

From Metal to Humanoid*

Nothing is more important than seeing the sources of invention which are, in my opinion, more interesting than the inventions themselves. —Leibnitz (1646-1716)

Since before the dawn of civilization, our forefathers pondered the elements (earth, water, fire, air, and emptiness) as the fundamental building blocks of all matter. Today, drawing from those same principles, we build systems that move, learn, and walk in our likeness. In this course, we begin not with philosophical inquiry, but with metal: a few magnets, copper wire, some torque, and a mounting plate. From this raw foundation, we build upward (through motion, behavior, and learning) toward a humanoid robot capable of navigating the world and standing beside us, side by side. Metal2Humanoid is a journey through embodied intelligence, where each layer (from mechanical grounding to ethical reflection) is both built and understood from first principles.

Prerequisite: The course is open to undergraduate and graduate students from all disciplines. The only prerequisite is an Introduction to Computer Science or Robotics, or equivalent experience. All core ideas (mechanics, electronics, control, machine learning) are introduced within the course. Students without hardware backgrounds are welcome; we will simulate most systems before building.

Overview: Metal to Humanoid is a builder's journey through robotics, combining engineering and ethics in a single, full-stack course. It starts with the fundamental motor unit (our "NAND gate") and builds toward a functioning humanoid with vision, behavior, and learning capability. Each week adds a new subsystem (mechanical, algorithmic, or cognitive) culminating in a working robot and a design manifesto. Along the way, we'll explore planning, control, imitation learning, and generative policies (like diffusion or VLA models), with reflection on embodiment and the future of human-robot coexistence. This course builds a complete humanoid robot: mechanically, cognitively, and philosophically. Each week yields a reusable module: from actuators to behavior to purpose. Inspired by NAND2Tetris, the learning is constructive, full-stack, and reflective.

When to Take This Course: If you've taken mechanical design, control, or machine learning courses, this class will unify them through concrete systems building. If you're newer to robotics, it offers a clear roadmap: from torque and motion through electronics to intelligence and purpose. Wherever you begin, you'll leave with an integrated perspective and a working humanoid.

Methodology: This is a hands-on, simulation-augmented journey. Each component (actuator, limb, policy, planner) is introduced with a functional specification and tested against physical and behavioral criteria. Students will build most systems in simulation (e.g. PyBullet, Isaac Gym), and implement at least one real-world artifact by the end.

Programming and Build Tools: Modeled by industry standards, mechanical and control systems are modeled via URDF and CAD. Simulations use Python APIs (e.g. Mujoco or Bullet). Learning modules can be implemented using PyTorch, RL libraries, or custom pipelines. Actuator builds use common microcontrollers or motor driver stacks (recommended kits provided).

Testing: Simulated components are evaluated using standard physics engines, performance plots, and scenario benchmarks. Real-world modules are tested for reliability, heat, repeatability, and task completion. Student demos and final projects are reviewed for robustness, creativity, and reproducibility.

Course Grade: 60% projects grades and 40% final examination. **Textbooks:** [Intro to Robotics \(John Craig\)](#), [Modern Robotics \(Lynch & Park\)](#), [Underactuated Robotics \(Russ Tedrake\)](#), [A.I. \(Russell & Norvig\)](#) (*none required*)

*This syllabus describes a one-semester course of ten approximately lectures.

Course Plan* (by week)

Part I: Earth — From Metal to Actuation

We begin with the fundamentals: physical matter, mechanical grounding, and torque. This part focuses on the building blocks of embodied systems, starting with motors and materials. Students will design and build an actuator, then simulate it and explore how components like fasteners, tolerances, gearing, and heat interact in real-world systems. By the end of this unit, you'll have a validated motor module (real or simulated) ready to serve as the “unit cell” for more complex limbs.

1. **Build a Motor:** In project 1 you will design, assemble, and characterize an actuator module.
2. **CAD + Simulation:** In project 2 you will create a digital twin of your motor and simulate performance.

Part II: Water — From Actuation to Motion

With building blocks in hand, we move toward motion. This part focuses on turning actuators into coordinated movement. Inspired by fluidity and adaptability, students will model full humanoids, implement inverse kinematics, and develop controllers that ensure stability and compliance. Using insights from underactuated robotics and trajectory generation, your robot will learn to move like a body in water, gracefully and safely.

