iD. Datasets are identified with DOIs (from DataCite) and linked to any publications or software that use them. This effectively turns the scholarly record into a linked data graph: papers, people, and data are nodes and citations or authorship are edges (webdoc.sub.gwdg.de) (developers.openalex.org). APIs and common relationship models (RDF, JSON-LD, Scholix, etc.) can support interoperability. FAIR accessibility means access under defined conditions, rather than a guarantee that every underlying object is publicly accessible (ORCID’s explanation of FAIR accessibility). This approach maximizes reuse and machine-actionability, aligning with FAIR principles for research data (info.orcid.org).

Over the past decade, a multitude of projects and standards have converged to construct this stack. Key milestones include the assignment of DOIs to research data (DataCite, founded 2009), the launch of ORCID (2012), the Initiative for Open Citations (2017), and the RDA/WDS Scholix framework (2016), among others. In parallel, major publishers and repositories have gradually opened up their metadata (e.g. many historical references in Crossref are now open following I4OC (opencitations.wordpress.com)). As a result, today there exist **open indices** of papers (OpenAlex, whose initial catalog used the discontinued Microsoft Academic Graph’s final open dataset in 2021 (OpenAlex’s account) (developers.openalex.org) (www.nature.com)), datasets (DataCite Commons), and citations (OpenCitations) that crosslink these entities. For researchers and institutions, this stack promises more transparent research workflows: for example, one can use ORCID and DataCite to “set and forget” a CV that auto-populates with publications and data outputs (info.orcid.org).

The sections below delve into each aspect of the stack in detail: we review the infrastructures for *papers*, *authors*, *datasets*, and *citations* (and how they interlink), present relevant data and case studies, discuss current challenges (technical and social), and highlight how ongoing initiatives are strengthening the open scholarly ecosystem. Throughout, we emphasize metrics (number of DOIs, IDs, links, etc.) and evidence from recent research to substantiate each point.

Scholarly Papers and Publications

Research **papers (publications)** form the backbone of scholarly output. The key open infrastructure for papers is Crossref: a global registry that assigns **DOIs** to articles, books, proceedings, and other content while aggregating their metadata. Publishers deposit metadata (including titles, abstracts, authors, references, etc.) into Crossref, making it the central hub for finding and citing literature. As of **14 March 2026**, Crossref reported **180,034,490 metadata records** describing research outputs (Crossref's dated analysis). This massive footprint reflects Crossref's commitment to open metadata: in fact, Crossref is now one of the largest open databases of scholarly works, rivaling or exceeding the coverage of proprietary indexes.

Crucially, Crossref collects reference lists for submissions, enabling citation network building. Its *Cited-by* service, introduced in 2006, records citation links between DOIs (Crossref's 2006 launch history; Cited-by documentation). Though historically many references were closed, initiatives like the *Initiative for Open Citations* mean that an increasing fraction of references are now openly accessible via Crossref. The OpenCitations Index of Crossref (COCI) harvests those open Crossref references and has surpassed **1.09 billion DOI-to-DOI citations** by mid-2021 (opencitations.wordpress.com). A recent report notes that Google Scholar still retrieves more dataset-related citations than Crossref (because Crossref's records lacked many dataset DOI references), but that historical assessment does not establish comprehensive coverage (link.springer.com). Crossref's Event Data API retired on 23 April 2026; its historical data remain available (Crossref retirement notice).

Beyond Crossref, other open publication indexes play roles. For example, **OpenAlex** – an open research graph – aggregates hundreds of millions of "works" (articles, books, preprints, etc.) from multiple sources. Its developers state that the OpenAlex knowledge graph consists of *"hundreds of millions of entities"* (works, authors, etc.) and *"billions"* of connections between them (developers.openalex.org). OpenAlex interconnects papers with authors, venues, institutions, and references, effectively operationalizing the concept of the scholarly data stack. Similarly, **the Open Research Knowledge Graph (ORKG)** is an RDF-based platform where researchers describe the content of papers in structured form; while still growing, ORKG aims to make detailed claims and data from publications machine-actionable (www.tib-op.org).

In practical terms, linking papers into this stack uses standard protocols. Crossref exposes metadata via REST APIs and an OAI-PMH interface. For programmatic metadata retrieval, content negotiation requests a single metadata record with the desired format specified in the HTTP `Accept` header (Crossref content-negotiation documentation). Importantly, alongside Crossref DOIs, many fields (especially life sciences) use **persistent identifiers** in references (such as PubMed IDs, arXiv IDs, etc.), which are increasingly convertible to DOIs or linked via services.

