HUMANEDGE ACADEMY™

Human-AI Supervisor Curriculum

A four-level pathway for responsible AI literacy, practice, supervision and leadership

I
Foundations

II
Practitioner

III
Supervisor

IV
AI Leadership

Human Authority → AI Assistance → Human Verification → Human Decision

---

Curriculum framework, lesson plans, assignments, assessments and certificate templates

A signature pathway for responsible human-AI work

The Human-AI Supervisor Curriculum turns one principle into teachable practice: **AI may assist with the work. Humans remain responsible for the judgment.**

What this document provides

- A four-level, stackable learning pathway
- Six program competencies and nine supervisory moves
- Sixteen module lesson plans
- Four applied Supervisor Challenges
- A common performance rubric
- Certification standards and printable templates
- Instructor routines, learner records and implementation guidance

Recommended default

80 guided hours + 32 applied hours

Each level includes four five-hour modules and an eight-hour applied challenge. Providers may extend practice time without changing the performance standard.

Credential model

Each level earns a stackable academy-issued certificate. Level III carries the title **Certified Human-AI Supervisor**.

> “Direct it. Check it. Question it. Decide.”

Status: Draft curriculum framework for institutional review and pilot delivery. This is an academy-issued, non-degree pathway; employers, schools and regulators determine how they recognize it.

Curriculum at a glance

Program foundation

1. Purpose and philosophy
2. Architecture and pacing
3. H.U.M.A.N.E.D.G.E. framework
4. Competencies and evidence

Four certification levels

5. Level I: Foundations
6. Level II: Practitioner
7. Level III: Human-AI Supervisor
8. Level IV: AI Leadership

Assessment and implementation

9. Assessment blueprint and rubric
10. Certificate pathway
11. Facilitator and audience guidance
12. Responsible-use policy
13. Printable certificates
14. Learner tools and rollout plan

Design principle

The pathway measures judgment made visible. Learners preserve their goals, instructions, revisions, checks, decisions and reflections so an assessor can see how they supervised the technology.

Use in multiple settings

Education
Middle school through college and career and technical education.

Workforce
Onboarding, upskilling, supervision and professional development.

Organizations
Government, schools, corporations, restaurants, small businesses and community groups.

Responsible collaboration—not competition or surrender

HumanEdge Academy prepares learners not simply to become users of artificial intelligence, but Human-AI Supervisors.

The human contribution

Humans establish purpose, values, ethics and acceptable risk. They bring context, empathy, creativity, lived experience and responsibility for consequences.

The AI contribution

AI systems may analyze information, generate ideas, recognize patterns, automate steps and assist people in producing useful work.

The supervisory obligation

Capability does not equal authority. A person or accountable institution must decide what may be delegated, what evidence is required and whether a result should be used.

The educational promise

Learners practice directing, checking, questioning and deciding until responsible supervision becomes a habit rather than an afterthought.

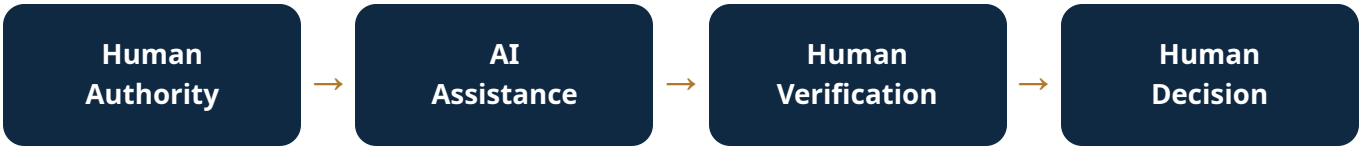

> Technology may become increasingly capable. Education must ensure that capability is accompanied by increasing human wisdom, responsibility and leadership.

Four stackable levels build one supervisory practice

Level	Learner shift	Guided / applied	Primary performance
I Foundations	From trusting output to checking it	20 / 8 hours	Reliable Brief
II Practitioner	From prompting to controlled workflow	20 / 8 hours	Responsible Work Product
III Supervisor	From individual use to operational oversight	20 / 8 hours	Operational Workflow
IV AI Leadership	From workflow oversight to institutional authority	20 / 8 hours	Operating Charter & Pilot

Standard module rhythm

30 min launch and diagnose
45 min concept and demonstration
60 min guided practice
120 min applied lab
45 min debrief and assessment

Flexible delivery

- Five consecutive days
- Weekly five-hour sessions
- Two 2.5-hour meetings per module
- Hybrid instruction plus supervised lab

Do not compress verification, feedback or authentic performance to preserve seat time.

Entry and progression

Level I has no prerequisite. Higher levels require the preceding credential or a documented prior-learning review using the same performance rubric. Learners may contextualize assignments to their sector while preserving the common evidence requirements.

H.U.M.A.N. begins with purpose and active supervision

H

Human Purpose

Define the goal, problem, values, stakeholders and desired outcome before engaging AI.

U

Understand the Task

Determine what requires human judgment and what can appropriately be delegated or assisted.

M

Match the AI

Select a suitable system for research, writing, analysis, design, coding, communication or automation.

A

Assign & Instruct

Provide objectives, context, constraints, standards, evidence and expected deliverables.

N

Navigate & Supervise

Monitor the work, question assumptions, preserve control and avoid automatic acceptance.

Instructional move: Ask learners to label every piece of submitted evidence H, U, M, A or N. If a step has no evidence, it may have been skipped.

E.D.G.E. converts output into accountable action

E

Evaluate

Check accuracy, sources, reasoning, bias, safety, quality, originality, privacy and ethical implications.

D

Decide

The human makes the final decision—not the AI—and records the basis for that decision.

G

Guide & Improve

Correct mistakes, refine instructions, strengthen controls and redirect the system when necessary.

