Agentic AI Bootcamp — Course Outline (Cohort 2)

krmaazha · 4 Weeks · 8 Live Sessions · Starts Monday, Sep 1

Week 1: Foundations

Session 1 — What Is an Agent?

- The four types of AI agents: reflex, model-based, goal-based, utility-based
- The "intelligence ladder" — how agent sophistication scales
- The Agent-Environment Loop: Perceive → Decide → Act → Feedback
- What actually makes something an "agent" (vs. a script or a chatbot wrapper)
- Centralized vs. distributed AI systems
- **Hands-on:** Setting up your dev environment — cloning a real open-source repo, reading docs in VS Code

Session 2 — Environment Setup & the Perception Module

- Hands-on debugging: fixing real-world errors (port conflicts, toolchain setup) — because real engineering starts with broken environments, not clean ones
- The Perception Module: how agents sense and collect data
- Sensors, data processing, and data cleaning
- Percept vs. Perception — one snapshot vs. the pattern built from many
- Physical vs. digital sensors

Week 2: Reasoning & Applied Code Review

Session 3 — The Reasoning Module

- How agents plan step-by-step
- Using logic and rules to make decisions
- Handling roadblocks and fallback strategies
- The Inference Engine: facts + rules → actions
- The feedback loop between perception and reasoning (it's not one-way)

Session 4 — Code & Documentation Review Methodology

- Reading GitHub issues and documentation like a reviewer, not just a user
- The 3-step review pattern: **Identify the Risk → Explain the Harm → Propose the Fix**
- Turning vague criticism into concrete, mergeable rewrite suggestions
- Applied practice: fixing real wording issues (privacy/consent claims, overpromising compatibility claims)
- Using neutral, factual language for legal/ethical topics (incl. real legal concepts like one-party vs. all-party consent)
- UI/UX bug awareness (modal overflow, edge cases like long filenames)

Week 3: Action & Engineering Practice

Session 5 — The Action Module

- How agents turn decisions into real commands
- Commands and effectors (software and hardware)
- Fine-grained vs. coarse-grained actions

- Coordinating multiple effectors
- Monitoring whether actions actually worked — "an action nobody checks isn't an action, it's a guess"
- Real-world vs. digital actions

Session 6 — Applied Engineering: Issues, Forks, and Fixes

- Understanding real issues and defining clear goals before writing code
- Forking a repository
- Creating and working on a feature branch
- Making code fixes and running tests
- Connecting the fork → branch → fix → test loop back to the monitor-and-adjust cycle from Session 5

Week 4: Git Workflow & Learning Systems

Session 7 — Git & GitHub Pull Request Workflow

- Branching in Git
- Committing and pushing changes
- Opening a pull request (PR)

- Handling review feedback like a professional engineer
- Applying the Session 4 review methodology to real PR feedback

Session 8 — The Learning Module & Course Wrap-Up

- How agents improve over time
- Upgrading perception, reasoning, and action based on feedback
- The parallel between "handling PR review feedback" and an agent's own learning loop — both take an external signal and use it to improve the *next* attempt, not just patch the current one
- Course recap: connecting all four modules (Perceive → Reason → Act → Learn) into one complete agent
- Portfolio review: what students walk away with (GitHub history, working agent code, applied review practice)

What students leave with

- A working understanding of all four core agent modules: Perception, Reasoning, Action, Learning

- Hands-on Git/GitHub experience: forking, branching, PRs, and professional code review
- A real, populated GitHub portfolio — not toy exercises
- Practical debugging experience on a real codebase, not a sanitized classroom example