# AI Summer Camp – Week 4

**Dates:** Aug 17th – Aug 21st (Mon-Fri)

**Instructor:** Prof. Shichao Pei, Assistant Professor, Department of Computer Science, University of Massachusetts Boston

**Format:** Hybrid — 2 in-person days (Mon & Fri) + 3 remote days (Tue–Thu)

## I. Course Overview

1.1 Course Description

This week introduces students to the foundations, creative power, and emerging research questions of Generative Artificial Intelligence (AI). Through a combination of conceptual lectures, hands-on labs, guided experimentation, and team-based projects, students will explore how Large Language Models generate text, how diffusion models synthesize images, and how neural speech systems produce audio. In addition to building creative artifacts, students will learn how to design small experiments, compare system behaviors, and evaluate the strengths and limitations of generative AI. Emphasis is placed on structured prompt design, responsible AI use, critical analysis, and research-minded exploration.

1.2 Overall Goals

By the end of the program, students will:

1. Understand the conceptual foundations of generative AI systems.
2. Apply structured prompt engineering techniques across text, image, and audio tasks.
3. Evaluate limitations of generative AI, including hallucinations, inconsistency, bias, and instability.
4. Design and conduct a small comparative AI experiment.
5. Interpret results using clear criteria, examples, and reflections.
6. Develop and present a multi-modal AI-assisted creative artifact with a research-inspired analysis component.
7. Have the opportunity to develop their projects into research posters or papers.

II. Logistics and Technology

2.1 Required Materials: Google account

2.2 Technologies Used During the Program

- Large Language Models (LLMs)
- Transformer-based architectures (Conceptual Overview)
- Diffusion models (Conceptual Overview)

- Text-to-Speech (TTS) systems
- Google NotebookLM (Lab Use)
- Prompt engineering frameworks

*Who Should Enroll: This program is designed for high school students interested in creativity, technology, media, and research in AI. No prior programming experience required.*

## Syllabus - August 17th, 2026 – In Person

### Day 1: Large Language Models and Text Generation

**Morning: Lecture**

- **Technologies Covered**: Large Language Models (LLMs), Transformer architecture (conceptual overview), Tokenization, and probabilistic text generation.
- **Topics Covered**: Foundations of generative AI, distinction between predictive and generative models, mechanism of next-token prediction, influence of training data, limitations of LLM systems.

**Morning: Lab (Using NotebookLM)**

- **Lab Focus**: Source-grounded text generation and structured summarization.
- **Activities**: Upload short instructor-provided readings, generate summaries and structured outlines, extract key themes and supporting evidence, and compare AI outputs to original texts.
- **Lab Deliverable**: Structured summary (300–400 words), organized outline, short reflection on output reliability.

**Afternoon: Project - AI-Assisted Research Brief + Mini Evaluation Study**

- **Requirement**: Students work in teams to produce a short research-based article and evaluate how prompt design affects the quality of the output.

- **Details**: Curate 3–5 topic-related sources, design at least two prompt strategies, generate a 400–600-word research article, compare outputs for clarity and factual grounding, and document which strategy worked better and why.
- **Project Deliverable**: Completed research brief, prompt documentation, short comparative analysis, and team presentation.

## Syllabus - August 18th, 2026 - Remote

### Day 2: Diffusion Models and Image Generation

**Morning: Lecture**

- **Technologies Covered:** Diffusion models, noise-to-image generation process, text embeddings in image synthesis, and latent space representation (high-level overview).
- **Topics Covered:** Digital representation of images, iterative denoising, role of textual conditioning, stylistic modifiers, coherence, and how prompt specificity influences output.

**Morning: Lab**

- **Lab Focus:** Designing and testing structured prompts for image generation.
- **Activities**: Develop detailed scene descriptions, add stylistic modifiers, compare variations in prompt specificity, refine visual concepts, and observe how small prompt changes alter the generated image.
- **Lab Deliverable**: Three refined image prompts and a written explanation of how prompt structure affected the visual output.

**Afternoon: Project - Themed Visual Portfolio + Prompt Comparison Study**

- **Requirement**: Students create a cohesive AI-generated visual collection while exploring how different prompt strategies influence semantic alignment and artistic consistency.
- **Details**: Produce five conceptually related images with a consistent theme and style, compare at least two prompt designs, and analyze which wording choices best improved coherence, detail, and faithfulness to the intended concept.
- **Project Deliverable**: Digital gallery (slides), prompt documentation, comparative reflection, and presentation.

