ME 4XX — Physical Intelligence and Embodied AI

Department of Mechanical Engineering Upper-level elective (3rd/4th year) Instructor: Prof. Malcolm A. MacIver (ME / BME; Center for Robotics and Biosystems) Tentative meeting times: 3:30pm-4:50pm MW

Catalog description (short)

Recent advances in artificial intelligence have been driven primarily by large-scale learning from static datasets. However, biological intelligence emerges through continuous interaction among brains, bodies, and environments. This course develops **physical intelligence**—the idea that morphology, mechanics, and the sensorimotor loop perform real computation—and connects it to modern **embodied AI**, where agents that act in physical or simulated environments learn very differently from disembodied ones. Building on the mechanics, dynamics, controls, and programming background of mechanical engineering students, the course moves from "brainless" mechanical intelligence (passive-dynamic walkers, compliant bodies) through active sensing and feedback to reinforcement learning, planning, and the evolutionary origins of cognition. Students implement embodied agents in simulation and build a final agent that senses, acts, and learns in a predator–prey arena modeled on the instructor's research.

Prerequisites: dynamics; programming (Python or equivalent); some exposure to feedback control recommended.

Description (extended, for the catalog committee)

A defining trend across robotics, AI, and the biological sciences is the recognition that intelligent behavior is *embodied*: it arises from the tight, closed loop among a nervous system (or controller), a body with particular mechanics, and a structured environment. A passive walker descends a slope with no actuation and no "brain," yet walks—its dynamics do the work that would otherwise require control. An animal that actively senses in all directions also has evolved a locomotory system that moves in all directions. A reinforcement-learning agent that has to *move to see* solves a fundamentally different problem than one handed a clean state vector.

This elective gives mechanical engineering students a principled, hands-on introduction into this territory. Building on the dynamics and computing skills ME students already have by their 3rd year, the course explores morphological computation, active sensing, reinforcement learning, planning under uncertainty, and the co-evolution of bodies and brains. Throughout, the course draws on case studies from bio-inspired robotics and computational neuroethology, including undulatory swimming and artificial electrosense, ergodic (information-maximizing) active sensing, and naturalistic paradigms that bridge robots and multiple mammalian species. Students leave able to read across the robotics/AI/biology literature, to model an agent as a brain–body–world system, and to build and train embodied agents in simulation.

Who this course is for

Third- and fourth-year ME undergraduates. No biology or machine-learning background is assumed. The course is engineered so that the student's existing strengths—**Newtonian and Lagrangian dynamics, basic fluid mechanics, classical feedback control, linear algebra/ODEs, and a working knowledge of Python**—are used to teach the new material.

Prerequisites

- **Required:** A dynamics course (e.g., ME 314 *Dynamics of Multibody Systems* or equivalent); ability to program (Python

preferred); linear algebra and differential equations.
- **Recommended:** A first course in feedback control (e.g., ME 390 / GEN_ENG 205-4 level) and/or fluid mechanics. Helpful but not required.
- **Not assumed:** Neuroscience, evolutionary biology, probability beyond the basics, or any prior machine learning.

Learning objectives

By the end of the course, students will be able to:

1. **Articulate the embodiment thesis** and explain, with concrete mechanical examples, what it means for a body or an environment to "compute" (morphological computation).
2. **Model an agent as a closed brain–body–world loop**, identifying where dynamics, sensing, control, and learning each enter, and predict how changing the body or environment changes required control.
3. **Analyze locomotion and active sensing** quantitatively—relating morphology and movement to thrust, sensing volume, and information gain—using tools from dynamics, fluids, and optimization.
4. **Formulate a natural embodied task as a reinforcement-learning problem** (MDP, reward, policy, value) and train an agent in a standard simulation framework.
5. **Reason about planning, prediction, and partial observability**, and explain why and when planning beats reactive control.
6. **Connect biological and engineered intelligence**, including the evolutionary pressures (e.g., sensory range and environmental complexity) that may select for planning.
7. **Design, implement, simulate, and document an embodied agent** end-to-end, and communicate the result in writing and a presentation.

Format and workload

- **Meetings:** Two 80-minute sessions per week. Most weeks blend lecture with a short in-class build/analysis. Several sessions are studio/lab days.
- **Computing:** Python with NumPy, Matplotlib, **Gymnasium** (the maintained successor to OpenAI Gym), **Stable-Baselines3**, and a lightweight physics engine (**PyBullet** or **MuJoCo**) or Webots.
- **Effort:** Roughly 3 short computational labs, final project, and AI-conducted oral exams on readings, midterm, and final exam. Expect ~6–8 hrs/week outside class.

Assessment

Component	Weight	Notes
Computational labs (3)	10%	Each builds one piece of the embodied-agent stack.
Oral exams on content	30%	Student examined by socratic AI agent seeded with course content
Midterm (concept exam)	15%	Course content and a set of papers needed for final project.
Final project	15%	Design + simulate an embodied agent; report + presentation.
Final oral exam	30%	Studio days and peer review.

Labs (the agent stack, one piece at a time): - **Lab 1 — The brainless walker.** Simulate a passive-dynamic (or minimally actuated) walker; show how slope, mass distribution, and leg compliance "control" gait. *(Dynamics.)* - **Lab 2 — The body as propulsor.** Model an undulatory swimmer / central-pattern-generator and relate body kinematics to thrust and maneuvering. *(Dynamics + fluids.)* - **Lab 3 — Moving to sense.** Implement an information-seeking (ergodic / proportional-betting) controller that moves a sensor to localize a hidden target under noise. *(Control + optimization.)* - **Lab 4 — Learning to evade.** Train a reinforcement-learning agent in a predator–prey gym environment; study sample efficiency and the effect of partial observability. *(RL.)*

Final project (weeks 7–10). Teams of 2–3 extend the agent stack into a small original investigation of their choosing, such as how obstacle density changes the value of planning, or a morphological-computation demonstration of the student's choosing. Deliverables: a working simulation, a 6–8 page report, and a 10-minute presentation.