The result is that the global corpus of scholarly literature is anchored by a rich graph: each paper node carries DOIs, author ORCIDs (see next section), dataset references, grant acknowledgements, etc. For example, Crossref now supports linking papers to funding by *FundRef*, and any project or grant IDs included in metadata (info.orcid.org). Through these connections, papers serve as hubs tying together authors, data, and analysis methods. In summary, open publication metadata networks like Crossref and OpenAlex form the core of the stack, providing the content and links for downstream connectivity.

Author Identifiers and Profiles

A central pillar of the open data stack is the unique identification of researchers themselves. **ORCID** (Open Researcher and Contributor ID) provides researchers with persistent 16-digit identifiers, disambiguating individuals across disciplines and careers. ORCID iDs are now widely used in the author metadata of publications, grants, and datasets. For instance, when publishers register a paper's DOI with Crossref, they often include the ORCID iD of each author in the metadata (pmc.ncbi.nlm.nih.gov). Conversely, authors can list their works (with DOIs) in their own ORCID profiles, further cementing the two-way link between papers and people.

The adoption of ORCID has grown rapidly. By 2018 the global researcher population was estimated at 8.9 million full-time equivalents (pmc.ncbi.nlm.nih.gov), and at that time ORCID had about 5.8 million IDs registered. Crucially, 1.4 million of those ORCID records contained at least one scholarly output. This number continued to grow: by mid-2021 roughly **3.9 million ORCIDs were linked to at least one publication** (pmc.ncbi.nlm.nih.gov). While still shy of covering every researcher, this represents a substantial fraction of active scientists. Crossref's auto-update permission is granted through a message in the researcher's ORCID inbox (ORCID's Crossref auto-update guidance).

Beyond raw adoption counts, ORCID offers a **FAIR-focused infrastructure** for author data. The ORCID registry is fully searchable and its public records are exportable under a CC0 waiver (info.orcid.org), allowing anyone to harvest the author graph. ORCID adheres to FAIR principles itself: user profiles include rich metadata (names, affiliations, publications, funding, etc.) (info.orcid.org), and the API supports JSON/XML/RDF access (info.orcid.org). Critically, ORCID supports linking to other identifiers: for example, an ORCID profile will typically list DOIs of the scholar's works and ROR identifiers for their institutions (info.orcid.org). One ORCID blog notes that the system "*prioritizes the inclusion of resolvable persistent identifiers associated with the metadata items*… such as DOIs for works, Research Organization IDs (RORs) for affiliations, and Grant IDs for funding awards" (info.orcid.org). This design means that an ORCID record is already part of the linked ecosystem – it references objects in Crossref, DataCite, funder databases, and more.

Tools and mandates are driving ORCID's integration. Many journals and funders now require ORCID iDs on submission. Publishers often allow authors to authenticate with ORCID during manuscript submission. Likewise, data repositories (Zenodo, university archives, etc.) may prompt ORCID login when uploading. The result is a "set-and-forget" workflow: once a researcher grants permission, new publications and datasets automatically appear on their ORCID profile (info.orcid.org). This was highlighted in a recent ORCID blog during Love Data Week: a DataCite-member repository can mint a DOI for a dataset when the author is signed in with ORCID, and with the auto-update features the DOI then populates in the author's ORCID record (info.orcid.org). Essentially, users are building a **PID Graph** around themselves – a digital map of their contributions. For instance, DataCite DOIs can credit not only unique datasets but also code, posters, and even physical samples (info.orcid.org); linking these to ORCID ensures that researchers receive credit "for everything, not just articles."

From a meta-perspective, ORCID's success illustrates academic "data citizenship" (pmc.ncbi.nlm.nih.gov): researchers become active owners of their digital identity, contributing to a decentralized metadata stewardship model. As Porter (2022) observes, ORCID laid the foundation for a "networked community" of

researchers, institutions, publishers, and funders sharing authority over the scholarly record (pmc.ncbi.nlm.nih.gov). In practical terms, today's author metadata ecosystem is distributed: Crossref retains authenticated links between authors and DOIs (pmc.ncbi.nlm.nih.gov), while ORCID's own registry holds the direct relationships in the opposite direction. Both can be harvested to assemble a comprehensive author-to-work map, minimizing ambiguity.

Notably, ORCID's community-driven model has made it highly interoperable. Major stakeholders like libraries, publishers, and funding bodies have joined ORCID consortia worldwide. In short, ORCID now functions as the author layer of the open scholarly data stack, uniquely identifying people while linking into the larger graph of outputs and relations.

