



Originally Published by EDRM and JD Supra

Editor's Note

This is the fourth article in a series that began with Hybrid Search and now turns to AI-Integrated Discovery. The full series [can be found here](#).

By: John Tredennick and Dr. William Webber

For decades, discovery and investigation teams faced the same empty search box. They conjured up keywords, nested terms in parentheses, and added wildcards and proximity operators. Even when they mastered Professor Boole's nineteenth-century logic, the reward was thousands of search hits. In every case, someone still had to painstakingly read the documents, decide what mattered, and determine where to go next.

In [***Keyword Search Isn't Dead. Searching With Keywords Alone Should Be***](#), we examined that starting point. Keyword search is indispensable when exact language matters, but it cannot find words the team does not know to use. Even a carefully built query may miss a



document that calls a problem by an unexpected name, while a broad query leaves people to sort through results that add little to the investigation.

In [*AI Search Can Find the Meaning. But It Can Still Miss the Evidence*](#), we added semantic search, which can find the underlying idea when documents use different words. Our Project Checkmate example showed the converse: an opaque project name may carry the crucial signal even though its ordinary meaning tells semantic search almost nothing. Combining keyword and semantic methods gave the team complementary ways to find the initial evidence.

In [*Search Finds the Initial Evidence. CAL Learns What Comes Next*](#), we explored what an AI-powered system should do after hybrid search returned its first 50 results. An LLM assessed those documents and identified the clearest positive and negative relevance judgments. We used those examples to train a continuous active learning system (CAL), which helped rank the results presented next. Across the next 100 results, that approach found 82 more responsive passages than continuing down the original search ranking. AI could learn from what it found and move the search forward.

That led us to a larger question:

If AI can learn from the documents and improve search results, why stop there? Why not let AI help us investigate what happened?

That is the idea behind **AI-Integrated Discovery**. AI becomes a member of the team from the first question through final resolution. It helps frame the inquiry, build searches, read the results, pursue leads, test competing explanations, and assemble an account based on the documents found. The human members of the team direct that work, examine critical evidence, challenge conclusions, and decide what the record supports.

Here is how that team could work in practice. Using Alchemy Intelligence, our document intelligence platform, we follow one investigation from the first question through a refined assignment, an initial report, a challenge to its conclusions, and a plan for what comes next. The example shows how quickly an AI teammate can help turn an initial question into an account the human team can inspect, test, and build upon.

1. Start with what you want to know

Imagine a team examining documents connected to the UK Post Office's Horizon accounting system. The team wants to investigate Fujitsu's knowledge of problems with Horizon and its communications with the Post Office. At the outset, the lawyer may not know the internal terminology, the key custodians, or the most important time periods. The question is nevertheless clear enough to begin:

What did Fujitsu know about problems with Horizon, when did it know, and what did it do in response?

That is an assignment, not a search string. It contains several issues: what counted as a problem, who learned of it, whether it could affect branch accounts, what was done internally, what was communicated externally, and what evidence points the other way. A system that immediately converts the question into a keyword or semantic query has missed its first opportunity to help.

Traditional search begins with a query. AI-Integrated Discovery begins with intent.

Our AI Prompt Helper can explore the collection before the team commits to a full search strategy. In this example, it reviewed an initial sample of roughly 30 documents in our Alchemy system. The following Prompt Suggestions show the issues it surfaced and the questions it put to the team.

Prompt Suggestions

I have reviewed a sample of 30 documents relating to your topic. Based on this review, there is a strong basis to investigate Fujitsu's knowledge through its incident-management systems, technical investigations, communications with Post Office, and responses to identified defects.

Key angles surfaced by the sample

The Horizon Post Office scandal

Horizon was the accounting system used in UK Post Office branches and operated by Fujitsu. Faults in the system could produce apparent shortfalls in branch accounts.

The Post Office held branch operators responsible for many shortfalls; some lost their livelihoods, and others were prosecuted and convicted.

