

Licensing research IP

From open source to the IP long tail - and the tools in between

● EOSA Summer School | 7 September 2026

Ed Maughfling & José Fernández | E-lucid Solutions Ltd.

Before we start: two questions

1

Have you produced code, a dataset or a trained model in the last 12 months?

2

Do you know who owns it?

The gap between those two answers is what today is about.

What we'll cover

2 hours, one break, questions welcome throughout

The licensing spectrum

**Who owns your code, data
and model?**

The tooling landscape

e-lucid

**Break
15:30**

The academic perspective

Live demo

Lessons learned


Questions

01

The licensing spectrum

Open source is one end of a line, not a separate world

Every research output sits somewhere on this line



Open-source licensing governs the left end.

Exclusive deals need lawyers.

The middle ground - where most research outputs lands - has had very few tools. That is the gap we will keep coming back to.

The IP long tail

A few big deals get all the attention. Everything else goes unlicensed.

The long tail: high volume, low value, predictable terms - and mostly not licensed at all



Assets ranked by deal value →

- Software, datasets, questionnaires, life science research tools, teaching materials.
- Standard, terms. Fixed or tiered price - often zero.
- Each one is too small to justify a lawyer's afternoon. So, it 'sits on a shelf', gets sent by email or via download button with no record.
- Nobody knows who used it, where, or why.
- Institutional IP owners & funders increasingly want to know!

Two kinds of licensing work. One gets all the resources.

The head: complex deals

- Exclusive licences, spin-outs, major partnerships
- Negotiated terms, due diligence, months of work
- Needs lawyers and experienced deal-makers
- Rightly gets the team's attention - cannot be automated

The tail: routine transactions

- Non-exclusive licences for software, data, materials, questionnaires
- Standard terms, fixed price (often €0), simple checks
- Managed by email / download buttons. Or not at all.
- But it could be hundreds a year - if the friction is removed

02

Who owns your code, data and model?

Who knows!?

It depends on where you are, who you are & who paid

General patterns - your institution's IP policy is what actually applies



Employee (staff, postdoc)

Work made in the course of employment usually belongs to the employer - the institution. You are the author, not the owner.



Student

Students are not employees. Many institutions let students keep their IP - unless a contract, a studentship or a funded project says otherwise.



Visiting or collaborating researcher

Two institutions, two policies. Joint ownership is common and messy. Agree it in writing before any IP is created.



Funded project

A grant agreement can override everything above: open access mandates, background/foreground IP rules, consortium agreements.

Funders have opinions

Open science mandates shape what you can do with the output

Horizon Europe

- Open access to **publications** is mandatory
- **Research data**: 'as open as possible, as closed as necessary'
- **FAIR** (Findable, Accessible, Interoperable and Reusable) data management plans
- Commercial **exploitation** is allowed — and expected — under the grant's IP rules

National funders (UKRI, AEI...)


- Open access policies for **publications**
- Increasing expectations on **software and data sharing**
- Expectation of commercial **exploitation** when reasonable
- **Impact** reporting: who used it, where, for what

What this means for you

- Open by default does not mean 'no licence' - it means choose an open one deliberately
- 'Closed where necessary' is a licence decision too
- 'Good' Impact reporting needs a record of every download or use. Emailing or anonymous downloads does not give you one.

The dependency trap

'We'll decide licensing later' IS a decision



- Permissive licences (MIT, BSD, Apache-2.0) let you re-license your combined work. Copyleft (GPL, AGPL) does not.
- Check before you import. A licence scanner takes minutes; a rewrite takes months.

Four things people forget until it is too late

Think about these from the moment you push.

Patents inside OSS

Releasing code under MIT does not license your patent. Apache-2.0 does. Decide which you mean.

Trademarks

The licence covers the code, not the name. Firefox is open source; the name is not.

Export control

Encryption, dual-use algorithms, HPC codes. 'Open' does not exempt you - publication rules do, sometimes.

Data & models

CC-BY, CC0, ODbL for data. OpenRAIL and bespoke terms for models. OSS licences fit weights badly.

03

The tooling landscape

What exists, what each tool is for, and where they all stop

The open end: tools for releasing

They tell you how to publish. None of them help you license the middle.

choosealicense.com

Pick a licence in five minutes. Plain-English summaries of permissions and conditions.