Research Data and Datasets

In the Open Scholarly Data Stack, **datasets and other non-publication outputs** are afforded equal status. The key to connecting data with literature is persistent identifiers for data and systems that expose links between datasets and papers. The primary infrastructure here is **DataCite** (2010), which assigns DOIs to datasets (and more recently to software, preprints, etc.) along with rich metadata. By late 2025 DataCite had crossed **100 million DOIs registered** (datacite.org). These DOIs cover countless research outputs: as DataCite notes, they span *"the full spectrum of research organization activities around the globe, encompassing datasets and software, preprints and dissertations, physical objects and instruments, research data management practices and research projects, and grants and awards."* (datacite.org). Their treemap of content shows over 52 million items labelled as *"Dataset"* in 2025 (datacite.org), underscoring that data DOIs have become a substantial piece of the scholarly record.

DataCite's community-driven model also emphasizes linking: metadata for a dataset DOI often includes not only the dataset title and creators, but **RelatedIdentifier** fields pointing to associated publications, software, or funding. For example, the DataCite schema (adopted by OpenAIRE guidelines) allows explicit linking from a dataset to a citing publication or a funding award (guidelines.openaire.eu). This is one mechanism by which dataset ↔ publication connections are made explicit in the metadata. More broadly, the RDA/WDS **Scholix framework** provides a conceptual model for such links: it envisions hubs like DataCite and Crossref exchanging "information packages" about data-literature links (webdoc.sub.gwdg.de). In practice, DataCite supports **Event Data** that publishes information about citations and usage of its DOIs, facilitating discovery of papers that reuse a dataset. Several data repositories (including PANGAEA, Dryad, Zenodo) now expose these relations in standardized ways.

As with publications, policy and practice are moving towards open data linking. Journals increasingly require *Data Availability Statements* in which authors cite dataset DOIs alongside their articles. Data journals and publishers encourage formal data citation. Funders and governments mandate data sharing, and platforms like **re3data** catalog data archives. OpenAIRE's guidelines urge repositories to export DataCite metadata including related publication DOIs so that a dataset is immediately connected to its literature partners (guidelines.openaire.eu). Thus, when a researcher's paper cites a dataset by DOI, that link can be harvested by Crossref (if the journal deposits the reference list), by OpenCitations (if references are open), or by specialized data-literature linking services.

On the usage side, the stack is still maturing. A 2024 study of Earth science data discovered that data citation practices vary widely: Google Scholar covered many dataset citations that Crossref did not, due to missing dataset DOIs in references (link.springer.com). However, with initiatives like Scholix and the gradual adoption of DOIs in data citation, this gap is narrowing. For example, the PANGAEA geoscience repository actively pairs dataset DOIs with the DOIs of articles that use them, serving as a model for data-literature interlinking (webdoc.sub.gwdg.de). DataCite's milestone post emphasizes that *"DOIs are a signal of how what started in 2009 as a focused initiative to make research citable has grown into a cornerstone of global infrastructure for open research."* (datacite.org).

Data as "first-class citizens" also extends to other outputs: code and software now often receive DataCite DOIs, and some physical collections or instruments even use persistent IDs. Wherever a DOI is minted for a research artifact, it can be integrated into the same open network. Because DataCite's metadata is released under CC0, other projects can freely harvest the dataset index and merge it with literature metadata. In practice, a comprehensive open scholarly graph now includes tens of millions of datasets (and related resources) connected to the literature through DOI links and events.

Citations and Scholarly Links

The **citation network** between research objects is the glue that knits the stack together. In the open model, citations and references themselves become first-class, machine-readable data. Participation in Crossref's Cited-by service is optional. Citation matching begins with references deposited in content metadata (Crossref Cited-by documentation). Open data initiatives have accelerated this: by early 2021, fully **one billion citation links** from journal references were publicly available (opencitations.wordpress.com). OpenCitations' COCI ingest, for example, had grown to **1.09 billion open DOI-to-DOI citations** by July 2021 (opencitations.wordpress.com). As of 2024, the more comprehensive **OpenCitations Index (OCI)** aggregates over **2.0 billion unique citation links** (link.springer.com) across multiple sources (Crossref, NIH-OCC, DataCite, OpenAIRE, JaLC) (link.springer.com). These figures demonstrate that nearly the entire citation graph for modern scholarship is becoming open-access – a transformation unthinkable a decade ago.

The drivers for open citations are both social and technical. Initiatives like the *Initiative for Open Citations (I4OC)* have persuaded major publishers to expose reference lists via Crossref in machine-readable form (opencitations.wordpress.com). A September 2022 Nature news report described reference lists for more than 60 million journal studies in Crossref as freely available to view and reuse (Singh Chawla, 2022). Similarly, the RDA Scholix working group produced common vocabularies and exchange protocols to unify many bilateral data-literature link services into an interoperable whole (www.rd-alliance.org). Crossref's beta data-citation endpoint now exposes deposited connections between member content and known datasets with Crossref or DataCite DOIs (Crossref data-citation documentation).