## Syllabus - August 19th, 2026 - Remote

### Day 3: Style Transformation

**Morning: Lecture**

- **Technologies Covered**: Conditional generation in LLMs, style adaptation mechanisms, conceptual overview of model behavior.
- **Topics Covered**: Definition of stylistic conditioning, creativity versus recombination, and the role of revision in AI-assisted writing.

**Morning: Lab (Using NotebookLM)**

- **Lab Focus**: Style transformation grounded in source material.
- **Activities**: Upload informational or literary text, rewrite content in alternative styles, apply formatting constraints, and compare how well different rewrites preserve meaning while changing tone or structure.
- **Lab Deliverable**: Three stylistic rewrites and a 250-word analytical reflection on style, fidelity, and creativity.

**Afternoon: Project - Multi-Modal Creative Artifact + Style Analysis**

- **Requirement**: Students integrate text and visuals into a cohesive product while examining how style choices influence interpretation.
- **Details**: Create an illustrated short story, educational comic, or concept proposal with visuals; generate text via LLM and images via diffusion tools; document how stylistic choices shaped the final artifact.
- **Project Deliverable**: Completed artifact, 3-slide explanation of workflow, and brief commentary on how style transformation changed audience experience.

## Syllabus - August 20th, 2026 - Remote

**Day 4: Generative Audio and Voice AI Systems**

**Morning: Lecture**

- **Technologies Covered**: Text-to-Speech (TTS) systems, neural speech synthesis, voice generation models, audio diffusion models (conceptual overview), multi-modal generative systems.
- **Topics Covered**: Digital representation of audio signals, neural speech generation processes, text-to-speech pipeline, and how vocal delivery shapes meaning and perception.

**Morning: Lab**

- **Lab Focus**: Producing and evaluating AI-generated spoken content.
- **Activities**: Generate a script using an LLM, refine the script for clarity and pacing, convert the script to audio using TTS, compare voice styles and tonal variations, and reflect on which version sounds most natural and effective.
- **Lab Deliverable**: Final script (200–400 words), generated audio file, and short reflection on voice synthesis quality.

**Afternoon: Project - Final Capstone Project: AI-Generated Multimedia Production**

- **Requirement**: Teams create a complete AI-generated media artifact integrating text, image, and audio, while documenting how iterative refinement improved the final result.
- **Details**: LLM-generated script, diffusion-generated visual element, TTS-generated audio, prompt strategy, and evidence of revision across multiple generations.

- **Project Deliverable:** Complete multimedia artifact, workflow notes, and a short explanation of what the team learned about multi-modal generation.

**Syllabus - August 21st, 2026 – In Person**

**Day 5: AI in Friday - Final Capstone Project**

**Final Capstone Project: AI-Generated Multimedia Production**

**Project: AI-Generated Multimedia Production**

- Teams create a complete AI-generated media artifact integrating text, image, and audio.

**Possible Formats**:

- Podcast introduction
- Audiobook excerpt with cover art
- Fictional news broadcast
- Story narration with visuals

**Requirements**:

- LLM-generated script
- Diffusion-generated visual element
- TTS-generated audio
- Documented prompt strategy
- Evidence of iterative refinement

**Student Discussion / Research Reflection Questions**:

- When was generative AI most useful this week?
- Where did it perform well, and where did it fail?
- What kinds of prompts or workflows produced the most reliable outputs?
- What surprised you most about how these systems behave?
- Can generative AI support human creativity without replacing it?
- If you were to continue this as a research project, what question would you investigate next?

**Capstone Project Presentation**

Student teams present their final multimedia projects, explain their design choices, and share one insight about how generative AI behaved during their experimentation. Parents and guests are invited to attend.

**What Makes This Program Distinct**

Unlike a standard AI creativity camp, this program encourages students to think like both creators and investigators. In addition to building text, image, and audio projects, students will test prompt strategies, compare outputs, examine system limitations, and reflect on how generative AI behaves across different tasks. The result is a learning experience that combines creativity, technical literacy, critical thinking, and an early introduction to research-minded AI exploration.