Module 3: Introduction to Large Language Models

Total Course Hours: 12 hours
Number of Sessions: 12
[Tuition assistance available for eligible students]
Session Length: 55 minutes
(35 min Lecture, 15 min Activity, 5 min Take-Home Lab Assignment Discussion)
Mode: In-person OR Synchronous Online Session options available.
One session Every week starting Sept 9th (Day/Time to be determined)

Session 1: NLP Basics and The Idea of Language Models

- **Lecture (35 min):**
  - 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 (15 min):**
  - 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?
- **Take-Home Lab Assignment (5 min):**
  - 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.

Session 2: From Simple Models to "Large" Language Models

- **Lecture (35 min):**
  - 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 (15 min):**
    - 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.
- **Take-Home Lab Assignment (5 min):**
    - Research one example of a language model or NLP application (e.g., ELIZA, SHRDLU). Briefly describe its capabilities and limitations.
    - Research Paper Review & implement coding for that research paper

Session 3: Understanding Large Language Models (LLMs)

- **Lecture (35 min):**
    - 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 (15 min):**
    - 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.
- **Take-Home Lab Assignment (5 min):**
    - Download Ollama
    - Using Ollama try asking it to summarize a short news article (2-3 paragraphs). Compare its summary to your own.

Session 4: 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.
        - **Role-Playing/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).
- **Activity (15 min):**
    - Prompt Improvement: Students are given a poorly written prompt and work in pairs to rewrite it using the basic principles discussed.
    - Discussion: Share improved prompts and explain the changes.
- **Take-Home Lab Assignment (5 min):**
    - Write three different prompts for an LLM to generate a short story about a specific topic (e.g., "a dog who can talk"). Experiment with different levels of detail and persona.

Session 5: The Art of Prompting: Advanced Techniques

- **Lecture (35 min):**
    - **Few-Shot Prompting:** Providing examples within the prompt to guide the LLM's response.
    - **Chain-of-Thought Prompting:** Encouraging the LLM to "think step-by-step" to solve complex problems.
    - **Iterative Prompting:** Refining responses by asking follow-up questions and providing feedback.
    - The importance of experimentation and iteration in prompting.
- **Activity (15 min):**
    - Complex Problem Prompting: Students are given a multi-step problem (e.g., "Explain how photosynthesis works to a 5th grader, then list 3 fun facts about plants"). They design a prompt using chain-of-thought or few-shot examples.
- **Take-Home Lab Assignment (5 min):**
    - Using an LLM (if available), try to solve a simple math word problem by asking it to "think step-by-step." Analyze if the step-by-step approach helped get a better answer.

Session 6: Computer Vision & its Applications

- **Lecture (35 min):**
    - 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 (15 min):**
    - To be Formalized
- **Take-Home Lab Assignment (5 min):**
    - 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.

Session 7: Diverse Applications of Large Language Models

- **Lecture (35 min):**
    - **Content Creation:** Generating articles, marketing copy, social media posts, creative writing.
    - **Customer Service:** Powering chatbots and virtual assistants for instant support.
    - **Education:** Personalized learning, tutoring, generating study materials.
    - **Coding Assistance:** Generating code snippets, debugging, explaining code.
    - **Research and Information Retrieval:** Summarizing documents, extracting key information.
- **Activity (15 min):**

    - "LLM for Film Making or other": Students choose a specific industry or area (e.g., healthcare, entertainment, law) and brainstorm 2-3 innovative ways LLMs could be used.
    - Group Share: Present ideas and discuss their feasibility and impact.
- **Take-Home Lab Assignment (5 min):**
    - Find an example of an LLM being used in a real-world product or service (e.g., a chatbot on a website, a writing assistant). Describe what it does and how it benefits users.

Session 8: Generative Adversarial Networks (GANs)

- **Lecture (40 min):**
    - 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.
- **Activity (10 min):**
    - "Real or Fake" Challenge: Show students a few real images and a few GAN-generated images (ensure they are appropriate). Ask them to identify which are real and discuss the subtle differences.
    - Discussion: What makes it hard to tell the difference?
- **Take-Home Lab Assignment (5 min):**
    - Research one interesting application of GANs (e.g., creating realistic landscapes, data augmentation for training other models). Briefly describe how it works.

Session 9: Introduction to Retrieval Augmented Generation

- **Lecture (40 min):**
    - 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 (10 min):**
    - "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.
- **Take-Home Lab Assignment (5 min):**
    - Imagine you're asking an LLM a question about a very recent event. Why might its answer be incomplete or inaccurate without RAG?

Session 10: How RAG Works (Conceptually)

- **Lecture (35 min):**
    - The RAG Pipeline (simplified):
        1. **Indexing:** Preparing the external knowledge base (e.g., splitting documents, creating embeddings).
        2. **Retrieval:** Finding the most relevant pieces of information based on the user's query.
        3. **Augmentation:** Passing the retrieved information along with the original query to the LLM.
        4. **Generation:** The LLM uses this context to generate a more accurate and informed response.
    - The role of "embeddings" and "vector databases" (very high-level, conceptual).
    - Benefits of RAG: Improved accuracy, reduced hallucinations, access to real-time data, source citation.
- **Activity (15 min):**
    - Information Retrieval Simulation: Students are given a small "knowledge base" (e.g., a few short paragraphs about a historical event). Given a specific question, they must "retrieve" the most relevant sentences before "generating" an answer.
- **Take-Home Lab Assignment (5 min):**
    - Think of a question you'd ask an LLM that would *require* it to look up information from a specific book or document to answer accurately. Describe how RAG would help.

Session 11: Ethical Considerations of Large Language Models

- **Lecture (35 min):**
  - **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.
- **Activity (15 min):**
  - Ethical Dilemma Discussion: Present a scenario where an LLM's output could have negative ethical implications (e.g., generating harmful stereotypes, providing incorrect medical advice). Students discuss the ethical issues involved.
- **Take-Home Lab Assignment (5 min):**
  - Find a news article or video discussing an ethical concern related to LLMs (e.g., bias, misinformation, copyright). Briefly summarize the issue.

Session 12: The Future of LLMs and Responsible Use

- **Lecture (35 min):**
  - Recap of core LLM concepts: What they are, how they work, their power, and their challenges.
  - Emerging trends: Multimodal LLMs (text + images/audio), smaller, more efficient models, personalized LLMs.
  - The importance of critical thinking and skepticism when using LLM outputs.
  - The role of individuals, developers, and policymakers in ensuring responsible LLM development and deployment.
  - Next steps for continued learning in NLP and LLMs.
- **Activity (15 min):**
  - "My LLM Vision": Students share one positive future application of LLMs they

envision, and one ethical safeguard they believe is most important.
    - Group Discussion: What are your biggest hopes and concerns for the future of LLMs?
- **Take-Home Lab Assignment (5 min):**
    - Identify one specific area related to LLMs (e.g., ethical AI, prompt engineering, a specific application) that you would like to explore further. Find one online resource (website, course, book) that could help you learn more about it.