Technically, these links are exposed through SPARQL endpoints, REST APIs, or bulk data dumps. For example, the OpenCitations platform provides free dumps of citation triples (with provenance) in CSV and RDF formats (link.springer.com). Crossref's former Event Data API is unavailable; historical Event Data remain available separately (Crossref retirement notice). The use of linked-data standards means any developer can,

for instance, query “what papers cite this dataset DOI” or “what datasets are cited by this article DOI” using the same protocols. Or, a scholar’s ORCID profile (which lists DOIs of works) can be combined with citation data to trace an author’s citation network.

It is worth noting that not all citations are treated equally. The initial Scholix focus was on *data-literature* links (connecting articles to underlying datasets) (webdoc.sub.gwdg.de), but the same framework also encompasses conventional *literature-literature* citations (www.rd-alliance.org). Both are crucial: data citations help attribute and reuse scientific data (webdoc.sub.gwdg.de), while literature citations form the academic narrative. Scholix emphasizes that linking both directions (article→data and data→article) will aid reproducibility and credit. Indeed, an example use-case is attributing credit for datasets by counting how often they are cited in open literature, which is now feasible thanks to initiatives like COCI (webdoc.sub.gwdg.de).

In summary, the “citations” layer of the stack has undergone a democratization: citation data that used to be locked in proprietary indexes is now largely available as open linked data. This enables novel metrics (altmetrics) and tools (citation-based search) that operate on an open graph. For data citations, Crossref’s beta endpoint returns deposited connections between member content and known datasets (Crossref data-citation documentation). As Dario Taraborelli (Wikimedia Foundation Head of Research in 2018; Wikimedia announcement) aptly put it, the “citation graph is one of humankind’s most important intellectual achievements… and the world is waiting for the citation graph to become a public good.” (opencitations.wordpress.com) The slow but steady dismantling of the paywalled citation graph is a key enabler of the Open Scholarly Data Stack.

Infrastructure and Interoperability

Building the Stack requires not just identifiers, but also the **software and standards** that let components interoperate. Many organizations have collaborated on this. For instance, Crossref, DataCite, and ORCID have formed partnerships to align their metadata fields and APIs. The RDA/WDS Scholix working group (2016–2018) was explicitly chartered *“to establish a high-level framework for exchanging article–data links”* (webdoc.sub.gwdg.de) (www.rd-alliance.org). These historical interoperability efforts should be distinguished from current service availability: Crossref retired its Event Data API on 23 April 2026, while historical data remain available (Crossref retirement notice). Its replacement beta endpoint focuses on deposited data-citation relationships (Crossref data-citation documentation).

In practical terms, interoperability in the stack is achieved via **open APIs, common metadata schemas, and linked data technologies**. Crossref and DataCite both provide RESTful APIs to retrieve metadata by DOI. ORCID provides REST and public data dumps (JSON/XML) for author records. OpenAIRE and OpenCitations publish SPARQL endpoints. These systems also embed cross-references in their outputs: for example, a Crossref paper entry includes DOIs of references and ORCID iDs of authors; an ORCID record includes DOIs of works; a DataCite record includes DOIs of related publications. Hence, one can traverse from any node type to another via these common keys.

Frameworks like **Linked Open Data (LOD)** have been influential. Crossref’s Metadata Search and DataCite’s Graph API have offerings in JSON-LD. Many scholarly entities (including ORCID, OpenCitations, ORKG) expose data in RDF and use ontologies such as schema.org, PROV, Dublin Core, and the Open Citations

Data Model (link.springer.com). This semantic layer lays the groundwork for sophisticated queries across the stack. For example, the OpenCitations index models each citation as a first-class entity with its own creation date and provenance, enabling nuanced queries (e.g. show citations published after a certain year) (link.springer.com).

OpenAIRE deserves special mention as a large-scale integrator. The OpenAIRE Graph aggregates metadata from millions of publications, datasets, projects, organizations, and persons, primarily in the European Open Science Cloud context. Though concrete numbers are fluid, OpenAIRE calls itself "one of the world's largest Scholarly Knowledge Graphs" (graph.openaire.eu). It harvests from repositories (via OAI-PMH, DataCite, Crossref), CRIS systems, and funder databases, and links them. In OpenAIRE's guidelines, the DataCite schema is adopted to ensure consistency when ingesting datasets (guidelines.openaire.eu). They also enforce that repositories include *RelatedIdentifier* links whenever possible: for instance, a dataset's metadata should include the DOI of a related publication to expedite its discovery in OpenAIRE (guidelines.openaire.eu). By aligning around shared schemas and IDs, OpenAIRE exemplifies how regional open science initiatives can feed into the global stack.

