

Prompt Engineering, ICL & more

LLM course, UW, Seattle
Feb 9th 2026

Dr. Karthik Mohan

Prompt Engineering use-cases

Prompt Engineering for information retrieval

Data Augmentation

Transfer Learning to smaller models

Prompt Engineering use-cases

Prompt Engineering for information retrieval

Data Augmentation

Transfer Learning to smaller models

Prompt Engineering use-cases

**Prompt Engineering for information
retrieval**

Data Augmentation


**Transfer Learning to smaller
models**

Prompt Engineering use-cases

Prompt Engineering for information retrieval

Data Augmentation

Transfer Learning to smaller models



Prompt Engineering use-cases

**Prompt Engineering for information
retrieval**

Data Augmentation

**Transfer Learning to smaller
models**

More use cases!

Prompt Engineering use-cases

**Prompt Engineering for information
retrieval**

Data Augmentation

**Transfer Learning to smaller
models**

Data Annotation

Prompt Engineering use-cases

Prompt Engineering for information retrieval

Data Augmentation

Transfer Learning to smaller models

Data Annotation

LLM as a Judge

Today's Focus

Prompt Engineering for information retrieval

Data Augmentation

Transfer Learning to smaller models

Data Annotation

LLM as a Judge

Prompt Engineering - Notebook Demo

Let's go take a look!

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Example from the notebook

The key word doesn't have to be present in the text. Also the key word shouldn't have a space in it.

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Example from the notebook

One question should be something a **five year old** would ask. Another second should be something a **mature adult** would ask.

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Example from the notebook

Generate 3 distinct key words that capture the most important topics in the text.

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Chain of Thought Prompting

Example from the notebook

Explain step by step. Use of <think> tags and <answer> tags

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Chain of Thought Prompting

In-Context Learning

ICL

Zero-shot

Few-shot

CoT

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Chain of Thought Prompting

In-Context Learning vs RAG Models

ICL vs RAG

In In-context Learning - All data is provided within the context window.

In RAG - First a database of docs is searched before generation from a single relevant doc. Search here is a vector search or semantic search

Pointers for effective prompt engineering

Clarity in Instructions, Goals

Providing context

Specificity/Conciseness

Chain of Thought Prompting

In-Context Learning vs RAG Models

ICL vs RAG

In In-context Learning - All data is provided within the context window.

In RAG - First a database of docs is searched before generation from a single relevant doc. Search here is a vector search or semantic search

Zero-shot vs ICL

in the following context, respond with one word. if for example the input is modi - respond with nehru jacket.

obama
trump
pope

suit
tie
cassock



butterfly

wings



kangaroo

pouch

ICL: Model adapts to the “local context” provided

Putting it all together...

Real-World Examples

Problem	Best approach	Why
Sentiment classification	Zero-shot	Well-learned task
Resume → JSON parsing	Few-shot	Format matters
New internal policy rules	In-context learning	No retraining, fast iteration
Math word problems	Chain of Thought	Reasoning critical
Company wiki Q&A	RAG	Knowledge outside model

Many industry use-cases

Industry Use-cases

Many industry use-cases can be addressed by a combination of **smart prompting (ICL with CoT), **evaluation** (LLM as a judge), **RAG** where needed and **good inferencing** (quantization, vllm, etc).
Fine-tuning LLMs are specialized use-cases where ICL or RAG didn't meet the needs.**

This leads us to.. ChatBot Design

ChatBot Design

A lot of **chatbot design** in the industry is about understanding the **business context**, **product requirements**, **engineering requirements** and then **creatively packaging** llms, tools, inferencing, agents to meet the requirements

Mini-Project 2

ChatBot Design & Evaluation

Use of RAG, ICL, ChatBot Design principles, LLM as a judge and putting it together into a web app. MP2 or MP3 will be part of a team final-project presentation on March 16 and March 18