

NEWS LETTER

Research — born,
community — sustained

November 2026

<https://core.ac.uk/governance>



1.Introductory Letter	4
2.CORE as an independent organisation.....	5
3.Services in pilot	5
3.1FreshFinds: from pilot to Sustaining Members.....	5
3.2FAIR Certification: what the pilots are teaching us.....	6
4.Repository tools and integrations.....	7
4.1ECORE in your repository platform: DSpace and DSpace-CRIS.....	7
5.Projects and partnerships.....	8
5.1EOSC AIsstant is under way.....	8
6. Community and collaboration.....	9
6.1Understanding the CORE Dashboard: the series continues.....	9
7.Events: webinars and conferences.....	9

Welcome to Our New
CORE Member 🎉



Delft University of Technology

Thank You for Renewing Your Membership



University of
Nottingham



Queen Mary
University of London

Our Thanks to Every CORE Member



UNIVERSITY OF
BIRMINGHAM



BUCKINGHAMSHIRE
NEW UNIVERSITY
EST. 1891



UNIVERSITY OF
BATH



University of
BRISTOL



Imperial College
London

US
UNIVERSITY
OF SUSSEX



White Rose
Libraries
Universities of Leeds, Sheffield & York

Lancaster
University



ual:

UNIVERSITY OF
OXFORD



Nottingham Trent
University



Manchester
Metropolitan
University

TIB

LEIBNIZ INFORMATION CENTRE
FOR SCIENCE AND TECHNOLOGY
UNIVERSITY LIBRARY



University of
Reading



University
of Exeter

CARDIFF
UNIVERSITY
PRIFYSGOL
CAERDYDD



UNIVERSITY OF
GREENWICH



University
of Glasgow



Durham
University



THE UNIVERSITY OF
CHICAGO



UNIVERSITY
OF HULL

Dear Members of the CORE Community,

Before anything else, thank you. CORE has reached this point because of the community that uses it and the members who fund it, and this year that support has mattered more than ever.

CORE: CONnecting REpositories is now an independent, not-for-profit organisation. After fifteen years of being built inside a university, it holds its own agreements, employs its own staff and sets its own priorities. What this means in practice for members and data service clients, including the transfer of agreements on their existing terms, is set out on the next page. Here I want to say something about why it matters.

[CORE began in 2010 as a PhD project in the Knowledge Media Institute](#), built on the idea that open research should be accessible to machines as well as to people. A Jisc-commissioned tender in 2013 made CORE the UK's national open access indexing service, and the tools repository teams use today, the Dashboard and the Recommender among them, grew out of those early years. CORE now indexes more than [450 million records from 16,000 data providers in more than 150 countries](#).

That work happened alongside colleagues at the Knowledge Media institute (KMi), and I am grateful to them for it.

An infrastructure of this size also needs to apply for funding in its own name and build financial reserves. Independence gives us that, and it is why this change was worth making.

What does not change is our [mission and our values](#). CORE exists to index open research worldwide and make it discoverable and reusable, giving repositories the visibility and reuse their content deserves. Income from membership and data licences will be reinvested in the service. We have recently amended the Articles to reflect that, formalising CORE as a not-for-profit organisation. The Board of Supporters and the Advisory Board continue in their existing form.

Independence also lets us progress further on commitments we have held for some time. We have been committed to the [Principles of Open Scholarly Infrastructure](#) since our first self-audit in 2022, and we can now meet more of what POSI asks, including holding financial reserves dedicated to the long-term sustainability of the service and setting out arrangements for the continuity of CORE's infrastructure. We will publish an updated self-assessment shortly.

In this issue, we specifically bring you FreshFinds and FAIR Certification. FreshFinds has spent the summer being built and tested with repository teams, and it will become available to [Sustaining Members](#) in November. I am grateful to the teams at Oxford, Cambridge, Glasgow and Buckinghamshire New University for the time they have given to their participation in our pilot. Repository teams shaping a service while it is still being written is how CORE has always built, and it remains the part of this work I enjoy most. [FAIR Certification](#) is going through the same process with pilot repositories now, and we expect to award the first certification before the end of the year.

There is also the question of how AI can be used to facilitate research workflows and how can this be done responsibly. [EOSC Assistant](#) brings together twelve organisations across seven countries, coordinated by TIB – Leibniz Information Centre for Science and Technology, to build research assistants for the European Open Science Cloud (EOSC). Alongside partners working in biodiversity, energy systems and materials science, CORE contributes the scholarly record and helps develop a research agent to support the creation of literature reviews. The project is built on the principle that these tools should reference and acknowledge the research they rely on, providing a complete provenance trail; with researchers, not the system, keeping control of scientific judgement.