E

Exercise Accountability

Accept responsibility for decisions and outcomes produced through human-AI collaboration.

> The AI can recommend, create, analyze and execute authorized tasks. The human remains responsible for deciding whether the result should be accepted and used.

Evidence of evaluation

Source checks, calculations, comparison tests, bias review, privacy screening, accessibility review and expert escalation.

Evidence of accountability

Named decision owner, signed approval, documented conditions, uncertainty statement, monitoring trigger and correction record.

Six competencies define AI literacy at HumanEdge Academy

Competency	Learners demonstrate	Typical evidence
AI Literacy	Explain capabilities, limits and uncertainty	Boundary map; concept explanation
AI Direction	Set effective goals, context, constraints and standards	Instruction brief; revision log
AI Evaluation	Detect errors, weak reasoning and unsupported claims	Verification table; test suite
AI Ethics	Consider fairness, privacy, transparency and consequences	Risk review; disclosure; control
AI Collaboration	Divide work appropriately between people and systems	Workflow map; checkpoint record
AI Leadership	Maintain authority, accountability and organizational learning	Decision memo; governance charter

Universal learning outcomes

- Define purpose before selecting technology.
- Use the minimum capability and access needed.
- Verify according to consequence, not convenience.
- Communicate uncertainty and AI involvement.
- Escalate when evidence or authority is insufficient.
- Own the final decision and its effects.

What mastery is not

Fast prompting, fluent output, tool loyalty, technical vocabulary or an impressive final product do not by themselves establish supervisory competence.

The learner's supervision record is the spine of every assignment

Every major task produces two things: a work product and a record of the human judgment that governed it.

1. Plan

Purpose, stakeholder, constraints and human-only decisions.

2. Direct

Tool choice, instructions, permissions and success criteria.

3. Verify

Checks, sources, tests, corrections and unresolved uncertainty.

4. Decide

Approve, modify, reject, escalate or stop—and explain why.

Required evidence habits

- Date and label meaningful AI interactions.
- Preserve prompts or instruction briefs.
- Identify which content came from AI.
- Record corrections and reasons.
- Protect personal, confidential and restricted data.
- Keep the human decision visible.

Authenticity checks

- Short oral defense
- Live revision or verification
- Scenario transfer task
- Peer challenge questions
- Instructor observation
- Reflection on a rejected output

Assessment rule: The grade is based primarily on how effectively the learner supervised the AI, not on what the AI produced.

Level I — Foundations

I

HUMAN-AI SUPERVISOR FOUNDATIONS CERTIFICATE

Use AI with purpose, healthy skepticism and a repeatable verification habit.

20 guided + 8 applied hours | Four modules | One Supervisor Challenge

Designed for

Beginning learners, middle/high school, career exploration, onboarding and community education

Level outcomes

- Explain what generative AI can and cannot reliably do.
- Separate human judgment from tasks that AI may assist.
- Write clear instructions with goals, limits and success criteria.
- Verify claims, disclose AI use and make a reasoned final decision.

Performance standard

Learners must produce observable evidence of purpose, direction, verification and decision. A polished AI output without a supervision record is incomplete.

Instructor checkpoint

Confirm tool access, data boundaries, accessibility needs and the local escalation route before instruction begins.

Module	Lesson focus	Assignment
F1	Human authority and AI literacy	AI Boundary Map
F2	Define purpose and decompose work	Purpose-and-Task Brief
F3	Direct AI with clear instructions	Three-Prompt Improvement Log
F4	Verify before use	Claim Verification Lab

Human authority and AI literacy

5 guided hours | Human Purpose + Exercise Accountability | Assignment: AI Boundary Map

Learning outcome

Learners can explain why AI output is a proposal, not a decision.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Compare two plausible AI responses and mark what must be checked.

45 MIN • MINI-LESSON & MODEL

Model patterns, probability, limitations, hallucinations, bias and the difference between capability and authority.

60 MIN • GUIDED PRACTICE

Sort twelve tasks into human-only, AI-assisted and context-dependent categories; defend two disputed choices.

120 MIN • SUPERVISOR LAB

Create an AI Boundary Map for a school, workplace or community scenario, naming decisions humans retain.

45 MIN • DEBRIEF & ASSESSMENT

Deliver a two-minute explanation using the HumanEdge Rule and answer one challenge question.

Instructor preparation

Prepare two contrasting AI outputs and a task-card set. Use no personal or confidential information.

Evidence to collect

Boundary Map, annotated outputs, reflection and exit explanation.

Specific assignment: AI Boundary Map

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Define purpose and decompose work

5 guided hours | Human Purpose + Understand the Task | Assignment: Purpose-and-Task Brief

Learning outcome

Learners can define success before opening an AI tool.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Diagnose why a vague request produces a polished but unusable answer.

45 MIN • MINI-LESSON & MODEL

Introduce purpose, stakeholder, values, constraints, risk and task decomposition.

60 MIN • GUIDED PRACTICE

Rewrite a broad goal into a human decision, AI-assistable steps and verification checkpoints.

120 MIN • SUPERVISOR LAB

Build a Purpose-and-Task Brief for a real class, job or community need.

45 MIN • DEBRIEF & ASSESSMENT

Peer-review briefs against clarity, boundaries and measurable success criteria.

Instructor preparation

Select three age- and sector-appropriate scenarios; prepare the brief template.

Evidence to collect

Completed brief, task map and peer-review notes.

Specific assignment: Purpose-and-Task Brief

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Direct AI with clear instructions

5 guided hours | Assign & Instruct + Navigate & Supervise | Assignment: Three-Prompt Improvement Log

Learning outcome

Learners can direct AI through context, constraints, standards and revision.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Run one underspecified prompt and identify what the system had to guess.