Subsequent litigation and appeals exposed serious problems with Horizon's reliability and with the investigation and disclosure of those problems.

The scandal raises a question central to this article's example: what did Fujitsu know about the faults, and what did it do with that knowledge?

- **Early and recurring fault knowledge:** Formal problem-management records existed from at least 2001, with known-error logs, root-cause investigations, and reports of branch cash mismatches and performance problems.
- **How knowledge moved within Fujitsu:** Support, reconciliation, development, testing, and release teams had distinct roles. Evidence suggests recurring issues were logged and escalated, but information was not always shared effectively across teams or with those supporting legal cases.
- **Specific Horizon Online issues:** Potential focal incidents include Oracle database bugs identified in March 2010; the Receipts/Payments Mismatch affecting about 40 branches; the Suspense Account bug affecting 12 branches; and incidents involving remote intervention or balancing transactions.
- **Response and disclosure:** Fujitsu investigated, developed workarounds and fixes, and sometimes provided written responses or material for the Post Office. The sample also raises questions about delay, incomplete identification of affected branches, audit trails, and the adequacy of assurances about system integrity.
- **Assurances and litigation support:** A particularly important line of inquiry is what Fujitsu told the Post Office about data integrity, remote access, auditability, and known defects—especially where information may have been relevant to investigations, prosecutions, or litigation.

Is your primary focus on (a) all Horizon history, (b) the Horizon Online migration period around 2009–13, or (c) later litigation and disclosure issues?

Do you want the investigation to concentrate on technical fault handling—logs, investigations, fixes and releases—or on what Fujitsu communicated to the Post Office and legal teams?

Should the output assess potential shortcomings, such as delayed escalation, incomplete branch identification, or inadequate assurances, or remain strictly descriptive?

Give us your thoughts on these questions or otherwise tell us how we can best help in revising and finalizing your prompt.

The lawyer can answer those questions and redirect the inquiry. Here is the initial prompt suggested by the Prompt Helper:

Prepare a comprehensive report on what Fujitsu knew about Horizon faults, defects, data-integrity concerns, reconciliation discrepancies, remote-access capabilities, and audit-trail limitations when it knew each matter; and what it did in response.

Cover the period from early rollout through Horizon Online and later reviews. Identify the relevant Fujitsu teams and personnel, the mechanisms by which issues were recorded or escalated, communications with Post Office, investigations, workarounds, fixes, releases, and any gaps or delays in identifying affected branches or disclosing information.

Distinguish documented facts, witness recollections, allegations, and unresolved issues.

In this way, the first encounter with the documents can help refine and clarify what the team is investigating.

2. Give the assignment more than one search method

Once the inquiry is sharper, AI can break it into investigative questions. Which records show that a potential defect was identified? Who assessed its effect on branch accounts? What did technical staff tell management? What did Fujitsu tell the Post Office? Were there reports that attributed a discrepancy to another cause?

As our first two articles showed, each question may require a different way to look for evidence. Keyword search is strong for names, ticket numbers, exact phrases, product terminology, and dates. Semantic search can find discussions of a problem even where nobody used the lawyer's words. Once documents have been assessed, CAL can learn the characteristics of responsive material and improve what the system presents next. The system can combine the methods and revise the mix as the investigation develops.

Here are two targeted examples of the kinds of queries the LLM might build for our topic:

Semantic search

Documents describing when Fujitsu learned of faults in the Post Office Horizon system, from early rollout through Horizon Online and later reviews, and how its staff investigated, escalated, communicated, and addressed them. Relevant records discuss data integrity, branch account reconciliation discrepancies, remote access to branch data, audit-trail limitations, fixes and workarounds, and delays or gaps in identifying affected branches or informing the Post Office.