Before the year ends, we would like to hear from you. Our [community survey](#) is included with this newsletter. It takes around ten minutes and asks how you use CORE, what is working, and where we should put our effort in the year ahead. What members tell us shapes what we build next, we would be grateful if you could complete our survey. I hope to see many of you at the Board of Supporters meeting on the 30th of October 2026, where there will be plenty to discuss. Thank you for building CORE with us.



Petr Knoth

Founder and Head of CORE

2. CORE as an independent legal entity

CORE has been spun-out from The Open University as a dedicated independent legal entity for the development and sustainability of the CORE digital service. The process was planned in a way to ensure that repositories and institutions experience continuity. Here is what you need to know.

Your agreements

If you hold an agreement with CORE, you will have received a novation notice and a letter from us confirming that we accept it. The terms of your agreement are unchanged.

Our services

All CORE services, including the CORE Dashboard, API, Dataset and FastSync, continue without interruption.

Questions? Contact us theteams@core.science.

Your contacts

CORE team email addresses now follow the format name.surname@core.science.

Our governance and commitments

CORE operates as a not-for-profit organisation, and income from membership and data licences is reinvested in the service. The Board of Supporters and Advisory Board continue in their roles. We will publish an updated POSI self-assessment shortly.

3. Services in pilot

3.1 FreshFinds: from pilot to Sustaining Members

When Jisc's Publications Router was switched off on 1 July, many repositories lost a feed they had relied on for nearly a decade. In May we set out how CORE would support deposit workflows after Router, followed it with a webinar alongside Jisc, and committed to building [FreshFinds](#) in the open.

The pilot began on 3 August with **Oxford, Cambridge, Glasgow** and more recently joined **Buckinghamshire New University**, a deliberate mix of large research-intensive repositories and a smaller team. Since then, the parts that matter most to repository workflows have been built and we worked first on ensuring a smooth transition from Jisc Publication router ensuring the same deposit workflow would work with Freshfind. We are now working on establishing a continuous deposit pipeline exploring different models that suits best the different realities

- Discovery that reaches beyond publisher notifications, building on CORE's indexing of preprint servers, Crossref and more than 16,000 data providers.
- Institutional matching that starts with identifiers such as ROR IDs, then goes further by reading affiliations in the full text of the paper.
- Duplicate detection that recognises a paper your repository already holds, even when its title or version differs.
- SWORD deposit, the same standard route Publications Router used, so FreshFinds fits into existing workflows.
- Export in RIS and BibTeX.
- We have been open about what is unfinished too. Matching needs tuning against real institutional output, and the pilot is where that tuning happens.



FreshFinds will be available to CORE **Sustaining Members** in the next few months. Contact us and the CORE team will set up access for FreshFinds in your dashboard.

membership@core.science

3.2 FAIR Certification: what the pilots are teaching us

FAIR Certification helps repositories evidence how well they align with the FAIR principles. The module assesses repositories against objective criteria for findability, accessibility, interoperability and reusability, and gives practical recommendations for improvement. Certification is offered at Bronze, Silver and Gold levels and is supported through the CORE Dashboard.

In April we committed to piloting the module with repositories. Since then, **Oxford and Glasgow** have been working through the assessment with us, and their feedback is shaping the criteria, the report and the Dashboard view.



FAIR	Seamless Dashboard Integration	We have added a dedicated "FAIR Certification" tab to the CORE Dashboard, where repositories can manage their applications and view results in real time.
	Hybrid data-driven assessment	The certification combines data-driven indicators (DD) (~80%) - which rapidly evaluate metadata richness, persistent identifier adoption, and full-text accessibility - with a qualitative review (~20%) covering open questions (OQ) on deposit management, take-down protocols, and long-term digital preservation.
	Review report and certification levels	Submissions are reviewed by the CORE team within 3 to 5 weeks, preparing Review report with specifying which criterias were met for Bronze, Silver, or Gold status to reflect FAIR maturity level. We also start a conversation with the repository to make sure all the aspects are covered and that we work together on a path to improvement.
	Feedback integration and continuous calibration	Insights gathered during pilot reviews are being used to refine the documentation and ensure the certification framework provides actionable guidance for repository managers.

We aim to award the first certification before the end of the year. At launch, certification will be available to Sustaining Members at no additional cost, opening to other repositories after that. Certification is a process rather than a single check, and repositories should expect it to take from several months up to 1 year. More information about FAIR Certification you can find [here](#).