45 MIN • MINI-LESSON & MODEL

Model an instruction brief: objective, context, audience, inputs, boundaries, format and quality standard.

60 MIN • GUIDED PRACTICE

Improve the same request across three rounds; change one instruction variable at a time.

120 MIN • SUPERVISOR LAB

Produce a useful draft and document why each revision improved or weakened it.

45 MIN • DEBRIEF & ASSESSMENT

Conduct a prompt conference in which the learner explains—not merely shows—the revisions.

Instructor preparation

Provide a safe sandbox task and examples of strong and weak instructions.

Evidence to collect

Three prompts, three outputs, comparison notes and final instruction brief.

Specific assignment: Three-Prompt Improvement Log

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Verify before use

5 guided hours | Evaluate + Decide | Assignment: Claim Verification Lab

Learning outcome

Learners can test accuracy, sources, bias and suitability before use.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Find confident language that lacks evidence in a sample output.

45 MIN • MINI-LESSON & MODEL

Teach lateral verification, source quality, calculations, date checks, quotation checks and bias questions.

60 MIN • GUIDED PRACTICE

Verify five claims using independent sources or direct calculations; label verified, uncertain or false.

120 MIN • SUPERVISOR LAB

Correct an AI-produced fact sheet and write an approval, revision or rejection decision.

45 MIN • DEBRIEF & ASSESSMENT

Use the verification checklist and require evidence for every status label.

Instructor preparation

Create a seeded fact sheet with accurate, outdated, unsupported and false claims.

Evidence to collect

Verification table, corrected fact sheet, sources and decision statement.

Specific assignment: Claim Verification Lab

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Supervisor Challenge — The Reliable Brief

Use AI to help produce a one-page workplace, school or community information brief. The final product matters, but the supervision record is the primary evidence.

Required deliverables

- Purpose-and-Task Brief
- AI-use and revision log
- Claim-by-claim verification table
- Final brief with AI-use disclosure
- Human decision memo: accept, modify or reject

What assessors look for

- Purpose and human authority are explicit
- Instructions include context and limits
- Every material claim is checked
- Final decision is explained
- No private data is entered

Eight-hour challenge sequence

HOUR 1 • SCOPE
Confirm purpose, stakeholder, authority boundary, success criteria and prohibited data.

HOURS 2–3 • DIRECT
Select the tool, document instructions and run the first supervised attempt.

HOURS 4–5 • VERIFY
Test claims, calculations, bias, privacy, usability and failure cases in proportion to risk.

HOURS 6–7 • IMPROVE
Correct defects, strengthen controls and complete the evidence package.

HOUR 8 • DECIDE & DEFEND
Approve, modify, reject, escalate or stop. Present the decision in a five- to eight-minute defense.

Scoring: Common rubric, 24 points. The work product itself is supporting evidence; it cannot compensate for unsafe practice or absent human judgment.

Level II — Practitioner

II

HUMAN-AI SUPERVISOR PRACTITIONER CERTIFICATE

Produce dependable work by selecting tools deliberately, directing iterations and documenting quality control.

20 guided + 8 applied hours | Four modules | One Supervisor Challenge

Designed for

High school, college, career and technical education, entry-level professionals and small-business teams

Level outcomes

- Match AI systems to task demands and risk.
- Iterate instructions using testable quality standards.
- Identify privacy, bias, intellectual-property and safety concerns.
- Complete an end-to-end supervised work product with an audit trail.

Performance standard

Learners must produce observable evidence of purpose, direction, verification and decision. A polished AI output without a supervision record is incomplete.

Instructor checkpoint

Confirm tool access, data boundaries, accessibility needs and the local escalation route before instruction begins.

Module	Lesson focus	Assignment
P1	Match tools to tasks	Tool Selection Memo
P2	Iterative direction and collaboration	Controlled Iteration Study
P3	Quality, bias, privacy and responsible use	Risk and Quality Review
P4	Build a supervised workflow	End-to-End Workflow Build

Match tools to tasks

5 guided hours | Understand the Task + Match the AI | Assignment: Tool Selection Memo

Learning outcome

Learners can choose a tool based on task fit rather than popularity.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Compare one task completed by two tool types and identify trade-offs.

45 MIN • MINI-LESSON & MODEL

Examine modality, context window, retrieval, privacy settings, cost, access, integration and failure modes.

60 MIN • GUIDED PRACTICE

Score candidate tools against a common requirement matrix without entering sensitive data.

120 MIN • SUPERVISOR LAB

Recommend a primary tool, fallback and human-only alternative for a defined workflow.

45 MIN • DEBRIEF & ASSESSMENT

Defend the recommendation in a five-minute review panel.

Instructor preparation

Prepare current tool capability sheets; verify local access and policy before class.

Evidence to collect

Requirement matrix, selection memo and oral defense notes.

Specific assignment: Tool Selection Memo

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Iterative direction and collaboration

5 guided hours | Assign & Instruct + Guide & Improve | Assignment: Controlled Iteration Study

Learning outcome

Learners can improve quality through deliberate, documented revisions.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Show how changing audience, constraints or criteria changes the result.

45 MIN • MINI-LESSON & MODEL

Teach decomposition, examples, critique loops, role framing, comparison and stopping rules.

60 MIN • GUIDED PRACTICE

Run controlled A/B instructions and score outputs with a shared quality checklist.

120 MIN • SUPERVISOR LAB

Complete three purposeful iterations and justify the stopping point.

45 MIN • DEBRIEF & ASSESSMENT

Evaluate whether revisions target actual defects rather than merely add length.

Instructor preparation

Choose a work product with visible quality dimensions and a fixed evidence set.

Evidence to collect

Iteration log, scored outputs, final selection and rationale.

