

AI KNOWLEDGE AGENT

400 papers in.

The questions you couldn't ask out.

*Researcher-built — gap analysis, hidden-gem surfacing,  
citation lineage mapping. Private corpus. UK-hosted.*

## THE PROBLEM

# Literature reviews are buried.

A PhD reading list is hundreds of papers deep. The thesis-shaped questions — the ones the field raised but never answered — are buried with them.

**3 weeks**

to skim 40 papers properly

**Months**

to spot a research gap by hand

**Hours**

per chapter mapping citations

*"The research assistant that doesn't get tired at paper 80, doesn't lose its place at paper 200, and doesn't forget what it read last Tuesday."*

# How it works.

*Six steps from first contact to a calibrated partner.*

**1**

## Upload

Drag-and-drop your reading list.

**2**

## Index

Every paper, every section, every citation.

**3**

## Ask

Natural-language Q&A grounded in your corpus.

**4**

## Cite

Paper, page, and passage with every answer.

**5**

## Map

Cross-paper citation lineage on demand.

**6**

## Surface

Hidden gems and unanswered questions.

---

## Built for researchers

### Azure UK South

University and NHS data handling.

### Private corpus

Not pooled. Not used to train public models.

### Citation-grounded

Every answer pinpoints paper / page / passage.

### GDPR Article 28

DPA template + sub-processor list public.

# Concrete savings.


Activity	Manual / typical tool	Knowledge Agent
Search 400+ papers for an argument	Days of skimming	Seconds
Identify gaps in literature	Months of cross-reading	Hours
Map citation lineage	Hours per chapter	Minutes
Find hidden questions	Often missed	Surfaced systematically
Onboarding new RA	Weeks	One-time corpus upload
Re-checking thesis defence facts	Re-reading every paper	Targeted queries

## Weeks → Hours

### literature review

*Cross-paper reference mapping that takes days runs in minutes.*

#### COMPARATIVE



From weeks to hours per literature review.

*Researcher-built differentiation vs. generic Q&A tools — gap analysis, hidden-gem surfacing, citation lineage, thesis drafting.*

BUILT FOR TRUST

# It's NotebookLM — built for researchers.

## DATA RESIDENCY

Azure UK South. UK university and NHS handling.

## PRIVATE CORPUS

Your reading list isn't pooled with anyone else's, ever.

## NO RETRAINING

Your corpus does not feed public foundation models.

## GDPR ARTICLE 28

Standard DPA template. Sub-processor list public.

## CITATION-GROUNDED

Every answer pinpoints paper / page / passage.

## WHY EXCOMS

# It's NotebookLM — built for researchers.

*Gap analysis, hidden-gem surfacing, cross-paper citation mapping, thesis drafting. Plus UK data residency and an institutional procurement path.*

NEXT STEP

# Ready to read 400 papers in seconds?

*We're working with UK university research offices and NHS knowledge teams. Talk to us about scoping a pilot for your institution.*

**ExComS.ai**

STUDIO · LUTON · UK

## CONTACT

**Request a demo**

[coding@excoms.com](mailto:coding@excoms.com)

**Talk to us about institutional pilots**

[coding@excoms.com](mailto:coding@excoms.com)

**Web**

[excoms.ai/products/knowledge-agent](https://excoms.ai/products/knowledge-agent)