

Connecticut AI Alliance (CAIA)

AI Computing Access Program

Sample Applications: General Education Integration Track

Award Period: Summer 2026 – Fall 2026

Note to Applicants: The following two sample applications are provided to illustrate the level of detail and creativity expected in a competitive submission to the General Education Integration track. All names, institutions, and contact information are fictional. These samples are not intended to be copied verbatim. Reviewers prioritize the quality and thoughtfulness of the AI integration plan over technical sophistication. You do not need prior AI expertise or a computer science background to submit a strong application. If you are unsure about computing requirements, the CAIA Program Office is available to help you develop your resource estimates.

Sample Application 1: AI Literacy in an Undergraduate Nursing Course

Scenario: A nursing professor at a four-year university wants to introduce a three-week AI module into a required clinical reasoning course. Students will use LLM interactions to explore AI-generated patient assessments, evaluate their accuracy, and debate the role of AI in clinical decision-making. The instructor has moderate personal experience with AI tools but has never used them in the classroom. No programming is involved.

Section 1: Instructor Information

1. Full Name

Dr. Karen Delgado

2. Email Address

kdelgado@lakeview.edu

3. CAIA Member Institution

Lakeview University

4. Department / Program

School of Nursing

5. Title / Rank

Associate Professor

6. Primary Discipline

- Humanities
- Social Sciences
- Business / Management / Accounting
- Health Sciences / Nursing / Public Health
- Education / Teacher Preparation
- Natural Sciences
- Engineering (non-CS)
- Arts / Music / Design / Communications
- Criminal Justice / Forensic Science / Legal Studies
- Mathematics / Statistics

- Other

Section 2: Course Information

7. Course Title

Clinical Reasoning and Evidence-Based Practice

8. Course Number

NURS 3200

9. Course Level

- Introductory / 100-level (Gen Ed / First-Year)
- Intermediate / 200–300 level
- Upper division / 400 level
- Graduate

10. Does this course fulfill a general education requirement?

- Yes
- No – but it serves a broad, non-major audience

NURS 3200 is a required course for all BSN students and is also taken as an elective by students in public health and health sciences programs. Approximately 30% of the enrollment consists of non-nursing majors.

11. When will this course be offered?

- Summer 2026
- Fall 2026
- Both Summer + Fall 2026

12. Expected Enrollment per Section

28 students

13. Number of Sections

2 (same instructor, same materials)

14. Course Description

NURS 3200 teaches students to synthesize patient data, apply clinical guidelines, and make evidence-based care decisions. The course uses case-based learning: students work through increasingly complex patient scenarios, identify relevant assessment findings, formulate nursing diagnoses, and develop care plans. Topics include diagnostic reasoning under uncertainty, integrating lab values and vital signs into clinical judgments, interprofessional communication, and patient safety. The course has not previously addressed AI or automated clinical decision support, but both the Connecticut State Board of Nursing and our accrediting body (CCNE) have recently emphasized that nursing graduates should be prepared to work alongside AI-assisted tools in clinical settings. The proposed AI module directly addresses this emerging competency.

Section 3: AI Integration Plan

15. Describe how you plan to integrate AI into this course.

I plan to add a three-week "AI in Clinical Reasoning" module (weeks 10–12 of a 15-week semester) that builds on the clinical case analysis skills students develop in the first nine weeks. The module consists of three interconnected activities:

Activity 1: "Can AI Write a Nursing Assessment?" (Week 10, 1 class session + homework).

Students receive a detailed patient scenario (the same type of case they have been working with all semester). In class, they first complete their own written assessment and preliminary care plan individually. Then, using a CAIA-provided chatbot interface, they give the same patient information to an LLM and ask it to produce a nursing assessment and care plan. The homework assignment asks each student to write a structured comparison (approximately 800 words) analyzing where the AI's output was clinically accurate, where it was wrong or incomplete, and where it provided information the student had not considered. Students submit both their original assessment and the comparison.

