Connecticut AI Alliance (CAIA)
AI Computing Access Program

Sample Applications: Capstone / Senior Design Projects Track

Award Period: Summer 2026 – Fall 2026

Note to Applicants: *The following two sample applications are provided to illustrate the level of detail and specificity expected in a competitive submission to the Capstone / Senior Design Projects track. All names, institutions, and contact information are fictional. These samples are representative of common application scenarios and are not intended to be copied verbatim. Reviewers value clear articulation of AI relevance, realistic resource estimates, and attention to responsible AI practices.*

Sample Application 1: Course-Wide Capstone Access for Multiple Teams

Scenario: *A CS faculty member supervises a summer capstone course with 22 students forming 4 teams. Projects are not yet finalized, so the application requests general-purpose GPU and LLM API access to support a range of anticipated AI/ML workloads.*

Section 1: Applicant Information

1. Full Name
Dr. Elena Marchetti

2. Email Address
emarchetti@silverfield.edu

3. CAIA Member Institution
Silverfield University

4. Department / Program
Department of Computer Science

5. Title / Rank
Associate Professor

6. Phone Number
(860) 555-0142

Section 2: Project Information

7. Number of Capstone/Senior Design Teams
4

8. Team Details

Team 1

a. Project Title
To Be Determined (AI/ML Application Area)

b. Project Description

Projects for this course section have not yet been finalized. Students will form teams during the first two weeks of the summer session and select projects in consultation with the instructor. All projects are required to incorporate AI or machine learning as a core technical component. Based on prior offerings, anticipated project areas include: natural language processing applications (e.g., document summarization, sentiment analysis), computer vision systems (e.g., object detection, image classification for accessibility), and LLM-powered tools (e.g., domain-specific chatbots, retrieval-augmented generation pipelines). Each team will be expected to produce a working prototype, a technical report, and a final presentation. Specific project goals, methods, and expected outcomes will be documented in individual project proposals submitted by Week 3 of the course.

c. Team Size

5–6 students

d. Course / Program Context

CS 490: Senior Capstone Project (required for B.S. in Computer Science)

e. Industry or Community Partner

To be determined per team

f. Expected Project Timeline

- ■ Summer 2026
- ■ Fall 2026
- ■ Both (Summer + Fall 2026)

Note: *Teams 2, 3, and 4 follow the same structure as Team 1. In a real submission, you may copy the team block for each team or simply note that all four teams share the same parameters (course, timeline, team size) and that project topics will be determined in the first weeks of the course. The form's "Add another team" feature will generate additional blocks.*

Section 3: Computing Resource Requirements

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

- ■ GPU Computing – Google Colab Pro/Pro+
- ■ GPU Computing – JupyterHub (managed cluster)
- ■ LLM API Access – Google Vertex AI

10. Describe the computational workloads your project(s) will require.

Because specific projects have not yet been selected, we are requesting a flexible allocation that can accommodate a range of AI/ML workloads typical of undergraduate capstone projects. Based on the last three offerings of this course, anticipated workloads include: fine-tuning pre-trained models (e.g., BERT, ResNet, or small Transformer-based architectures) on moderate-sized datasets (10K–500K samples); training CNNs or vision Transformers for image classification or object detection tasks; running inference with LLMs via API for NLP-focused projects (summarization, classification, conversational agents); and iterative experimentation during model selection and hyperparameter tuning. Most workloads will require T4 or equivalent GPU instances with 16 GB VRAM. None of the anticipated projects require A100-class hardware. Peak usage is expected during weeks 6–10 of the course when teams move from prototyping to final model training.

11a. Which models do you plan to use?

Gemini Pro (primary); PaLM 2 for text classification tasks; possibly Gemini Ultra for a limited number of evaluation runs if one team pursues a multimodal project.

11b. Estimated number of API calls or tokens per month

Approximately 2–3 million input tokens and 500K–1M output tokens per month across all four teams. This estimate assumes moderate usage during development (weeks 1–5) and heavier usage during integration and evaluation (weeks 6–10).

11c. Use case description

RAG pipelines for domain-specific Q&A; chatbot prototyping; text classification and summarization; code generation assistance for teams exploring automated testing or documentation tools.

12. Estimated total cloud credit budget requested

$4,500 ($2,500 for GPU compute via Colab Pro+ across 22 students for ~10 weeks; $2,000 for Vertex AI API credits across 4 team projects)

Section 4: Educational Impact & Learning Outcomes

13. How will access to these computing resources enhance student learning outcomes?

Access to cloud GPU resources and LLM APIs will directly address a persistent bottleneck in our capstone course: students have historically been limited to CPU-only training on personal laptops, which forces them to work with toy-sized datasets and prevents meaningful experimentation with modern deep learning architectures. With Colab Pro+ access, teams will be able to fine-tune pre-trained models on realistic datasets, iterate on hyperparameters within practical timeframes, and gain hands-on experience with the GPU-accelerated workflows used in industry. LLM API access will enable teams working on NLP projects to build functional prototypes that interact with state-of-the-art models, rather than relying on small local models that poorly represent production capabilities. These resources will prepare students to work with cloud AI services upon entering the workforce, and they will practice cost-aware resource management, a skill rarely addressed in coursework.

