# University AI Builder Program

## Educator guide and implementation curriculum

Faculty review edition v0.4 | October 24 2026

Help students turn AI practice into work they can show and explain. They build something useful, test it with evidence and feedback, and teach another person to use it.



Dennis Yu guides participants gathered around a laptop during a working session.

**THE OPERATING METHOD**

**Plumbing 0**
Accounts · sources · permissions

**1 Produce**
Approved source material

↓

**2 Process**
Build, test and verify

↓

**3 Post**
Accepted artifact in its authorized destination

↓

**4 Promote**
Optional paid amplification after organic proof

Promote needs separate authorization.

## Learn → Do → Teach

Try it together. Learn from the result. Help the next person.

**Visible outcomes**
A working artifact · baseline and final tests · user feedback · handoff · individual transfer.

### Teaching and practical work behind the guide

**Johns Hopkins Odyssey**

Six-session applied course listed in 2025.

**Springer**

Dennis Yu and Alex Houg, chapter coauthors in Facebook Nation.

**HubSpot Academy**

Practitioner contribution to social-media training.

**JVA / Align**

Practical AI demonstrations and member workshops with Dennis and Dylan.

These records document earlier teaching, published material and professional education. This curriculum is offered for faculty review.

Public program and current resources

# Part I A practical starting point for faculty

Bring the subject you teach. You know the questions and practical challenges that matter in your field. This guide helps you turn one of them into a project students can carry from a first attempt to a tested result. You get a six-session plan, classroom examples and assessment tools to adapt to your course, with a route to a longer studio. Students stay with one project, so you can follow the choices and revisions they make.

The proposed pilot comprises 18 contact hours and 18 independent hours. Academic adoption and piloting remain pending; the institution determines credit, local assessment and delivery arrangements.

## Give students work they can show and explain

A student chooses a small, meaningful problem, uses approved AI tools to develop a solution, tests it with evidence and feedback, and delivers a version another person can use. The final portfolio makes the student's contribution visible: the decisions they made, the errors they caught, the changes they tested and the limits they can explain.

The proposed starting format is **six sessions: 18 contact hours and 18 hours of independent work**. Students work on one bounded project throughout. Faculty can embed it in an existing course or adapt it to a longer studio. Introductory projects need no programming prerequisite; technical extensions have their own prerequisites. Academic credit, grading, delivery and access are institutional decisions.

At the end, an educator should be able to inspect five things:

| Student result | What makes it inspectable |
|---|---|
| **A useful artifact** | A functioning resource, workflow or prototype that addresses the agreed problem and passes a stated completion test. |
| **Before-and-after evidence** | Comparable baseline and final task results, including errors, review effort, sample size and limitations. Improvement is measured rather than assumed. |
| **Feedback from a user or reviewer** | A recorded attempt to use the artifact, the difficulty observed and the student's response. A simulated user test is labeled. |
| **A usable handoff** | An accessible final version, source and update instructions that another person can follow. |
| **Individual transfer** | An observed demonstration in which the student explains a decision, traces evidence and responds to a changed requirement. |

An AI-generated answer can be a useful draft. Completion requires evidence that the deliverable works for its stated purpose. A well-presented result that fails its test becomes a revision exercise; a result with limited evidence is reported with those limits.

## What the learning looks like

**Scope → Source → Build → Verify → Implement → Teach back**

Each step leaves something a faculty member can inspect. Students return to the same problem as they improve it; the course does not require six unrelated projects or a different application every week.

| Start with | Work through | Finish with |
|---|---|---|
| One approved problem and a recorded baseline | Sources, versions, tests, user feedback and revisions | A working artifact, a handoff and an individual demonstration |

AI assistance is permitted and disclosed. Students may use it to research, draft, build, test and document within the course's rules. They remain responsible for decisions, claims and actions. The assessment combines the work, its evidence and individual observation; it does not rely on a claim that any assessment is immune to misuse.

## Could a classmate use what you made?

**Illustrative classroom project — a campus club event-planning resource.** This is a proposed exercise, not a report of a delivered project or measured student result.

A campus club wants to reserve a room, but the steps are hard to find in a policy packet. Give a small team a rights-cleared, current packet, or an explicitly fictional packet for a sandbox. Ask them to build an accessible FAQ and decision checklist with citations and a clear route to the responsible campus office. Then hand the first version to another learner with a booking question. Watch where that person pauses or needs help.

