SAKHI

A Memory-Augmented Robotic Lab Assistant for Guided Electronics Instruction and Campus Companionship

1. Project Type

Hardware + Software Design Practicum Project (Embodied AI, Human-Robot Interaction, Edge AI, Long-Term Memory Systems).

2. Domain

Educational and assistive robotics, embodied conversational AI, human-robot interaction, edge AI/LLM systems, retrieval-augmented tutoring, campus community support infrastructure.

3. Background

Every incoming IIT Mandi student passes through hands-on electronics/instrumentation labs (breadboarding, multimeter use, basic circuit debugging) with a small number of Teaching Assistants stretched across many benches at once. A student who gets stuck on something small — a reversed LED, a misread resistor color code, an incorrectly probed multimeter — often waits several minutes for help, or quietly falls behind without asking. Static lab manuals do not adapt to what a specific student got wrong, and once a TA moves to the next bench, whatever guidance was given is not retained anywhere: the same student can make the same mistake the following week with no record of it.

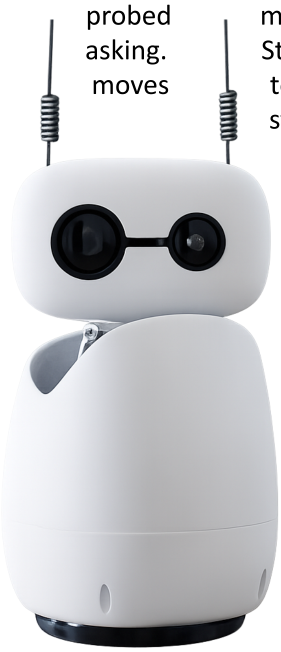

Separately, commercial desktop AI companions (including Hugging Face's **Reachy Mini**) demonstrate that a small expressive robot is an effective, engaging interaction medium — but their own demo applications are stateless: the robot has no memory of who it is talking to beyond the current conversation, and none are grounded in a specific institution's actual course material.

There is a clear gap for a low-cost, desk-resident robotic assistant that combines physical embodiment (so it can literally look at and gesture toward the equipment in front of a student), grounded subject-matter guidance (so its answers match the actual lab manual, not generic internet advice), and persistent per-student memory (so it can recognize a returning student and adapt to their specific history) — deployable directly at an IIT Mandi lab bench, and, as a lighter secondary role, at a hostel common desk as a general campus companion.

4. Problem Statement

Design and build a low-cost, desk-mounted, memory-augmented robotic assistant capable of recognizing an individual student at a shared lab bench, recalling their prior session's progress and common errors, and providing spoken and on-screen step-by-step guidance for basic electronics experiments — while remaining simple and safe enough to sit unattended at an active lab bench — within an estimated budget of INR 55,000–65,000.

5. Target Users

- First- and second-year students in Basic Electronics/Instrumentation labs (SCEE/SMME/ECE-equivalent courses)
- Teaching Assistants and lab instructors currently stretched across many benches per session
- Lab coordinators who currently have no structured record of which concepts/steps students commonly struggle with
- Hostel common areas or personal hostel desks, as a secondary general-purpose memory-augmented companion outside lab hours
- IIT Mandi's Design Practicum and academic-support ecosystem more broadly, as an IIT Mandi Community-track project

6. Core Idea

A small expressive desktop robot — a rotating base, a pan/tilt head, and animated antennas — sits at a shared electronics lab bench. A student identifies themselves with a simple ID-card tap (RFID), and the robot retrieves that student's persistent memory profile: skill level, experiments completed, and specific past mistakes. Using a speech pipeline and an LLM-based tutoring agent grounded in the actual lab manual (via retrieval-augmented generation, so its guidance matches the syllabus rather than generic advice), the robot walks the student through the current step verbally and on a small display, physically turning its head/gaze toward the relevant part of the breadboard or instrument as it explains. A camera-based check can verify simple things like LED orientation or a resistor's color bands before letting the student proceed. At session end, the robot updates that student's memory record and logs an anonymized entry to a lightweight instructor dashboard, so TAs can see, across the whole class, which steps are most commonly getting students stuck. Outside lab hours, the same hardware can be redeployed as a general memory-augmented companion at a hostel common desk.

