Tutorial 4

Week 8

March 18, 2026

1 Foundations

1. What is a large language model fundamentally trained to do? What type of machine learning problem is it?
2. What is LLM hallucination? Why does it happen? Give examples to hallucinations you encounter in your life.
3. What is RAG? Where do you encounter RAG?
4. What is reasoning in LLM systems?
5. What is an AI Agent? Which AI agents do you use? Which AI Agents do you wish you have existed.

2 ElmasAI

You are designing a fiction writing assistant named ElmasAI, which is an AI Agent with a Retrieval-Augmented Generation system that integrates with a writer's drafting environment, for example a manuscript editor. It has access to local files on the user's computer. It employs an LLM through API calls.

ElmasAI provides real time feedback on narrative elements such as grammar, clarity, character consistency, and plot coherence, and can generate or revise story content upon request. It should perform well on large writing projects with multiple interconnected documents stored locally, such as chapters, character profiles, timelines, and setting descriptions.

1. What is the machine learning problem ElmasAI solves? Is it descriptive, predictive, or prescriptive?
2. Describe the potential data & data sources to train and fine tune ElmasAI.
3. Apart from the data that is used to train and fine tune ElmasAI, what documents should ElmasAI use when running user commands?

4. Suppose the user asks ElmasAI to check whether the main character's face is described consistently in the current chapter and to fix any inconsistencies. Describe the sequence of high level steps ElmasAI should follow to handle this request effectively, from receiving the user's instruction to delivering the final response.

5. What are the security and privacy concerns with ElmasAI? How can these be mitigated? What are the solution's drawbacks?