

Course Syllabus: Artificial Intelligence, Robotics, and Music

1. Course Overview

Unit	Weeks	Main Focus	Module Title	Total Hours
Unit 1	Week 1	Foundations of AI & Intelligent Agents	AI Methodology, Intelligent Agents, and Robotics	5 hours
Unit 2	Week 2	Generative AI	Generative Models and Prompt Engineering	5 hours
Unit 3	Weeks 3-4	Audio Signal Foundations & Analysis	Signals, Sound Analysis, and Synthesis	10 hours
Unit 4	Weeks 5-6	Digital Music and AI	Synthesis, Digital Audio, and Computational Analysis	12 hours
Unit 5	Weeks 7-8	Final Integration	Project Work: AI-Driven Music System	7 hours
TOTAL				39 hours

Unit 1 – Foundations of AI, Intelligent Agents, and Robotics

Timeline: Week 1 | **Title:** AI Foundations and Intelligent Agents

Delivery: Synchronous & Asynchronous MOOC

Main Focus

A theoretical introduction to Artificial Intelligence, ranging from classical systems to the frontiers of Generative AI, with a specific focus on autonomous agent architecture and their interaction with the physical and robotic world.

Learning Objectives

- Distinguish between Weak AI, Strong AI, and AGI.
- Understand the *Perception-Reasoning-Action* cycle in intelligent agents.
- Analyze reinforcement learning paradigms applied to robotics.
- Evaluate Human-Robot Interaction (HRI) dynamics.

Lessons

- **Lesson 1 – Intro to Generative AI** (1h video lecture / 1h self-study): AI vs. Intelligent Agents; Practical Reasoning; Planning; Commitment Strategies.
- **Lesson 2 – Multi-Agent Interaction** (1h video lecture / 2h self-study): Utility and preferences; Nash Equilibrium; Game Theory (Prisoner's Dilemma); Mechanism Design; Auctions; Negotiation.
- **Lesson 3 – Reinforcement Learning** (1h video lecture / 2h self-study): Reward functions; Exploration vs. Exploitation; Autonomous driving and control applications.

Unit 2 – Generative Models and Prompt Engineering

Timeline: Week 2 | **Title:** Generative AI and the Art of Prompting

Delivery: Synchronous (Lecture/Lab) & Asynchronous MOOC

Main Focus

Exploration of modern generative architectures (LLMs and Diffusion models) and development of advanced Prompt Engineering skills.

Learning Objectives

- Understand Transformer-based models and *Attention mechanisms*.
- Master prompting techniques (*Zero-shot, Few-shot, Chain-of-Thought*).
- Analyze "hallucinations" and the ethical-technical limits of generative models.

Lessons

- **Lesson 1 – Architectures** (1h video lecture / 1h self-study): Large Language Models (LLMs); Self-Attention; Discriminative vs. Generative AI.
- **Lesson 2 – Prompt Engineering** (1h video lecture / 2h self-study): Anatomy of a prompt; Context; Input data; Reducing model ambiguity.
- **Lesson 3 – Computer Vision and Diffusion** (1h video lecture / 2h self-study): Latent space and denoising; Multimodal pipeline (Text-to-Image).

Unit 3 – Audio Signal Foundations and Analysis

Timeline: Weeks 3–4 | **Title:** Signals, Sound Analysis, and Synthesis **Delivery:** Asynchronous MOOC | **Total Hours:** 10

3.1 – Audio Signal Fundamentals

- **Lesson 1 (Intro):** Physical nature of sound; Electroacoustic chain (Source → Mic → Preamp → Conversion).
- **Lesson 2 (Sampling):** Digital audio basics; Nyquist–Shannon theorem; Aliasing; Bit depth and quantization noise.

3.2 – Signal Analysis and Processing (6 hours)

- **Lesson 3 (Frequency):** Time domain vs. Spectral domain; Fourier Transform; FFT.
- **Lesson 4 (STFT):** Short-Time Fourier Transform; Windowing; Lossless vs. Lossy compression (PCM, etc.).
- **Lesson 5 & 6 (Filters):** FIR/IIR filters; Z-transform; Low-pass, High-pass, Band-pass, Band-stop.
- **Lesson 7 (Advanced Filters):** Notch, Comb, Peak, All-pass, Shelving filters.
- **Lesson 8 (Effects):** Delay, Reverb, Chorus; Dynamic processors (Compressor, Limiter, Gate).

3.3 – Sound Synthesis Techniques (2 hours)

- **Lesson 9:** Additive, Subtractive, Granular, AM/FM Synthesis.
- **Lesson 10:** Physical Modeling; **Karplus–Strong algorithm**; Hybrid synthesis.

Unit 4 – Digital Music and Artificial Intelligence

Timeline: Weeks 5–6 | **Title:** Synthesis, Digital Audio, and Computational Analysis **Delivery:** Asynchronous MOOC | **Total Hours:** 12

Main Focus

AI-assisted composition, virtual orchestration, and audio-video multimedia projects using professional software.

Lessons

- **Lesson 1 – Assisted Composition:** Evolution from algorithmic systems to AI; Creative control.
- **Lesson 2 – AI as Workflow Support:** Prototyping and material generation.
- **Lesson 3 – Virtual Orchestration:** Orchestral libraries; Expressive control and articulations.
- **Lesson 4 – Frame-by-Frame Method:** Music-to-image synchronization; Markers and cues.
- **Lesson 5 – DAW Workflow:** Nuendo, Cubase, and Logic Pro management; Timecode and delivery.
- **Lesson 6 – AI-based Grading (“TuttiBot”):** Using AI for technical and stylistic feedback on audio tracks.
- **Lesson 7 – Human-Robot Interaction (HRI):** Multimodal interfaces; AI as a cognitive extension of the performer.

Unit 5 – Final Integration

Timeline: Weeks 7-8 | **Title:** Project Work: AI-Driven Music System **Delivery:** Intensive Lab (Synchronous) & Mentoring | **Total Hours:** 7

Project Work Phases

1. **Ideation & Design (2h):** Design Thinking applied to music; Technical Rider drafting.
2. **AI Core & Integration (3h):** Implementing agent logic and advanced prompts; DAW data flow.
3. **Refinement & Rendering (1h):** Mixing, finalization, and AI model fine-tuning.
4. **Presentation (1h):** Live demo and Q&A session.

Possible Project Tracks

- **The Improvising Agent:** Real-time FFT analysis responding with an AI orchestral counterpart.
- **The Semantic Synthesizer:** A text-to-patch interface (e.g., "Create a metallic granular sound").
- **Generative Visual Music:** Video parameters evolving based on audio descriptors (pitch, brightness).

Final Exam Requirements

- **Prototype Delivery:** Functional technical system.
- **Technical Report:** Max 5 pages explaining algorithmic (AI) choices and signal processing.
- **Final Pitch:** Oral presentation of the cross-unit integration.