Activity 2: "AI Under Pressure" (Week 11, 1 class session). Working in groups of 4, students receive three time-pressured clinical scenarios of increasing complexity (a stable post-op patient, a patient with ambiguous symptoms, and a rapidly deteriorating patient). For each scenario, the group queries the LLM for a suggested assessment and then must decide, under a 10-minute time constraint, whether to follow, modify, or override the AI's recommendation. Each group documents their decision rationale. The full class then debriefs: Which scenarios did the AI handle well? Where was clinical judgment essential? What are the risks of over-reliance? This activity directly addresses the course's learning objective on reasoning under uncertainty.

Activity 3: "AI Policy Position Paper" (Week 12, homework + class debate). Each student writes a 1,000-word position paper arguing for or against a specific policy proposal (assigned randomly): "Hospitals should be required to disclose when AI tools are used in patient care decisions." Students must draw on their experiences from Activities 1 and 2, at least two peer-reviewed articles from the nursing or health informatics literature, and their own clinical practicum observations. In class, students participate in a structured debate, followed by a reflection discussion on how their views evolved.

Throughout all three activities, the LLM is used as a foil and a provocation, not as a teaching tool. The pedagogical goal is not to teach students to use AI but to develop their ability to critically evaluate AI outputs in a domain where errors have real consequences. No programming or technical skills are required; all LLM interaction happens through a web-based chat interface.

16. What type of AI integration are you proposing?

- AI Literacy Module
- LLM Interaction – students use chatbots or language models for disciplinary tasks
- AI-Assisted Analysis
- Critical Evaluation – students assess AI-generated outputs for accuracy, bias, etc.
- Creative AI
- AI-Powered Research
- Build a Simple AI Tool
- Ethical / Policy Debate
- Other

17. How many class sessions will involve AI-related activities?

- 1–2 sessions
- 3–5 sessions (a significant multi-week unit)
- 6+ sessions
- Ongoing / embedded throughout

18. What is your level of prior experience with AI tools?

- Extensive
- Moderate – I have experimented with AI tools but haven't used them in a course
- Beginner
- None

Section 4: Computing Resource Requirements

19. What type(s) of computing resources are you requesting?

- LLM API Access – Google Vertex AI
- GPU Computing – Google Colab
- Not sure

20. How will students interact with the models?

- Through a CAIA-provided web interface or chatbot (no coding required)
- Through Jupyter/Colab notebooks with guided prompts
- Through a custom application built by the instructor or TA
- Directly via API
- Not sure

21a. Which models do you anticipate using?

Gemini Pro. The activities require a model that can produce detailed, fluent clinical text so that the comparison and evaluation exercises are substantive. If CAIA recommends an alternative model for health-related scenarios, I am open to guidance.

21b. Estimated usage per student

Activity 1: approximately 3–5 chatbot interactions per student, each involving a detailed patient scenario prompt (~500 words) and a model response (~400 words). Estimated: ~8K–12K input tokens and ~5K–8K output tokens per student. Activity 2: approximately 6–9 group queries (3 scenarios x 2–3 queries per scenario), shared across groups of 4. Estimated per student: ~4K–6K input tokens and ~3K–5K output tokens. Total per student over the module: approximately 12K–18K input tokens and 8K–13K output tokens. Across 56 students (2 sections): roughly 1M input tokens and 700K output tokens total.

22. Colab GPU access

Not applicable.

23. Estimated total cloud credit budget requested

\$450 (LLM API access only, for 56 students across 2 sections). This is a conservative estimate based on Gemini Pro pricing; the actual cost may be lower. I am requesting a small buffer (~20%) to allow for additional student experimentation beyond the structured activities.