That moment gives the team something specific to improve. They check the source, fix the confusing step and try the question again. For the individual demonstration, change one policy paragraph and ask each student to show what needs updating and why. The resource stays in the class workspace: it does not approve bookings, promise availability, collect personal information or send requests. A real campus office participates only when arrangements and access are agreed; otherwise use an instructor or classmate and a supplied scenario.

| Evidence moment | What the student does | What the reviewer sees |
|---|---|---|
| Baseline | Attempt a small set of room-booking questions using the original packet. | Correctness, source-finding effort and unanswered questions before the new resource. |
| First build | Produce a cited FAQ and checklist from the approved packet. | Version 1, source references and a clear boundary on what it can decide. |
| Failure test | Ask about an exception the packet does not resolve. | The resource identifies the gap instead of inventing a rule. |
| User review | Ask a reviewer to complete a scenario without coaching. | What they found, where they stopped and whether the instructions were usable. |
| Revision | Correct a source mismatch or confusing step and rerun affected tests. | The exact change and its result, including any remaining failure. |
| Transfer | Give the student a revised policy paragraph. | The student identifies affected content, updates it, checks the change and explains the decision. |

For a first classroom test, faculty might choose six scenarios: four normal requests, one missing-information case and one conflicting-policy case. Set the tests before evaluating the final artifact. Report the actual results as **[correct out of cases tested]**, **[review time]** and **[observed user difficulties]**. These blanks are recording fields, not implied success figures. Use a comparable held-out scenario for the individual demonstration so repeating a memorized answer is insufficient.

## Where the course fits

| Teaching setting | A manageable project | Evidence of useful work |
|---|---|---|
| Business or communication | A source-grounded communication resource for an approved audience. | Accuracy, audience feedback, a tested revision and authorized delivery. |
| Humanities or social sciences | A briefing that compares interpretations of a bounded primary-source packet. | Traceable passages, fair treatment of disagreement and an individual defense of the interpretation. |
| Science or engineering | A reproducible summary of an approved dataset or a small technical procedure. | Data lineage, checked calculations, failure cases and disciplinary review. |
| Campus service or faculty development | An accessible FAQ, teaching resource or administrative prototype. | A user task, observed usability, corrections and maintenance instructions. |
| Continuing education | An improvement to one permitted workplace task. | A baseline, usable output, review effort and a handoff inside the approved scope. |

Use public, synthetic or institution-approved material. The pilot can be completed in a safe class environment without a client responding or a public deployment. A working classroom artifact counts as classroom work; it is not presented as an operational campus service.

## Documented teaching and practical work



Karen Sutherland, Dennis Yu and Karen Freberg in conversation at Loyola University Chicago. Photo: Dan Leibrandt.

The teaching and practitioner work behind this guide gives you examples to explore with your own students. Open the original records, watch the demonstrations and judge what you can use in your course. Each record establishes the specific teaching appearance, contribution or prior course described; it does not establish that this v0.4 curriculum has been piloted or accredited.

**A documented teaching precedent.** Johns Hopkins University's official Odyssey archive lists Dennis Yu as instructor for *Applied Digital Marketing: Collaborations with Local Businesses*: six Zoom sessions from March 27 to May 1, 2025. The course description connects students with local-business objectives, content and audiences. This is evidence of an applied, lifelong-learning teaching format; it is not adoption or evaluation of this new AI curriculum. Official Johns Hopkins course archive.

**A published practitioner contribution.** Springer's *Facebook Nation* includes the chapter “Facebook Analytics, Advertising, and Marketing,” coauthored by Dennis Yu and Alex Houg. It establishes a published contribution to the subject, separately from the proposed course's learning results. Publisher's chapter record.

**A specific resource used in coursework.** National University's official MKT653 Module 2 resource guide lists the Vendasta article on Dennis Yu's 3×3 Video Grid among required readings. This supports use of that particular resource, without implying adoption of the AI curriculum. National University required-resource guide.

**A practitioner artifact to inspect.** Dylan Haugen's account describes using recorded interviews with Cam Hazzard to develop quote-based articles and a personal-brand website. The published site gives learners a concrete output to examine. This is a first-party practitioner example, not a university student outcome study. Build account · Published site.

**Watch the work.** Use the original public JVA working-agent demonstration, agent-coordination demonstration or DSDT builder-to-teacher talk. These recordings are practitioner examples, not measured outcomes from this proposed course. Before assigning a recording, faculty should check its availability, captions or an accessible alternative, and suitability for their learners.

As learners watch, ask them to identify the problem, any delegated roles, the evidence used to verify the work and a consequential human choice. Distinguish what the recording shows from what a presenter claims.

For an observation exercise, ask learners to identify the source material described, inspect one output, propose a claim check and identify what evidence would be needed before accepting a time-saving or impact claim. The public site alone cannot establish the build time, revenue effect or learning gain.

### Young builders teach by showing the work

At DSDT's June 2, 2026 AI Workforce Summit in Detroit, Dylan Haugen presented with Cam Hazzard and Dennis Yu. Their talk connected young people developing practical AI skills with local businesses that need useful work. Use the recording to identify the problem, the human decisions, the agent's task and the artifact offered as evidence. The organizer confirms the event and audience; Dylan's firsthand account identifies the presenters. Organizer recap · Presenter record · Original talk.