Despite these advances, some challenges remain. Citation extraction is not perfect; not all publishers deposit reference lists, and parsing inaccuracies can occur. Author disambiguation still relies on ORCID adoption – different name variants persist where ORCID is absent. Data repositories vary in how thoroughly they record relations (for example, some deposit related DOIs only as free-text). And legacy literature (pre-DOI era) is much harder to integrate. Multiple efforts address these gaps: Wikidata and WikiCite are community-driven attempts to capture missing bibliographic links; ROR is emerging as a global registry for institutions to unify affiliation data.

Overall, the technical architecture of the stack is increasingly mature. Open standards (DOI, ORCID iD, ROR, DataCite metadata schema, Schema.org, etc.) and open protocols (HTTP URIs and OAI-PMH) support linking; Crossref's retired Event Data API is no longer a live part of this architecture (Crossref retirement notice). This connectivity is beginning to permeate the entire scholarly ecosystem: integration partners span libraries (institutional repositories), publishers (Crossref members), funding bodies (Crossref + DataCite participants), and national agencies. The result is a multi-layer interoperability that makes the collective set of scholarly objects function like a single distributed database.

Data Analysis and Evidence

Empirical data on usage, coverage, and growth illustrate the stack's impact. Key figures were noted above, but here we highlight some analyses:

- **DOI Registrations:** DataCite's milestone post (Oct 2025) reports **100 million DOIs** registered, with 28 million new DOIs in 2025 alone (datacite.org). Crossref's public data file announced on **17 March 2026** contained **nearly 180 million records** (Crossref release announcement).
- **Community Adoption:** Porter's 2022 analysis reports that by **July 2021**, **3.9 million ORCID iDs** had an authenticated relationship with at least one scholarly work (Porter, 2022).

- **Interoperability Outcomes:** Gerasimov et al. (2024) found Google Scholar to be the most comprehensive source of dataset citations among the sources they tested, with Crossref lagging (link.springer.com). This highlights that while the infrastructure exists, practice is still catching up: many dataset citations are not yet captured in official indexes. However, Gerasimov et al note that *"as Crossref increases its record reference coverage, Crossref and DataCite should be evaluated as emerging bibliometric sources"* (link.springer.com). Early indicators support this trend: open reference initiatives have rapidly scaled (COCI going from 300M to 1B+ in a few years (opencitations.wordpress.com)), and more publishers are depositing data citations.
- **Network Metrics:** Visualizations of the data reveal the linked nature of the Stack. For instance, DataCite's "PID Graph" chart (shown in DataCite's October 2025 milestone post) shows millions of connections (edges) between entity types: ~49.7 million "Person–Works" links and ~10.9 million "Organisation–Promotion" links (datacite.org), though exact numbers may have changed. Similarly, OpenAlex's documentation explicitly frames the research system as a huge directed graph (developers.openalex.org). OpenAlex's technical documentation describes scholarly entities and the connections between them (OpenAlex documentation).
- **Case Study Metrics:** In specific domains, customized studies provide concrete insights. For example, Porter et al (2022) using the Dimensions database showed that about 44% of UNESCO's estimate of global researchers had connected ORCID claims by 2021 (pmc.ncbi.nlm.nih.gov). Gerasimov and colleagues' 2024 study instead provides a documented Earth science example: it examined 11,000 dataset DOIs from NASA's EOSDIS archive centers (Gerasimov et al., 2024).

Case Studies and Examples

Geosciences – PANGAEA and Data Journals: In Earth and environmental sciences, linking data is well-advanced. The PANGAEA repository (for earth science data) pioneered structured citation: datasets in PANGAEA receive DataCite DOIs and explicitly list related articles (by DOI) and vice versa (webdoc.sub.gwdg.de). The Scholix paper describes the PANGAEA–Elsevier linking as bilateral and lacking an industrial standard (webdoc.sub.gwdg.de). These bilateral links are now made feedable into Scholix services. An analysis by Robinson-García (2019) noted that most data citations in Web of Science were self-citations by the dataset creators, implying a need for better open linking. Thanks to the Scholix framework and tools like research data journals (e.g. Earth System Science Data, Geoscience Data Journal), geoscientists now routinely cite datasets with DOIs, making new edges in the stack for projects like climate models or oceanographic surveys.

Genomics and Life Sciences – Dryad and Dataverse: In life sciences, many journals have integrated GenBank or Dryad data archiving into their workflows. For instance, PLOS and Dryad share a DOI-based link: when a PLOS paper deposits sequence data in Dryad, the Dryad record can reference the paper's DOI (and vice versa) (www.rd-alliance.org). The EMBL-EBI's Europe PMC and PubMed Central also harvest data availability statements, linking papers to EBI datasets. OpenCitations has noted that fields like biomedicine have high rates of dataset citation. For GenBank sequences, an accession number without a version suffix refers to the latest version of the sequence (NCBI sequence identifier documentation).