Section 5: Learning Outcomes & Broader Impact

24. What AI-related learning outcomes will students achieve?

- Understand what AI and large language models are at a conceptual level
- Critically evaluate AI-generated content for accuracy, bias, and limitations
- Use AI tools effectively and responsibly in a disciplinary context
- Articulate ethical, social, and policy implications of AI in their field
- Recognize how AI is transforming their discipline or profession
- Produce work that demonstrates meaningful engagement with AI tools
- Develop basic technical fluency
- Other

25. How does this AI integration align with your institution's general education goals?

Lakeview University's 2025–2030 Strategic Plan identifies "AI readiness across all programs" as a Tier 1 priority, and the Provost's Office has encouraged faculty across disciplines to pilot AI integration in existing courses. The School of Nursing's 2026 CCNE self-study explicitly calls for preparing students to evaluate AI-assisted clinical decision support tools, a competency that has no current home in our curriculum. This module provides a concrete, assessable response to that gap. It also aligns with the

university's general education goal of critical thinking, as students must evaluate the reliability and limitations of an information source (in this case, an LLM) in a high-stakes professional context.

26. Will course materials be shareable?

- Yes – I will contribute my AI module/assignments to CAIA's shared resource library
- Yes – I am willing to present my approach at a CAIA workshop or faculty development event
- Yes – the module is designed to be adaptable across disciplines
- Possibly
- No

The three-activity structure (produce and compare, evaluate under pressure, argue a policy position) is deliberately discipline-agnostic in its pedagogy. The clinical scenarios can be swapped for legal cases, business strategy problems, social work assessments, or any other discipline where AI outputs need expert evaluation. I intend to document the module with instructor notes and a "how to adapt" guide for other disciplines.

27. How will you assess whether AI integration improved student engagement or learning?

Students will complete a brief pre-module survey (before Week 10) and post-module survey (after Week 12) measuring their confidence in evaluating AI-generated clinical information, their attitudes toward AI in healthcare, and their self-reported critical thinking about AI outputs. The position paper will be scored on a rubric that includes a "depth of AI engagement" criterion. I will also compare student performance on the clinical reasoning case exam (Week 14) to prior semesters as an informal benchmark.

Section 6: Support & Faculty Development Needs

28. What support would be most helpful?

- One-on-one onboarding session with CAIA technical staff
- Faculty workshop on AI tools for non-CS instructors
- Sample assignments or module templates for my discipline
- A TA or student assistant trained in AI tools
- Technical documentation / quick-start guides
- A peer mentor (another faculty member who has integrated AI in a similar course)
- Guest lecture from an AI researcher
- No additional support needed

29. Would you participate in a CAIA community of practice?

- Yes – very interested
- Maybe
- No

Section 7: Responsible Use & Academic Integrity

30. How will you address academic integrity concerns?

The course syllabus will include an explicit AI use policy stating that LLM tools may only be used during designated course activities and not for completing other assignments (e.g., writing the position paper itself). The comparison assignment (Activity 1) requires students to submit both their own pre-AI assessment and the AI-generated version side by side, making the distinction between human and AI work transparent. The position paper rubric awards points for original argumentation and personal clinical experience, which cannot be replicated by an LLM. All AI interactions occur through the CAIA-provided interface, which logs session transcripts; students are informed that transcripts are available to the instructor for academic integrity purposes.

31. How will you help students think critically about ethical dimensions of AI?

The entire module is structured around critical evaluation. Activity 1 asks students to identify clinical errors in AI output, directly confronting the assumption that AI-generated text is authoritative. Activity 2 forces students to make real-time decisions about when to trust vs. override AI recommendations, simulating the ethical tensions they will face in clinical practice. Activity 3 requires students to take and defend a policy position on AI transparency in healthcare, drawing on published evidence and their own clinical observations. I will also assign two readings from the nursing ethics literature on AI in healthcare (Topol, 2023; Ronquillo et al., 2021) to ground the discussion in disciplinary scholarship.

32. Will students work with sensitive or real-world data?

- ☐ No – only synthetic, public, or curated academic datasets
- ☐ Yes

All patient scenarios are fictional cases created by the instructor. No real patient data, EHR records, or identifiable clinical information is used at any point.

