EdSutra

Curriculum

Grade 12

About age 17–18 · Year 13 in the UK · Class 12 in India · 27 Units · about 16 hours

Every EdSutra Unit rated for Grade 12, in the suggested order. Each Unit is one simulation in which the learner takes on a real job, and each lists what learners will be able to do at the end.

- Three stages in a suggested order. Module 1 comes first; after it, any Module can be opened.
- Every Unit is one simulation. The learner takes a job, climbs Bloom's steps from Remember to Evaluate, and ends on an In real life card. Create happens in the browser-based Labs.
- "I can..." lines are each Unit's learning outcomes. "Ask the class" is a question to open discussion.
- Module 1 is free with no account. Everything else is one $99 purchase.

Stage 1 · Beginner

Foundations: How does a machine learn?

How a machine learns from examples, and why the data matters.

Module 01 · AI Foundations Free

The words AI people use, and how a machine learns from examples and makes a call.

1 Units · about 1 hours · Labs: Computer Vision Lab, Neural Playground

Unit 01 · 60 min · All grades · Role: Codebook Keeper

AI Alphabet A to Z Words

- I can understand 26 foundational AI words in simple language
- I can connect each AI term to an everyday example
- I can use AI vocabulary in short, everyday sentences of their own
- I can distinguish correct from incorrect uses of AI terms in real tools and conversations

Ask the class If an AI tool gets something wrong, who should check it: the tool, the company, or you?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-alphabet.pdf

Module 04 · Data & Knowledge

Where the data comes from, what it leaves out and how it shapes what a model learns.

2 Units · about 1 hours · Labs: Data Explorer Lab

Unit 02 · 25 min · Grades 8–12 · Role: Claims Checker

Validation & Trust The Truth Engine

- I can explain why a verifier gives you an observation, not the truth
- I can tell independent sources apart from copies of one original
- I can choose the right check for a claim, and say when the evidence cannot settle it
- I can judge a verification policy on its observed errors, delays, and effort

Ask the class How much checking is enough before you share something, and when is waiting worth it?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-validation-lab.pdf

Unit 03 · 30 min · Grades 8–12 · Role: Loop Supervisor

Human-in-the-Loop The Feedback Circuit

- I can trace how a person's choice becomes training data for a reward model
- I can predict how a reward model's weights change after an expected or a surprising pick
- I can train a reward model with your own preferences and test it on unseen cases
- I can explain why loops trained by different choosers learn different values
- I can choose where human review belongs in a feedback loop, with evidence and an honest limit

Ask the class If your likes trained an app for a year, what would it learn about you?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-human-in-the-loop.pdf

Stage 2 · Intermediate

Capabilities: What can AI do with pictures, sound and words?

How AI sees, hears, reads, writes and makes new things.

Module 02 · Seeing & Sensing

How a machine turns a picture or a video into something it can recognize.

1 Units · about 1 hours · Labs: Computer Vision Lab

Unit 03 · 45 min · Grades 8–12 · Role: Fusion Officer

Multimodal AI The Fusion Deck

- I can name the kinds of input a multimodal AI can combine: image, text, audio, video and context
- I can combine clues in the Fusion Deck and see which ones an answer needs
- I can trace each part of an answer back to the clue it came from, and spot a clue that added nothing
- I can judge when combining signals helps and when one signal is enough for the task
- I can design a multimodal pipeline: its senses, its goal, and what each sense could get wrong

Ask the class When you're unsure what's happening, which do you check first: what you see, hear or read?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-multimodal-ai.pdf

Module 03 · Language & Text

How a language model reads a sentence, guesses the next word and translates.

3 Units · about 2 hours · Labs: AI Story Canvas Lab

Unit 06 · 40 min · Grades 8–12 · Role: Dial Operator

Sampling & Decoding The Dial Room

- I can explain how language models generate text token by token
- I can describe what temperature controls and its effect on output diversity
- I can compare top-K and top-P (nucleus) cuts, and use a repetition penalty to break loops
- I can design sampling configurations for different use cases (creative vs. factual)

Ask the class For homework facts, would you want a predictable AI or a surprising one? What about story ideas?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-sampling-studio.pdf

Unit 07 · 40 min · Grades 6–12 · Role: Request Writer

Prompting AI The Signal Tuner

- I can name the four features that make a prompt clear: clarity, specificity, role and audience, guardrails
- I can rewrite a vague prompt until it reaches an expert score
- I can predict and test which kinds of noise hurt a prompt most
- I can back your best change with evidence, and explain what a heuristic prompt score can't see