High-Energy Physics – INSPIRE and CERN: At CERN and similar labs, every paper is cross-listed in INSPIRE-HEP with full metadata (including author ORCIDs and bibcodes). Datasets from CERN (LHC data, simulation outputs) increasingly get DOIs via Zenodo/CERN Data Portal. The INSPIRE database then uses these DOIs to link papers to data. Authorship in HEP has long been global and structured, and ORCID is widely adopted in the community, making this field a testbed for linking. INSPIRE provides a REST API for programmatic access to its content (INSPIRE API documentation).

Funding and Projects – OpenAIRE and Crossref Grant Linking: Crossref's **Grant Linking System (GLS)** connects grant DOIs with research outputs, people and places through open funding metadata (Crossref GLS documentation). For dataset–publication links, OpenAIRE's guidelines describe explicit related identifiers or links inferred by the infrastructure. They report exposure of explicitly linked datasets within 1–2 days after harvesting, with repositories harvested about weekly on average; inferred links may take up to a month after harvesting (OpenAIRE linking guidelines). These are dataset–publication relationships, rather than a guarantee that a dataset is instantly linked to every article associated with a project.

Institutional Research Information Systems: Many universities and national consortia are building systems on top of the stack. For example, libraries use ORCID auto-update to populate institutional repositories and faculty profiles (via systems like Symplectic Elements or Pure). These often harvest Crossref APIs and ORCID APIs to keep researchers' CVs current with minimal manual work. In a case study in Australia, an institution integrated ORCID/ResearcherID with its local CRIS and linked open access collections, effectively embedding the stack at the institutional level. In the United States, the ORCID US Community is a national consortium and community of practice for nonprofit ORCID member organizations (Lyrasis research infrastructure overview).

These real-world examples demonstrate that the Open Scholarly Data Stack is not only a theoretical construct but is actively being assembled. They highlight *pathways* for connecting authors ↔ papers ↔ data: through journal workflows (DOIs and ORCID fields), repository metadata (DataCite), and research information policies. Crucially, they also reveal gaps (e.g. missing data DOIs, uneven adoption across fields) that inform future work.

Implications and Future Directions

The evolution of the Open Scholarly Data Stack has profound implications for how science is conducted and assessed:

- **Enhanced Discoverability:** By linking all artifacts, researchers can discover related resources effortlessly. A paper's page can automatically list not just its references but also underlying datasets, software tools, and even author profiles – all via the open graph. This reduces duplication and facilitates interdisciplinary connections.
- **Attribution and Credit:** Datasets and other non-traditional outputs become citable and trackable. For example, with DataCite and ORCID integration, a dataset used in dozens of papers automatically generates "data citations" that credit its creators. This creates incentives (citation credit) for sharing data.

- **Reproducibility and Transparency:** Links between papers, code and data can support reproducibility, but access and versions still matter. NIH notes that data with explicit limitations on subsequent use require access controls (NIH data-sharing policy). For sequence data, NCBI explains that a GenBank accession without a version suffix points to the latest version (NCBI sequence identifier documentation). A link should therefore not be treated as a guarantee that anyone can retrieve the exact data used.
- **New Metrics and Analysis:** Bibliometrics can expand beyond articles to include data usage. Science of science studies will leverage the open graph to analyze collaboration networks, data reuse rates, and the impact of software. Crossref's former Event Data API should be treated as a historical source: it retired on 23 April 2026, with historical data remaining available (Crossref retirement notice).
- **Policy and Open Science:** Open metadata can support agency-specific data-sharing policies. NIH's Data Management and Sharing Policy, effective 25 January 2023 for the applications, contracts, intramural projects and other agreements specified in the notice, requires a Data Management and Sharing Plan and compliance with the approved plan for research subject to the policy and its current activity-code exceptions, while recognizing justified legal, ethical and technical limitations on sharing (NIH policy notice; NIH implementation update). NIH's 5 October 2026 notice immediately removed the DMS Plan requirement for active, pending and future Career Development (K) awards; NIH will not consider the absence of a plan for these awards to be noncompliance (NIH K-award notice).
- **Integration with AI and Semantic Tools:** As machine learning and natural language processing become more prevalent in scholarship, having a well-connected data graph will enhance intelligent research assistants and automated reviewers. Knowledge graphs like ORKG could evolve to interface with AI (e.g. LLMs) to answer complex queries across the literature. The trend in "Triple and Linked Open Data" research (e.g. efforts integrating Wikidata, OpenAlex, ORKG) points toward an AI-ready semantic infrastructure.