Specific assignment: Controlled Iteration Study

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Quality, bias, privacy and responsible use

5 guided hours | Evaluate + Exercise Accountability | Assignment: Risk and Quality Review

Learning outcome

Learners can identify material risks and apply proportionate controls.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Trace how a harmless-looking prompt can expose confidential or personal data.

45 MIN • MINI-LESSON & MODEL

Cover data minimization, consent, bias, accessibility, copyright awareness, transparency and sector rules.

60 MIN • GUIDED PRACTICE

Use a risk matrix to rate impact and likelihood; add a control and escalation path.

120 MIN • SUPERVISOR LAB

Audit a simulated AI-enabled process and redesign its riskiest step.

45 MIN • DEBRIEF & ASSESSMENT

Require a justified go, revise, escalate or stop recommendation.

Instructor preparation

Adapt scenarios to the audience; distinguish instruction from legal or compliance advice.

Evidence to collect

Risk matrix, revised process, disclosure language and decision.

Specific assignment: Risk and Quality Review

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Build a supervised workflow

5 guided hours | Navigate & Supervise + Decide | Assignment: End-to-End Workflow Build

Learning outcome

Learners can place human checkpoints around AI-assisted work.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Map a familiar task and locate where errors could become consequences.

45 MIN • MINI-LESSON & MODEL

Introduce inputs, transformations, outputs, owners, checkpoints, evidence and approval gates.

60 MIN • GUIDED PRACTICE

Storyboard a workflow; test one normal case and one failure case.

120 MIN • SUPERVISOR LAB

Run the workflow, preserve an audit trail and revise the weak checkpoint.

45 MIN • DEBRIEF & ASSESSMENT

Observe performance and review the evidence package, not just the final artifact.

Instructor preparation

Provide a workflow canvas and a safe set of source materials.

Evidence to collect

Workflow map, execution log, test results, corrected output and approval record.

Specific assignment: End-to-End Workflow Build

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Supervisor Challenge — The Responsible Work Product

Create a useful professional deliverable—such as a customer guide, training aid, proposal or communication plan—through a documented human-AI workflow.

Required deliverables

- Tool selection memo
- Instruction and iteration record
- Risk and privacy review
- Verified final work product
- Decision and disclosure statement

What assessors look for

- Tool choice fits the task
- Revisions respond to evidence
- Risks receive controls
- Sources and calculations are checked
- Human approval is explicit

Eight-hour challenge sequence

HOUR 1 • SCOPE
Confirm purpose, stakeholder, authority boundary, success criteria and prohibited data.

HOURS 2–3 • DIRECT
Select the tool, document instructions and run the first supervised attempt.

HOURS 4–5 • VERIFY
Test claims, calculations, bias, privacy, usability and failure cases in proportion to risk.

HOURS 6–7 • IMPROVE
Correct defects, strengthen controls and complete the evidence package.

HOUR 8 • DECIDE & DEFEND
Approve, modify, reject, escalate or stop. Present the decision in a five- to eight-minute defense.

Scoring: Common rubric, 24 points. The work product itself is supporting evidence; it cannot compensate for unsafe practice or absent human judgment.

Level III — Human-AI Supervisor

III

CERTIFIED HUMAN-AI SUPERVISOR

Design and oversee reliable workflows with clear authority, permissions, checkpoints and escalation.

20 guided + 8 applied hours | Four modules | One Supervisor Challenge

Designed for

College and CTE learners, managers, team leads, trainers and professionals who oversee AI-enabled work

Level outcomes

- Design multi-step workflows with accountable human ownership.
- Set permissions and escalation rules for assistants and agents.
- Build quality assurance tests and audit-ready evidence.
- Lead incident response and corrective improvement.

Performance standard

Learners must produce observable evidence of purpose, direction, verification and decision. A polished AI output without a supervision record is incomplete.

Instructor checkpoint

Confirm tool access, data boundaries, accessibility needs and the local escalation route before instruction begins.

Module	Lesson focus	Assignment
S1	Workflow architecture and human checkpoints	Authority-and-Checkpoint Map
S2	Agents, permissions and escalation	Agent Permission Plan
S3	Quality assurance and auditability	AI Quality Test Suite
S4	Team supervision and incident response	Incident Tabletop

Workflow architecture and human checkpoints

5 guided hours | Human Purpose through Human Decision | Assignment: Authority-and-Checkpoint Map

Learning outcome

Learners can assign ownership across a multi-step AI workflow.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Analyze a workflow where each step works but the overall decision lacks an owner.

45 MIN • MINI-LESSON & MODEL

Teach decision rights, handoffs, segregation of duties, checkpoints, reversibility and traceability.

60 MIN • GUIDED PRACTICE

Create a responsibility map naming who directs, verifies, approves and can stop work.

120 MIN • SUPERVISOR LAB

Redesign a workflow so high-consequence steps require stronger human review.

45 MIN • DEBRIEF & ASSESSMENT

Run a tabletop walkthrough and challenge every unowned decision.

Instructor preparation

Prepare a sector scenario and define non-negotiable policy boundaries.

Evidence to collect

Workflow diagram, authority matrix, checkpoint criteria and walkthrough notes.

Specific assignment: Authority-and-Checkpoint Map

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Agents, permissions and escalation

5 guided hours | Match the AI + Navigate & Supervise | Assignment: Agent Permission Plan

Learning outcome

Learners can limit automated action to an appropriate authority envelope.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Contrast drafting assistance with an agent that can send, purchase or change records.

45 MIN • MINI-LESSON & MODEL

Introduce least privilege, scoped access, confirmation gates, rate limits, sandboxing, monitoring and shutdown.

60 MIN • GUIDED PRACTICE

Write permitted, prohibited and approval-required actions for a simulated agent.

