



AI-RESILIENT ASSIGNMENT DESIGN

Master Framework for Learning and Assessment
in the Age of Generative AI

VERSION 1.0 • FALL 2026

by

Taner Onsay, PhD & Sebnem Onsay

Course-agnostic framework for homework, laboratories, projects,
programming assignments, reports, computational activities, and related assessments

DESIGN ASSUMPTION

Students have unrestricted access to capable AI systems.

*Published by AI First Engineering (AIFE) LLC • Version 1.0 • Fall 2026
Developed from cross-course work in engineering and computing education.*



1. Purpose and Status

This document is a course-agnostic design authority for creating and auditing assignments in an era when students can readily access powerful generative AI. It is intended to guide instructors and AI assistants during assignment development without forcing different disciplines or assignment types into a single format.

The framework applies to homework, laboratories, projects, programming assignments, reports, computational investigations, design activities, and other graded or formative work. Course-specific master plans remain authoritative for content, sequencing, disciplinary methods, and local policies. This document governs the higher-level question: what should an assignment require so that meaningful evidence of learning remains visible when AI is available?

Living Framework: Version 1.0 is intentionally revisable. New course-design documents, classroom observations, student outcomes, and advances in AI capability should be checked against this framework and incorporated when they materially improve it.

2. Foundational Purpose

The purpose is not simply to incorporate AI into education. It is to adapt education to the AI evolution for the benefit of students - preserving and strengthening human reasoning, judgment, creativity, responsibility, and disciplinary understanding while teaching students to use AI as a powerful intellectual and professional partner.

As AI makes answers easier to obtain, education must place greater value on asking good questions, doing meaningful work, making sound decisions, validating results, and understanding what the answers mean.

3. Core Design Position

Do not attempt to make AI unusable. Assume students can ask AI to solve, write, code, explain, revise, and critique. Instead, design the assessment so that the intended learning cannot be represented solely by a polished final artifact.

AI may assist. The student must still provide meaningful evidence of learning.

The practical goal is AI-resilient assessment, not AI-proof assessment. No prompt, agent instruction, or prohibition can reliably prevent delegation. The durable strategy is to make student participation, reasoning, evidence, validation, adaptation, and accountability part of what is assessed.

4. What Counts as Evidence of Learning?

A useful assignment distinguishes the final artifact from the learning that produced it. Depending on the discipline, evidence may include:

- student-generated measurements, observations, experimental records, traces, or datasets;
- problem formulation, assumptions, representations, model selection, or algorithm selection;
- predictions made before computation or before an AI-generated answer is revealed;
- intermediate calculations, code behavior, debugging evidence, plots, parameter choices, or iteration history;
- comparison of alternatives, tradeoffs, sensitivity, uncertainty, or edge cases;



- validation against physics, known cases, independent methods, tests, data, or expected behavior;
- interpretation, critique, engineering/computational judgment, and justified decisions;
- ability to explain, reproduce, modify, transfer, or defend the submitted work.

5. The Eight Universal Design Principles

1. Require Authentic Student-Generated Evidence

Whenever the learning objective permits, anchor the assignment in something the student or team actually does: measures, observes, records, chooses, builds, runs, tests, debugs, compares, or decides. The evidence need not be elaborate, but it should connect the submission to a real learning activity rather than only to a prompt that could be delegated.

2. Assess the Process, Not Only the Final Answer

A correct number, polished report, or functioning program is not sufficient evidence by itself. Reward relevant intermediate reasoning: assumptions, representations, model/algorithm choice, test cases, iterations, debugging, calculations, plots, and decisions. The process evidence should be proportional to the learning objective; do not create paperwork merely to prove that work occurred.

3. Require Interpretation and Disciplinary Judgment

Move beyond 'calculate,' 'code,' or 'report.' Ask what the result means, whether it is reasonable, which assumptions matter, what tradeoff is present, what should be changed, or what conclusion is justified. In engineering this may be physical or design judgment; in computing it may be representation, complexity, correctness, invariant, edge-case, or algorithmic judgment.

4. Design Validation Into the Assignment

Validation is part of the task, not an afterthought. Require an appropriate check: units, signs, limiting behavior, physical plausibility, known cases, residuals, independent calculations, experimental comparison, alternative implementation, test cases, boundary/edge cases, or AI-assisted critique followed by human verification.

5. Introduce Variation, Choice, Iteration, or Uncertainty

When appropriate, avoid assignments that collapse to one deterministic prompt and one deterministic answer. Use measured or personalized data, parameter choices, alternative methods, changed conditions, sensitivity studies, debugging branches, design decisions, uncertainty, or iterative improvement. Variation should serve learning rather than merely obstruct copying.