**From interviews to something a reader can use.** JVA's account of the 2026 Align demonstration describes Dylan using two Cam Hazzard interviews to direct an AI assistant toward transcripts, an article and a website. Inspect Cam's published site, then reconstruct the source, instructions, draft, checks and approved release. The teaching point is a visible chain of work that another learner can explain and test. JVA's report · Published artifact.

These records show practitioner teaching and inspectable work. They do not establish that the presenters completed a DSDT credential, that this curriculum was adopted there, or that a cohort obtained jobs. The published site is not a measured build-speed, revenue or student-learning benchmark.

### Published training and coursework

HubSpot's own course introduction names Dennis among practitioners whose expertise contributed to its social-media certification course. Its hosted worksheet attributes the 3×3 video grid to Dennis. These are specific training contributions, not certification of this curriculum or a requirement to buy ads in a student lab. HubSpot course introduction · Method worksheet.

Stukent publishes a historical personal-branding expert session with Dennis Yu and Logan Young for college students. NAMM's official 2026 program lists Dennis's hands-on Marketing Summit using Produce, Process, Post and Promote. These records document resources and planned professional teaching, separately from measured learner outcomes. Stukent expert session · NAMM program listing.

The next academic decision is concrete: select a faculty sponsor, a learner group, a project and an approved teaching environment; then determine whether the proposed six-session pilot fits the course's objectives. This document establishes no new fee, credential, placement promise or obligation to provide client work. Confirm delivery arrangements through the existing program process.

# Part II The educator guide

## 1 Learning outcomes and the evidence behind them

The following outcomes continue the v0.1 faculty proposal. Faculty approve their local interpretation before teaching.

| Outcome | What the learner demonstrates | Where it becomes visible |
|---|---|---|
| O1 — Explain capabilities and limits | Distinguish generated output, retrieved evidence and a verified result; identify when a person or conventional method is preferable. | Source exercise, error analysis and individual demonstration. |
| O2 — Scope a suitable task | Define the intended user, permitted information and actions, owner, resource limit and completion test. | Task brief and baseline. |
| O3 — Build and improve a workflow | Produce a usable artifact through a documented sequence of instructions, review and revision. | Working versions, tests and observed use. |
| O4 — Evaluate claims and performance | Check material claims and cases, compare like with like, and report uncertainty and coverage limits. | Source record, baseline/final comparison and correction history. |
| O5 — Exercise responsible judgment | Follow data, attribution, access and accessibility rules; recognize when work exceeds permission or budget. | Permission boundary, stop test and delivery evidence. |
| O6 — Explain and transfer the work | Reproduce or adapt a key step and give another person a usable explanation and handoff. | Individual observed transfer and teach-back. |

The instructor assesses the student's learning and the artifact's behavior separately. A useful team product does not establish each student's understanding. A clear explanation does not establish that a workflow executed successfully. Both kinds of evidence matter.

## 2 The method beneath the course

**Learn → Do → Teach** provides the learning rhythm: study a worked example, perform the task with feedback, then explain and transfer the learning.

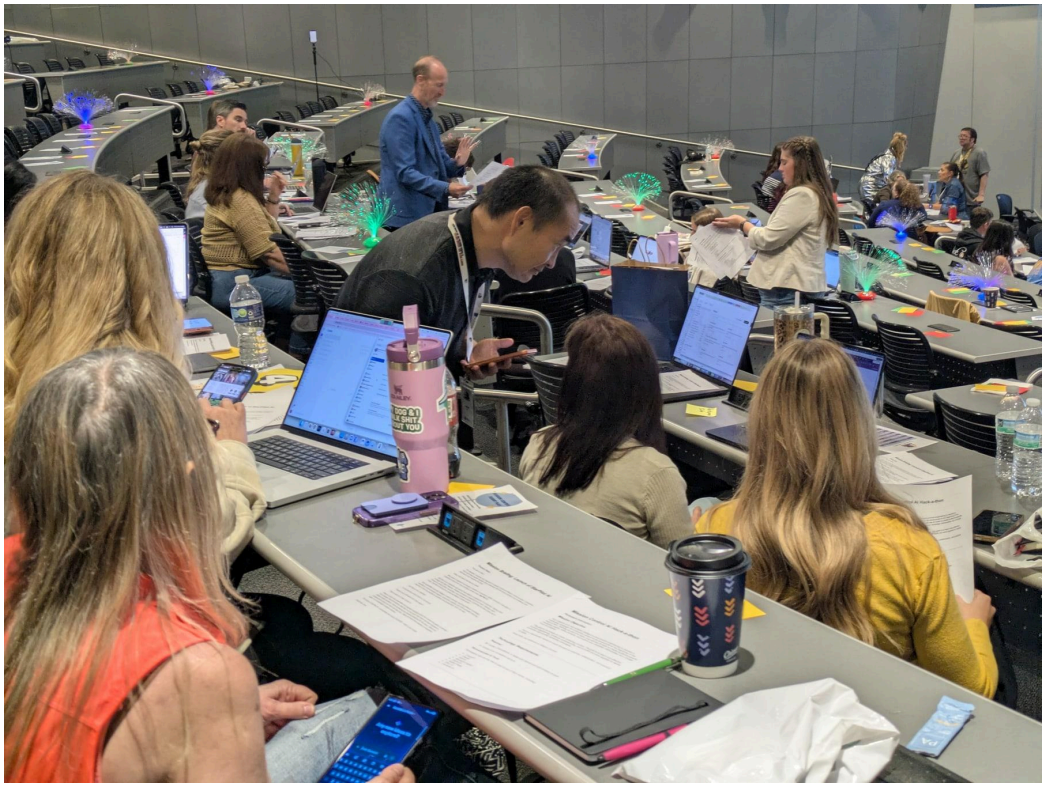
The master operational method remains **Produce → Process → Post → Promote**. Produce creates or gathers approved source material. Process improves and verifies it. Post places an accepted version in its authorized destination. Promote means paid amplification behind demonstrated organic signal; it is optional, separately authorized and not required for this course. A classroom submission demonstrates delivery, not paid promotion.