120 MIN • SUPERVISOR LAB

Test the permission plan against normal, ambiguous and adversarial scenarios.

45 MIN • DEBRIEF & ASSESSMENT

Require safe handling of uncertainty and a clear stop/escalate path.

Instructor preparation

Use simulations only; do not grant live account access for the exercise.

Evidence to collect

Permission matrix, escalation tree, test log and revision memo.

Specific assignment: Agent Permission Plan

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Quality assurance and auditability

5 guided hours | Evaluate + Guide & Improve | Assignment: AI Quality Test Suite

Learning outcome

Learners can define tests that reveal failure before deployment.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Show why a single successful example does not establish reliability.

45 MIN • MINI-LESSON & MODEL

Teach acceptance criteria, representative cases, edge cases, red-teaming, sampling, error taxonomy and audit logs.

60 MIN • GUIDED PRACTICE

Build and run a ten-case test set against a prototype workflow.

120 MIN • SUPERVISOR LAB

Classify defects by severity, correct the workflow and re-test.

45 MIN • DEBRIEF & ASSESSMENT

Judge whether tests match real risk and whether evidence supports approval.

Instructor preparation

Provide benchmark inputs, expected outcomes and protected test data.

Evidence to collect

Test plan, results, defect log, corrective action and release decision.

Specific assignment: AI Quality Test Suite

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Team supervision and incident response

5 guided hours | Decide + Exercise Accountability | Assignment: Incident Tabletop

Learning outcome

Learners can respond when AI-enabled work causes or nearly causes harm.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Present a near-miss and ask who acts in the first fifteen minutes.

45 MIN • MINI-LESSON & MODEL

Cover containment, preservation of evidence, notification, impact review, correction, root cause and lessons learned.

60 MIN • GUIDED PRACTICE

Assign incident roles and rehearse an escalating scenario.

120 MIN • SUPERVISOR LAB

Lead a tabletop, document decisions and propose preventive controls.

45 MIN • DEBRIEF & ASSESSMENT

Evaluate speed, transparency, evidence preservation and accountability.

Instructor preparation

Set psychological-safety norms; use a fictional incident without real personal data.

Evidence to collect

Incident timeline, decision log, stakeholder message and corrective-action plan.

Specific assignment: Incident Tabletop

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Supervisor Challenge — The Operational Workflow

Supervise a simulated AI-enabled operation from intake through delivery. The workflow must survive routine, ambiguous and failure cases without surrendering human authority.

Required deliverables

- Purpose and authority charter
- Workflow and permission maps
- Quality test suite and results
- Incident tabletop record
- Final release or no-release decision

What assessors look for

- Every decision has an owner
- Permissions follow least privilege
- Tests reflect meaningful risk
- Escalation works under uncertainty
- Approval is evidence-based

Eight-hour challenge sequence

HOUR 1 • SCOPE
Confirm purpose, stakeholder, authority boundary, success criteria and prohibited data.

HOURS 2–3 • DIRECT
Select the tool, document instructions and run the first supervised attempt.

HOURS 4–5 • VERIFY
Test claims, calculations, bias, privacy, usability and failure cases in proportion to risk.

HOURS 6–7 • IMPROVE
Correct defects, strengthen controls and complete the evidence package.

HOUR 8 • DECIDE & DEFEND
Approve, modify, reject, escalate or stop. Present the decision in a five- to eight-minute defense.

Scoring: Common rubric, 24 points. The work product itself is supporting evidence; it cannot compensate for unsafe practice or absent human judgment.

Level IV — AI Leadership

IV

HUMANEDGE AI LEADERSHIP CERTIFICATE

Set the conditions under which AI creates value without outrunning human values, governance or accountability.

20 guided + 8 applied hours | Four modules | One Supervisor Challenge

Designed for

Executives, directors, public leaders, educators, program owners and senior workforce professionals

Level outcomes

- Prioritize AI opportunities against mission, evidence and risk.
- Establish governance, authority and oversight mechanisms.
- Lead workforce redesign with participation, fairness and capability building.
- Approve, pause or stop an AI initiative through accountable executive judgment.

Performance standard

Learners must produce observable evidence of purpose, direction, verification and decision. A polished AI output without a supervision record is incomplete.

Instructor checkpoint

Confirm tool access, data boundaries, accessibility needs and the local escalation route before instruction begins.

Module	Lesson focus	Assignment
L1	AI strategy and value discipline	Opportunity-Risk Portfolio
L2	Governance, policy and authority	AI Governance Charter
L3	Workforce design and change leadership	Workforce Transition Plan
L4	Executive accountability and decision	Executive Decision Briefing

AI strategy and value discipline

5 guided hours | Human Purpose + Understand the Task | Assignment: Opportunity-Risk Portfolio

Learning outcome

Leaders can distinguish valuable use cases from technology-driven activity.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Review an attractive AI proposal that lacks a defined problem or owner.

45 MIN • MINI-LESSON & MODEL

Connect mission, stakeholder value, process readiness, data fitness, risk, cost and measurable outcomes.

60 MIN • GUIDED PRACTICE

Score a portfolio of use cases and identify what should proceed, wait or stop.

120 MIN • SUPERVISOR LAB

Write a one-page investment thesis for one use case and a rejection memo for another.

45 MIN • DEBRIEF & ASSESSMENT

Test whether expected value and human accountability are concrete.

Instructor preparation

Prepare a balanced portfolio; avoid vendor-specific claims without current evidence.

Evidence to collect

Prioritization matrix, investment thesis, rejection memo and metrics.

Specific assignment: Opportunity-Risk Portfolio

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Governance, policy and authority

5 guided hours | Match the AI through Exercise Accountability | Assignment: AI Governance Charter

Learning outcome

Leaders can establish enforceable decision rights and oversight.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Trace what happens when ownership is spread across technology, legal and operations.