Section 8: Prior Experience

33. Have you previously integrated AI tools into any of your courses?

- ☐ Yes
- ☐ No – this would be my first time

34. Have you previously received computing resources through CAIA?

- ☐ Yes
- ☐ No

Section 9: Applicant Certification

35. Electronic Signature

Karen Delgado

36. Date

May 20, 2026

Sample Application 2: AI-Augmented Primary Source Analysis in a First-Year History Course

***Scenario:** A history professor at a community college teaches a large first-year survey course that fulfills a general education humanities requirement. She wants to use LLM interactions to help students practice analyzing primary source documents and to build information literacy skills by comparing AI-generated historical interpretations to scholarly sources. The instructor has beginner-level experience with AI tools and requests onboarding support from CAIA.*

Section 1: Instructor Information

1. Full Name

Prof. Rebecca Stanton

2. Email Address

rstanton@millbrook.edu

3. CAIA Member Institution

Millbrook Community College

4. Department / Program

Department of History and Social Sciences

5. Title / Rank

Assistant Professor

6. Primary Discipline

- ☐ Humanities (English, History, Philosophy, Languages, etc.)
- ☐ Social Sciences
- ☐ Business / Management / Accounting
- ☐ Health Sciences / Nursing / Public Health
- ☐ Education / Teacher Preparation
- ☐ Natural Sciences
- ☐ Engineering (non-CS)
- ☐ Arts / Music / Design / Communications
- ☐ Criminal Justice / Forensic Science / Legal Studies
- ☐ Mathematics / Statistics
- ☐ Other

Section 2: Course Information

7. Course Title

United States History to 1877

8. Course Number

HIST 1110

9. Course Level

- ☐ Introductory / 100-level (Gen Ed / First-Year)
- ☐ Intermediate / 200–300 level
- ☐ Upper division / 400 level
- ☐ Graduate

10. Does this course fulfill a general education requirement?

- Yes – Humanities distribution requirement and Information Literacy competency
- No

11. When will this course be offered?

- Summer 2026
- Fall 2026
- Both Summer + Fall 2026

12. Expected Enrollment per Section

40 students

13. Number of Sections

1

14. Course Description

HIST 1110 is a first-year survey of American history from indigenous civilizations through Reconstruction. The course fulfills the college's Humanities distribution requirement and contributes to the Information Literacy general education competency. Students come from all programs across the college; the majority are not history majors. The course emphasizes primary source analysis as its central skill: students read, contextualize, and interpret historical documents (letters, speeches, legal texts, images) in every unit. A recurring challenge is that first-year students often struggle to distinguish between credible and unreliable interpretations of historical evidence. The proposed AI integration leverages this challenge as a teaching opportunity.

Section 3: AI Integration Plan

15. Describe how you plan to integrate AI into this course.

I am proposing a semester-long AI thread woven into the course's existing primary source analysis assignments, plus a culminating reflective exercise. The AI is not a separate module; it is integrated into the work students are already doing.

"AI as Historian" Exercises (5 assignments, weeks 3, 5, 8, 11, and 13). The course already includes five primary source analysis assignments spread across the semester (one per major unit: Colonial Era, Revolution, Early Republic, Antebellum, and Civil War/Reconstruction). I propose adding an AI component to each. In the current format, students read a primary source, answer guided analysis questions, and write a short interpretive paragraph. In the revised format, students will complete their own analysis first (as before), then pose the same analysis questions to an LLM using the CAIA chatbot interface. They will then write an additional paragraph (200–300 words) evaluating the AI's interpretation: Did the AI identify the correct historical context? Did it introduce anachronisms, factual errors, or unsupported claims? Did it miss nuances that a historian would catch? How does the AI's analysis compare to what we discussed in class?

Each assignment is designed so that the AI's response will contain a mix of accurate and problematic elements. For example, when asked to interpret an excerpt from a Federalist Paper, an LLM might provide a plausible summary but miss the rhetorical strategy, misattribute authorship to the wrong Federalist author, or impose a modern political framework on an 18th-century debate. These are exactly the kinds of errors that teach students what good historical analysis requires.