Ask the class Can a score tell you whether a piece of writing is good? What might it miss?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-prompting-ai.pdf

Unit 09 · 25 min · Grades 7–12 · Role: Reading Assessor

Evaluating AI Writing The Response Judge

- I can name the four judging lenses and the failure each one catches
- I can explain how a response's verdict is counted from evidence, not guessed
- I can tag claims, tone cues, loaded words, and clarity signals in a fresh response
- I can diagnose which lens a broken response fails, and cite the exact span that proves it
- I can weigh two rival responses and defend a ship / revise / reject call with evidence

Ask the class Have you read something that sounded certain but turned out wrong? How could you have checked?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-eval-writing.pdf

Module 07 · Creative AI

How generative AI makes pictures, music and text, and whether that counts as creativity.

1 Units · about 1 hours · Labs: Sound and Music Lab

Unit 01 · 5 min · All grades · Role: Studio Technician

Generative AI Start Here

- I can recognize generative AI as producing an output from learned patterns
- I can name examples of generated outputs: words, images, music, and code
- I can recall that learning patterns happens before a prompt steers generation

Ask the class If one prompt can make many different outputs, who decides which one is good?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-generative-ai.pdf

Stage 3 · Advanced

Responsibility: Who answers for what AI does?

How to build with AI, set rules for it and keep a person in charge.

Module 05 · Ethics & Society

Bias, accountability and the rules society sets, worked through as cases to decide.

6 Units · about 4 hours

Unit 01 · 50 min · Grades 8–12 · Role: Standards Examiner

AI Ethics & Bias The Fairness Scale

- I can recall that bias enters through data, algorithm design, and deployment — not just intent
- I can predict which category caused a real documented AI-bias case, then check against the record
- I can operate the fairness scale and read the tilt as a real sum, not a mood ring
- I can defend which mitigations would cut a given bias fastest, citing case evidence

Ask the class Is a tool fair if it's right most of the time but wrong more often for one group?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-ethics.pdf

Unit 02 · 45 min · Grades 7–12 · Role: Duty Registrar

AI Accountability The Accountability Grid

- I can differentiate the responsibilities of developers, companies, users, and regulators
- I can apply transparency, fairness, privacy, and safety principles to specific stakeholder actions
- I can identify distractor 'responsibility-washing' actions that deflect rather than allocate accountability
- I can defend a complete accountability map with reasoning grounded in stakeholder capability

Ask the class If an AI tool made an unfair decision about you, who would you want to answer for it?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-accountability.pdf

Unit 03 · 40 min · Grades 7–12 · Role: Energy Accountant

Environmental Ethics One Message · Two Pathways

- I can explain why an AI chat belongs to a material and technical lifecycle, not only a screen
- I can distinguish recurring use-phase demand from embodied hardware and upstream impacts
- I can compare cloud and on-device pathways across energy, network, cooling, hardware, and task capability
- I can apply proportionality by choosing a responsible pathway without hiding lifecycle trade-offs

Ask the class Should apps show the energy cost of using AI? What would you want to know?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-carbon-cost.pdf

Unit 04 · 35 min · Grades 7–12 · Role: Duty Advisor

Responsible AI Use The Use-Case Compass

- I can name the four value axes a use-case decision can be judged against
- I can explain why the same dilemma can produce different scores for different learners
- I can calibrate personal value weights and apply them to six real dilemmas
- I can compare a categorical gut judgment against a numeric alignment score and locate the exact source of disagreement
- I can defend a specific use-case judgment using the dilemma, the relevant value axis, and context

Ask the class Which task would you never hand to AI, and which would you happily share with it?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-responsible-use.pdf

Unit 05 · 25 min · Grades 7–12 · Role: Independent Review Officer

AI in Society The Civic Ripple

- I can identify the five links between an AI signal and a social outcome
- I can explain why institutions, operating rules, and feedback loops shape AI's impact
- I can calculate how reach, observed miss rate, and human review change unresolved misses
- I can compare deployment plans and locate which assumption creates the difference
- I can defend a launch posture with safeguards matched to the stakes

Ask the class Should a new AI tool start as a small pilot before everyone uses it? Why?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-society.pdf

Unit 06 · 30 min · Grades 8–12 · Role: Governance Route Detective

AI & Regulation The Rule Router

- I can distinguish law, regulation, standard, policy, and voluntary guidance
- I can explain why the same AI capability can follow different governance routes in different contexts
- I can route three authored cases through inspectable conditions and attach relevant dated sources
- I can diagnose missing evidence across sectors and jurisdictions
- I can defend an oversight plan using duties, safeguards, costs, and uncertainty
- I can create a printable governance brief that records unresolved questions rather than claiming legal certainty

Ask the class What's one AI rule you think every country should have? Who would enforce it?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-regulation.pdf

Module 06 · Building with AI

Plan an AI tool, chain its parts into a pipeline and find out what an agent can and cannot do.