45 MIN • MINI-LESSON & MODEL

Cover governance bodies, risk tiers, use-case intake, documentation, review cadence, exceptions, audits and sunset decisions.

60 MIN • GUIDED PRACTICE

Draft a risk-tier model and approval route for four use cases.

120 MIN • SUPERVISOR LAB

Produce a governance charter with roles, evidence requirements and escalation.

45 MIN • DEBRIEF & ASSESSMENT

Conduct a mock governance review of a contested proposal.

Instructor preparation

Align examples with applicable organizational policy; note when specialist advice is required.

Evidence to collect

Governance charter, risk tiers, review checklist and meeting record.

Specific assignment: AI Governance Charter

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Workforce design and change leadership

5 guided hours | Human Purpose + Guide & Improve | Assignment: Workforce Transition Plan

Learning outcome

Leaders can redesign work without treating people as obstacles to automation.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Compare a headcount-only automation plan with a capability-centered redesign.

45 MIN • MINI-LESSON & MODEL

Address task analysis, role redesign, participation, training, accessibility, incentives, workload, fairness and feedback.

60 MIN • GUIDED PRACTICE

Map tasks that are automated, augmented, retained or newly created.

120 MIN • SUPERVISOR LAB

Build a ninety-day transition plan with communications, training and worker feedback.

45 MIN • DEBRIEF & ASSESSMENT

Use an employee-impact review and require mitigation for unequal burdens.

Instructor preparation

Create a fictional organization or use approved internal information only.

Evidence to collect

Task map, role changes, learning plan, communication and impact review.

Specific assignment: Workforce Transition Plan

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Executive accountability and decision

5 guided hours | Evaluate + Decide + Exercise Accountability | Assignment: Executive Decision Briefing

Learning outcome

Leaders can make and defend a consequential AI decision under uncertainty.

Lesson plan

30 MIN • LAUNCH & DIAGNOSE

Present conflicting evidence from performance, risk and stakeholder perspectives.

45 MIN • MINI-LESSON & MODEL

Model decision standards, uncertainty statements, dissent, conditions of approval, monitoring triggers and stop rules.

60 MIN • GUIDED PRACTICE

Hold a decision hearing with assigned stakeholder roles.

120 MIN • SUPERVISOR LAB

Deliver a board-style briefing and issue an approve, conditional-approve, pause or stop decision.

45 MIN • DEBRIEF & ASSESSMENT

Judge evidence use, trade-off reasoning, accountability and clarity of conditions.

Instructor preparation

Provide a complete case packet with unresolved uncertainty; avoid a predetermined answer.

Evidence to collect

Decision brief, oral briefing, conditions, monitoring plan and reflection.

Specific assignment: Executive Decision Briefing

Submit the listed evidence with a brief statement answering: **What did the AI contribute, what did I verify, and what decision remained mine?**

Supervisor Challenge — The AI Operating Charter and Pilot

Design a responsible AI pilot for an organization or community institution, including the case for action, authority boundaries, workforce plan, evidence standard and leadership decision.

Required deliverables

- Use-case investment thesis
- AI governance and authority charter
- Human-AI operating model
- Workforce transition and learning plan
- Executive decision briefing and pilot scorecard

What assessors look for

- Mission value is measurable
- Authority boundaries are enforceable
- Workers and stakeholders are included
- Monitoring has triggers and stop rules
- Leadership accepts final accountability

Eight-hour challenge sequence

HOUR 1 • SCOPE
Confirm purpose, stakeholder, authority boundary, success criteria and prohibited data.

HOURS 2–3 • DIRECT
Select the tool, document instructions and run the first supervised attempt.

HOURS 4–5 • VERIFY
Test claims, calculations, bias, privacy, usability and failure cases in proportion to risk.

HOURS 6–7 • IMPROVE
Correct defects, strengthen controls and complete the evidence package.

HOUR 8 • DECIDE & DEFEND
Approve, modify, reject, escalate or stop. Present the decision in a five- to eight-minute defense.

Scoring: Common rubric, 24 points. The work product itself is supporting evidence; it cannot compensate for unsafe practice or absent human judgment.

Performance, not polish, determines mastery

Recommended level weighting

Evidence	Weight
Four module assignments	40%
Supervisor Challenge	40%
Oral defense / live check	10%
Reflection and professionalism	10%

Non-negotiable gates

- No reckless or prohibited data handling
- No fabricated verification evidence
- AI involvement is disclosed
- A human decision is identifiable
- Required assignments are complete

A learner may revise after feedback. A serious integrity or safety violation requires remediation before certification.

Recommended passing standards

Tier	Rubric average	Additional requirement
Foundations	2.0 / 4	No safety domain below 2
Practitioner	2.5 / 4	Verified work product
Supervisor	3.0 / 4	Operational defense
AI Leadership	3.25 / 4	Governance and pilot defense

Attendance recommendation: at least 80% of guided time plus all assessed performances. Local providers may set a higher standard, but should not lower the safety gates.

Rubric A — Purpose, direction and verification

Use this rubric for every module assignment and Supervisor Challenge.

Domain	1 • Emerging	2 • Developing	3 • Proficient	4 • Exemplary
Purpose & task judgment	Goal, stakeholders and human-only decisions are unclear.	Goal is stated; some boundaries or consequences are missing.	Purpose, success criteria, boundaries and human decisions are clear.	Reframes the problem insightfully and anticipates competing values or consequences.
Direction & tool match	Tool is chosen by habit; instructions are vague.	Tool generally fits; instructions include some context and criteria.	Tool choice is justified; instructions specify context, limits and quality.	Compares alternatives, controls permissions and adapts direction with precision.
Verification & evidence	Accepts output or performs superficial checks.	Checks obvious claims but misses important uncertainty or sources.	Uses independent, proportionate checks and records corrections.	Designs robust tests, triangulates evidence and defines stopping conditions.