6. Preserve Student Accountability for AI-Assisted Work

Students remain responsible for understanding and defending what they submit. AI may tutor, explain, code-assist, analyze, critique, edit, or review, but the student should be able to explain important reasoning, reproduce essential steps, identify limitations, and respond when conditions change.

7. Verify Transfer of Understanding

Include opportunities to adapt knowledge rather than only repeat it. Change a parameter, condition, dataset, constraint, representation, test case, or problem feature and ask the student to predict, modify, diagnose, or explain the consequence. Transfer is stronger evidence of understanding than reproduction.



8. Balance AI-Augmented Capability with Independent Evidence

Not every activity must prohibit AI, and not every activity should depend on AI. Across a course, deliberately collect evidence of both: (a) the student's ability to work effectively with AI and modern tools, and (b) the student's independent conceptual and procedural understanding. Short in-class, oral, paper-based, or otherwise controlled checks can complement open AI-enabled assignments.



6. Three Capabilities to Assess: KNOW - DO - DIRECT

AI-era assessment benefits from distinguishing three related capabilities. An assignment need not assess all three equally, but the instructor should know which one is being targeted.

Capability	Core Question	Typical Evidence	AI Relationship
KNOW	Does the student understand the concept?	Explanation, prediction, trace, interpretation, short calculation, conceptual distinction.	May be checked with or without AI depending on purpose.
DO	Can the student perform the disciplinary work?	Model, calculate, program, measure, debug, analyze, validate, design, test.	AI may assist, but evidence must show the student can execute and interpret.
DIRECT	Can the student effectively direct AI/tools while retaining responsibility?	Prompt/context quality, decomposition, critique, verification, accepted/rejected suggestions, final judgment.	AI use is part of the competency being assessed.

7. The Assignment Development Protocol

Use this sequence when creating a new assignment. An instructor may perform it directly or instruct an AI assistant to follow it.

1. Define the learning objective first. State what students should understand, be able to do, and/or be able to direct by the end of the assignment.
2. Identify the authentic disciplinary activity. What should the student actually do - observe, model, calculate, code, measure, investigate, compare, debug, design, critique, or decide?
3. Specify the evidence of learning. Decide what artifacts or behaviors will demonstrate the intended learning. Avoid equating polish with mastery.
4. Choose the intended AI role. Examples include Tutor, Coding/Problem-Solving Companion, Analyst, Modeling Partner, Critic/Validator, Workflow Orchestrator, Pre-Submission Reviewer, or Agentic Assistant.
5. Define the human-responsibility boundary. Identify what the student must personally understand, decide, validate, or defend even if AI assists elsewhere.
6. Build validation into the task. Decide what independent evidence will support trust in the result.
7. Add transfer, variation, or iteration where it improves the objective. Do not add artificial complexity solely to frustrate AI.
8. Align the rubric with the evidence. Award meaningful credit to reasoning, evidence, validation, interpretation, and decisions when those are the learning targets.
9. Run the AI Substitution Test and redesign weak components.
10. Run a clarity/workload/fairness audit. Confirm that the assignment is feasible, accessible, appropriately scoped, and that students know what evidence is expected.
11. Prepare the student-facing AI guidance. State permitted/encouraged uses, accountability expectations, and whether any portion requires independent work.
12. After implementation, collect observations and revise the framework or assignment design as needed.



8. The AI Substitution Test

Give the completed assignment instructions to a capable AI system and assume the student is willing to delegate as much work as possible. Which graded components could AI complete with little or no student engagement in the intended learning?

For each vulnerable component, classify the result:

- **ACCEPT** - AI completion is compatible with the learning objective; the skill being assessed lies elsewhere.
- **MODIFY** - AI assistance is appropriate, but stronger student evidence, validation, transfer, or explanation is needed.
- **REDESIGN** - AI can substitute for the core learning activity, so the task should be changed while preserving the learning objective.

The goal of this test is not to defeat AI. It is to reveal where the assignment is accidentally grading AI-accessible product generation instead of the intended human learning.

9. Designing the Rubric

The rubric should measure the learning evidence the assignment was designed to produce. Avoid giving disproportionate credit to presentation quality when the intended learning is modeling, reasoning, computation, experimentation, debugging, validation, or judgment.

- Separate correctness from validation when both matter.
- Reward evidence of appropriate model/algorithm/representation choice, not merely execution.
- Reward interpretation and justified decisions when they are explicit objectives.
- Where relevant, give credit for identifying and correcting failed approaches rather than requiring a sanitized reasoning trail.
- Do not require excessive process documentation that creates workload without improving evidence of learning.
- Use the rubric as a specification for a student pre-submission review agent: the agent may identify missing evidence, but should not be treated as the official grader.