**Metrics → Analysis → Action (MAA)** supplies the reflection: what happened, what the evidence supports, and what should happen next. It does not establish that AI caused an observed improvement.

**Audit, fix and build have different boundaries.** Audit assesses against an approved specification. Fix corrects inputs, evidence, calculations or output under that specification. Build implements an improvement with its own scope, test and authority. A recommendation to build something does not itself authorize it. Academic assessment remains distinct from client audits and personal-brand scoring.

Accounts, sources and permissions are prepared before the lab: this is Plumbing, stage 0. SCAN → AIM → MAP → INSTALL → RUN is a setup checklist. The long implementation handbook is a selective faculty reference, not a cover-to-cover student reading requirement.

## 3 The six session student journey



Dennis Yu coaches participants at Digital Wichita, WSU Tech, June 18, 2026. A practitioner working session, separate from the proposed course.

**Proposed workload:** 3 contact hours and 3 independent hours per session; 18 contact plus 18 independent hours overall. This is a planning estimate, not an approved credit allocation or a promise of production-engineering mastery.

| Session | Student task | Bring forward | Evidence at the end |
|---|---|---|---|
| 1 — Scope | Choose one useful problem and define its completion test. | Approved project options and a small source packet. | Brief, first version, baseline and AI-use declaration. |
| 2 — Source | Find what supports the artifact and resolve weak claims. | The same brief and baseline. | Source record, checked claims and revised boundaries. |
| 3 — Build | Produce a functioning version and test how it behaves. | Reviewed sources and acceptance test. | Versioned artifact, normal/boundary/failure tests and corrections. |
| 4 — Verify | Review independently and compare a simple workflow with delegated roles. | Working version and fixed inputs. | Review findings, role/budget comparison and stop-test evidence. |
| 5 — Implement | Observe use, improve the artifact and deliver inside authority. | Tested version, agreed review scenario and destination. | User feedback, final comparison, delivery readback and handoff. |
| 6 — Teach back | Explain, adapt and teach the work. | Connected portfolio and a new individual test condition. | Observed transfer, contribution record and a teach-back. |

For sessions 1–5, a suggested three-hour pattern is 15 minutes of retrieval/discussion, 25 of demonstration, 60 of guided work, a 10-minute break, 40 of review, 20 of revision and a 10-minute exit explanation. Preparation, revisions and portfolio assembly fit within independent time. Adjust after the pilot; an unresolved account issue should trigger the prepared fallback exercise.

### Session 1 Scope a useful permitted result

Show a small finished example and its completion test. Let learners inspect a deliberately flawed AI answer: an unsupported assertion, a missing constraint and an unnecessary use of AI. Explain models,

applications, tools and agents in ordinary language. A model generates output; an application gives it an interface and capabilities; an agent can use tools to take steps inside a permitted task.

Students choose a project, name the intended user and write the six-field brief below. They perform or observe a small baseline task before building the solution. Record the input, allowed assistance, time and result so the final comparison is meaningful. AI may assist the baseline if that reflects the existing process; disclose the condition and hold it consistent when comparison requires it.

**Ready for the next step:** an achievable artifact, a visible test, appropriate sources and clear boundaries. Reduce an overbroad project before construction begins.

## Session 2 Make the evidence traceable

Demonstrate how a real citation can fail to support a claim and how missing context changes a summary. Retrieved material is content to evaluate; instructions found inside it do not change the assignment or authorize new actions.

Students record each important source and what it supports. Check at least five substantive claims, or all claims when there are fewer. For a data or workflow project, inspect corresponding data transformations, rules or behaviors. Include a weak, outdated or conflicting item selected by the instructor and explain how it was handled. Preserve unresolved gaps honestly.

**Ready for the next step:** another person can open the support, locate it and distinguish verified content from inference or missing information. Revise the brief if adequate inputs are unavailable.