R(2.1) Do you provide access to full texts of research papers for your content?



Yes

Number of papers with access to full text:

123.61K

Percentage of papers with access to full text:

65%

Example of requirement from FAIR Certification report.

4. Repository tools and integrations

4.1 CORE in your repository platform: DSpace and DSpace-CRIS

Earlier this year we began working with 4Science to bring CORE functionality directly into DSpace and DSpace-CRIS, and presented the work at Open Repositories 2026. The first phase focuses on four integrations: importing metadata from CORE into new submissions through the CORE Search API; adding OAI identifiers during submission to prefill metadata; enriching records with SDG metadata identified by CORE, with repository administrators deciding what is added; and showing the CORE Recommender on item pages.

What We Have Done: Implemented Features (pilot)

We have successfully developed three core integration use cases linking CORE and DSpace/DSpace-CRIS platforms:



CORE Recommender Integration	Streamlined Deposit Workflow	COAR Notify Integration for SDG Classification
Displays recommended open-access publications directly on DSpace item display pages, expanding research discovery directly within institutional repositories.	Integrated CORE as a lookup source during the deposit process, enabling researchers and repository managers to search CORE's global index to automatically pre-fill metadata fields (including author names, titles, publication dates, and DOIs).	Implemented metadata enrichment using the COAR Notify protocol. CORE automatically detects Sustainable Development Goal (SDG) classifications for repository items and sends enrichment notifications back to DSpace. A confidence score mechanism routes high-confidence matches for automatic metadata updates and medium-confidence matches to a review queue for staff approval (<i>planned release - the end of the year</i>).

Towards the integration of CORE into DSpace (OR2026 Video)



Watch video



University of
Nottingham

Upcoming Pilot: University of Nottingham

To validate these integrations in a real-world repository environment, CORE and 4Science are launching a pilot with the University of Nottingham.

This pilot will evaluate the newly implemented features across active institutional workflows. Insights gathered during this phase will help refine the functionality prior to a broader release across all 4Science repositories and CORE Member institutions in an upcoming software update

5. Projects and partnerships

6.1 EOSC Alssistant is under way

EOSC Alssistant brings together twelve organisations across seven countries to build an agentic research assistant for the European Open Science Cloud. The project is coordinated by TIB, the Leibniz Information Centre for Science and Technology, who are also CORE Sustaining Members, and runs for 36 months.

The consortium met for the first time in Hannover on 1 and 2 October, hosted with the support of consortium partner Leibniz Universität Hannover. Partners presented their work package plans and set the roadmap for the first phase, which covers FAIR data foundations, the first operational research agents, and the project's governance and communication structures.

From CORE's perspective, the first phase will see us particularly involved in shaping the project's scholarly data foundations, developing the literature-survey agent, and contributing to the benchmarking and evaluation of research agents. We will also help ensure that the agents remain grounded in, and properly attribute, the research literature they use.

CORE contributes the scholarly record and helps develop the agents that reason over it, alongside partners working in biodiversity, energy systems and materials science. The assistant is built on a human-in-the-loop model, which keeps researchers in control of scientific judgement, and on the principle that AI tools should reference and acknowledge the research they rely on. For repositories contributing content to CORE, it means their research is well placed to be found by the next generation of research tools built for EOSC.



TIB – Leibniz Information Centre for Science and Technology

COORDINATOR

DE Germany



UNIVERSITÀ DEGLI STUDI DI CAGLIARI

Università degli Studi di Cagliari (UOC)

PARTNER

IT Italy



Leibniz Universität Hannover (LUH)

PARTNER

DE Germany



Stichting Naturalis Biodiversity Center

PARTNER

NL Netherlands



Technische Universiteit Eindhoven (TUE)

PARTNER

NL Netherlands



National and Kapodistrian University of Athens

PARTNER

GR Greece



Instytut Chemii Bioorganicznej PAN (PSNC)

PARTNER

PL Poland



Institut Mihajlo Pupin (PUPIN)

PARTNER

RS Serbia



CORE: Connecting REpositories

PARTNER

UK United Kingdom



OFFIS e.V.

PARTNER



The University of Manchester (UNIMAN)

PARTNER

UK United Kingdom



LAION gemeinnütziger e.V.

PARTNER

DE Germany

Follow the project: eosc-aissant.eu

EOSC Alssistant is funded by the European Union under Horizon Europe, grant agreement 101292566. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

6. Community and collaboration

7.1 Understanding the CORE Dashboard: the series continues

Understanding the CORE Dashboard

Since March, the educational series has covered the **Indexing Status tab**, managing content with the **Content tab**, and checking open access compliance for REF 2029.

