

cape^{start}

CASE STUDY




How **Conversational Search** Accelerates Systematic Literature Review For Medical Affairs Teams With An AI-Enabled Chat Assistant

Problem

Medical affairs and HEOR teams conducting systematic literature reviews work across thousands of scientific articles per project. The answers they need — prevalence figures, reported efficacy, safety signals — exist in the corpus, but there was no way to simply ask for them. Teams needed a solution that could overcome:

- Keyword searches that surface large volumes of articles with no way to refine intent, forcing analysts to sift results by hand
- Answers locked behind manual re-reading — data points such as disease prevalence or reported efficacy retrievable only by opening and skimming every included article
- Summaries drafted from scratch, consuming hours of skilled analyst time per deliverable
- Information silos — charts and artifacts built in third-party tools through repeated uploading and downloading
- Limited visibility for managers and directors into the questions asked and artifacts generated by their teams



Solution

An AI conversational search assistant, deployed as a layer over the indexed literature corpus — whether that corpus is structured by a document-intelligence pipeline or lives in an existing literature review platform. **The AI-enabled solution provides:**



Grounded, RAG-based answers — natural-language questions returned with citations and source-page highlighting for one-click verification



Intelligent query refinement — the assistant asks clarifying questions (disease subcategory, date ranges) before surfacing the most relevant articles



Project-aware context — distinguishes screening phases (title & abstract vs. full text) and recognizes human include/exclude decisions



On-demand answers and draft summaries — delivered as text, tables, images, or graphs



An orchestration layer that routes each request to the best-suited LLM for the task



Regulatory context on demand — users attach country-specific guidelines for the AI to validate responses against

To effectively develop and deploy the platform, the CapeStart team worked with the client to implement the following strategies:

- Built on a conversational search engine proven in a global enterprise deployment indexing 10,000 technical documents
- Approximately 70% reusable components, enabling a three-month development cycle from build to production
- Human-in-the-loop validation workflow keeping subject-matter experts in control of final outputs
- Full audit trails and session records, with super-admin oversight of every question asked and artifact generated
- Hosted on AWS and built to life sciences standards — SOC 2 Type II, HIPAA, and GDPR, with data-residency support
- A familiar chat interface, consistent with leading AI assistants, for near-zero onboarding burden
- Available within the platform or via build-operate-transfer into an organization's in-house literature review systems

Results

The development and implementation of the conversational search assistant delivers significant improvements for literature review teams:

92% Annual ROI

46 of 50 user-reviewed queries validated as correct in the proven enterprise deployment.

Human-in-the-loop validation before responses are finalized

70% Less Manual Effort

Draft summaries pre-generated — analysts ideate and edit instead of starting from scratch.

Time-to-answer: expert search sessions cut from 30 to 10 minutes in the proven deployment — answers 3x faster.

8x+ Annual ROI

Estimated 8x+ yearly return on run cost in the proven enterprise deployment, compounding over a multi-year term. Capacity freed for higher-value evidence synthesis and review oversight

Scalability

Proven at 10,000 documents in a single deployment. Multi-tenant architecture processing thousands of documents simultaneously across concurrent literature reviews platformwide.

Pairs with CapeStart Document Intelligence — Document Intelligence structures the corpus; Conversational Search turns it into answers. Together they take a document estate from raw files to instant, cited answers.

RAG + Citations

Query Refinement

Human-in-the-Loop

Audit Trails

AWS

SOC 2 Type II

HIPAA

GDPR

Build-Operate-Transfer

CapeStart Conversational Search is available for life-sciences and medtech organizations — within the platform or via buildoperate-transfer.

For more information, contact hello@capestart.com · www.capestart.com

capestart
Production AI for regulated life sciences