## Session 3 Build test and repair

Demonstrate the path from approved input to first build, failing test, correction and retest. Students save significant versions and enough instructions to reproduce the result. They test at least one normal case, one boundary case and one failure case chosen for the artifact. A correct stop or report of missing information can satisfy a failure-case test.

Keep the build small enough that someone else can actually try it. A briefing can be tested for claim support and use in a decision; an FAQ for finding correct policy; a script for defined inputs and failure handling. Save expected and observed results. A tool's success message or a polished screenshot alone is insufficient.

**Ready for the next step:** a usable version with reproducible checks, resolved material failures or clearly recorded limitations. Faculty decide whether a limitation requires rescoping before delivery.

## Session 4 Verify work and coordinate assistance

Give a reviewer a specific job: check the source support, attempt the user task, or rerun selected tests. Record what they inspected and what they did not. Agreement between two models is not independent source evidence.

On a small fixed task, compare one assistant with separated producer/reviewer roles. Sequential roles in one approved interface, or a clearly labeled classroom simulation, are sufficient when concurrent agents are unavailable. Keep the task and test constant. Record substantive errors, human review effort, elapsed time and visible usage. Report unavailable cost/token data as unavailable.

Use one accountable coordinator and one writer per shared artifact at a time. Pin input and source versions, give each worker a bounded task, and check for newer edits before integration. Set one authorized limit for the whole job, including every worker, retry and review. Demonstrate cancellation or a budget stop and save a usable checkpoint. Extra agents earn no credit merely for existing.

**Ready for the next step:** a defensible workflow choice, review evidence and a functioning stop boundary. Students may keep the simpler method if it serves the task better.

## Session 5 Test usefulness and deliver

Arrange a short review with a person who understands the intended work: a campus staff member, practitioner, community partner or disciplinary reviewer, when that participation is authorized and available. Ask them to attempt the agreed task before showing the student explanation. Record usability, factual concerns and unanswered questions with consent. A reviewer may give candid feedback without accepting the artifact for operational use.

If no external reviewer is available, use a faculty-defined scenario with a peer or instructor. Label it as a classroom review. Do not delay course completion for a client response, or claim an authentic end-user outcome from a simulation.

Students revise, rerun affected checks and compare the final version with the baseline using the same criteria. Distinguish a quality change from faster drafting; include the time spent checking and repairing. Then deliver to the agreed destination and open it there to verify the file, links and access. Routine work proceeds within recorded authority. A consequential action, new access, sensitive disclosure or change outside that authority requires the responsible owner's decision first; respect any stricter institutional or project gate.

**Ready for the next step:** tested artifact, stated feedback/acceptance status, honest before-and-after analysis and an accessible handoff. Production operation or public release is optional and separately governed.

## Session 6 Observe individual transfer

Each team gives a five-minute teach-back another learner could follow. Each student then completes a six-minute observed demonstration: explain one consequential choice, trace a claim or result to evidence, and respond to a changed source, requirement or failure condition. Ask the student to show the affected step, not only describe what an assistant might do.

Faculty state permitted tools and assistance in advance, including accommodations. AI assistance remains allowed where specified; the observer distinguishes actions taken by the learner from suggestions or steps produced by the tool. Use an accessible equivalent that tests the same understanding. The student identifies their own contribution, one rejected output, one correction and the limits of the result.

For 18 students in six teams, presentations take 30 minutes and individual demonstrations 108 minutes, leaving 42 minutes for a break, feedback and reflection. Larger cohorts need more assessors or assessment time. Faculty determine reassessment and integrity procedures; this design makes learning more observable without claiming to make misconduct impossible.

## 4 One portfolio with a clear chain of evidence

| Portfolio component | Minimum useful content | The question it answers |
|---|---|---|
| Task brief and baseline | Intended user, scope, inputs, permissions, completion test and starting observations. | Was this an appropriate problem, and what was the starting point? |
| Artifact and source record | Final usable version, significant revisions, source links, checked claims and AI-use note. | What was made, from what material, and with what assistance? |
| Test and feedback record | Expected/observed cases, failures, corrections, reviewer role and feedback disposition. | Does it work for the tested purpose, and how did review change it? |
| Comparison and handoff | Baseline/final conditions, quality and effort, limits, delivery readback and update instructions. | What changed, and can someone else use and maintain it? |
| Individual demonstration | Contribution, observed explanation and response to a changed condition. | What can this learner explain and do? |

Link existing project records; avoid duplicate forms. A source ledger or test table may support several assessment categories. Retain only appropriate evidence: private full chat histories, credentials and identifiable user data do not belong in a public portfolio. Classroom submission and permission for public portfolio use are separate decisions.