10. Student Pre-Submission Review Agent

A pre-submission review agent is an encouraged application of AI when the course permits it. Its purpose is quality assurance: compare the student's completed work against the assignment requirements and rubric, identify omissions or inconsistencies, and point the student back to evidence.

Recommended operating rule: Audit my submission against every requirement. Identify missing, incomplete, inconsistent, or unsupported items and tell me where you found the evidence. Do not assume that a requirement is satisfied merely because the writing sounds plausible.

The agent does not need to be tamper-proof. Students can reconfigure AI systems. The assessment itself should therefore remain meaningful even if a student chooses to ask AI for more extensive help. The student remains accountable for the submitted work.

11. Assignment-Type Adaptations

Assignment Type	Strong Evidence	Useful AI Roles	Common Vulnerability
Homework	Prediction; model/representation; personalized or chosen parameters; calculations/code; interpretation; validation; changed-case response.	Tutor; coding/problem-solving companion; critic; pre-submission reviewer.	A conventional numerical or coding prompt can be delegated end-to-end.
Laboratory	Prediction; measurements/observations; raw/intermediate data; uncertainty; model-data comparison; troubleshooting; validation; conclusion.	Tutor; analyst; coding companion; critic; report reviewer.	A polished report can hide weak participation or weak understanding.
Project	Problem definition; design choices; decomposition; iterations; evidence trail; tradeoffs; testing; validation; recommendation/defense.	Modeling partner; workflow orchestrator; agentic assistant; critic.	AI can generate a sophisticated artifact while obscuring ownership of key decisions.
Programming / CS	Representation choice; invariants; traces; tests; edge cases; debugging; complexity/correctness reasoning; adaptation.	Coding companion; debugger; critic; test-case generator.	Working code can be generated without algorithmic understanding.
Report / Technical Communication	Accurate representation of completed work; traceable evidence; interpretation; consistency; limitations; conclusions.	Editor; reviewer; requirements auditor.	Writing quality may be mistaken for disciplinary mastery.
Quiz / Checkpoint	Independent prediction, conceptual reasoning, short analysis, critique, transfer.	Often restricted or absent when independent evidence is the purpose.	Overreliance on open AI-enabled work can leave independent understanding unmeasured.



12. Cross-Disciplinary Examples

Engineering Computation

Instead of only asking for a numerical result, ask students to define or approve the model, make a prediction, compute using an appropriate tool, compare with an independent method or measured result, interpret discrepancies, and change one condition to test transfer.

Engineering Dynamics

Anchor work in physical motion where appropriate. Require system/model definition, assumptions, coordinate or method choice, solution, physical interpretation, and validation. Evidence may include video/data, plots, hand checks, limiting cases, or comparison of alternative dynamics methods.

Computer Science / Algorithms

Do not equate generated code with learning. Ask students to predict behavior, select representations, identify invariants, analyze runtime, create or critique test cases, diagnose failures, compare approaches, and adapt an algorithm when constraints or inputs change.

13. AI Roles: A Deliberate Progression

The role of AI should be selected intentionally rather than left implicit. Across a course, responsibility may progressively shift from instructor/AI scaffolding toward student-directed AI-augmented work.

Tutor: Explains concepts, asks questions, provides alternative explanations.

Coding / Problem-Solving Companion: Helps with syntax, algebra, implementation, debugging strategies, or organization while the student owns the underlying method.

Analyst: Helps interpret results, compare cases, identify trends, or suggest useful views of data.

Modeling Partner: Discusses assumptions, representations, simplifications, candidate models, or algorithms; the student approves the model.

Critic / Validator: Challenges assumptions, units, edge cases, correctness, physical plausibility, tests, or conclusions.

Pre-Submission Reviewer: Audits completed work against instructions and rubric; identifies missing or inconsistent evidence.

Workflow Orchestrator: Helps decompose and coordinate multi-step work across tools or artifacts.

Agentic Assistant: Executes larger delegated workflows with explicit human checkpoints, verification, and final responsibility.

14. Independent Evidence Is Complementary, Not Anti-AI

An AI-enabled course can still include controlled individual evidence. Short paper-based checkpoints, oral explanations, in-class predictions, concept questions, live modifications, or other independent checks can establish that students possess the foundational understanding needed to supervise AI well.

Some activities should teach students to work with AI. Other activities should verify what the student can understand and do without AI. A coherent course needs both.