Keyword search

Fujitsu ICL Pathway Post Office POL Horizon Horizon Online faults defects bugs data integrity reconciliation discrepancies branch accounts remote access remote alteration audit trail audit logs SSC Software Support Centre PEAK KEL known error log ARQ incident reports emails correspondence investigations escalations workarounds fixes patches releases affected branches disclosure reviews witness statements

Under the hood, AI builds both keyword and semantic searches and runs them independently against the larger population. The two search engines rank documents across the population for relevance and send them to the system in relevance order. Alchemy then merges the results and presents the first 100 most promising documents or sections.

3. Let AI read and report what the first search finds

The first 100 results are a starting point, not the deliverable. AI reads them, assesses how each relates to the assignment, prepares targeted summaries, and draws together an initial report. Here is an excerpt from Alchemy's report on the Horizon inquiry:

Executive Summary: Fujitsu's Knowledge of and Response to Horizon Issues

Fujitsu had established mechanisms for recording and escalating Horizon faults. These included PinICL (replaced by PEAK in 2003), the Known Error Log, change control and the Release Management Forum, supported by teams such as the Service Support Centre, the Management Support Unit and the Fraud and Litigation Support Office ([LONDON_00000331 §2 ¶11-24](#); [LONDON_00000840 §3 ¶23-34](#)). Fujitsu's corporate evidence accepted that 29 identified bugs, errors and defects operated between November

1999 and May 2018, and that some could and did affect branch-account integrity ([LONDON_00000025 §5 ¶46-47](#); [LONDON_00000118 §5 ¶50-51](#)).

Documented examples include the Callendar Square issue, the Craigpark/CABSPProcess locking issue, ARQ 447 duplication, duplicate JSNs and Sysman3 errors during the Horizon Online pilot, and a receipts-and-payments mismatch affecting 62 branches. In each case the summaries show that Fujitsu identified the problem and introduced a fix. They show much less about whether every affected branch was identified, whether fixes were verified across the estate, and whether historic accounts were corrected.

The most serious gap concerns the link between Fujitsu's technical knowledge and the evidence it supplied for prosecutions and litigation. Fujitsu accepted that it could filter, insert or amend ARQ data, and that supplied ARQs did not exactly reflect branch transactions ([LONDON_00000025 §5 ¶53-57](#)). Its 2005 prosecution-support policy did not require a check of the Known Error Log ([LONDON_00000025 §7 ¶66-78](#)). Its corporate witness accepted that defects were generally absent from Fujitsu witness statements and that Fujitsu failed in its duty to provide accurate and complete data ([LONDON_00000025 §5 ¶47-52](#)). A Horizon Online bug was not disclosed in the Misra proceedings because the witness regarded Online as separate from Legacy Horizon ([LONDON_00000138 §6 ¶60-69](#)). Subpostmaster representatives allege that information was delayed, downplayed or misrepresented. These are submissions, not findings ([LONDON_00000039 §2 ¶17-22](#)).

On remote access, the evidence establishes that Fujitsu staff held privileged access and could amend branch data. Witness accounts differ on what was represented to Post Office about auditability. A Fujitsu executive recalled that developers could deploy code to live HNG-X in 2010 and that this persisted into 2011. Ernst & Young raised privileged-access concerns in 2011 ([LONDON_00001080 §7 ¶72-77](#)). In 2024, Fujitsu told police that remote access and known or undetected defects required case-specific investigation ([LONDON_00000141 §21 ¶212-222](#)). Scrutiny remains incomplete. A 2024 review reportedly could not assess privileged-access controls because Fujitsu did not cooperate ([LONDON_00001092 §2 ¶19-24](#)). Much of the remedial account also rests on the recollections of those responsible for the remediation, rather than on contemporaneous records.

The detailed findings, and five recommended areas for further investigation, are set out in the separate memo to the team leader: what "fixed" actually meant; identification of affected branches and cases; the gap between technical knowledge and litigation evidence; ARQ provenance and reliability; and remote-access capability, use and accountability.

See the [full report by clicking here](#).

Within Alchemy, team members can move from a finding to its supporting summaries and documents, test the conclusion, and direct the next inquiry. They need not accept an AI report any more readily than an associate's first memo. But they begin with an organized account of the evidence and specific questions to pursue, rather than a list of search hits.