4 Units · about 3 hours · Labs: ML Pipeline Builder

Unit 01 · 55 min · Grades 9–12 · Role: Build Lead

Build Your AI The Pipeline Lab

- I can sequence the six pipeline stages and name a failure mode for each
- I can change a deterministic pipeline simulation and trace cause to effect
- I can diagnose why a simulated run underperforms and which stage caused it
- I can judge whether evidence is enough to deploy, revise, or run a limited trial

Ask the class What should a school test before using AI to spot cheating? Who could a mistake hurt?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-build-your-ai.pdf

Unit 02 · 40 min · Grades 9–12 · Role: Autonomy Supervisor

Agentic AI The Autonomy Lab

- I can name the four parts of the agent action loop — Perceive, Reason, Plan, Act
- I can predict how a reactive agent and a planning agent respond when a plan fails
- I can construct a bounded agent with a scoped tool and an explicit stopping condition
- I can trace an agent failure to the part of the loop that broke, citing transcript evidence
- I can choose an autonomy level and guardrail for a deployment and defend the trade-off

Ask the class What's one job you'd let an AI agent do alone, and one it should always ask about?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-agentic-ai.pdf

Unit 03 · 50 min · Grades 9–12 · Role: Flight Director

AI Pipelines Mission Control

- I can name what each pipeline specialist contributes: index, embed, retrieve, ground, generate
- I can predict which failures retrieval can fix — and which it cannot guarantee
- I can operate a RAG system: tune TOP K, inspect scores, and build a cited evidence brief
- I can trace a bad reply to the first broken stage, citing the artifact that proves it
- I can choose and defend a sequential, parallel, or conditional architecture under constraints

Ask the class Would you trust an answer more if it showed its sources? What would you still check?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-pipelines.pdf

Unit 04 · 45 min · Grades 9–12 · Role: Field Engineer

On-Device AI The Edge Runner

- I can name the four parts of a local AI pipeline: camera, compact model, device chip, result
- I can explain why lower-precision models use less memory — and that this can affect quality
- I can run a simulated offline translation and identify what made it possible without a connection
- I can diagnose a failed local run as a memory, compute, or task-fit problem, and choose a valid repair
- I can choose and defend where an AI task should run — on-device, cloud, or hybrid — using named constraints

Ask the class What would you rather keep on your own device instead of sending online, and why?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-on-device-ai.pdf

Module 08 · The Future of AI

AGI, AI safety and what people should keep for themselves as machines improve.

5 Units · about 2.5 hours

Unit 01 · 30 min · Grades 9–12 · Role: Reward Designer

Reinforcement Learning The Reward Lab

- I can name the agent, state, action, reward and policy in a reinforcement learning system
- I can predict how exploration, discounting and learned values shape what an agent does
- I can write a reward function and observe the behavior it actually trains
- I can diagnose reward hacking by comparing what a rule pays with what the designer meant
- I can design a reward function that closes its loopholes and transfers to a new environment

Ask the class Think of a rule at school or home that people "game". How would you rewrite it?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-reinforcement-learning.pdf

Unit 02 · 30 min · Grades 7–12 · Role: Lifecycle Assessor

AI & Climate The Climate Ledger

- I can distinguish power, energy, emissions, water, and materials, and read each with its correct unit
- I can explain how a system boundary, a grid, and the definition of “one prompt” change an estimate
- I can calculate and compare two deployment scenarios from visible assumptions, carrying ranges
- I can trace how rebound, utilization, or a boundary change reverses an apparent advantage
- I can decide whether an AI use is climate-justified against a bounded non-AI baseline, and specify a lower-impact plan

Ask the class When is an AI task worth its energy, compared with doing it another way?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-climate.pdf

Unit 03 · 30 min · Grades 7–12 · Role: Safety Case Reviewer

AI Safety The Safety Case Lab

- I can match hazard, risk, alignment, misuse, robustness, monitoring, containment, and recovery to plain-language examples
- I can predict where a failure path should be interrupted and explain how the four control layers differ
- I can assemble a safety case by linking each claim to supporting evidence and a named control owner
- I can run failure injections and locate the first missing or failed control
- I can decide deploy / pilot / revise / do-not-deploy under explicit constraints using residual paths, evidence gaps, and reversibility
- I can redesign an architecture or operating policy to close a hazardous path and document the remaining uncertainty

