

A life-sciences chatbot that answers 65% of product and technical queries — grounded only in approved documentation, with a citation on every answer

Cytiva — 8 weeks

We built a grounded AI chatbot for a life-sciences business that now handles an estimated 65% of product and technical questions — answering strictly from the company's approved documentation, citing the source on every response, and deferring to a human whenever it lacks a grounded answer.

THE CHALLENGE

A life-sciences company wanted an assistant that could answer detailed product and technical questions — but in healthcare, a chatbot that hallucinates is worse than none at all. Answers had to be grounded in approved source material, traceable to a citation, and safe within a regulated context. An off-the-shelf general-purpose chatbot, with no grounding and no audit of what it told users, carried unacceptable accuracy and compliance risk.

WHAT WE BUILT

- Indexed the company's approved documentation into a retrieval-augmented (RAG) knowledge base so answers are grounded in real, controlled sources
- Engineered the assistant to cite its sources, so every answer is traceable back to an approved document
- Added guardrails and fallbacks — the assistant defers or escalates to a human rather than guessing when it lacks a grounded answer
- Built a human review and feedback loop so subject-matter experts correct and improve responses over time, with interactions logged for oversight

OUTCOMES

- The chatbot now handles an estimated 65% of product and technical queries, grounded only in approved documentation with a citation on every answer and zero ungrounded responses reaching users
- Every response traceable to an approved document — critical for a regulated industry
- Hallucination risk contained through retrieval grounding, guardrails and human escalation