4. Let the first results open new lines of inquiry.

In Article Three, CAL used clear relevance judgments to bring more responsive documents into the next batch. But AI can also notice when the evidence changes the question. A record saying a defect was "fixed," for example, may leave open whether the fix was tested, whether earlier accounts were corrected, and whether affected branches were told.

Information in one set of documents can lead to a whole new investigation pathway and generate new searches to find more relevant information.

We asked Alchemy where its initial Horizon report needed further investigation. It proposed five areas for the team to pursue:

Five Areas for Further Investigation into Fujitsu's Handling of Horizon Issues

1. What "fixed" actually meant

A recurring pattern runs through the evidence. A defect is identified, a release or patch follows, and a witness says they understood the problem did not recur. What happened in between is rarely documented. At Callendar Square, the S90 release was presented as the final fix, but a Fujitsu employee questioned whether the multi-site problem had truly been resolved and proposed estate-wide checks. No confirmation of that follow-up has emerged ([LONDON_00001179 \\$7 ¶¶68-76](#)). The key question is whether "fixed" meant only that code was deployed, or also that recurrence was ruled out, historic accounts were corrected and affected postmasters were told.

2. Identifying who was affected

Fujitsu and Post Office produced figures for affected branches: two discrepancy occasions at Craigpark, at least four branches for Callendar Square and 62 branches for the receipts-and-payments mismatch. The method behind these figures is largely unexplained (

[LONDON_00000895 §12 ¶114-116](#); [LONDON_00000145 §6 ¶59-65](#)). One witness could not recall any consideration of wider implications, or whether the information reached those handling dismissals ([LONDON_00000145 §6 ¶63-65](#)). Understanding how each affected group was defined, and whether it was checked against prosecutions, recoveries and dismissals, may be central to assessing delay and disclosure failures.

3. The gap between technical knowledge and litigation evidence

Fujitsu's corporate witness accepted that defect information should have been included in litigation evidence. He also accepted that it was generally missing from Fujitsu witness statements and that Fujitsu had failed in its duty to provide accurate and complete data ([LONDON_00000025 §5 ¶47-52](#)). Fujitsu's 2005 prosecution-support policy did not require a check of the Known Error Log ([LONDON_00000025 §7 ¶66-78](#)). In one prosecution, a Horizon Online bug was not disclosed because the witness regarded Online as separate from Legacy Horizon ([LONDON_00000138 §6 ¶60-69](#)). Further work is needed to establish whether these omissions stemmed from how the policy was designed, from individual judgment, from drafting practices or from outside instruction.

4. The reliability of ARQ audit data

Fujitsu accepted that it could filter, insert or amend ARQ data, and that the ARQs it supplied did not exactly reflect branch transactions. It maintained that the underlying data did ([LONDON_00000025 §5 ¶53-57](#)). Practical limitations compounded this. Late-collected audit files were not flagged, and until around November 2010 some ARQs lacked reporting of duplicate or missing transactions ([LONDON_00001413 §3 ¶24-25, ¶32-35](#)). Case-level testing of what was changed, who authorised it and whether recipients were told would show whether these weaknesses affected evidence used against individuals.

5. Remote access: capability, use and accountability

The remote-access evidence is fragmented and partly contradictory. One witness recalled that branch data could be amended with Post Office authorisation. Another recalled Fujitsu representing that there were no unauditible access points. A Fujitsu executive recalled that developers could deploy code to live systems in 2010 and that this continued into 2011 ([LONDON_00001184 §5 ¶54](#); [LONDON_00001505 §10 ¶101](#); [LONDON_00001080 §7 ¶74-77](#)). In 2024, a review reportedly could not assess privileged-access controls because Fujitsu did not cooperate ([LONDON_00001092 §2 ¶19-24](#)). That suggests the question of scrutiny is ongoing rather than purely historical. A reconciled chronology should separate what Fujitsu could do, what it actually did, and whether the logs can prove either.

