


Physical AI Education with the SV Robotics AMR Robot Arm

AMR Robot Arm — an educational robot: a self-driving wheeled robot (AMR) carrying a 5-axis robot arm (gripper included) and a depth camera. In the order imitation learning → reinforcement learning and Sim-to-Real → VLA, students train AI themselves in the “RL training lab” on GPU PCs and move it onto the real robot.



AMR Robot Arm — 5-axis arm, gripper, depth camera, omnidirectional wheels

The heart of the university course — reinforcement learning and Sim-to-Real

Students go through the whole process of moving an AI trained in a virtual world onto the real robot themselves. The key is to compare a hand-coded method (rules and calculation) with the AI method on the same task, and to find and close the gap between virtual and real.

Follows the order of robot-learning courses at leading universities abroad (Stanford, Berkeley, ETH and others): imitation learning → reinforcement learning and Sim-to-Real → VLA.

In the reinforcement-learning sessions, students train directly in the RL training lab and compare their models with the reference model we provide.

From a virtual world to the real robot — six steps for students

1 Build the digital twin

Move the real robot's model (URDF), camera position and blocks into the virtual world (Isaac Sim) exactly as they are.

3 Train thousands at once

In Isaac Lab, thousands of virtual robots practice at the same time. Students read progress from learning curves and change methods when it goes badly — for example, first having the AI imitate a hand-coded method and then refining it.

5 Deploy and close the gap

Run the trained AI inside the robot with no internet. Find, record and close the differences between virtual and real, such as movement while the camera tilts down and the arrival criterion.

2 Design the task and the reward

Students design what the AI sees (input), what it decides (output) and what earns it points when done well.

4 Decide which part to give the AI


“Calculation where calculation can solve it, learning where goals conflict” — the approach segment, which must be fast without missing the block, goes to the AI, while search and picking connect to the calculation-based method.

6 Compare under the same conditions

With the same layout, change only the decision method and compare the hand-coded and AI methods. Failed sessions are recorded too — cause and fix — and presented.



Virtual world built with Isaac Sim (digital twin)



Thousands of robots training at once in Isaac Lab

University course — 20-session curriculum

After learning the robot's body, eyes and lidar navigation, students teach the arm and wheels by demonstration (imitation learning), by reward (reinforcement learning) and by words (VLA), and move the AI trained in a virtual world onto the real robot.

Follows the order of robot-learning courses at leading universities abroad (Stanford, Berkeley, ETH and others): imitation learning → reinforcement learning and Sim-to-Real → VLA.

Session	Course	What students learn	Hands-on
1	Foundation Sessions 1–10	What is Physical AI? Starting with the final mission demo Physical AI technology landscape (RL, IL, VLA)	Technology overview; safety rules
2		Robot structure and ROS 2 ROS 2 nodes and topics	Write nodes and topics yourself
3		Arm coordinates and motion Forward and inverse kinematics, coordinate frames	Implement forward and inverse kinematics
4		Measuring distance with the depth camera RGB-D, pixel → 3D coordinates	Implement pixel-to-3D conversion
5		Mapping with lidar (SLAM) Lidar SLAM mapping	Build a SLAM map and tune settings
6		Lidar navigation: reaching a destination Navigation (localization, path planning)	Tune path planning; analyze driving logs
7		Intro to the virtual world (Isaac Sim, digital twin) Digital twin, robot model (URDF)	Modify the robot model and environment
8		First steps in imitation learning: an arm that copies demonstrations Imitation learning (behavior cloning) concepts	Behavior cloning principles; demonstration data structure RL training lab · Reference model
9		First steps in reinforcement learning: the arm learns to reach (virtual → real) RL (state, action, reward); arm-reaching Sim-to-Real	Reward design and training experiments; measure real-world error RL training lab · Reference model
10		AI eyes: finding objects and their 3D position Pretrained object detection + depth-camera 3D position	Small-scale extra training; measure accuracy
11	Advanced Applied AI + Sim-to-Real Sessions 11–20	Naming the target in words (vision-language AI) Open-vocabulary recognition, VLM	Analyze failure cases; prompt design
12		Imitation learning ①: collecting demonstrations by moving the arm by hand Demonstration data collection (kinesthetic teaching)	Collect and organize demonstration data RL training lab · Reference model
13		Imitation learning ②: learning from demonstrations to pick by itself Behavior cloning, ACT policy training	Compare results by number of demonstrations RL training lab · Reference model
14		Sim-to-Real ①: shaking up the virtual world for robust learning Domain randomization (approaching a block)	Set up randomization and train RL training lab · Reference model
15		Sim-to-Real ②: lidar obstacle avoidance with RL (virtual → real) Lidar-based obstacle avoidance policy	Train the avoidance policy and deploy it · Optional add the depth camera to avoid low and high obstacles RL training lab · Reference model
16		Sim-to-Real ③: real-robot block-picking test + imitation vs reinforcement learning Real-robot deployment and evaluation; method comparison	Classify failure types and report RL training lab · Reference model
17		Task planning with a language model LLM task planning, tool calling	Design functions; implement tool calling
18		VLA ①: an AI that understands words and acts VLA (vision-language-action) model inference	Set up the inference server and client RL training lab · Reference model
19		VLA ②: fine-tuning for our robot and comparing three methods VLA fine-tuning; imitation vs RL vs VLA comparison	Fine-tune on the class's own demonstration data RL training lab · Reference model
20		Final mission: a tidy-up robot you instruct by voice — presentation and assessment Integrating LLM planning, perception, navigation and learned skills; presenting results	Experiments to justify each technology choice; short results write-up in conference-presentation format


