

Q&A

One of the most common use-cases for an LLM application is to answer questions about a set of documents. LlamaIndex has rich support for many forms of question & answering.

Types of question answering use cases

Q&A has all sorts of sub-types, such as:

What to do

- **Semantic search**: finding data that matches not just your query terms, but your intent and the meaning behind your question. This is sometimes known as "top k" search.
- **Summarization**: condensing a large amount of data into a short summary relevant to your current question
- **Structured Analytics**: translating your query into a SQL statement that can be run against a database (e.g. Postgres, MySQL, Snowflake). These types of questions usually include analytical quantities like max, sum, averages (e.g. "Tell me the average order volume in 2021").

Where to search

- **Over documents**: LlamaIndex can pull in unstructured text, PDFs, Notion and Slack documents and more and index the data within them.
- **Over structured data**: if your data already exists in a SQL database, as JSON or as any number of other structured formats, LlamaIndex can query the data in these sources.

How to search

- **Combine multiple sources**: is some of your data in Slack, some in PDFs, some in unstructured text? LlamaIndex can combine queries across an arbitrary number of sources and combine them.
- **Route across multiple sources**: given multiple data sources, your application can first pick the best source and then "route" the question to that source.
- **Multi-document queries**: some questions have partial answers in multiple data sources which need to be questioned separately before they can be combined

Examples

For examples of all of these types of Q&A, check out [\[Q&A\]/understanding/putting_it_all_together/q_and_a.md](#)) under "Putting it all together".