Culminating Reflection (Week 14). At the end of the semester, students write a 600–800 word reflective essay responding to the question: "What can AI do well as a historian, and what can it not do? What did evaluating AI outputs teach you about how historians think?" This essay asks students to draw on all five AI exercises and articulate what they have learned about source analysis, historical reasoning, and information reliability. The reflection counts as the course's Information Literacy competency assessment.

In-Class Demonstration (Week 2). Before the first AI exercise, I will conduct a 30-minute in-class demonstration showing students how LLMs work at a non-technical level (what training data is, why models "hallucinate," why they sometimes sound authoritative when wrong). This sets expectations and gives students a framework for the critical evaluation they will practice all semester.

16. What type of AI integration are you proposing?

- ☐ AI Literacy Module
- ☐ LLM Interaction
- ☐ AI-Assisted Analysis
- ☐ Critical Evaluation
- ☐ Creative AI
- ☐ AI-Powered Research
- ☐ Build a Simple AI Tool
- ☐ Ethical / Policy Debate
- ☐ Other

17. How many class sessions will involve AI-related activities?

- ☐ 1–2 sessions
- ☐ 3–5 sessions
- ☐ 6+ sessions
- ☐ Ongoing / embedded throughout course activities

18. What is your level of prior experience with AI tools?

- ☐ Extensive
- ☐ Moderate
- ☐ Beginner – I have limited experience but am eager to learn
- ☐ None

Section 4: Computing Resource Requirements

19. What type(s) of computing resources are you requesting?

- ☐ LLM API Access – Google Vertex AI
- ☐ GPU Computing – Google Colab
- ☐ Not sure

20. How will students interact with the models?

- ☐ Through a CAIA-provided web interface or chatbot (no coding required)
- ☐ Through Jupyter/Colab notebooks
- ☐ Through a custom application
- ☐ Directly via API
- ☐ Not sure

21a. Which models do you anticipate using?

Gemini Pro. I need a model that generates fluent, detailed prose so that the "errors" students must find are subtle and intellectually engaging rather than obviously nonsensical. If CAIA staff recommend a different model for this type of humanities exercise, I am open to that guidance.

21b. Estimated usage per student

Each of the 5 AI exercises involves 2–4 chatbot interactions per student: students paste the primary source excerpt and analysis questions (~300–500 words input), receive a response (~300–500 words output), and may ask 1–2 follow-up questions. The in-class demonstration adds a small amount of shared usage. Estimated per student over the semester: approximately 10K–15K input tokens and 8K–12K output tokens. Across 40 students: roughly 400K–600K input tokens and 320K–480K output tokens total.

22. Colab GPU access

Not applicable.

23. Estimated total cloud credit budget requested

\$350 (LLM API access only, for 40 students in 1 section). Usage is light because each interaction involves short prompts and responses. I am requesting a 25% buffer to accommodate students who want to experiment further on their own.

Section 5: Learning Outcomes & Broader Impact

24. What AI-related learning outcomes will students achieve?

- Understand what AI and large language models are at a conceptual level
- Critically evaluate AI-generated content for accuracy, bias, and limitations
- Use AI tools effectively and responsibly in a disciplinary context
- Articulate ethical, social, and policy implications of AI in their field
- Recognize how AI is transforming their discipline or profession
- Produce work that demonstrates meaningful engagement with AI tools
- Develop basic technical fluency
- Other

25. How does this AI integration align with your institution's general education goals?

Millbrook Community College's general education framework includes an Information Literacy competency that requires students to "evaluate the credibility, relevance, and accuracy of information sources." This competency was written with traditional sources in mind (academic journals, news media, websites), but AI-generated text is rapidly becoming one of the most common information sources students encounter. By asking students to evaluate LLM-generated historical interpretations against scholarly standards, this module updates the Information Literacy competency for the AI era. The college's Digital Strategy Committee has identified "AI across the curriculum" as a 2026–2027 priority, and this course would serve as a pilot for the humanities division.