Each area can become a new search assignment. The team might begin with the first: find the S90 release records, proposed estate-wide checks, later reports of the Callendar Square problem, and records addressing affected accounts and branches. AI can assess

what those searches find and revise the working account. The team decides which lead to pursue and whether the new evidence answers the question.

5. Put an AI team to work on the next questions

Alchemy's initial report identified five areas for further investigation. In an agentic workflow, an **AI team leader** could turn each area into a defined assignment and dispatch agents to pursue them. Each agent would build and run searches, read the results, test competing evidence, and report back with document references and questions still open.

To show what one assignment could produce, we asked Alchemy to investigate what "fixed" meant for Callendar Square. It examined release records, proposed estate-wide checks, later incidents, and records concerning affected branches. Here is an excerpt from its follow-up report:

Executive Summary: Callendar Square and the S90 "Fix"

When Fujitsu said the Callendar Square (Riposte lock) problem was fixed after the S90 release, the records show it meant only that a code change was delivered in S90, rolled out to counters in March–April 2006, and followed by fewer lock and timeout events. The PEAK closure said the fix "appeared to have worked", and one witness accepted that this was less definitive than saying the problem was fixed.

The contemporaneous records do not show verified resolution:

- No test results. No S90 test results are available.
- No estate-wide check. There is no evidence that the proposed estate-wide check was carried out.
- No complete list of affected branches. A spreadsheet of affected branches existed but appears not to have been disclosed.
- No estate-wide account corrections. There is no evidence of a retrospective correction of branch accounts across the estate.
- No notice to postmasters. There is no record that the Post Office or affected postmasters were told this was a wider problem.

Later evidence contradicts a clean fix. A 2010 PEAK records a similar incident, and the High Court found the bug continued until 2010. The claim that the problem was "fixed during release S90" was nonetheless repeated in the Castleton and Misra cases without verification.

The full chronology, supporting citations, contrary evidence, missing records and recommended next searches are set out in the separate report, [*Callendar Square and the S90 "Fix": What "Fixed" Meant.*](#)

Other agents could investigate how affected branches were identified, the gap between technical knowledge and litigation evidence, the reliability of ARQ data, and remote-access capability and use. The AI team leader would compare their reports, connect findings across issues, and assign narrower follow-up work when new names, terms, events, or contradictions emerged.

The AI team leader would maintain a working record of the issues, people, terminology, significant documents, conflicting accounts, unanswered questions, and searches already run. The human team should be able to inspect and correct that record. An uncertain identification should not quietly become the premise for every later search or report.

6. An AI teammate should investigate what could prove its first answer wrong.

An investigator who finds evidence supporting a theory has not finished testing it. The next assignment should seek records that could undermine that account and test competing explanations. The Callendar Square report raised a consequential question: did the S90 release resolve the underlying defect, or did it only appear to do so? We asked Alchemy to search for evidence that could change its first answer.

Challenge the Callendar Square report's assessment of the S90 fix. Find the strongest evidence that the original defect was resolved, including release testing, post-release monitoring, incident records, and communications.

Determine whether later Riposte lock incidents involved the same underlying fault or a distinct problem. Present the evidence for each explanation, identify missing records, and cite the documents behind your assessment.

Alchemy ran new searches, examined the results, and prepared a second report testing the first. In particular, it looked at whether similar symptoms in later records established that the original defect persisted, and whether records outside the initial search supported a verified fix.

Executive report: Challenging the Callendar Report

Fujitsu's best case is that S90 fixed the identifiable Callendar Square failure, not that it eliminated every Riposte lock. Escher's correction was implemented in S90; a migration report recorded 99.9% counter-release completion by 22 March 2006; and Fujitsu recorded on 27 March that timeout locks had "gone right down". This is affirmative evidence of deployment and improvement, although the supplied summaries contain no fix-specific S90 test results ([LONDON_00000331 §19 ¶217–¶219](#)).

