Large Language Models

NLP Basics and The Idea of Language Models

- What is Natural Language Processing (NLP)? Why is it important for computers to understand human language?
- Brief overview of how computers process text (tokens, words, sentences).
- Introduction to the concept of a "language model": predicting the next word in a sequence.
- Simple examples of language models in action (e.g., autocomplete on your phone).

- **Activity**
  - Word Prediction Game: Students are given the start of a sentence and asked to predict the next 3-5 words. Discuss the different possibilities and why certain words are more likely.
  - Discussion: What makes some predictions more "natural" than others?
  - Observe text prediction/autocomplete on your phone or a website. Write down 3 examples where it was helpful and 1 example where it was incorrect or amusing.

From Simple Models to "Large" Language Models

- Brief history of language models: From simple statistical models (n-grams) to early neural networks.
- The breakthrough of the Transformer architecture (conceptual, focusing on its ability to understand context).
- What makes LLMs "Large"? Scale of data, number of parameters, and computational power.
- The concept of "pre-training" on vast amounts of text data.

- **Activity**
  - Contextual Understanding: Provide students with a short passage of text. Ask them to identify words that have different meanings based on context (e.g., "bank" of a river vs. financial "bank"). Discuss how a computer might learn this.

Understanding Large Language Models (LLMs)

- Defining LLMs: Powerful generative AI models capable of understanding and generating human-like text.
- Key capabilities: Text generation, summarization, translation, question answering, creative writing.
- How LLMs are "fine-tuned" for specific tasks after pre-training.
- Introduction to popular LLMs (e.g., ChatGPT, Gemini, Claude - mentioning they are examples of this technology).

- **Activity:**
  - LLM Capabilities Brainstorm: In groups, students brainstorm 3-5 ways they could use an LLM for schoolwork, creative projects, or daily tasks.
  - Group Share: Discuss the most interesting or surprising ideas.

The Art of Prompting: Basic Techniques

- **Lecture (35 min):**
  - What is a "prompt"? The instructions given to an LLM.
  - Why effective prompting is crucial for getting good results.
  - Key principles of basic prompting:
    - **Clarity and Conciseness:** Be direct and avoid ambiguity.
    - **Specific Instructions:** Tell the LLM exactly what you want.
    - **Persona:** Ask the LLM to act as a specific character or expert.
    - **Output Format:** Specify how you want the answer (e.g., bullet points, paragraph, table).

Computer Vision & its Applications

- Introduction to Computer Vision: How computers "see" and interpret images.
- What are Convolutional Neural Networks (CNNs)? Their role in image recognition.
- Conceptual understanding of convolution (feature detection, filters) and pooling (downsampling).
- Why CNNs are effective for images (spatial hierarchy, parameter sharing).

    - Examples of CNN applications: Image classification, object detection (brief mention).
- **Python Coding Activity :**
    - Feature Detection Game: Students are given a simple image (e.g., a drawing of a face) and asked to identify basic "features" (edges, corners, circles) that a computer might look for.
    - Discussion: How do these features help identify what's in the image?
    - Find an app or website that uses image recognition (e.g., Google Lens, a photo tagging app). Describe what it does and how you think a CNN might be involved.

Generative Adversarial Networks (GANs)

- What are Generative Adversarial Networks (GANs)? Creating new, realistic data (images, audio, etc.).
- The "Generator" vs. "Discriminator" concept: A "fake artist" and a "critic" learning from each other.
- How GANs learn to generate realistic outputs through competition.
- Examples of GAN applications: Generating realistic faces, style transfer, creating art.

Introduction to Retrieval Augmented Generation

- Limitations of "pure" LLMs: Hallucinations (making up facts), outdated information, lack of domain-specific knowledge.
- The problem of "knowledge cut-off" in pre-trained models.
- What is Retrieval Augmented Generation (RAG)? Combining LLMs with external, up-to-date knowledge bases.
- Conceptual explanation: The LLM "looks up" relevant information before generating a response.

- **Activity:**
    - "Fact-Checking Challenge": Provide students with a few statements (some true, some false, some debatable). Ask them to identify which ones would be hard for a "pure" LLM to answer accurately without external information.

Ethical Considerations of Large Language Models

- **Bias and Fairness:** How biases in training data can lead to biased or discriminatory outputs from LLMs.
- **Misinformation and Disinformation:** The potential for LLMs to generate convincing but false information.
- **Copyright and Intellectual Property:** Issues around LLMs being trained on copyrighted material and generating new content.
- **Privacy Concerns:** LLMs potentially memorizing and revealing sensitive information from training data.
- The "black box" problem revisited: Difficulty in understanding *why* an LLM made a particular decision.