3. **Design a Humanoid:** In project 3 you will build a full-body URDF/CAD model of a humanoid.
4. **Move It:** In project 4 you will implement inverse kinematics and trajectory following to control body parts.

Part III: Fire — From Motion to Behavior

Motion becomes behavior when it's structured toward goals. In this part, we focus on task definition, planning, and robustness. Drawing from classical AI and behavior trees, students will design reusable action primitives (like grasp or place) and compose them into full tasks. You'll explore the transition from control loops to intentional action, from force to fire.

5. **Manipulation:** In project 5 you will implement a manipulation skill and demonstrate it repeatedly.
6. **Planning:** In project 6 you will compose skills into a multi-step behavior with task and motion planning.

Part IV: Air — From Behavior to Learning

Air is expansion: the system must now go beyond what we've hand-coded. In this part, we introduce learning from data, through demonstrations, reinforcement, and generative policies. Students will explore behavioral cloning, policy training, and instruction following with VLA or diffusion-based systems. The robot begins to generalize, improving with data and handling situations it wasn't explicitly designed for.

7. **Learning to Walk:** In project 7 you will train a policy for locomotion using RL or imitation techniques.
8. **Diffusion & VLAs:** In project 8 you will use generative or vision-language models to produce behavior.

Part V: Void — From Learning to Humanoid

What happens when machines reflect us? The final part is philosophical: we consider the ethical, social, and symbolic implications of humanoids. We read speculative fiction, robotics manifestos, and labor critiques to explore embodiment, trust, safety, and the post-labor future. The course culminates in a full build or presentation that brings together all parts: body, behavior, learning, and intent.

9. **Reflection:** In project 9 you will write a builder's manifesto; discuss meaning, and the future of labor.
10. **Final Build:** In project 10 you will work as a team to assemble the humanoid and have it take its first steps.

*This syllabus describes a one-semester course of ten approximately lectures.

Resources:

1. Basics:
 - a. <https://mitsoul.org/>
1. Project 1 - Metal to Motor (\$500 a motor)
 - a. <https://youtu.be/GakFB8Tdd98?si=5NrgeUVQFeOJjJPH>
 - i. <https://github.com/aaedmusa/Internal-Cycloidal-Actuator>
 - b. <https://www.youtube.com/watch?v=GakFB8Tdd98&t=277s>
 - c. <https://www.youtube.com/watch?v=OZarwftUh8w>
 - d. <https://www.youtube.com/watch?v=e-tAsM0K7Sg>
 - e. <https://www.youtube.com/watch?v=m4KiFf9MHiw>
 - f. Refer to [Feynmann Lectures on Physics](#) for physics
2. Project 2 - Motor to Actuation
 - a. Pybullet, Mujoco, Issac Sim,CAD tutorials: **TBD**
3. Project 3 - Actuation to Design
 - a. Refer to [HBRN OSH Atlas](#) for designs
4. Project 4 - Design to Motion
 - a. Refer to [Modern Robotics \(Lynch & Park\)](#) for IK
5. Project 5 - Motion to Manipulation
 - a. Refer to [Modern Robotics \(Lynch & Park\)](#) for motion planning
6. Project 6 - Manipulation to Behavior
 - a. Refer to [AI \(Russel and Norvig\)](#) for planning
7. Project 7 - Behavior to Models
 - a. Refer to [Underactuated Robotics \(Tedrake\)](#) for dynamics models
8. Project 8 - Models to Learning
 - a. Refer to [Underactuated Robotics \(Tedrake\)](#) + [TRI](#) + [Pi0](#) for policies
 - b. <https://diffusion.csail.mit.edu/2025/index.html>
9. Project 9 - Learning to Ideas
 - a. "I, Robot" by Isaac Asimov
 - b. "Human Compatible: Artificial Intelligence and the Problem of Control" by Stuart Russell
 - c. "Superintelligence: Paths, Dangers, Strategies" by Nick Bostrom
 - d. "The Technological Singularity" by Murray Shanahan
 - e. "Mind Design II" by collection of authors
 - f. "The Culture" by Iain M. Banks
 - g. "The Hanish Cycle" by Ursula K. Le Guin
 - h. Movies/Pop-Culture:
 - i. Movies: Ex Machina, Terminator, Wall-E, Wild Robot, Iron Giant
 - ii. Themes: Solarpunk/Ecopunk vs Cyberpunk
10. Project 10 - Ideas to Humanoid
 - a. Refer to [HBRN UMOJA](#) for checklist

*This syllabus describes a one-semester course of ten approximately lectures.