26. Will course materials be shareable?

- Yes – I will contribute my AI module/assignments to CAIA's shared resource library
- Yes – I am willing to present my approach at a CAIA workshop
- Yes – the module is designed to be adaptable across disciplines
- Possibly
- No

The "AI as Historian" exercise structure is directly transferable to any discipline that involves evaluating source material: literature courses could have AI analyze a poem; political science courses could have AI interpret a Supreme Court opinion; sociology courses could have AI analyze interview data. I will document the framework with a discipline-adaptation guide and sample prompts for at least three other fields. I would be glad to present this approach at a CAIA workshop or the Connecticut Community College faculty development day.

27. How will you assess whether AI integration improved student engagement or learning?

The culminating reflection essay (Week 14) will be scored on a rubric assessing depth of critical engagement with AI outputs, quality of meta-cognitive reasoning about historical analysis, and specificity

of examples drawn from the five exercises. I will compare scores on the primary source analysis assignments to prior semesters (same instructor, same prompts, without the AI component) as an informal measure. End-of-semester evaluations will include three questions about the AI exercises. I will also track whether the AI component affects course completion rates, which is an institutional priority at Millbrook.

Section 6: Support & Faculty Development Needs

28. What support would be most helpful?

- One-on-one onboarding session with CAIA technical staff
- Faculty workshop on AI tools for non-CS instructors
- Sample assignments or module templates for my discipline
- A TA or student assistant trained in AI tools
- Technical documentation / quick-start guides
- A peer mentor (another faculty member who has integrated AI in a similar course)
- Guest lecture from an AI researcher for one class session
- No additional support needed

As a beginner with AI tools, I would especially value a one-on-one onboarding session to test the chatbot interface with my actual course materials before the semester begins, and a peer mentor who has done something similar in a humanities or social sciences context. I would also welcome a guest lecture from an AI researcher during the Week 2 demonstration session, as a domain expert could answer student questions about how LLMs work more effectively than I can.

29. Would you participate in a CAIA community of practice?

- Yes – very interested
- Maybe
- No

Section 7: Responsible Use & Academic Integrity

30. How will you address academic integrity concerns?

This is a critical concern for a history course where writing is the primary mode of assessment. My approach is to make AI use visible and structured rather than prohibited. The AI exercises explicitly require students to submit both their own analysis AND the AI's output, so there is no incentive to secretly use the AI for the human-written portion. The evaluation paragraph asks students to critique the AI, which requires genuine understanding of the material. The position paper and reflection essay are graded on personal voice, specific references to in-class discussion, and clinical use of evidence, all of which are difficult to replicate with an LLM. The syllabus will state clearly: "In this course, you will use AI in specific, structured assignments. Using AI to complete other coursework without authorization is an academic integrity violation." I will also discuss AI and academic honesty directly in the Week 2 demonstration.

31. How will you help students think critically about ethical dimensions of AI?

The entire semester-long AI thread is built around critical evaluation. By repeatedly encountering AI outputs that are confidently wrong, factually incomplete, or anachronistically framed, students develop a healthy skepticism about AI-generated content that transfers far beyond this course. The Week 14 reflection essay requires students to articulate what AI cannot do well in historical analysis and why. The in-class demonstration (Week 2) includes a brief discussion of training data bias: whose voices are overrepresented in the data that LLMs learn from, and how might that affect AI interpretations of history? This connects directly to the course's existing themes about whose perspectives are preserved and privileged in the historical record.

32. Will students work with sensitive or real-world data?

- ☐ No – only synthetic, public, or curated academic datasets
- ☐ Yes

All primary source documents are drawn from published, public-domain historical collections (Library of Congress, National Archives, Avalon Project). No student data or personal information is used.

Section 8: Prior Experience

33. Have you previously integrated AI tools into any of your courses?

- ☐ Yes
- ☐ No – this would be my first time

34. Have you previously received computing resources through CAIA?

- ☐ Yes
- ☐ No

Section 9: Applicant Certification

35. Electronic Signature

Rebecca Stanton

36. Date

May 22, 2026

These sample applications are provided by the CAIA Program Office for illustrative purposes only. All names, institutions, and contact information are fictional. Questions? Contact caia@newhaven.edu.