Scoring note: Assess the quality of human supervision shown in the evidence. Do not infer mastery from the fluency, visual quality or confidence of AI-generated output.

Rubric B — Responsibility, improvement and decision

Score each domain from 1 to 4. Add all six domains for a 24-point total.

Domain	1 • Emerging	2 • Developing	3 • Proficient	4 • Exemplary
Ethics, safety & privacy	Misses material risks or mishandles protected information.	Identifies common risks; controls are incomplete.	Protects data, tests bias/accessibility and applies suitable controls.	Anticipates systemic impacts and designs preventive, monitored safeguards.
Collaboration & improvement	Treats AI output as final or revises without a clear reason.	Makes revisions but evidence of learning is uneven.	Uses feedback and defect evidence to guide purposeful improvement.	Builds efficient feedback loops and transfers learning across contexts.
Decision & accountability	Final decision is missing, delegated to AI or unsupported.	Makes a decision with limited rationale or ownership.	Makes, explains and owns an evidence-based human decision.	Handles uncertainty, dissent and consequences with transparent accountability.

Scoring note: Assess the quality of human supervision shown in the evidence. Do not infer mastery from the fluency, visual quality or confidence of AI-generated output.

Feedback must point to a stronger supervisory move

Feedback protocol

1. **Name the evidence.** Quote or point to the learner's supervision record.
2. **Name the risk or gap.** Explain what could go wrong and why it matters.
3. **Name the next move.** Connect the revision to H.U.M.A.N.E.D.G.E.
4. **Require new evidence.** A revised claim must be rechecked, not merely rewritten.

Three useful questions

- What did you decide before involving AI?
- What evidence changed your mind?
- What would make you stop or escalate?

Reassessment

Learners may resubmit after targeted practice. Preserve both versions so growth is visible.

Authenticity interview prompts

Prompt	Evidence it reveals
Show one AI suggestion you rejected. Why?	Independent judgment
Recheck this claim in front of me.	Verification skill
What did your tool choice make easier—and riskier?	Tool matching
Who could be affected by this decision?	Ethical awareness
What would you do if the system acted outside scope?	Escalation and accountability

Four credentials make growth visible and stackable

Credential	Recipient can credibly say	Required evidence
Foundations	I use AI with a clear purpose, check its work and own my decision.	4 assignments + Reliable Brief
Practitioner	I select and direct AI to produce verified work through a documented process.	4 assignments + Responsible Work Product
Certified Human-AI Supervisor	I oversee people, AI systems and controls across an operational workflow.	4 assignments + Operational Workflow
AI Leadership	I set strategy, authority and governance for responsible organizational AI.	4 assignments + Operating Charter & Pilot

Issuance checklist

- ☐ Identity verified by provider
- ☐ Passing score confirmed
- ☐ Safety gates satisfied
- ☐ Portfolio retained per policy
- ☐ Completion date recorded
- ☐ Credential ID assigned

Certificate integrity

Maintain an internal credential register containing recipient name, tier, completion date, credential ID, issuing instructor and status. Publish or share verification only with appropriate consent.

Renewal recommendation

For workforce credentials, offer a two-year update or continuing-evidence review because tools, regulations and organizational risks change.

Model the supervision you expect learners to practice

Before the course

- Confirm approved tools and age/access rules.
- Set data, privacy and account boundaries.
- Prepare accessible alternatives.
- Test scenarios without real confidential data.
- Identify the escalation contact.

During each module

- Make uncertainty visible.
- Think aloud while verifying.
- Reward a justified rejection.
- Ask for the human decision.
- Separate speed from quality.

After each module

- Collect the supervision record.
- Score evidence with the common rubric.
- Return one strength and one next move.
- Log safety or access issues.
- Adjust examples, not standards.

Instructor stance

The instructor is not required to know every AI tool. The instructor must be able to maintain boundaries, ask for evidence, model verification and recognize when specialist guidance is needed.

Facilitation principle: Never demonstrate a prohibited practice merely to make the lesson more realistic. Simulate access, data and consequences safely.

The framework stays constant while context and independence change

Audience	Instructional emphasis	Appropriate evidence
Middle school	AI is a tool; checking is a habit; no personal data	Annotated output, check list, oral explanation
High school	Direction, source evaluation, disclosure and agency	Prompt log, source checks, decision paragraph
College / CTE	Professional standards, workflow, quality and role boundaries	Work product, test record, supervisor defense
Workforce	Organizational policy, customer impact and escalation	Process map, risk control, approval record
Leadership	Authority, governance, workforce impacts and stop rules	Charter, scorecard, decision hearing

Accessibility

- Offer keyboard, speech and screen-reader compatible tools.
- Provide non-AI or partner alternatives when access differs.
- Assess judgment, not typing speed or paid-tool access.
- Caption demonstrations and describe visuals.

Sector contextualization

Replace scenarios with local tasks, standards and risks. Keep the same competency language, evidence set, rubric and human-approval requirement so credentials remain comparable.

Safe practice is a condition of participation

Learners must

- Use only authorized accounts and tools.
- Follow age, school, workplace and sector rules.
- Minimize data and remove identifiers.
- Disclose material AI assistance.
- Verify before submitting or acting.
- Escalate uncertainty and potential harm.

Learners must not

- Enter restricted, confidential or personal data without authorization.
- Misrepresent AI work as entirely human.
- Fabricate sources, checks or approvals.
- Bypass safeguards or permissions.
- Let AI make a final high-consequence decision.

