

AI Summer Camp – Week 3
Speech and Natural Language Processing
Dates: August 10-14

Course Overview

Format: Hybrid (2 in-person days + 3 remote days)

Duration: 5 days, 3hr lecture + 3hr lab per day

Instructor: Dr. Xiaohui Liang, Associate Professor and Associate Chair of Computer Science, UMass Boston

Lab Instructor: Graduate student assistant

Schedule:

	Monday	Tuesday	Wednesday	Thursday	Friday
9 am	Introduction	Warm-up	Warm-up	Warm-up	Warm-up
10 am	Warm-up Preparation	Lecture	Labs	Lecture	Labs
11 am					
12 pm					
1 pm	<i>Break</i>	<i>Break</i>	<i>Break</i>	<i>Break</i>	<i>Break</i>
2 pm	Lecture	Labs	Lecture	Labs	Lecture Presentation
3 pm					
4 pm					
5 pm	Recap	Recap	Recap	Recap	Recap

Detailed Schedule:

Lecture Sessions (Dr. Liang)

Day 1 Theme: How Computers "Read" (Foundations & Tokenization)

Lecture:

- Introduction to Natural Language Processing (NLP). Why is human language hard for computers? (Sarcasm, slang, and context).
- Tokenization & Vectors. Learning how words are turned into numbers.

Labs:

- "The Human Tokenizer" – Students manually break down complex sentences into tokens to see how meanings change. Write a program to do this.

Day 2 Theme: Teaching Machines to Feel & Translate (Sentiment & Translation)

Lecture:

- Sentiment Analysis. How social media companies detect "toxic" comments or happy reviews.
- Machine Translation. From the "Rosetta Stone" to Google Translate. How Neural Networks bridge the gap between English and other languages.

Labs:

- Building a simple "Mood Detector" logic flow and implement a prototype.

Day 3 Theme: The Era of Giants (LLMs & ChatGPT)

Lecture:

- Transformers & GPT. What does the "G," "P," and "T" actually stand for? A look at how GPT-4 predicts the next word.
- Prompt Engineering. Learning how to talk to AI to get the best results (Question Answering & Summarization).

Labs:

- "AI vs. Human" – A competition to see if students can distinguish between a student-written essay and an AI-summarized text.

Day 4 Theme: Machines with Voices (Speech Processing)

Lecture:

- ASR & TTS. How Automatic Speech Recognition (listening) and Text-to-Speech (speaking) work.
- Speaker Identification. The science of "Voice Prints." How your bank knows it's you on the phone.

Labs:

- Experimenting with free TTS tools to see how changing "pitch" and "speed" affects the "emotion" of a robot voice.

Day 5 Theme: The "Deepfake" Debate & Final Showcase (Ethics & Trends)

Lecture:

- Ethics & Societal Impact. Bias in AI, Deepfakes, and the "End of Homework." Who owns the words an AI writes?
- The Horizon. What's next? Multimodal AI (AI that sees, hears, and speaks).

Presentations:

- Students present a 10-minute project for an AI tool that solves a specific problem

Outcomes:

- Demystify AI: Understand that Large Language Models are sophisticated "prediction machines," not magic.
- Think Critically: Evaluate the ethical dangers of voice cloning and algorithmic bias.
- Gain Technical Literacy: Explain the basic pipeline of how a voice command becomes an action.
- Hands-on activity: Design and implement an AI project.

Required Tools

- Laptops with Internet access.
- Access to free-tier AI tools (ChatGPT, Google Gemini, or Hugging Face demos).
- No prior coding experience is required, though a "logic-first" mindset is encouraged!