PART I Earth: From Metal to Actuation (≈4 hrs)

- **Lecture 1: Build a Motor (≈2 hrs)**
 - 1.1 What a motor really is (electromagnetism intuition) — 15 min
 - 1.2 Torque, speed, power (engineering intuition) — 20 min
 - 1.3 Motor types (DC, BLDC, stepper) — 20 min
 - 1.4 Heat, efficiency, failure modes — 15 min
 - 1.5 Project 1 Overview: build + characterize — 20 min
 - 1.6 Perspective: why actuators are not the only bottleneck — 10 min
 - ≈120 min
- **Lecture 2: What is an Actuator (≈2 hrs)**
 - 2.1 Gearboxes & transmission — 20 min
 - 2.2 Bearings, shafts, tolerances — 15 min
 - 2.3 Encoders & sensing — 15 min
 - 2.4 Drivers, ESCs, wiring — 20 min
 - 2.5 CAD + digital twins — 20 min
 - 2.6 Project 2 Overview — 20 min
 - ≈120 min

PART II Water: From Actuation to Motion (≈4 hrs)

- **Lecture 3: Design a Humanoid Body (≈2 hrs)**
 - 3.1 Kinematic chains & joints — 20 min
 - 3.2 DOFs and morphology tradeoffs — 20 min
 - 3.3 URDF / CAD modeling — 25 min
 - 3.4 Mass, inertia, balance — 15 min
 - 3.5 Project 3 Overview — 20 min
 - ≈120 min
- **Lecture 4: Motion & Control (≈2 hrs)**
 - 4.1 Forward & inverse kinematics — 20 min
 - 4.2 Trajectory generation — 20 min
 - 4.3 PID & impedance control — 25 min
 - 4.4 Stability & compliance — 15 min
 - 4.5 Project 4 Overview — 20 min
 - ≈120 min

PART III Fire: From Motion to Behavior (≈4 hrs)

- **Lecture 5: Manipulation Skills (≈2 hrs)**
 - 5.1 What is a “skill”? — 15 min
 - 5.2 Grasping & contact basics — 20 min
 - 5.3 State machines / behavior trees — 20 min
 - 5.4 Robustness & retries — 15 min
 - 5.5 Project 5 Overview — 20 min
 - ≈120 min
- **Lecture 6: Task and Motion Planning (≈2 hrs)**
 - 6.1 Classical planning intuition — 20 min
 - 6.2 Task & motion planning — 25 min
 - 6.3 Symbolic vs continuous — 15 min
 - 6.4 Failure recovery — 15 min

- 6.5 Project 6 Overview — 20 min
- ≈120 min

PART IV Air: From Behavior to Learning (≈4 hrs)

- **Lecture 7: Learning to Walk (≈2 hrs)**
 - 7.1 Why learning? limits of hand coding — 15 min
 - 7.2 Imitation learning — 20 min
 - 7.3 Reinforcement learning basics — 25 min
 - 7.4 Simulation → reality gap — 15 min
 - 7.5 Project 7 Overview — 20 min
 - ≈120 min
- **Lecture 8: Generative & VLA Policies (≈2 hrs)**
 - 8.1 From policies to distributions — 20 min
 - 8.2 Diffusion policies intuition — 25 min
 - 8.3 Vision-language-action models — 20 min
 - 8.4 Instruction following & grounding — 15 min
 - 8.5 Project 8 Overview — 20 min
 - ≈120 min

PART V Void: From Learning to Humanoid (≈4 hrs)

- **Lecture 9: Reflection & Meaning (≈2 hrs)**
 - 9.1 Why humanoids? — 20 min
 - 9.2 Labor, automation, power — 25 min
 - 9.3 Trust, safety, alignment — 20 min
 - 9.4 Sci-fi as design space — 15 min
 - 9.5 Manifesto Overview — 20 min
 - ≈120 min
- **Lecture 10: Final Integration (≈2 hrs)**
 - 10.1 System integration — 25 min
 - 10.2 Debugging embodied systems — 20 min
 - 10.3 Cost, reproducibility, openness — 20 min
 - 10.4 Demo & evaluation criteria — 15 min
 - 10.5 Final build overview — 20 min
 - ≈120 min