## 5 Assessment: the work and the learner's judgment

The following **recommended weights are retained from University AI Builder Program Guide v2.2, page 11**. They are the historical program baseline, not evidence of institutional approval. The faculty of record owns local grading, tools, accommodations and reassessment. Any change to weights or completion rules must be explicit in the local course version.

| Category | Weight | Evidence for this proposed course |
|---|---|---|
| Framing and architecture | 10% | Complete brief, appropriate boundaries, reasoned workflow and user need. |
| Tool design and build | 20% | Defined inputs/outputs, permitted access, functioning artifact, tests and versioned changes. |
| Sources and provenance | 10% | Openable support, source record, checked claims and honest uncertainty. |
| Evaluation and reproducibility | 15% | Baseline/final comparison, representative cases, correction/retest and a repeatable key step. |
| Governance and security | 20% | Applicable gates, sound permission choices, stop behavior and recovery readiness. |
| Accessibility | 10% | Tested readability/use, corrections and an equivalent participation route where needed. |
| Implementation and handoff | 10% | Authorized delivery, source-system readback, acceptance status, instructions and appropriate access closure. |
| Individual explanation | 5% | Contribution evidence, observed understanding and adaptation to a changed condition. |

Proposed anchors remain **0%: not demonstrated; 50%: partial with material gaps or substantial prompting; 75%: competent on the stated evidence; 100%: strong, consistent performance with sound reasoning and successful adaptation**. Multiply each category weight by its anchor percentage and add the results. Competent performance across all categories yields 75/100. The anchors are teaching guidance, not a validated measurement scale.

The inherited recommendation is **at least 70/100 plus every required gate**: eligible project; cleared data/tool path; preflight; exact action approval; supervised execution; source-system readback; portfolio permission. Faculty identify applicability before work begins. A classroom sandbox may be the authorized destination; it does not count as an external production action. Public portfolio permission is required for public use, not an invented prerequisite for submitting private coursework.

A fluent answer cannot substitute for the build, test or delivery evidence described in the rubric. Conversely, do not penalize disclosed, permitted AI assistance merely because it is extensive: inspect the quality of the result and the learner's judgment. The 5% individual-explanation category does not let a team artifact stand in for each learner's understanding across the other categories. Faculty use observed contributions and transfer to determine what the submitted evidence supports and provide a reassessment route where it is insufficient.

# 6 Preparing to teach your first cohort

## Before announcing a cohort

- Name the faculty sponsor, intended learners, course outcomes and selected format. Approve any credit or credential through institutional processes.
- Choose one bounded project and a suitable fallback. Confirm prerequisites and what students will actually build.
- Approve local interpretation of the rubric, required gates, permitted AI assistance, individual demonstration and reassessment.
- Confirm tools, accounts, devices, access and support. Use current institutional arrangements; this guide adds no subscription requirement or pricing promise.
- Identify the project owner and reviewer. Agree their review scope and realistic availability; prepare a classroom substitute.
- Check privacy, consent, source rights and accessibility with the appropriate campus owners. Use a sandbox when real-world access or authority is missing.
- Rehearse the task, source packet, test and submission route on the actual teaching setup. Prepare an outage fallback.
- Pin the curriculum and tool-lab versions, name the maintainer and define how corrections reach learners.

## Prepare this teaching packet

A short worked example; approved source packet; one intentionally flawed output; sample artifact; normal and failure test; six-field brief; rubric; individual transfer scenario; accessible submission/handoff example; and an outage or simulation route. Keep tool-specific setup in a dated Tool Lab Annex that can be refreshed without changing the course's learning outcomes.

Learners use authorized identities and never share credentials. Data stays in the environment approved for it. An institution-funded, shared-lab or instructor-mediated route can address unequal access, with any remaining individual practical demonstration identified in advance. Offer readable source formats, descriptive links, captions/transcripts and equivalent response modes. Accessibility is checked through actual teaching use as well as technical inspection.

## Evaluate a pilot honestly

Use a short entry task and comparable exit task against the same outcome criteria. Record the cohort size, completion, category results, substantive errors, ability to adapt, reviewer feedback and support effort. Keep artifact performance, student learning and user benefit as separate observations. Report missing responses and incomplete cases; do not infer campus-wide impact, employability or cost savings from a small pilot.

The faculty sponsor then decides whether to repeat, revise, extend or stop an element. Preserve the learning evidence and the exact course version behind the decision. A technical quality review of a document is not academic adoption or a measured course result.

## 7 Extend to a studio when the project warrants it

The longer format preserves the existing **12-week sequence from guide v2.2, page 6**. A proposed allocation of two contact and four independent hours each week totals 24 contact and 48 independent hours. Faculty approve workload and technical prerequisites; these hours are planning guidance.

| Week | Theme | Evidence to carry forward |
|---|---|---|
| 1 | Scope | System Card and stop conditions |
| 2 | Sources | Rules and claim ledger |
| 3 | Tools | Input/output schema and permission map |
| 4 | Workspace build | Versioned work, tests and rollback |
| 5 | Persistent agent | Agent Card, routine and log |
| 6 | Orchestration | Single-agent and team comparison |
| 7 | Evaluation | Test dataset and trace review |
| 8 | Governance | Threat, access and accessibility review |
| 9 | Preflight | Baseline and exact action authority |
| 10 | Act | Approved change and readback |
| 11 | Measure | Metrics, analysis and action |

| Week | Theme | Evidence to carry forward |
|---|---|---|
| 12 | Teach | Case, operating instructions and defense |

The extra time supports repeated tests, deeper disciplinary feedback and a stronger handoff. A persistent agent can remain in a sandbox. The studio does not require an unattended production service. A 10-week adaptation combines selected blocks only after faculty recalculate time and preserve necessary evidence and individual assessment. A 60–90-minute guest lab uses a prepared environment for one small task and teach-back; it does not establish the full course outcomes.