14. Will this project produce any shareable deliverables?

- ■ Open-source code repository (e.g., GitHub)
- ■ Technical report or paper
- ■ Poster or presentation (e.g., at CAIA Research Symposium)
- ■ Deployed application or demo
- ■ Dataset or benchmark
- ■ Other

15. Does this project involve collaboration with another CAIA institution?

- ■ Yes
- ■ No

Section 5: Responsible AI & Data Practices

16. Does your project involve any of the following?

- ■ Human subjects data (requires IRB approval)
- ■ Protected health information (PHI / HIPAA)
- ■ Personally identifiable information (PII)
- ■ Biometric data
- ■ Synthetic media generation
- ■ None of the above

17. Safeguards and approvals

N/A. Should any team's finalized project involve sensitive data categories, the instructor will require IRB consultation before data collection begins. A responsible AI checklist is built into the project proposal template.

18. How will your team ensure responsible and ethical use of the AI computing resources?

All students complete a mandatory Responsible AI module during Week 1 of the course, covering bias awareness, data governance, and the ethical implications of AI system deployment. Teams are required to include a "Limitations and Ethical Considerations" section in their final reports. API keys and cloud credentials are managed centrally by the instructor through a shared project account with per-team budget alerts to prevent overuse. Students sign an acceptable use agreement at the start of the course. Any project involving user-facing AI outputs must include appropriate disclaimers and transparency measures.

Section 6: Prior Support & Additional Funding

19. Have you previously received computing resources through CAIA?

- ■ Yes
- ■ No

20. Is this project supported by any other funding sources?

- ■ Institutional funding
- ■ Federal grant
- ■ Industry sponsorship
- ■ No other funding
- ■ Other

Section 7: Department Chair / Program Director Acknowledgment

21. Department Chair / Program Director Name

Dr. Robert Anand

22. Department Chair / Program Director Email

ranand@silverfield.edu

Section 8: Applicant Certification

23. Electronic Signature

Elena Marchetti

24. Date

May 5, 2026

Sample Application 2: Targeted Senior Design Project with Industry Partner

Scenario: *A mechanical engineering faculty member sponsors a single senior design team of 4 students building an AI-powered predictive maintenance system for a regional manufacturing partner. The project spans both summer and fall semesters and has well-defined computing needs.*

Section 1: Applicant Information

1. Full Name

Dr. James Okafor

2. Email Address

jokafor@ridgeport.edu

3. CAIA Member Institution

Ridgeport State University

4. Department / Program

Department of Mechanical Engineering

5. Title / Rank

Assistant Professor

6. Phone Number

(203) 555-0278

Section 2: Project Information

7. Number of Capstone/Senior Design Teams

1

Team 1

a. Project Title

AI-Driven Predictive Maintenance for CNC Machining Operations

b. Project Description

This senior design project, conducted in partnership with Precision Metalworks LLC (Bridgeport, CT), aims to develop a machine learning system that predicts tool wear and potential failures in CNC milling operations using time-series vibration and acoustic emission sensor data. The team will collect labeled sensor data from the partner's production floor (approximately 200K samples across normal, degraded, and failure states), then train and evaluate deep learning models including 1D-CNNs, LSTM networks, and Transformer-based time-series classifiers. The system will output tool condition predictions and maintenance alerts through a lightweight dashboard. A secondary objective is to use an LLM-based interface that allows shop floor operators to query the system in natural language (e.g., "Which tools are due for replacement this week?"). Expected outcomes include a trained model achieving at least 90% classification accuracy on held-out test data, a functional dashboard prototype, and a technical report suitable for conference submission.

c. Team Size

4 students

d. Course / Program Context

ME 495/496: Senior Design I & II (two-semester capstone sequence required for B.S. in Mechanical Engineering, ABET-accredited program)

e. Industry or Community Partner

Precision Metalworks LLC (Bridgeport, CT). The partner is providing sensor data from their CNC milling machines, access to their production environment for validation testing, and a dedicated industry mentor (their Director of Operations) who meets with the team biweekly.

f. Expected Project Timeline

- ■ Summer 2026
- ■ Fall 2026
- ■ Both (Summer + Fall 2026)

Section 3: Computing Resource Requirements

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

- ■ GPU Computing – Google Colab Pro/Pro+
- ■ GPU Computing – JupyterHub (managed cluster)
- ■ LLM API Access – Google Vertex AI

10. Describe the computational workloads your project(s) will require.

The primary computational workload involves training deep learning models on time-series sensor data. The dataset comprises approximately 200K labeled vibration and acoustic emission samples, each consisting of multi-channel signals at 10 kHz sampling rate over 2-second windows. We plan to train and compare three model architectures: a 1D-CNN (approximately 2M parameters), a stacked LSTM network (approximately 1.5M parameters), and a lightweight Transformer encoder (approximately 5M parameters). Each training run takes approximately 30–45 minutes on a T4 GPU. We estimate 40–60 training runs over the project duration for hyperparameter sweeps, ablation studies, and cross-validation experiments. Peak GPU usage will occur during August (model development) and October–November (final evaluation and ablation studies). A T4 GPU with 16 GB VRAM is sufficient for all anticipated workloads.

