

DocScope

RAG over Papers — Ask Your Documents Anything

Grounded answers · cited sources · human-reviewed

Domain packs: Legal · Clinical · Financial · Regulatory

ERC2 RAG Challenge Winner · Livedata Limited — Female-led

August 2026

The Problem: Knowledge Buried in Documents

Thousands of Pages, Limited Time

- Case law, guidelines, contracts and reports span tens of thousands of pages
- Manual reading is slow, inconsistent and error-prone
- Key answers are buried in passages nobody can recall
- Each query becomes a new search across multiple systems

What this means for you: Valuable knowledge sits unused because finding the right passage takes too long.

The Cost of Slow Answers

- Hours of research for every review, decision or response
- Risk of citing out-of-date or wrong sources
- Knowledge locked in individual officers' heads
- No audit trail of where an answer came from

In practice: Teams need to ask questions in plain English and get answers that cite the exact source — fast, consistent and auditable.

What DocScope Does

Ask in Plain English

- Ingest a corpus: papers, case law, guidelines, reports
- Ask questions in natural language
- Every answer cites its exact source passages
- Domain-tuned: Legal, Clinical, Financial, Regulatory

What this means for you: Your team asks questions in their own words — and sees exactly which document, page and paragraph the answer came from.

Trustworthy by Design

- Human-in-the-loop — officer reviews before anything is issued
- Full audit logging of every query and draft
- Source citation on every answer
- Document isolation between corpora

In practice: Answers are safe for professional use because every claim points back to evidence your team can verify.

Screen: Ask Your Documents

DocScope | RAG over Papers

ERC2 RAG Challenge Winner

Ask your documents anything

Grounded answers with citations · domain packs: Legal, Clinical, Financial, Regulatory

Corpora

Police Legal Services — case law Legal

48,214 documents · 1.2M chunks · pgvector

Clinical guidelines — evidence review Clinical

12,369 documents · Qdrant

Company annual reports — due diligence Financial

ERC2 challenge corpus · FAISS + BM25 hybrid

Ask a question over the corpus...

Ask

Answer with evidence

Q: What grounds for judicial review apply to a planning consent granted without a flood risk assessment?

Under the Flood & Water Management Act 2010 (s.23), an LLFA must be consulted before a planning decision affecting flood risk. Failure to consult, or to have regard to the LLFA response, may render the consent **ultra vires** and subject to judicial review on grounds of procedural impropriety and failure to take into account a relevant consideration. ^{[1][2]}

[1] Flood & Water Management Act 2010, s.23

... the authority must be consulted ... (p. 14)

[2] R (on the application of X) v Planning Authority [2022] EWHC 1159

... failure to have regard to flood risk ... (para 45)

[3] LLFA consultation guidance (2024)

... procedural steps for consultation ... (p. 7)

Corpus-first interface:


- Select a corpus (Legal, Clinical, Financial, Regulatory)
- Ask a question in plain English
- Grounded answer with inline citation markers

Evidence panel:

- Each citation opens the exact source passage
- Statutes, case law and guidance shown with page/para refs
- Highlighted passages make verification instant

What this means for you: From question to verified source in seconds — with the audit trail to prove it.

Screen: Ingestion & Retrieval Pipeline



Proven RAG pipeline:

- Ingest (RAGCrawler) → chunk → embed → store → retrieve → answer
- Hybrid FAISS + BM25 retrieval with ranked results
- Vector stores: pgvector, Qdrant, FAISS as required

Corpus health visible:

- Documents, chunks and embeddings tracked
- Daily incremental indexing keeps answers current

Guardrails:

- Human-in-the-loop · audit logging · citation on every answer · per-corpus isolation

Who DocScope Is For

Priority	Segment	What they use it for
1	Police legal services & MOJ	Case and document search with semantic vector retrieval
2	Health evidence & clinical	Evidence review and clinical guidelines (NICE, NHS trusts, MHRA)
3	Law firms & in-house legal	Case law, contracts and due diligence Q&A
4	Research & academic	Evidence synthesis across publications and reports
5	Financial & corporate	Annual reports and due diligence (ERC2 domain)

Why It Works — Proof Points

Claim	Evidence
Independently validated algorithm	Winner of the ERC2 RAG Challenge — Q&A over company annual reports
Built for the public sector	Legal corpus search for police legal services (legal_cases, pgvector)
Production RAG pipelines	Qdrant ingestion (RAGCrawler), hybrid FAISS + BM25 retrieval
Live NHS experience	NHS DSPT (Standards Met) — data protection capability
Compliance-ready	JOSCAR-accredited, ICO-registered

The algorithm won a challenge; the deployment is designed for the public sector’s trust requirements.

Commercial Model

Tier	What's included	Indicative
Pilot	First corpus ingested free	£0
Per-seat SaaS	Web interface, per user	£99–£249/mo
Enterprise	Per-corpus, domain packs	By corpus size
White-label	Embedded in partner platforms	Annual licence
Bare server	Deployed on your own hardware — all data saved locally	From £5,000

All pricing in GBP. Pricing quoted only when asked in formal engagement. First corpus ingested free to prove value. With the bare-server option, documents and indexes never leave your infrastructure.

How We Deliver

Pilot First

- First corpus ingested free (your documents, up to a defined size)
- Live demo: questions over your material with citations
- Showcase ERC2 win, hybrid retrieval, domain packs
- Set up within days, not months

What this means for you: You see DocScope answer questions from your own documents before any commitment.

Then Scale

- Choose per-seat SaaS, enterprise per-corpus, or **bare-server** (your hardware)
- With bare-server, documents and indexes stay entirely on your infrastructure
- Add domain packs as your needs grow
- Same evidenced RAG capability Livedata already runs
- Human-in-the-loop and audit logging throughout

In practice: From first ingestion to daily use in days — with answers your team can verify, cite and defend.

Next Steps

1. **15-minute call** — demo on a sample of your documents
2. **Pilot** — first corpus ingested free
3. **Review** — accuracy, citations, and speed on your material
4. **Choose a tier** — per-seat SaaS or enterprise per-corpus

Thank You

Livedata Limited

Female-led | UK Public Sector | SaaS Solutions

William Turbayevsky — CTO · admin@livedata.link · 07479 189835 · livedata.link