7. Objectives

1. Build a 3-DOF expressive desktop robot chassis (pan/tilt head, base rotation, animated antennas) sized to sit on a lab benchtop without displacing equipment.
2. Implement a student-identification layer (RFID/ID-card tap, primary) to reliably distinguish which student is at the bench.
3. Implement a persistent, per-student memory system (Mem0-style architecture) storing skill progress, past errors, and preferences across sessions, weeks apart.
4. Implement a voice-and-display guidance pipeline: ASR(Auto. Speech Recognition) → RAG-grounded LLM tutoring agent → TTS(Text-to-Speech) + on-board display, with head/gaze motion directed toward the relevant equipment.
5. Implement at least one camera-based perception check for a common beginner mistake (e.g., LED polarity or resistor color-band reading).
6. Build a lightweight instructor-facing dashboard summarizing session logs and common error patterns per experiment, across all recorded sessions.
7. Demonstrate a complete guided session in which a real beginner student completes a basic circuit with the robot's help, followed by a second session days later in which the robot correctly recalls their prior mistakes and adapts its guidance.

8. Proposed System Architecture

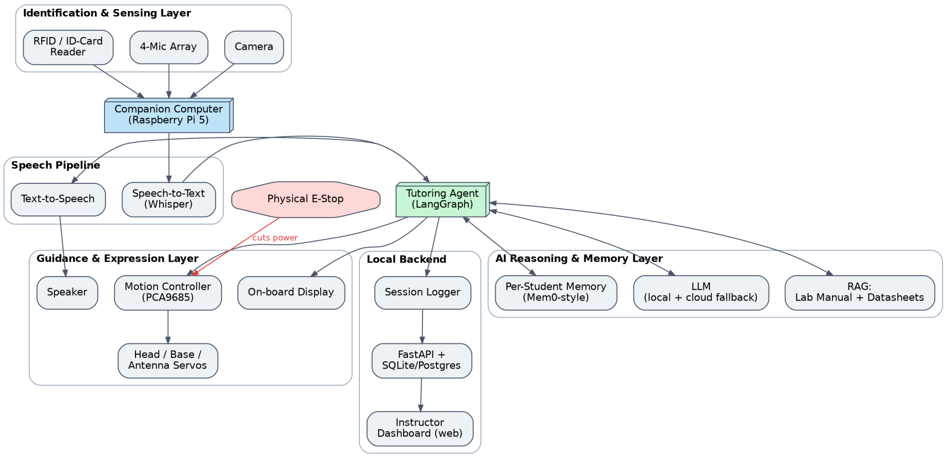

8.1 Mechanical Layer

- Compact rotating base with a pan/tilt head assembly and two animated antennas — footprint kept small so it sits on a lab bench without displacing student equipment.
- Frame 3D-printed in PLA for rapid iteration and near-zero incremental material cost.

8.2 Control Layer

- A PWM driver (I2C-Communication-Protocol) offloads servo timing from the Raspberry Pi and provides a separate, isolated power rail for the servos.
- A physical emergency-stop button cuts servo power immediately
- Status LEDs indicate robot state (idle / listening / thinking / speaking) so students always know when it's ready for input.

8.3 Identification & Sensing Layer

- RFID based ID-card reader as the primary and default student-identification method hich stands as a reliable, near-zero false-positive rate, and avoids the privacy/consent complexity of biometric identification for a Minimum Viable Product.
- 4-microphone array for reliable voice pickup in a noisy shared lab environment.
- Camera (Raspberry Pi Camera Module 3, wide) for perceiving the student's breadboard/instrument setup.

8.4 AI Reasoning & Memory Layer

- Whisper for speech-to-text; a small local LLM (e.g., Qwen/Llama class) for routine responses, with a cloud LLM as fallback for harder tutoring questions.
- The tutoring agent's answers are grounded via retrieval-augmented generation over the actual course lab manual and relevant component datasheets — not generic internet knowledge — so guidance is guaranteed to be syllabus-consistent.
- A Mem0/MemGPT-style persistent memory store, keyed per student ID, records: experiments completed, specific mistakes made, an estimated proficiency level, and preferred pacing — retrieved and updated every session.