11a. Which models do you plan to use?

Gemini Pro for the natural language query interface and for generating plain-language summaries of maintenance predictions.

11b. Estimated number of API calls or tokens per month

Approximately 500K–800K input tokens and 200K–300K output tokens per month. Usage will be concentrated in Fall 2026 when the dashboard and NL interface are under active development.

11c. Use case description

A retrieval-augmented generation (RAG) pipeline that allows operators to query the predictive maintenance system in natural language. The LLM interprets user questions, retrieves relevant prediction data and maintenance logs from a vector database, and generates concise, actionable responses (e.g., tool replacement schedules, anomaly summaries).

12. Estimated total cloud credit budget requested

$2,800 ($1,800 for GPU compute via Colab Pro+ for 4 students over ~20 weeks; $1,000 for Vertex AI API credits for the NL query interface)

Section 4: Educational Impact & Learning Outcomes

13. How will access to these computing resources enhance student learning outcomes?

This project sits at the intersection of mechanical engineering and AI, giving students direct experience applying deep learning to a real-world industrial problem. Without GPU resources, training even moderately sized neural networks on time-series sensor data would take prohibitively long on personal hardware, forcing the team to work with downsampled data or oversimplified models. Cloud computing

access will allow students to train on the full dataset and conduct meaningful comparisons between architectures. The LLM API component will expose students to modern NLP pipeline design (RAG, prompt engineering, output evaluation) in a domain-relevant context. Importantly, this project prepares students for the growing demand in Connecticut's manufacturing sector for engineers who can bridge domain expertise with AI capabilities. The industry partnership ensures that student work is grounded in real operational constraints and data quality challenges.

14. Will this project produce any shareable deliverables?

- ■ Open-source code repository (e.g., GitHub)
- ■ Technical report or paper
- ■ Poster or presentation (e.g., at CAIA Research Symposium)
- ■ Deployed application or demo
- ■ Dataset or benchmark
- ■ Other

15. Does this project involve collaboration with another CAIA institution?

- ■ Yes
- ■ No

The team will consult with Dr. Priya Nair in the Data Science department at Millhaven Community College, who has expertise in time-series anomaly detection and has agreed to serve as an external advisor. Two data science students from Millhaven may contribute to the data preprocessing pipeline as part of their own coursework.

Section 5: Responsible AI & Data Practices

16. Does your project involve any of the following?

- ■ Human subjects data (requires IRB approval)
- ■ Protected health information (PHI / HIPAA)
- ■ Personally identifiable information (PII)
- ■ Biometric data
- ■ Synthetic media generation
- ■ None of the above

17. Safeguards and approvals

N/A. The sensor data used in this project is machine-generated (vibration, acoustic emission, spindle load) and does not contain any personally identifiable or human subjects information. A data sharing agreement between Ridgeport State University and Precision Metalworks LLC is in place, covering data handling, IP protections, and publication rights.

18. How will your team ensure responsible and ethical use of the AI computing resources?

The team will follow Ridgeport State University's IT security policies for cloud resource management. All API keys and cloud credentials will be stored securely using environment variables and will not be committed to version control. The industry partner's proprietary data will be stored in an access-controlled cloud storage bucket with encryption at rest. Students will document model limitations and potential failure modes in their technical report, including conditions under which the predictive maintenance system should not be relied upon for safety-critical decisions. The team will also complete CAIA's Responsible AI training module.

Section 6: Prior Support & Additional Funding

19. Have you previously received computing resources through CAIA?

■ Yes

■ No

20. Is this project supported by any other funding sources?

■ Institutional funding

■ Federal grant

■ Industry sponsorship

■ No other funding

■ Other

Precision Metalworks LLC is providing in-kind support (sensor hardware, data access, and staff time for mentorship) valued at approximately $8,000. No direct financial support has been provided to the university for this project.

Section 7: Department Chair / Program Director Acknowledgment

21. Department Chair / Program Director Name

Dr. Linda Ferraro

22. Department Chair / Program Director Email

lferraro@ridgeport.edu

Section 8: Applicant Certification

23. Electronic Signature

James Okafor

24. Date

May 12, 2026
