

# DDS HR Chatbot: A Deep Dive into the Code

- A Retrieval-Augmented Generation (RAG) System for HR Queries
- Core Technologies:
  - OpenAI: For language understanding, embedding, and answer generation.
  - Pinecone: For efficient, scalable similarity search (Vector Database).
  - Gradio: For creating a simple, interactive web interface.

# High-Level Architecture (The RAG Flow)

**Indexing (One-time setup):** HR documents are converted into numerical representations (embeddings) and stored in a Pinecone vector database.

**User Query:** An employee asks a question through the Gradio web UI.

**Retrieve:** The system finds the most relevant document snippets from the knowledge base.

**Augment & Generate:** The retrieved snippets (context) and the original question are sent to an OpenAI model.

**Answer:** The model generates a final answer based on the provided context and sends it back to the user.

# Code Deep Dive: Setup & Initialization

Purpose: To securely load API keys and configure the clients for OpenAI and Pinecone.

Key Functions/Libraries:

```
import os, openai, pinecone
```

```
dotenv.load_dotenv(): Loads secrets from a .env file.
```

Code Snippet:

```
```python
# Load environment variables
OPENAI_API_KEY = os.getenv("OPENAI_API_KEY")
PINECONE_API_KEY = os.getenv("PINECONE_API_KEY")

# Initialize OpenAI client
client = OpenAI(api_key=OPENAI_API_KEY)

# Initialize Pinecone index
pc = Pinecone(api_key=PINECONE_API_KEY)
```
```

# Code Deep Dive: Loading & Indexing the Knowledge Base

Purpose: To read text files and store them as vector embeddings in Pinecone for fast searching.

Key Functions: `load_documents()` & `index_documents()`

Process:

1. Load: Recursively find all .txt files and read their content.
2. Embed: Convert document text into numerical vectors using OpenAI's text-embedding-ada-002 model.
3. Upsert: Upload these vectors and their metadata to the Pinecone index.

Code Snippet:

```
```python
def index_documents(index, docs: List[dict]):
    embeddings = client.embeddings.create(
        input=texts,
        model="text-embedding-ada-002"
    )
    index.upsert(vectors)
```
```

# Code Deep Dive: Retrieving Relevant Information

Purpose: To search the database for documents that are most relevant to the user's question.

Key Function: `retrieve()`

Process:

1. The user's question is also converted into a vector embedding.
2. Pinecone is queried to find the top\_k (e.g., top 5) document vectors with the highest similarity score.
3. The original text from these matching documents is returned.

Code Snippet:

```
```python
def retrieve(query: str, index, k: int = 5):
    embed = client.embeddings.create(
        input=[query],
        model="text-embedding-ada-002"
    ).data[0].embedding

    res = index.query(vector=embed, top_k=k, ...)
    return [m["metadata"]["text"] for m in res["matches"]]
```
```

# Code Deep Dive: Generating the Final Answer

- Purpose: To use a Large Language Model (LLM) to create a clear, concise answer from the retrieved text.
- Key Function: `generate_answer()`
- Process:
  1. A System Prompt instructs the AI to act as an HR assistant and only use the provided information.
  2. The retrieved documents (context) and the user's query are combined into a single message.
  3. This message is sent to OpenAI's gpt-3.5-turbo model.
- Code Snippet:
- ```
```python
def generate_answer(query: str, docs: List[str]):
    system_prompt = (
        "You are a helpful HR assistant. Use the provided context..."
    )
    response = client.chat.completions.create(
        model="gpt-3.5-turbo", messages=messages
    )
    return response.choices[0].message.content
```
```

# Bringing it to Life: The Gradio User Interface

Purpose: To provide a simple and interactive web app for users.

Key Library: gradio

UI Components:

`gr.Blocks()`: A flexible layout for organizing components.

`gr.Image`, `gr.Markdown`: For branding and titles.

`gr.Button`: For category selection.

`gr.Dropdown`: For example questions.

`gr.Textbox`: For user input and the final answer.

Event Handling:

`.click()` and `.change()` methods link UI actions to backend Python functions.

# Summary & Key Takeaways

The application uses a Retrieval-Augmented Generation (RAG) architecture to provide accurate, context-aware answers.

Pinecone acts as the specialized long-term memory for fast similarity search.

OpenAI provides the intelligence for both understanding text and generating answers.

Gradio makes it easy to build and share an interactive demo.

This architecture is highly effective for building factual, domain-specific chatbots.