Next up: RIOXX Metadata Validator Tab.

→ **A Closer Look at the Overview Tab**

[Read more](#)

→ **Managing Repository Content with the Content Tab**

[Read more](#)

→ **Understanding your repository's indexing status in CORE**

[Read more](#)


→ **Checking Open Access Compliance for REF2029**

[Read more](#)

→ **RIOXX Metadata Validator Tab**

Coming soon

7. Events: webinars and conferences



Warsaw

12 and 13 May

The Big Migration

CORE took part in the DSpace community's conference at the University of Warsaw Library, which drew more than 120 participants from across Central and Eastern Europe. Petr Knoth presented Connecting Repositories, an overview of CORE as open scholarly infrastructure, and he and Catherine Kuliavets ran a hands-on session, Unlock the potential of your repository with the CORE Dashboard, for repository managers and librarians.

[FIND OUT MORE](#)



May

CORE FreshFinds webinar with Jisc

Following the retirement of Publications Router, CORE and Jisc hosted a webinar on the development of FreshFinds as a solution for populating repositories. The recording is available for anyone who missed it.

[FIND OUT MORE](#)



COAR Annual Conference

Braga

1 to 3 June

The COAR Annual Conference took place at the University of Minho in Braga, bringing together 50 participants from around the world under the theme of maintaining openness and collaboration in times of uncertainty. Matteo Cancellieri represented CORE across three sessions. He introduced CORE's new services and developments to the network services strand, presented the CORE and DSpace integration work alongside Giuseppe Digilio of 4Science during the COAR Notify workshop, and joined Lautaro Matas of LA Referencia to set out emerging solutions for dealing with aggressive bots. That last topic ran through the whole meeting, with COAR announcing plans to work with industry partners on more frictionless interactions between repositories and crawlers, and presenting a new pilot project idea to better support machine crawling of repositories.

[FIND OUT MORE](#)



Singapore

2 to 5 June

FORCE2026

Petr Knoth presented Managing AI Bot Access to Open Scholarly Infrastructures, drawing on CORE's own harvesting data to show how widespread the problem has become. Of 9,129 data providers CORE attempted to harvest, 4,510 returned errors, and once routine protocol faults are set aside, between 35 and 40 per cent of harvesting activity appears to be affected by bot protection rather than by genuine technical failure. Much of it is invisible to the repository: timeouts and dropped connections rather than error messages. The argument he put to the community is that a repository blocking machines indiscriminately cannot be FAIR, because content that cannot be indexed is neither findable nor reusable. CORE's response is practical: new guidance on machine access in the CORE Data Provider's guide, a Repository Indexability Diagnostics tool that tests whether your repository is blocking indexing systems, and participation in COAR's Task Force on Dealing with Bad Bots, where work is under way on a community whitelist of trustworthy agents.

OPEN REPOSITORIES 2026

Online conference | 8-11 June 2026

Online

8 to 11 June

Open Repositories 2026

The 21st Open Repositories conference ran fully online, with two contributions from CORE. Matteo Cancellieri presented Towards the integration of CORE into DSpace, work with Viktoriia Pavlenko and colleagues at 4Science setting out where CORE services can be built directly into DSpace and DSpace-CRIS. David Pride presented Large Language Models for Software Mention Extraction in the AI and software preservation session, work carried out with colleagues at the University of Bologna and Brno University of Technology. Research software is often mentioned only in passing in the text of a paper, which makes it hard to find, credit or preserve. The work shows that large language models can match current state of the art approaches using zero-shot methods, meaning no specialist training is needed first, which brings automatic identification of software mentions within practical reach for repositories. A recording of the software mentions session is available.



The CORE team

Poznań

June 2026

In June the CORE team met in Poznań for our annual meeting, the one week in the year when a team spread across several countries works in the same room. The agenda was shaped by what the community has been telling us: where the Dashboard needs to be clearer, what repositories need after the retirement of Publications Router, and which pieces of work deserve our effort over the next year. Not everyone could travel, and colleagues joining online were part of every session and every decision. The service you use is built by this group, and the plans in this newsletter were largely made that week.



Online

14 September

Webinar for the Open Access Australasia community

CORE joined Open Access Australasia for a webinar introducing CORE and its services to repository teams across Australia and Aotearoa New Zealand, covering compliance monitoring, metadata validation, and the tools that support interoperability and discoverability. The session was aimed at anyone whose work touches repository content, not only those who manage repositories, and much of the time went to questions about what would be most useful to institutions in the region.

[FIND OUT MORE](#)