One session = about 3 hours, adjustable to the school schedule (by consultation). In the reinforcement-learning, imitation-learning and VLA sessions, students train directly in the [RL training lab](#) on GPU PCs and compare performance with the [reference model](#) we provide. The second line under "What students learn" gives the technical terms.

From a virtual trial to real-world use, proven by records

Students put the AI they trained in the RL training lab onto the real robot, compare it with the hand-coded method and the reference model under the same conditions, and finish with a results write-up and presentation.



Virtual trial — AMR Robot Arm in the virtual world



Real-world use — the real AMR Robot Arm

Records students keep — from virtual trial to real-world use

1 Training record

Note the learning curves and why you changed the training settings.

3 Same-condition comparison

Repeat a comparison experiment that changes only one thing, 3 or more times.

2 Virtual-real gap

Find and write down differences in movement and arrival criteria that did not exist in the virtual world.

4 Record failures too

Present the cause and the fix for sessions that failed.

Sample student experiment table — hand-coded vs AI

Stage	Hand-coded Mean of 3 runs (range)	AI method · reinforcement learning Mean of 3 runs (range)	Reading
Find	6.0 s	6.2 s	Same code → no difference
Approach	19.4 s (18.4–19.9)	13.2 s (13.1–13.3)	AI used here · ranges do not overlap → real effect
Align	11.4 s (9.4–13.6)	10.3 s (9.7–11.1)	Ranges overlap → not claimed as an effect
Total	63.3 s	56.3 s	About 88% of the 7.0 s gap comes from Approach

Example built from real experiment records (3 runs each under the same conditions; preliminary comparison).

Evaluation criteria

1. **Fair conditions** — did you change only one thing?
2. **Repeated measurement** — 3 or more runs; did you note the range?
3. **Interpretation** — does the difference exceed the range?
4. **State limits** — how many conditions were there?

Students train the AI themselves on GPU PCs

Students train reinforcement learning, imitation learning and VLA themselves in the “RL training lab” and compare performance with the reference model. Design → virtual test → real-world deployment → data comparison.

University course — train in the lab and compare with the reference model

RL training lab (GPU PC, required for the university course)

Students train the AI themselves with thousands of virtual robots in the “RL training lab” (Isaac Sim · Isaac Lab), then load it onto the real robot.

Reference model (provided for each session)

A benchmark model for comparing performance on the same task as the models students build themselves. Records show where training design makes the difference.

Everything from the robot to the PC, from SV Robotics

AMR Robot Arm

Classroom robot · 2 students per robot (pairs)

GPU PC for Sim-to-Real practice

Required for the university course; specification and number sized to the class plan (consultation)

Textbook · teacher's guide (PDF)

20-session student textbook and instructor's guide

RL class materials

Per-session explainer videos and slides

RL training lab

The environment where students train in the virtual world themselves (details by consultation)

Reference model

A benchmark for comparing performance with the models students build (per session)

Teacher training

The instructor practices with the robot and the virtual world before classes

Technical support · after-sales service

Problem solving and service during classes

The number of robots and GPU PCs is sized to the number of students and the class plan (2 students per robot); details, schedule and cost are agreed by consultation to suit each school. Some AI practice requires the internet and external AI services.

Final mission — a tidy-up robot you instruct by voice

“Put the red block in the blue box.”

① Understand

The AI understands the words and decides what to do

② Approach

Finds the block with the camera and drives to it


③ Pick with the AI arm

Picks it up with the chosen method: imitation, RL or VLA

④ Place

Carries it to the box and drops it in

Designed, demonstrated and presented as an end-of-semester team challenge.



What students can do after one semester

✓ Design the AI that moves a robot

Connect seeing, thinking and moving in code

✓ Apply AI trained in the virtual world to the real robot

Find and close the gap between virtual and real (Sim-to-Real)
Every lesson is verified on the real robot first.

✓ Teach the arm by demonstration, reward and words

Teach the same task with imitation learning, RL and VLA, and compare

✓ Compare with records and present

Explain the hand-coded and AI methods by comparing data

Getting started

① Consultation

Hear the school's situation and goals

② Course plan

Set sessions, level and quantities

③ Robot & PC setup

Install and check in the classroom

④ Teacher training

The instructor practices first

⑤ First class

With technical support