


GENERATIVE AI SERIES: DAY 21

The Generation Step: Synthesizing the Final Answer

We've found the right puzzle pieces (our context). Now, let's watch how LLM expertly puts them together to form a perfect, human-like answer! 🎯



We've Retrieved... Now We Generate!

So far, our RAG system has:

- Indexed our knowledge base (Day 19).
- Retrieved the most relevant chunks based on a user's question (Day 20).
- But the LLM itself hasn't seen this context yet!
- The Generation Step is the bridge that connects our retrieved evidence to the LLM's reasoning power.

@learn.machinelearning




It's All About the Prompt!

We don't just ask the LLM the original question. We create a new, much more powerful "Mega-Prompt" that includes:

- Clear Instructions:
- The Retrieved Context:
- The User's Question:

This combined prompt gives the LLM everything it needs.

@learn.machinelearning




Method 1: Stuffing (The Simple Way)

- This is the most straightforward approach. You literally "stuff" all the retrieved context chunks and the question into a single prompt.
- Example: "Use the following context to answer then.
CONTEXT: [chunk 1 text] [chunk 2 text]...
QUESTION: What were the key findings?"

- Pros: Easy and fast

- Cons:
Fails if the total text exceeds LLM's context window limit.




@learn.machinelearning

Method 2: Map-Reduce (For Big Docs)

- What if you have too much context to 'stuff'?
- Map-Reduce is a clever solution.
- Map: Send each context chunk to the LLM individually and ask it to answer question or summarize it.
- Reduce: Take all the individual answers/summaries and send the LLM in final in a final call to synthesize a single, cohesive answer.

@learn.machinelearning



How to Ensure Faithfulness & Trust

- Providing context is only half a battle. Your instructions are what guarantee a good answer. Always tell LLM:
Always tell LLM:
- ✓ 'Answer ONLY from the provided context.'
(Prevents making things up)
- ✓ If the answer is not in the context, say 'I don't know.'
(Builds trust)
- ✓ 'Cite your sources from the context.'
(Allows users to verify)

@learn.machinelearning



Use Context Only
Don't Hallucinate
Cite Sources

GENERATIVE AI SERIES: DAY 21

Day 21 Takeaway:

Prompting Turns Context into Answers!

The Generation step is where the magic happens. By carefully constructing with clear instructions and the right retrieved evidence, we guide the LLM to move beyond its general knowledge and become a true expert on your specific data.

Next Up (Day 22): Advanced Retrieval Strategies: Leveling Up Your RAG!



@learn.machinelearning



Perfect Answer

[1]

[2]

What's the #1 rule you'd give AI to make its answers more trustworthy? Let me know! ↓
Follow @learn.machinelearning for the rest of our RAG Deep Dive!
SAVE this post to remember how to build a better AI prompt!