8.5 Guidance & Expression Layer

- TTS(Text-to-speech) output through an onboard speaker for spoken step-by-step guidance.
- A small onboard display (OLED) shows the current circuit diagram or step text in sync with what the robot is saying.

- Head/gaze servos physically orient the robot toward the relevant breadboard region or instrument as it explains a step — a direct embodiment advantage over a flat tablet tutorial.

8.6 Perception-Based Error Checking

- At least one camera-based check in the MVP (e.g., verifying LED orientation or reading a resistor's color bands) before letting a student proceed to the next step

8.7 Local Backend & Instructor Dashboard

- A lightweight FastAPI service, running locally on the Raspberry Pi, logs each session (student ID, timestamp, steps completed, errors observed, time-to-completion) to a local SQLite/Postgres database.
- A simple web dashboard aggregates this data for TAs/lab coordinators: which steps most commonly trip students up, average time per experiment, and per-student progress — genuinely useful pedagogical data.

9. Guided Session Logic

1. Student taps their ID card at the bench.
2. Robot retrieves the student's memory profile: proficiency level, last experiment, and past mistakes.

3. Robot greets the student by name and briefly recaps where they left off (e.g., “Last time you had trouble reading resistor color codes — want a quick refresher before we start?”).
4. Student states or selects which experiment/step they are working on.
5. The RAG-grounded tutoring agent retrieves the relevant manual section and explains the current step, orienting its head/gaze toward the relevant equipment.
6. Where a camera-based check is available, the robot verifies the circuit/instrument state before advancing the student, flagging likely errors conversationally rather than just stating “wrong.”
7. At session end, the robot writes an updated memory entry (steps completed, new mistakes/successes) and logs an anonymized entry to the instructor dashboard.
8. In any later session — even weeks apart — the robot recalls this history and adapts its pacing and tone accordingly.

10. Evaluation Metrics

- Voice interaction latency (wake-to-response), target under approximately 3 seconds.
- Student identification reliability (RFID tap success rate).
- percentage of prior-session facts/mistakes correctly recalled in a follow-up session.
- Guidance correctness: percentage of tutoring responses verified consistent with the actual lab manual (a measure of RAG grounding quality).
- false positive/negative rate against a small test set of deliberately mis-wired circuits.

12. Estimated Bill of Materials

Component	Estimated Cost (INR)
Raspberry Pi 5 (8GB) + official power supply + microSD	10,400
Camera module (Pi Camera 3, wide)	2,500
4-mic array (ReSpeaker-class)	3,500
Speaker + I2S amplifier	700
Servos (base rotation, head pan/tilt, antennas) + spares	1,600
PCA9685 servo driver board	450
RFID/ID-card reader module	600
Onboard OLED/TFT display	2,000
Physical E-stop button + cutoff circuit	800
Lazy-susan bearing + 3D-printed chassis (in-house filament)	900
Desk-mount base plate, fasteners, cable management	1,300
Cooling (heatsink/fan for Pi 5)	400
Wiring, connectors, perfboard, misc.	700
Cloud LLM API credits (semester budget)	4,000
Fabrication and testing contingency buffer	3,150

Estimated Total (single unit): ≈ INR 33,000.

13. Novelty

- Most tutoring-robot demos are either purely software (a chatbot on a tablet) or purely mechanical (a generic desktop robot with scripted responses). This project combines a physically embodied, gaze-directing tutor with genuine persistent per-student memory and RAG-grounding in the actual course lab manual — a combination not present in commercial desktop robots such as Reachy Mini, whose own demo applications are stateless.
- The instructor-dashboard layer turns an assistive robot into a pedagogical data-collection tool of direct practical value to lab coordinators, beyond the individual student interaction.
- the robot is designed to live at a shared campus resource (the lab bench) rather than being a personal gadget, with a clear, low-cost path to a secondary deployment as a hostel common-area companion.