

OpenText Knowledge Discovery delivering off cloud virtual assistance

Answer natural language questions from enterprise data through the power of AI



Benefits

- **Get to the best answer faster** through direct questioning
- **Ensure correct answers** by matching questions to curated responses
- **Generate informed answers** from structured or unstructured data

The deluge of information facing customers and support staff is threatening to drown them. The amount of time it takes to sift through to find the correct answers (and the risk of missing them entirely) can do real damage to efficiency and customer and employee satisfaction.

OpenText™ Knowledge Discovery (IDOL) uses natural language question answering (NLQA) technology to allow natural conversations between users and a virtual assistant. Staff and customers can perform queries against curated (FAQ) responses (answer bank), data tables and data services (fact bank), and unstructured documents (passage extraction), following configured processes to facilitate knowledge discovery.

Built on proven technology and innovations, such as machine learning and deep neural networks, this solution unlocks hidden insights by revealing trends, patterns, and relationships. Develop an in-depth understanding of user profiles and actions to personalize knowledge delivery. NLQA transforms the end-user experience for greater efficiency and satisfaction.

With modular offerings of hundreds of advanced analytics functions, as well as an open and scalable architecture for easy embedding and third-party integration, OpenText Knowledge Discovery lends itself to supporting diverse use cases spanning a broad spectrum of industries.

Question answering

Streamlining search means getting to the point faster. If a user enters a direct question, for example, it is generally more appropriate to respond with a direct answer rather than a document or documents that may contain the answer. The technology must be able to understand, process, and answer direct questions to streamline the retrieval process, which allows for more convenient and user-friendly information retrieval. Moreover, it should offer an interface to allow configuration of that process to ensure existing human knowledge is incorporated to train the system to answer questions optimally.

There are typically three independent steps involved in accomplishing this:

Step 1: Answer bank

Many administrators of support or user-help systems have an existing set of frequently asked questions that human support agents are trained on and help pages are populated with.

The system should be loaded with these answers and its natural language processing (NLP) should ensure that differently worded variants of the same questions are all directed to the relevant answer. This knowledge base is generally referred to as the answer bank, where OpenText Knowledge Discovery's NLP is used to match a user's question to one of these curated questions and return the relevant answer.

Step 2: Fact bank

Users frequently have tables of information that are commonly matched via structured (SQL) queries. The system should be able to exploit such databases of information to allow natural language questions to receive direct factual answers. This is achieved by processing and understanding the question and mapping it to an appropriate structured query that will in turn respond with the desired answer.

In addition, the system should have the ability to understand and extract information (for example names and numbers) from unstructured data such as free-form documents. This means that databases can be automatically populated with a rich set of structured data from a corpus of unstructured documents. For example, ingesting a set of corporate annual reports could create tables of financial data so that a question such as, "What was BP's revenue in North America in the second half of 2023?" will return the correct value pulled from such a table.

Step 3: Passage extraction

In many cases, the information requested is simply not present in either an FAQ data set or a structured database, and an extended approach is required. At this point, the system must be able to process human information effectively.

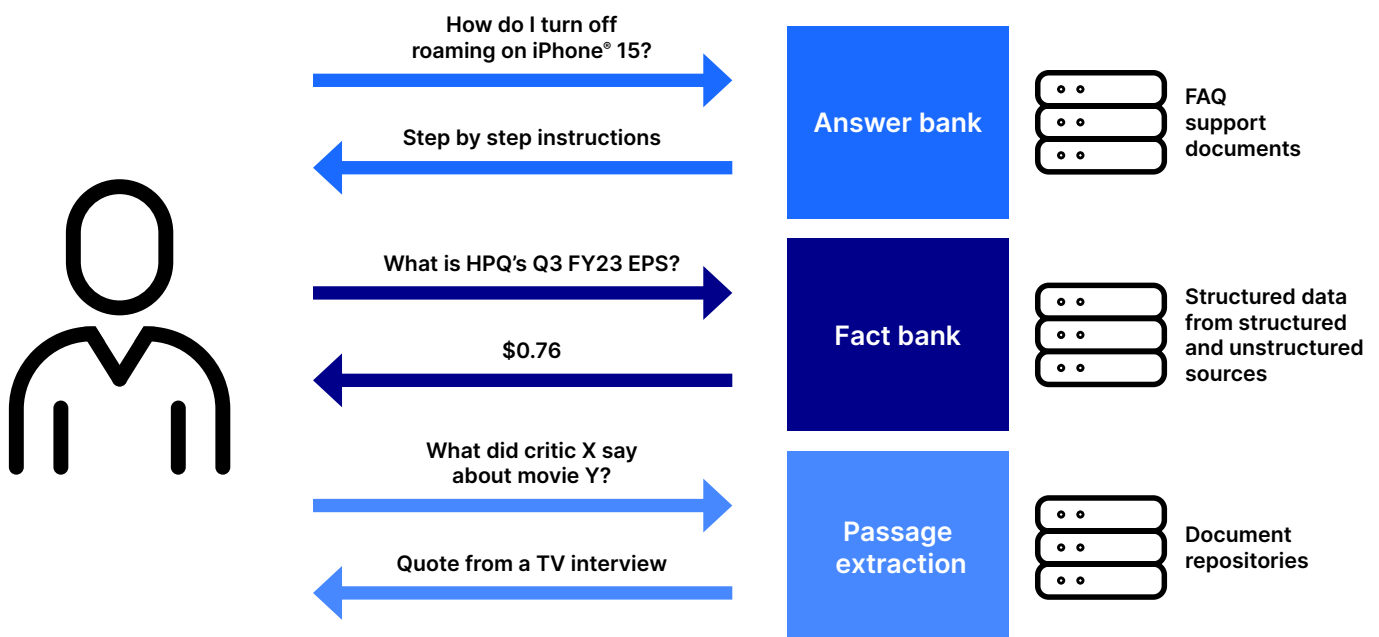
OpenText Knowledge Discovery integrates with your choice of on or off cloud LLM services to generate answers from your indexed documents. This integration mode is commonly referred to as "Retrieval Augmented Generation" or "RAG."