The defence must remain qualified. Fujitsu's 2019 analysis says it is unclear why a similar lock incident was reported in 2010. The summaries do not establish whether that event reproduced the original defect. They do identify a *different, specifically attributed* later end-of-day lock issue arising from a 2007 CABSPProcess code change, which supports the proposition that a later Riposte lock need not be a recurrence of Callendar Square ([LONDON_00000488 ¶8–¶10](#); [LONDON_00001226 §6 ¶65–¶68](#)).

Recommended position: challenge any inference that the 2010 PEAK alone disproves S90's effectiveness, while avoiding a claim of verified eradication. The decisive gaps are the S90 test results, complete post-release monitoring and incident records, the affected-branch list, and a technical comparison of the 2010 event with the original fault. Until those are obtained, the evidence supports an apparently effective fix with unresolved later-incident questions—not a conclusive finding either way ([LONDON_00000488 ¶5, ¶9–¶10](#); [LONDON_00000331 §19 ¶219](#)).

The full chronology, supporting citations, contrary evidence, missing records and recommended next searches are set out in the separate report, [Challenging the Callendar Report](#)

The challenge report narrows the initial account. It finds evidence that S90 was deployed and lock timeouts declined, while leaving open whether the 2010 incident involved the same defect and whether the original fix was fully validated. The AI team leader could carry forward that more precise account, assign searches for the missing test and incident records, and bring both reports and the critical documents to the human team for review.

7. Know when there is enough for the task

An AI team could keep finding relevant documents long after it has enough to help with the immediate assignment. It could also produce a persuasive report while important



evidence remains undiscovered. The team needs a way to judge whether further work is likely to change its understanding.

Alchemy's **Adaptive Sampling** system samples the ranked document set to estimate how many relevant documents are in the collection. An LLM reviews the sample and distinguishes documents that add important new information from those that are relevant but repeat what the team already knows. Together, those measures help the team see whether deeper review is still producing new evidence or mostly cumulative material.

For a Callendar Square witness interview, the chronology, competing reports, and principal documents may be enough to prepare focused questions. If Adaptive Sampling shows that deeper ranks are still producing new facts or contrary evidence, the team has reason to keep investigating. If they yield mainly cumulative records, the team may have enough for that interview while continuing targeted searches for the missing S90 test and monitoring records. A disclosure decision or a conclusion about all affected branches may require a different depth of review.

The AI team leader can present those estimates alongside the open questions, searches completed, and lines of inquiry deferred. But in all cases, the human team decides how deep to go for the task at hand.

8. Let AI suggest next steps for the investigation

After the initial report and the Callendar Square follow-up, we asked Alchemy to recommend the next phase of the broader Horizon investigation. It identified eight steps, from mapping defects to affected cases through recovering missing records and testing the evidence behind each conclusion.

Here is a summary of Alchemy's recommendations for next steps:

Recommended next steps

Map every defect against every affected case. Link each defect's operative period to prosecutions, civil recoveries, suspensions, dismissals and transaction corrections. Include the roughly 2,400 ARQs never tested against known reliability incidents, and identify cases that cannot be cleared because data no longer exists ([LONDON_00000025 §5 ¶46-47](#); [LONDON_00000025 §10 ¶102-105](#)).

Establish what "fixed" meant for each defect. Determine whether each fix was verified across the estate and whether branch accounts were corrected, rather than accepting closure "in the next

release" ([LONDON_00000164 §39 ¶412-416](#); [LONDON_00000112 §12 ¶123-127](#)).

Investigate the editing of witness evidence. Obtain every draft and final Fujitsu witness statement with its edit trail. Identify who proposed and approved the removal of the Riposte lock paragraphs, and whether other statements were edited in the same way ([LONDON_00000165 §2 ¶17-22](#); [LONDON_00000025 §5 ¶49-50](#)).