SPDX + REUSE

Machine-readable licence identifiers in every file. REUSE (FSFE) makes a repo's licensing unambiguous.

Scanners

ScanCode, FOSSology: find every licence in your dependency tree before someone else does.

CLA / DCO

Contributor License Agreement and Developer Certificate of Origin (signed-off-by). They keep relicensing possible. Without them, dual licensing is off the table.

All of these work on the left end of the spectrum. Now the middle.

IP management databases

Wellspring Sophia / Evolve · Inteum Minuet

What they do

- The tech transfer office's (TTOs) system of record
- Invention disclosures, patent families, deadlines, inventors
- Agreement records, royalty tracking, internal workflows
- Reporting to the institution and to funders

What they don't do

- Face the buyer. They are internal tools - a licensee never logs in
- Publish a catalogue, a price or the terms
- Let anyone self-serve: every licence still starts with an email to a person
- Take a payment or countersign anything

Generic e-commerce platforms

Shopify · WooCommerce · Etc

What they do

- Buyer-facing, self-service, 24/7
- Catalogue, cart, card payment, order history
- Cheap to start, huge ecosystem of plugins
- Several TTOs have tried them. e.g. UniQuest licensed the Geriatric Anxiety Inventory (clinical questionnaire) on BigCommerce for years

What they don't do

- Know what a licence is. You are selling a contract, not a t-shirt
- Ask qualifying questions and route to approval based on the answers
- Handle free-of-charge and paid transactions in the same flow
- Generate and countersign an agreement, or report who licensed what and why

E-signature software

DocuSign · Adobe Sign

What they do

- Execute an agreement legally and quickly, anywhere
- Templates, signing order, audit trail
- Integrate with almost everything
- Genuinely solved the 'print, sign, scan' problem

What they don't do

- Anything before the signature: discovery, pricing, qualification, approval
- Anything after it: delivery of the asset, renewal, impact reporting
- Payment
- Scale without a human preparing every contract

Discovery and matching platforms

In-Part · Lucidhub · Wellspring Flintbox · Inteum Techpublisher

What they do

- Buyer-facing and IP-specialised — the right corner of the map
- Help industry find research it would never have searched for
- In-Part: matches company R&D needs with university technologies
- Lucidhub: aggregates institutional catalogues; AI search across them (EU / UKRI co-funded)

What they don't do

- Transact. Discovery ends with an introduction, not a licence
- Set price or terms
- Take payment or deliver anything
- Lucidhub is the exception by design: it hands the buyer to an e-lucid store to complete the licence


04

e-lucid

What it is, where it came from, and what it is not

Where each tool sits

Position is about role in the workflow, not quality




Ten steps in a licence. Who covers which?

Simplified – up for discussion in the Q&A

Step	IP database	E-commerce	E-signature	Discovery	e-lucid
List the asset publicly	—	✓	—	✓	✓
Publish terms and price	—	✓	—	—	✓
Qualify the buyer (who, what for)	—	—	—	—	✓
Route for approval	✓	—	—	—	✓
Take payment — card, invoice or £0	—	✓	—	—	✓
Generate the agreement	—	—	✓	—	✓
Countersign	—	—	✓	—	✓
Deliver the asset and the agreement	—	✓	—	—	✓
Record the licence in one place	✓	—	—	—	✓
Report impact: who, where, why	✓	—	—	—	✓

How a university ended up building a licensing shop



What it looks like at scale

2025 figures across e-lucid stores

11,500+

licences processed per year
(2025)

€1M+

licensing revenue
(2025)

35+

institutions across 6 countries
(UK, US, FR, DK, NO, AUS)

€0

price of most* licences issued
*approx. 90%

Two kinds of store, both valuable.

Some run thousands of free academic licences a year and produce impact evidence.
Others make €200,000+ a year from fewer than 15 products.

A typical UK store: ~30 products, ~600 licences a year.

What e-lucid is not

Not a contract lifecycle tool

It does not negotiate or redline. If the terms change per deal, it is the wrong tool. Standard terms are the point.

Not an open-source compliance tool

It will not scan your dependencies or tell you which licence to pick. Use the tools from part 3 for that.

Not for exclusive deals

Spin-outs and exclusive licences need people. e-lucid frees those people up; it does not replace them.

