



No-Code RAG Pipeline Generator

Simplifying RAG Pipelines: Just Specify
Models, Get Production-Ready Code


The Challenge of Building RAG Pipelines

- **Complex Setup:** Building a RAG pipeline requires expertise in multiple areas: data indexing, retrieval, embedding, and LLM integration.
- **Time-Consuming Process:** Customizing each component and testing end-to-end functionality takes days or weeks, limiting quick prototyping.
- **Specialized Skills Required:** Limited accessibility for those without backend or ML expertise, raising barriers to adoption.

Who Can Benefit from the No-Code RAG Generator?



Data Scientists & Engineers: Speed up experimentation and focus on optimization.



Developers: Quickly integrate RAG into applications without complex setup.



Hackathon Teams: Save time, showcase RAG solutions fast.



Enterprises: Deploy RAG capabilities without hiring specialized experts.

Key Features of the RAG Pipeline Generator

Model Selection: Supports various embedding and LLM models for different project needs.

Automated Code Generation: Produces end-to-end RAG pipeline code, including indexing and retrieval.

Integration with Vector Databases: Compatible with Pinecone, FAISS, and Chroma for efficient storage and retrieval.

Flexible Retrieval Settings: Customize retrieval parameters, document chunking, and database configurations.

Model Architecture

We created 2 distinct Architectures:

- Prompting based
- Rag based