Stop-and-escalate conditions

- Possible harm to a person or group
- Unclear authority to act
- Security or privacy exposure
- High-stakes legal, medical, financial, safety or employment impact
- Evidence conflict that cannot be resolved
- System behavior outside the assigned scope

This curriculum supports responsible practice; it does not replace legal, compliance, cybersecurity or subject-matter advice.

F

HUMANEDGE ACADEMY™

Certificate of Completion

This certifies that

Recipient Name

has successfully completed the requirements for the

Human-AI Supervisor Foundations Certificate

demonstrating foundational ability to direct, check, question and decide when using artificial intelligence.

> Human Authority → AI Assistance → Human Verification → Human Decision

Completion date	Credential ID

Authorized Academy Representative

Instructor / Program Director

Academy-issued, non-degree credential • Verification subject to HumanEdge Academy records

P

HUMANEDGE ACADEMY™

Certificate of Completion

This certifies that

Recipient Name

has successfully completed the requirements for the

Human-AI Supervisor Practitioner Certificate

demonstrating the ability to select, direct and evaluate AI within a documented professional workflow.

Human Authority → AI Assistance → Human Verification → Human Decision

Completion date	Credential ID

Authorized Academy Representative

Instructor / Program Director

Academy-issued, non-degree credential • Verification subject to HumanEdge Academy records

S

HUMANEDGE ACADEMY™

Certificate of Completion

This certifies that

Recipient Name

has successfully completed the requirements for the

Certified Human-AI Supervisor

demonstrating the ability to supervise AI-enabled operations through clear authority, controls, verification and accountable decision-making.

Human Authority → AI Assistance → Human Verification → Human Decision

Completion date	Credential ID

Authorized Academy Representative

Instructor / Program Director

Academy-issued, non-degree credential • Verification subject to HumanEdge Academy records

L

HUMANEDGE ACADEMY™

Certificate of Completion

This certifies that

Recipient Name

has successfully completed the requirements for the

HumanEdge AI Leadership Certificate

demonstrating the ability to lead AI strategy, governance, workforce change and responsible organizational decision-making.

Human Authority → AI Assistance → Human Verification → Human Decision

Completion date	Credential ID

Authorized Academy Representative

Instructor / Program Director

Academy-issued, non-degree credential • Verification subject to HumanEdge Academy records

Human-AI Supervision Record

Complete for every major assignment. Attach prompts, outputs and source evidence as needed.

1. Human Purpose

Goal / problem:

Stakeholder and desired outcome:

Values, constraints and prohibited actions:

2. Understand & Match

Human-only decisions:

AI-assistable tasks:

Tool selected and why:

3. Assign & Navigate

Instruction summary:

Permissions / access:

What I monitored or changed:

4. Evaluate & Decide

Checks performed:

Corrections / remaining uncertainty:

Decision: approve / modify / reject / escalate / stop

Accountability statement: I understand that AI assisted with this work and that I remain responsible for the decision to submit, share or use the result.

Signature: _____ Date: _____

Verification and risk checklist

Accuracy and evidence

- ☐ Material claims checked independently
- ☐ Sources opened and matched to claims
- ☐ Dates and versions confirmed
- ☐ Calculations repeated independently
- ☐ Quotes and names confirmed
- ☐ Uncertainty clearly labeled

Quality and suitability

- ☐ Meets audience and purpose
- ☐ Follows required format and standards
- ☐ Accessible and understandable
- ☐ No invented citations or details
- ☐ Originality and reuse reviewed

Ethics, safety and privacy

- ☐ No unauthorized personal data
- ☐ Bias and unequal impact considered
- ☐ AI use disclosed where material
- ☐ Permissions match the task
- ☐ High-consequence issues escalated
- ☐ Stop rule identified

Final human decision

☐ Approve ☐ Modify ☐ Reject ☐ Escalate ☐ Stop

Basis for decision:

_____

_____

Decision owner: _____

Date: _____

Rule: The strength of verification should rise with the consequence of error.

A ninety-day pilot can test the model without lowering the standard

DAYS 1–15 • DESIGN
Name a program owner; choose the first audience; select one level; confirm policies, tools, instructor readiness and accessibility.

DAYS 16–30 • PREPARE
Contextualize scenarios, build safe case materials, calibrate assessors on sample evidence and establish the credential register.

DAYS 31–60 • DELIVER
Run the modules, collect supervision records, hold weekly instructor reviews and address access or safety issues promptly.

DAYS 61–75 • ASSESS
Conduct the Supervisor Challenge, oral defenses and rubric moderation. Offer targeted reassessment where appropriate.

DAYS 76–90 • IMPROVE
Review completion, rubric consistency, learner confidence, incident logs, instructor observations and stakeholder feedback. Revise before scaling.

Pilot evidence

- Enrollment and completion
- Rubric scores by domain
- Reassessment patterns
- Safety and access issues
- Learner and employer feedback

Scale decision

Expand only when assessors score consistently, learners can defend their decisions and the operating environment supports safe, equitable access.

Core glossary

AI assistant

A system that supports a person's task under human direction.

AI agent

A system that can take multi-step actions within granted tools, permissions and goals.

Authority envelope

The defined actions, data, limits and conditions within which a system may operate.

Checkpoint

A required moment of review before work proceeds.

Human-AI Supervisor

A person who defines, directs, evaluates and remains accountable for AI-assisted work.

Verification

Independent checking of claims, reasoning, calculations, sources, quality and suitability.

Escalation

Moving a decision to a person with the expertise or authority to resolve it.

Stop rule

A pre-defined condition that pauses or ends AI activity.

Supervision record

Evidence of purpose, instructions, checks, corrections and human decision.

Supervisor Challenge

An authentic performance task scored on how effectively the learner supervises AI.

> Direct it. Check it. Question it. Decide.