Required and recommended texts

- **Primary text:** Pfeifer & Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT Press). Selected chapters.
- **RL reference (selected sections):** Sutton & Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (freely available online).
- **Primary literature:** A reader of ~12–15 papers (provided), drawn from bio-inspired robotics, computational neuroethology, active sensing, embodied AI, and evolution of cognition. Several are from the instructor's group and anchor the weekly case studies.

Weekly schedule

Each week pairs a **concept**, an **ME on-ramp** (the engineering tool used to teach it), and a **case study**.

Week 1 — What is physical intelligence?

Concept: The embodiment thesis; intelligence as a brain–body–world loop; 4E cognition (embodied, embedded, enactive, extended); a short history from cybernetics to embodied AI. **ME on-ramp:** Re-reading a familiar mechanical system as a "controller." **Case study:** McGeer's passive-dynamic walker—stable walking with no actuation and no controller. **Out:** Lab 1 assigned. **Reading:** Pfeifer & Bongard ch. 1–2.

Week 2 — Bodies that compute (morphological computation)

Concept: What it means to offload computation onto mechanics; compliance, passive dynamics, tensegrity, mechanical feedback; the trade between brain and body. **ME on-ramp:** Dynamics and compliance; impedance. **Case study:** Undulatory bodies and compliant structures; how morphology pre-shapes the control problem. **Due:** Lab 1. **Out:** Lab 2.

Week 3 — Locomotion as embodied intelligence

Concept: Movement as the primary output of physical intelligence; gaits, central pattern generators, thrust generation; fluid vs. terrestrial regimes. **ME on-ramp:** Fluid mechanics and rigid-body dynamics; CFD intuition. **Case study:** Ribbon-fin / undulatory swimming—how counter-propagating waves cancel fore–aft thrust to produce heave and precise positioning (MacIver lab, bio-inspired robotics). **Studio day:** swimmer/CPG modeling.

Week 4 — Sensorimotor loops and closed-loop control

Concept: The action–perception loop; reafference; "sensing volumes" and "motor volumes"; closed-loop tracking vs. ballistic strikes. **ME on-ramp:** Classical feedback control; stability of the loop. **Case study:** Predator–prey sensing in weakly electric fish—how the body acts as a sensory antenna and prey capture is adaptive tracking, not a ballistic strike. **Due:** Lab 2. **Out:** Lab 3.

Week 5 — Active sensing: moving to sense

Concept: Sensing as a controllable, energetic act; information vs. energy (bits vs. joules); proportional betting on expected information gain. **ME on-ramp:** Optimization and optimal/ergodic control. **Case study:** Artificial electrosense and **Ergodic Information Harvesting**—predicting how the amplitude and frequency of search movements must change with signal-to-noise ratio. **Studio day:** information-seeking controller.

Week 6 — Learning in embodied agents I: reinforcement learning

Concept: Markov decision processes; reward, policy, value; model-free control; why RL is a natural language for embodied tasks. **ME on-ramp:** From feedback control and cost functions to value functions and policies. **Case study:** The Cellworld **botEvade** task ported to Gymnasium / Stable-Baselines3. **Due:** Lab 3. **Final project teams formed; proposals begin.**

Week 7 — Learning in embodied agents II: the intelligence gap

Concept: Sample efficiency; why animals learn in tens of trials and RL agents need millions; exploration; reward shaping and its hazards. **ME on-ramp:** Hyperparameters as design variables; experiment design. **Case study:** Real mice vs. RL agents in the same arena ("Of Mice and Machines"); "digital PTSD"—over-representing capture events in the replay buffer—and how it produces more animal-like, cautious trajectories. **Final project proposals due.**

Week 8 — Planning, prediction, and partial observability

Concept: Model-based vs. model-free; planning as internal simulation; partial observability; episodic projection (imagining future paths). **ME on-ramp:** State estimation and prediction; what an internal model buys you. **Case study:** Spatial planning and escape from visual predators; why partial observability (predator hidden behind obstacles) creates the highest demand for planning.

Week 9 — Evolution and the origins of physical intelligence

Concept: Why bodies and brains co-evolve; sensory ecology as a driver of cognition; when does planning pay off? **ME on-ramp:** Parameter sweeps and survival landscapes—planning value as a function of visual range and environmental complexity. **Case study:** The **Buena Vista Hypothesis**—the million-fold expansion of visual range at the water-to-land transition (~380 Mya) as a possible cause of the origin of planning. **Studio day:** final-project work + check-ins.

Week 10 — Synthesis and frontiers

Concept: From biology to robots and back; embodied AI in practice; humans-in-the-loop (VR / clinical translation); responsible translation and limits. **Case study:** Cellworld across robots, rodents, and humans. **Final project presentations. Reports due end of finals week.**

Why this course, taught by this instructor

The course is built directly on the instructor's research program at the intersection of evolutionary biology, neuroscience, robotics, and embodied AI—undulatory and electrosensory robotics, the Ergodic Information Harvesting theory of active sensing, the Buena Vista Hypothesis for the origin of planning, and the Cellworld paradigm linking RL agents, mice, and humans. That gives every abstract idea in the syllabus a concrete, first-hand engineering case study, and it positions ME students at a frontier the school has named in its strategic plan (Embodied AI). It also reflects the instructor's prior teaching across ME core, design, and graduate courses, including team-taught and studio formats.