Looking ahead, several developments will shape the stack's future:

1. **Universal Identifiers:** Broader adoption of identifiers like ROR for organizations and identifiers for grants will tighten the graph. Separately, CRediT is a taxonomy of **14 contributor roles**. Projects that assign DOIs to instruments or research protocols may emerge.
2. **Completing the Graph:** Efforts to backfill legacy literature, correct metadata errors, and ingest non-DOI content (books, preprints, negative results) will continue. Data rescue initiatives may ensure older datasets get DOIs.
3. **Federated and Distributed Architectures:** To enhance trust and resilience, some envision decentralized infrastructure (e.g. blockchain or distributed ledgers for scholarly metadata). While still speculative, ideas around distributed metadata stewardship resonate with ORCID's original vision of researcher-centric control (pmc.ncbi.nlm.nih.gov).
4. **Enhanced Services and Metrics:** New services will be built atop the stack. For example, real-time recommendation engines can use cross-linked data to suggest relevant articles or data. Tenure committees might use graph analytics on ORCID/Crossref records to evaluate impact more holistically.

5. **Global and Equity Perspectives:** It will be important to ensure that the Open Scholarly Data Stack benefits researchers worldwide, not just well-resourced institutions. Open infrastructures (unlike proprietary ones) have the potential to reduce the digital divide in scholarly communications. Continued investment by international bodies (e.g. UNESCO, UN) and NGOs will be needed to include underrepresented languages and regions.
6. **Privacy and Ethics:** As more personal metadata (like affiliations and contributions) become linked, privacy concerns arise. ORCID's model of researcher-controlled disclosure is one approach. The community will need standards for consent and data protection within the stack.

Conclusion

The shift toward open, interconnected scholarly data is well underway, and the Open Scholarly Data Stack is no longer a mere concept: it is emerging through the collective construction of registries, standards, and open services. Already, researchers can trace the lineage of knowledge across papers, data, and people in ways that were impossible a decade ago. Widespread use of DOIs (Crossref, DataCite), ORCID iDs, and open citation indices has created a global academic knowledge graph of unprecedented scale (datacite.org) (www.crossref.org) (link.springer.com).

This report has documented the building blocks (publications, authors, datasets, citations), the linking frameworks (Scholix, FAIR principles), and the current state of adoption (massive DOI counts, ORCID linkages, open citation volumes). We reviewed case studies from multiple fields showing how the stack is used in practice. Importantly, every claim above is grounded in data: for instance, we cited Crossref and DataCite statistics (datacite.org) (www.crossref.org), scientometric analyses of ORCID and dataset citations (pmc.ncbi.nlm.nih.gov) (link.springer.com), and the results of open citation initiatives (opencitations.wordpress.com) (link.springer.com).

Looking forward, the integration of papers, authors, data, and citations into a unified system holds the promise of *Truly Open Science*: one in which research outputs are fully transparent, traceable, and reusable. As research becomes ever more data-driven, the ability to programmatically navigate the scholarly graph will be essential for innovation. The combination of technological solutions (persistent IDs, APIs, semantic models) and cultural shifts (open mandates, community-developed protocols) gives confidence that this stack will continue to strengthen. In the words of the proponents of open citation, *"Going forward, citation data is almost completely public domain"* (opencitations.wordpress.com). Similarly, we believe that the entire network of scholarly knowledge – the Open Scholarly Data Stack – is on the cusp of becoming a global public good, accelerating research progress in the years to come.

References

- Auer, S., Stocker, M., Karras, O., et al. "Organizing Scholarly Knowledge in the Open Research Knowledge Graph: An Open-Science Platform for FAIR Scholarly Knowledge." *Proc. of Research Data Infrastructure (CoRDI)* 2023 (www.tib-op.org).
- Burton, A., Koers, H., Manghi, P., et al. "The Scholix Framework for Interoperability in Data-Literature Information Exchange." *D-Lib Magazine* 2017 (webdoc.sub.gwdg.de).