Not a replacement for the TTO

Ownership, policy and judgement stay with the institution. The platform runs the process they define.

What we are trying to change

1

Zero-friction licensing

A licence should take the buyer five minutes and the institution zero. Every email removed is a licence that actually happens.

2

Measure and monetise the tail

Every use recorded: who, where, what for. Free licences produce impact evidence. Paid ones fund the rest.

3

Discovery across institutions

A buyer should not need to know which university built the product they need. Lucidhub searches across all of them.

Break

Back in 5 minutes.

Next: Conversation with tech transfer expert, Dr Weng Sie Wong.

05

The academic perspective

How does e-lucid work for academics? Q&A with Dr Weng Sie Wong from UCL Ventures.



Dr Weng Sie Wong

Senior Business Manager, UCL Ventures

Commercialises medical technology, healthcare engineering, medical imaging and AI-enabled healthcare innovations from UCL research — forming multiple spinouts and serving on their boards.

Previously a senior technology transfer manager at Oxford University Innovation for 10 years. MBA, Edinburgh; PhD in pharmacology, UCL.

06

Live demo

How does it look like + live exercise

Live demo

Switching to a real store — this slide is here to come back to



Watch for: the buyer never emails anyone · the free and paid paths are the same checkout · the qualifying answers drive the approval rule · the agreement is generated and countersigned without a human · the admin sees who, where and why.

Exercise: license this

Eight minutes. Answer in the chat.

The asset

An HPC numerical library built in a Horizon Europe project. Three groups contributed. It imports two GPL dependencies. A commercial CFD vendor has asked to embed it. Two universities want it for teaching.

- 1 Who owns it - and can you even decide?
- 2 Which licence, or which two?
- 3 What pricing model for the vendor? For the universities?
- 4 Which qualifying questions? Who approves?

07

Lessons learned

Fifteen years of watching what licences perform best

Key lessons

Make the best of your dual licensing strategy

1 Standard terms beat bespoke

Negotiation kills the economic rationale of any low value licence. If a buyer wants different terms, the answer is a different product or licence - not an email thread.

2 Free licences create the impact data

UniQuest's Geriatric Anxiety Inventory: 216 licences, 43 countries, 23+ languages in 17 months. 93% free. That is a funder report nobody could have produced from disseminating assets via email.

3 Price the tail low and don't negotiate

£250 that closes today beats £5,000 that never does. The highest-revenue stores make £200k+ a year from under 15 products - at fixed prices.

Impact and income go hand in hand

Users interact with a shopping portal, not with a tech transfer office.

Impact

- Free-of-charge licences to academics, students, clinicians, non-profits
- Every one recorded: who, where, what for
- The report funders ask for, generated instead of reconstructed
- Ventura's CPAP breathing aid: licensed 'in volumes not seen before' while keeping the conditions enforced

Income

- Fee-bearing commercial licences on the same product
- Different price for academic, commercial and government buyers
- Fixed, variable or price-on-application (POA)
- Pay by Card or Invoice
- Pays for the platform - and frees the team for the deals that need input

Lucidhub: search IP across institutions

Co-funded by the European Union and UKRI under Horizon Europe

What it is

- One search across the catalogues of many research institutions
- AI identifies & presents relevant research outputs - software, data, materials, tools
- For buyers: one place to look instead of forty university websites
- For institutions: reach beyond your own audience with no extra work

How it connects

- Institutions already on e-lucid are listed automatically
- A buyer finds the asset on Lucidhub and completes the licence on the institution's own store
- Discovery & transaction in one path - the hollow dot and the filled dot from the map
- lucidhub.ai

How to get involved

1

Find your IP policy

Know who owns what you built before anyone asks. Then check your dependencies with a scanner.

2

Talk to your tech transfer office

Bring the asset and this line: 'I want it licensed with standard terms and a record of every use.' Point them to lucidhub.ai and e-lucid.com.

3

Pilot a store or list on Lucidhub

A store starts at ten products. Lucidhub listing extends the reach. We can walk a group through it in an hour.

e-lucid.com/demo · lucidhub.ai · [jose@e-lucid](mailto:jose@e-lucid.com) & ed@e-lucid.com

Questions

The tool map, the ten steps, the dependency trap - argue with any of it.

e-lucid.com · lucidhub.ai · jose@e-lucid.com & ed@e-lucid.com