Ask the class Who should decide whether an AI tool is safe enough to use at school?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-safety.pdf

Unit 04 · 30 min · Grades 7–12 · Role: Commons Steward

AI & Humanity The Future Commons

- I can distinguish automation, augmentation, substitution, agency, dependency, access, and recourse
- I can predict how one AI deployment changes a task and a human role, then compare stakeholder perspectives
- I can fill a consequence matrix for an authored scenario using provided evidence and explicit assumptions
- I can change value weights and identify which recommendation flips, why, and who bears the cost
- I can deliberate among plausible futures and defend adoption conditions using evidence, distribution, and reversibility
- I can draft a Human Choice Charter: protected decisions, disclosure, opt-out, skill preservation, access, appeal, review triggers

Ask the class Which decisions about your life should always stay with a person, never a machine?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-ai-humanity.pdf

Unit 05 · 30 min · Grades 7–12 · Role: Capability Auditor

AGI & Superintelligence The Capability Observatory

- I can distinguish narrow AI, generality, transfer, autonomy, AGI, and superintelligence, and separate definitions from predictions
- I can predict whether a high score in one domain transfers to a changed task
- I can run authored systems through a transparent multi-domain test battery and record pass / partial / fail with evidence
- I can inspect spiky capability profiles and diagnose memorization, tool dependence, brittle transfer, or missing embodiment/context
- I can audit a public-style 'approaching AGI' claim for definition, evidence breadth, counterexamples, and uncertainty
- I can design a capability test suite with task diversity, hidden transfer cases, failure criteria, and a limits statement

Ask the class What would an AI have to do to convince you it was truly general? What wouldn't count?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-agi.pdf

Module 09 · Intelligence Management

Match AI tools to the work they do best, check their results and keep a person responsible for the outcome.

4 Units · about 2 hours · Labs: Open Project Canvas

Unit 01 · 35 min · Grades 8–12 · Role: Duty Controller

What Is Intelligence Management? The Intelligence Control Room

- I can identify five useful AI specialist roles and the limits of each one
- I can match a task to a specialist by the output and evidence the task requires
- I can assign specialists to a provided multi-stage workflow by its required handoffs
- I can trace a failed workflow to a capability mismatch, lost context, or skipped verification gate
- I can judge whether an AI handoff should be approved, returned for revision, or escalated to a person

Ask the class If you used three AI tools on one project, how would you know where a mistake came from?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-what-is-intelligence-management.pdf

Unit 02 · 25 min · Grades 8–12 · Role: Override Officer

Human Oversight The Override Station

- I can distinguish human in, on, and out of the loop — and review, escalation, override, and appeal
- I can judge whether a reviewer has the information, time, context, authority, and recourse to change an outcome
- I can configure an escalation policy and route a full case stream through a review queue
- I can trace a missed case, a late review, and a rubber-stamp back to the rule or interface that caused it
- I can defend a policy on its exact missed, late, and unnecessary counts rather than on intention

Ask the class Is it always safer to have a person check everything? What could go wrong?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-human-oversight.pdf

Unit 03 · 25 min · Grades 7–12 · Role: Handoff Checker

Delegation & Verification The Proof Gate

- I can name the parts of a handoff: goal, task, context, constraint, deliverable, acceptance check, evidence and escalation
- I can predict what an agent does with a blank line on its job card, and which check (if any) catches it
- I can write a complete job card, and make the human call on briefs that pass every automatic check
- I can diagnose false acceptance, false rejection, gaming the test, stale context and a fake score on a real gate
- I can choose which job card to trust, backed by facts from your own runs

Ask the class What's something you'd hand to a helper, and how would you check it was done well?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-delegation-verification.pdf

Unit 04 · 30 min · Grades 7–12 · Role: Mission Director

Agent Orchestration The Orchestration Room

- I can identify specialist, task, dependency, handoff, shared context, coordinator, and verification gate
- I can predict which sequential, parallel, or conditional workflow shape completes a mission reliably
- I can decompose an authored mission, assign specialists, define context packets, and execute the deterministic workflow
- I can inspect failed traces and diagnose capability mismatch, missing context, stale version, dependency race, or skipped verification
- I can compare orchestration plans for correctness, latency, cost/rework, observability, and human escalation
- I can design and successfully run an orchestration plan for a fresh mission with explicit interfaces and recovery paths

Ask the class In a group project, what goes wrong when nobody coordinates? Would an AI coordinator help?

Teaching guide: edsutra.com/downloads/curriculum/teaching-guides/edsutra-teaching-guide-agent-orchestration.pdf