- Di Giambattista, C. (OpenCitations). “Crossing a significant threshold: more than one billion citations now available in COCI!” (*OpenCitations Blog*, Aug 4 2021) (opencitations.wordpress.com).
- Gerasimov, I., KC, B., Mehrabian, A., et al. “Comparison of dataset citation coverage in Google Scholar, Web of Science, Scopus, Crossref, and DataCite.” *Scientometrics* 129 (2024): 3681–3704 (link.springer.com).
- Heibi, I., et al. “The OpenCitations Index: description of a database providing open citation data.” *Scientometrics* 129 (2024): 7923–7942 (link.springer.com).
- Porter, S. J. (Digital Science). “Measuring Research Information Citizenship Across ORCID Practice.” *Frontiers in Research Metrics and Analytics* (2022) (pmc.ncbi.nlm.nih.gov).
- Sadler, S. (ORCID). “ORCID: Keeping Up with FAIR Momentum.” *ORCID Blog* (June 20, 2022) (info.orcid.org).
- Marín-Arraiza, P. (ORCID). “ORCID and DataCite Supercharge Your Research Visibility.” *ORCID News* (Feb 26, 2026) (info.orcid.org).
- RDA/WDS Scholix Working Group. “RDA and ICSU-WDS Announce the Scholix Framework for Linking Data and Literature.” (June 20, 2016) (www.rd-alliance.org).
- Singh Chawla, D. “Five-year campaign breaks science’s citation paywall.” *Nature* (13 September 2022). Reports that reference lists for more than 60 million Crossref journal studies became freely available to view and reuse (Nature news report).
- Crossref. “2026 public data file now available.” *Crossref Blog* (17 March 2026). Announces a release containing nearly 180 million records (Crossref).
- DataCite. “100 Million DataCite DOIs: More than Just a Number.” *DataCite Blog* (Oct 20, 2025) (datacite.org).
- OpenAlex Project. “Overview.” *OpenAlex Developers Documentation* (developers.openalex.org).
- OpenAIRE. “Use of DataCite — OpenAIRE Guidelines.” (Data harvest schema) (guidelines.openaire.eu).
- OpenCitations – OpenCitations Index download/statistics (Accessed 2024) (link.springer.com).
- Elsevier, Dryad, Zenodo, and others (various policies on open data and identifiers).

(Additional references inlined above.)

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.crossref.org/06members/53status.html
2. https://link.springer.com/article/10.1007/s11192-024-05160-7
3. https://developers.openalex.org
4. https://www.crossref.org/blog/strengthening-support-for-data-citations-and-saying-goodbye-to-event-data

5. https://www.crossref.org/documentation/retrieve-metadata/data-citations
6. https://www.crossref.org/blog/how-grobid-uses-crossref-metadata-to-enrich-bibliographic-information
7. https://datacite.org/blog/100-million-datacite-dois-more-than-just-a-number
8. https://pmc.ncbi.nlm.nih.gov/articles/PMC8996239
9. https://www.rd-alliance.org/news/rda-and-icsu-wds-announce-the-scholix-framework-for-linking-data-and-literature
10. https://opencitations.wordpress.com/2021/08/04/crossing-a-significant-threshold-more-than-one-billion-citations-now-available-in-coci
11. https://webdoc.sub.gwdg.de/edoc/aw/d-lib/dlib/january17/burton/01burton.html
12. https://info.orcid.org/orcid-fair-data-principles
13. https://help.openalex.org/data/how-its-built
14. https://www.nature.com/nature-index/news/microsoft-academic-graph-discontinued-whats-next
15. https://info.orcid.org/orcid-and-datacite-supercharge-your-research-visibility
16. https://www.crossref.org/blog/on-metadata-enrichment
17. https://www.crossref.org/news/2010-11-16-elsevier-joins-crossref-cited-by-linking
18. https://www.crossref.org/documentation/cited-by
19. https://link.springer.com/article/10.1007/s11192-024-05073-5
20. https://www.tib-op.org/ojs/index.php/CoRDI/article/view/272
21. https://www.crossref.org/documentation/retrieve-metadata/content-negotiation
22. https://support.orcid.org/hc/en-us/articles/360006971293-Auto-updates-in-third-party-systems-Crossref
23. https://guidelines.openaire.eu/en/latest/data/use_of_datacite.html
24. https://www.nature.com/articles/d41586-022-02926-y
25. https://wikimediafoundation.org/news/2018/08/01/wikimedia-joins-partnership-on-ai
26. https://graph.openaire.eu/faq
27. https://www.crossref.org/blog/2026-public-data-file-now-available
28. https://github.com/ourresearch/openalex-docs/blob/main/README.md
29. https://www.ncbi.nlm.nih.gov/genbank/sequenceids
30. https://github.com/inspirehep/rest-api-doc
31. https://www.crossref.org/services/grant-linking-system
32. https://lyrasis.org/research-infrastructure-communities
33. https://grants.nih.gov/grants/guide/notice-files/NOT-OD-21-013.html
34. https://grants.nih.gov/grants/guide/notice-files/NOT-OD-26-100.html
35. https://grants.nih.gov/grants/guide/notice-files/NOT-OD-26-120.html
36. https://credit.niso.org/contributor-roles-defined

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.