15. Final Assignment QA Checklist

- ☐ Learning objective is explicit and appropriate.
- ☐ The assignment asks students to perform a meaningful disciplinary activity.
- ☐ Required evidence of learning is identifiable.
- ☐ The intended AI role is explicit or intentionally unrestricted.
- ☐ Human-responsibility boundaries are clear.
- ☐ Validation is built into the task where appropriate.
- ☐ Interpretation/judgment is required where appropriate.
- ☐ Transfer, variation, iteration, uncertainty, or edge cases are used when they strengthen learning.
- ☐ The rubric rewards the intended evidence rather than merely polish.
- ☐ The AI Substitution Test has been performed.
- ☐ Any independent-evidence component is purposeful and proportionate.
- ☐ Instructions tell students what they must submit and what they must be able to explain or defend.
- ☐ Workload and documentation requirements are proportional to learning value.
- ☐ The assignment is aligned with the course-specific master plan and policies.

16. Instructions for an AI Assignment-Development Assistant

Before creating or revising an assignment, read this framework and the relevant course-specific master plan. Treat both as design constraints. Preserve the instructor's learning objectives and disciplinary standards. Do not merely add AI language to a conventional assignment.

The AI assistant should:

- identify the intended KNOW, DO, and/or DIRECT outcomes;
- propose authentic evidence of learning;
- identify where AI could substitute for the intended learning;
- recommend redesigns that preserve the objective rather than adding artificial obstacles;
- build validation and transfer into the task where appropriate;
- align rubric categories with evidence;
- flag excessive workload, redundant documentation, or requirements that reward polish over learning;
- distinguish course-specific requirements from this universal framework;
- perform a final AI Substitution Test and report remaining vulnerabilities to the instructor.

17. Version-Control Protocol

When a new master document, course policy, classroom observation, or assessment result becomes available, do not automatically append it. Compare it against the current framework:

- Does it introduce a genuinely new universal principle?
- Does it refine or clarify an existing principle?
- Is it course-specific and therefore better kept in the course master?
- Does it contradict an existing principle? If so, what evidence or context explains the difference?



- Would adopting it make assignment design more meaningful, fair, operational, or sustainable?

Minor clarifications should increment the version (for example, 1.0 to 1.1). Substantial changes to the design philosophy or architecture should produce a major revision.

18. Source Basis for Version 1.0

Version 1.0 was synthesized from existing Fall 2026 course-development materials and the accompanying cross-course design discussion. These sources informed the framework but were not copied as a combined template.

- ENGR 200 - Fall 2026 Master Plan with Cohort Calibration Addendum: engineering computation workflow; progressive AI roles; validation; debugging; human responsibility; KNOW/DO/DIRECT; laboratory, homework, and project philosophy.
- ME 345 - AI-Augmented Engineering Dynamics Master Plan: SEE-THINK-MODEL-SOLVE/ANALYZE-VALIDATE-REPORT; phenomenon-first learning; physical interpretation; model choice; validation; AI as tutor/coding companion/critic; homework evidence and AI-compatible design.
- CSE 331 Fall 2026 Redesign Thread Summary (provisional source): computational judgment; Predict-Discuss-Vote-Debug-Reflect; micro-investigations; AI-enabled collaborative learning paired with independent workflow checkpoints.
- ENGR 200 grading/feedback workflow documents supplied with the source set: considered as operational context for the principle that grading and feedback should trace to explicit evidence and authoritative rubrics/records.

19. Master Design Statement

Assume students have access to capable AI.

Design assignments so that AI can support learning without substituting for the evidence of learning.

Require students to participate, reason, validate, adapt, and remain accountable.

AI ASSISTS • HUMAN UNDERSTANDS • STUDENT VALIDATES • PROFESSIONAL DECIDES

Publication, Copyright & Educational Use

Publisher / Owner: AI First Engineering (AIFE) LLC

Copyright: © 2026 AI First Engineering LLC. All rights reserved.

This resource is part of AI First Engineering educational and AI-assisted instructional materials. Unless otherwise stated, it is intended for personal learning and appropriate noncommercial academic and educational use.

Generally permitted with attribution: personal study and self-directed learning; student educational use; individual engineer self-study; appropriate noncommercial academic use; limited classroom use with proper attribution; and linking students or colleagues directly to AI First Engineering resources.

Permission or licensing is generally required for: commercial or organizational training; company-wide or institution-wide redistribution; uploading materials to internal training portals; commercial education or paid



training programs; commercial consulting use; republishing substantial portions; hosting AI First Engineering materials on another website or repository; modification and redistribution; incorporation into commercial products or services; and resale or sublicensing.

Attribution: Recommended attribution: AI First Engineering LLC, “ASSIGNMENT PRE-SUB REVIEW AGENT” / “AI-Resilient Assignment Design,” Version 1.0 (2026), aifirstengineering.org.

Current terms: Use of this resource is subject to the current AI First Engineering Terms of Use and Educational Use & Copyright guidance published at aifirstengineering.org. If a resource-specific license or notice conflicts with this summary, the resource-specific notice controls.

Educational use permitted with attribution • Commercial or corporate use requires permission