The solution acts as a secure retrieval system by following this process:

1. It forms an understanding of the intent of a question and securely retrieves relevant sections of text from your documents.
2. It passes your original question to an LLM with that reference material and prompt engineering instructions to provide an answer from it.
3. The returned answer from the LLM is then compared to the reference content to provide a validation confidence level.

The result is concise, validated answers to your questions from your own data.

The above steps combine to form a powerful data analytics system for processing questions of any type. The typical workflow is to first check whether the query is appropriate for question processing, as many queries are not of this type. If it is, then the system matches the question against the answer bank to see if a pre-defined answer has been set. If none is found, it will then try the fact bank to see whether a precise answer is present. If an answer is still not found, it will pass the query to the system's unstructured index to see if an appropriate passage can be found that answers the question.



Question answering in natural language

Dialogue systems

For many users contacting a company for help with a problem or an inquiry, a search box is insufficient. Even if the system is capable of understanding and processing direct questions, the user is often unable to accurately describe the problem or need. They can also fail to provide enough information.

The natural response in such situations is to want to talk to a human operator who will guide them to the correct solution or answer. However, large teams of human operators are costly, and clients are often frustrated by having to wait to talk to someone. So, how can we accelerate the right information to the right user in the right context?

NLQA can be extended to create a more natural system of customer services. Rather than a single search box, an IM-style interface allows a user to enter an initial question or problem description. In many cases, the problem can be answered directly and a single answer or solution provided. However, in many other cases, a conversation between the system and the user is required to gather more information on how to diagnose and solve the problem.

Resources

[Request a demo ›](#)

[Knowledge Discovery
product page ›](#)

[What's new—rolling blog ›](#)

This can be achieved by training the system with classes of problems and the type of information that is needed in each case. The system will then respond with follow-up questions until it believes it has enough knowledge to suggest a solution.

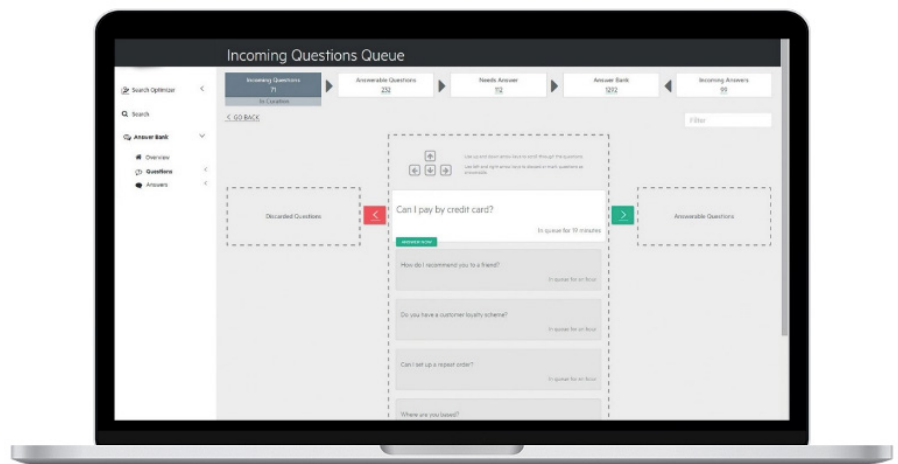
Here is a very simple example:

1. The user enters, “Cheapest ticket from London to Boston.”
2. The system responds, “Quantity, date and time, preferred airline?”
3. The user responds with answers.
4. The system presents best options optimized for different parameters: schedule, cost, and mileage program.

Simple curation of question/answer pairs in answer bank

While it is natural to expect a manually-intensive process for curating question/answer pairs (FAQs), OpenText Knowledge Discovery comes with an intuitive GUI tool to simplify the process. This tool allows administrators to easily refine the answers’ availability, accuracy, and relevancy with a quick click.

The growing amount of content in existing systems and business applications threatens to make it near impossible for users alone to understand and gain actionable insight. OpenText Knowledge Discovery’s NLQA transforms the end-user experience. It unlocks hidden insights by revealing trends, patterns, and relationships, delivering a user-centric answering service that overcomes the main barrier to any technology success: user adoption. Gain an in-depth understanding of user profiles and actions and personalize knowledge delivery.



Question/Answer pairs in answer bank