## Studio artifacts that make progress visible

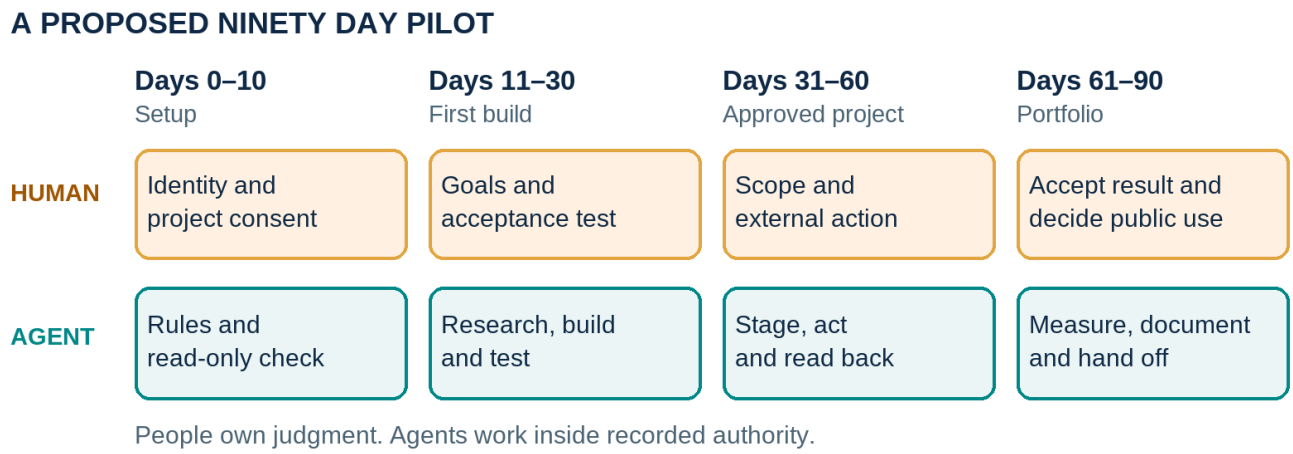
The detailed studio carries the original twelve-week progression into concrete records. Use these alongside the six-session lessons when the faculty-approved format allows deeper technical work. Link existing project evidence rather than creating duplicate paperwork.

| Weeks | Practical work | Visible learner artifact |
|---|---|---|
| 1–3 | Accountable scope; source and rule ledger; tool inputs, errors and permissions. | System Card; source ledger; Tool Card and edge-case test. |
| 4–6 | Supported workspace; persistent-agent routine; single-agent versus bounded team comparison. | Versioned build; Agent Card and action log; comparison of quality, time, cost and integration risk. |
| 7–9 | Representative evaluation cases; governance and accessibility; approved project discovery and preflight. | Test set; risk and accessibility record; exact implementation and recovery brief. |
| 10–12 | Authorized action and live readback; Metrics → Analysis → Action; portfolio, handoff and individual explanation. | Action and readback receipt; MAA comparison; case, operating instructions and teach-back. |

A **System Card** records the user, task, permitted inputs and actions, owner, completion test and stop condition. A **Tool Card** records inputs, outputs, errors and permissions. An **Agent Card** adds the agent's bounded role, source versions, routine, budget, review and handoff. These are records of how the work operates, not new grading categories. Faculty can use an equivalent local format.

## A proposed ninety day pilot

The program guide's ninety-day plan separates human accountability from work an agent may perform. People own identity, relationships, judgment and acceptance. Agents can handle repeatable steps within the recorded scope. The amount delegated is a design choice to test, not a promised productivity percentage.




| Phase | Human decisions and ownership | Agent work within authority |
|---|---|---|
| Days 0–10 Setup | Own identity, accounts, MFA and project consent. | Prepare rules, workspace, setup ledger, authority map and a read-only check. |
| Days 11–30 First build | Supply goals, source material, relationships and the acceptance test. | Research, organize sources, build, test and prepare recovery steps. |
| Days 31–60 Approved project | Approve scope, data, risk and the exact external action. | Stage the bounded change, compare it, execute only within authority and read back. |
| Days 61–90 Portfolio and handoff | Explain contribution, accept the result and decide portfolio use. | Measure, rerun checks and package the case, operating instructions and handoff. |

**Day 10:** safe access and a read-only check. **Day 30:** first accepted artifact, tests, source and process log, and individual explanation. **Day 90:** approved implementation or a clearly labeled simulation, handoff and portfolio decision. A delayed permission does not authorize acting; use the faculty-approved simulation route. These are proposed checkpoints, not a promise that every student gets a client in ninety days.