Test ARQ integrity at case level. Record what was filtered, inserted or amended, who authorised it, whether it was logged, and whether recipients were told ([LONDON_00000025 §5 ¶53-57](#)).

Build a remote-access chronology. Set out capabilities, approvals and logging for each period, and test every assurance given to Post Office, Second Sight and courts against it ([LONDON_00000112 §12 ¶116-120](#); [LONDON_00001080 §7 ¶72-77](#)).

Trace the Fujitsu–Post Office handoff. Follow BIMS reports through to error notices and establish whether action was taken against subpostmasters after a system-caused loss had been identified ([LONDON_00000087 §18 ¶205-210](#)).

Recover core documents. These include S90 release and test records; branch-level deployment data; pre- and post-S90 event data; the 195-ARQ review and 400,000-ARQ results; the 40- and 62-branch mismatch lists with repayment records; BIMS reports and error notices; privileged-access logs; and Ernst & Young and Deloitte findings.

Apply investigative discipline. Grade every finding as documented fact, recollection, corporate reconstruction or allegation, noting that Fujitsu's corporate admissions rest on documents selected by its legal team ([LONDON_00000025 §2 ¶14-16](#)).

You can see the full report here [Next Steps for the Investigation](#)

The full report gives the team its next assignments. It can ask Alchemy to trace a witness statement through its drafts, build a witness kit for Anne Chambers, or investigate how a particular defect affected individual cases. Each request can draw on the work already done, run new searches where needed, and update the account as the evidence develops. The human team sets the priorities, inspects the critical documents, and decides what the findings support.

A Better Way to Discover What Happened

The progression across our series now comes into focus. Keywords find language.

Semantic search finds meaning. CAL learns relevance from assessed examples.

AI-Integrated Discovery carries that learning through the matter: helping the team frame the inquiry, examine the evidence, pursue new leads, challenge its conclusions, and decide where to look next.



The division of labor matters. **AI investigates broadly. Humans direct, challenge and decide.** AI can search, read, learn and synthesize across a large collection. The human team examines critical evidence, tests the answers, and exercises judgment about what the matter requires.

For decades, discovery systems waited for us to supply the next query. AI-Integrated Discovery can help us frame the question, follow the evidence, test competing explanations, and see what remains unknown. It gives every member of that team a more effective way to understand what happened and a more efficient way to act on what they learn.

That matters beyond the discovery team. When evidence can be found and tested sooner at lower cost, parties can assess their positions earlier, resolve disputes faster, and pursue or defend claims that prolonged review might otherwise put out of reach.



About the Authors

John Tredennick

CEO & FOUNDER, MERLIN SEARCH TECHNOLOGIES

John Tredennick is CEO and founder of Merlin Search Technologies, a company pioneering AI-powered document intelligence for legal professionals. He spent the first 20 years of his career as a trial lawyer and senior litigation partner at Holland & Hart LLP, where he became one of the first CTOs of a major law firm. In 2000, he founded Catalyst Repository Systems, an international ediscovery technology company acquired by OpenText in 2019. He has authored or edited eight books and dozens of articles on legal technology and AI, has spoken on five continents, and served as Chair of the ABA's Law Practice Management Section. He has been recognized by The American Lawyer as one of the six top ediscovery pioneers.

jt@merlin.tech



Dr. William Webber

CHIEF DATA SCIENTIST, MERLIN SEARCH TECHNOLOGIES

Dr. William Webber is Chief Data Scientist of Merlin Search Technologies. He completed his PhD in measurement in information retrieval evaluation at the University of Melbourne under Professors Alistair Moffat and Justin Zobel, and his post-doctoral research at the E-Discovery Lab of the University of Maryland under Professor Doug Oard. With more than 30 peer-reviewed publications in information retrieval, statistical evaluation and machine learning, he is a leading authority on AI and statistical measurement for information retrieval and ediscovery. At Merlin he leads the data science behind the Alchemy™ platform, including its multi-LLM architecture and the validation methods used to measure AI performance.

wwebber@merlin.tech

