



AI chatbot feature overview

Last Modified on 06/25/2025 10:54 am EDT

Use our AI Chatbot to answer your readers' questions with information from your knowledge base. No separate accounts, complex integrations, or data import required.

The quick summary

Our AI Chatbot is built entirely within KnowledgeOwl's infrastructure. Your information is never sent anywhere else and isn't used to train models.

Your readers ask a question of the chatbot. It analyzes information in your knowledge base--and ONLY in your knowledge base--and returns an answer using the tone and audience info you provide in your setup. It will include images from your documentation if they're relevant to answering the question.

The chatbot cites the sources it used, encouraging your readers to learn more in your knowledge base and establishing trust and authority.

And it admits when it doesn't know.

The more detailed process

Our AI Chatbot is built entirely within KnowledgeOwl's infrastructure using [Anthropic](#) AI models. Your information is never sent to Anthropic or used to train their models. We use a static copy of one of their models and restrict it to your knowledge base.

When you first set up the chatbot, we generate a separate search index--called a semantic search index--for your knowledge base optimized for returning search results for questions. This index, like the normal keyword search index, is updated any time you save an article or category.

When a reader asks a question, we query that semantic search index. We send the top results from that search into the AI model and ask it to return an answer.

The AI model analyzes the question and information we gave it, uses the descriptions you've provided for your company/audience and the tone you'd like, and generates an answer.

It cites which of the results it used to create that answer as hyperlinks so your readers know where to go to learn more and can assess whether it feels like the chatbot properly understood the question.

If our initial semantic search doesn't return any good results, the chatbot admits that it can't answer the question.

Accuracy and security safeguards

You've probably heard a lot about Generative AI Large Language Models (LLM) offering incorrect information--known as hallucinations. So have we. We're firm believers that the information you store in KnowledgeOwl is *yours*. We know how much trust you place in us to be good custodians of that information, and we take that responsibility seriously.

We focused on designing a chatbot we felt confident using ourselves, one that could answer questions effectively and accurately while keeping our information secure.

Here's what that means in practice:

- The chatbot doesn't try to "learn". We don't send information back to Anthropic to train their models and we don't try to further train the model.
 - We don't store a history of the questions you asked. The question exists on our servers only as long as it takes to process the search query and feed the question and relevant results into the AI model.
 - Once the AI model generates its response, it flushes all history: the question that was asked, the search results we provided, and the response it generated. As a result, each question asked--even when a reader asks multiple questions within the same session--is treated as a discrete, standalone question.
 - The AI model only accesses the relevant results from your knowledge base for the question it's processing. Your information isn't shared back to the model in any way and will never be used as information for another knowledge base's chatbot.
 - Every step in this process is handled entirely within the KnowledgeOwl infrastructure. We don't share it back to Anthropic. We don't send it anywhere else. Your knowledge base information is as secure in this process as it's always been.
-