# 8 Reusable learner prompts

## Six field task brief

1. **Problem and user:** Who needs what result, and what is the starting point?
2. **Ownership and authority:** Who is accountable, and which actions are permitted?
3. **Inputs and evidence:** Which sources/data may be used, and which version?
4. **Deliverable and test:** What will be made, and how will someone check it?
5. **Resources and stop:** What time, calls or spending is allowed across the whole job, and when must it stop?
6. **Version and handoff:** Where is the editable source, who reviews it, and how will another person use it?

## AI use and verification note

“I used [tool and visible model/version, or unavailable] for [tasks]. My consequential decisions were [decisions]. I checked [claims/tests] against [evidence]. I rejected or corrected [example]. The remaining limits are [limits].” Include representative records sufficient for assessment; avoid unnecessary private transcripts.

## User review note

“[Reviewer role; name only with consent] attempted [task] using version [version]. This was [real-user review / disciplinary review / classroom simulation]. They observed [difficulty or success]. I changed [change] and checked [test]. Operational acceptance is [accepted / declined / pending / not requested].”

## MAA and handoff

**Metrics:** starting and final results under stated conditions, including quality and review effort. **Analysis:** what the comparison supports and cannot establish. **Action:** continue, revise or stop, with an owner. Link the final artifact, editable source, tests and update instructions; state the recipient's acceptance separately from successful delivery.

# 9 Optional advanced study models and agent teams

Keep this discussion after the core course. Compare tools on the same task, inputs and acceptance test. Record factual quality, completion, reproducibility, accessibility, elapsed time, human review effort and observable resource use. A personal preference suggests something to test; it is not a universal model ranking.

Distinguish the underlying model, application, account plan, enabled tools and permissions. Use campus-approved options and date the tool-lab comparison. Do not grade access to a particular vendor, subscription or local-model machine. Advanced students may add a supervised connector, retrieval, automated evaluation or local model when the project and authorization justify it.

A team of agents still needs a human owner, pinned evidence, one writer per shared artifact, a bounded reviewer, a whole-job budget and a verified handoff. Coordination is a skill to demonstrate; multiple simultaneous agents are not an enrollment requirement.

# Sources lineage and release notes

The v0.4 revision makes the faculty welcome, classroom example and teaching prompts more practical and personal. It preserves v0.3's detailed lessons, visual program material, operational method, learning loop, assessment framework and source qualifications. Teaching records describe their specific relationships; they do not establish endorsement or adoption of this curriculum.

The v0.3 foundation unified the complete v0.2 curriculum with selected visual and practical material from the University AI Builder Program Guide v2.2, developing the faculty companion v0.1. This v0.4 edition keeps that foundation and the recommended assessment weights and gates. The master AI Implementation Guide remains the operational source; its R4 cross-reference points to the related curriculum. This educator guide is not a new operational master, institutional endorsement or evidence that a cohort has completed the course.

The v0.2 foundation remains outcome-first framing, an explicitly illustrative project, a visible student journey, user-review evidence, clearer baseline/final comparison and individual transfer, and a more actionable faculty adoption plan. Course hours, detailed lessons, learning outcomes and rubric anchors remain proposals for faculty review.

Public program and current resources. Public AI implementation reference.

## Educational design references

The following sources inform the teaching design. They do not endorse this course or establish institutional approval.

- **UNESCO student competencies 2024.** The human centred, ethical, technical and system design dimensions inform the breadth of the proposed outcomes. This framework concerns school curricula and is not a university accreditation standard. AI competency framework for students.
- **UNESCO teacher competencies 2024.** The treatment of pedagogy and professional learning informs faculty preparation and the staff development adaptation. AI competency framework for teachers.
- **TEQSA assessment guidance 2025.** Varied evidence and assurance of learning inform the combination of portfolio work and observed individual demonstration. This educational resource is not a claim that the course meets a jurisdiction's regulatory requirements. Enacting assessment reform in a time of artificial intelligence.
- **NIST Generative AI Profile 2024.** Evaluation and documentation practices inform claim verification, testing and stated limitations. AI 600-1 publication.
- **CAST UDL Guidelines 3.0.** Engagement, representation, and action and expression inform accessible materials and equivalent demonstrations of learning. CAST guidelines.

Editable curriculum source for team members: University AI Builder editable source. Team access is required. The public program page above is the reader entry point.

Each institutional adaptation records its parent version, academic owner, changed outcomes/examples, local tool rules and approval state. Update the editable source first, export a versioned PDF and preserve the prior release. Use the institution's normal records for academic decisions. Team-only operating links, client information and unpublished evidence stay out of the faculty-facing release.