

AI Course Project Tracks

ECE 57000, Purdue University

David I. Inouye

Overview and Introduction

The course project is a significant component of this course, designed to provide you with a hands-on opportunity to delve deeply into an area of AI/ML that excites you. It is your chance to apply the theoretical concepts from lectures to a substantial, practical challenge. While discussing with fellow classmates about your project is encouraged, *this project should be done **individually** and will be graded individually.*

To accommodate a wide range of interests and career goals—from fundamental research to applied engineering—we offer three distinct project tracks. Whether you’re passionate about verifying scientific claims, building real-world applications, or pushing the boundaries of research, there’s a track for you. The table below provides a high-level comparison to help you choose the track that best aligns with your goals:

Feature	TinyReproductions	ProductPrototype	NovelResearch
Primary Goal	Verify a known claim	Solve a real problem	Propose a new claim
Core Challenge	Simplification & Distillation	System Integration & Design	Novelty & Rigor
Focus	Scientific Insight	Engineering Application	Research Contribution
Paper Reading	Deep dive into one paper	Broad search for applicable tools	Deep dive into a subfield
Ideal Outcome	An educational demo	A functional prototype	A submittable workshop paper

Deliverables & Evaluation

While the focus of your work will differ between tracks, the final submission deliverables are the same for everyone. Project grades will be determined through a multi-faceted evalua-

tion process, which will involve a combination of AI-assisted tools, manual evaluation by the teaching staff, and possibly a peer-review component.

Deliverables & Evaluation

While the focus and content of your work will differ significantly between tracks, the **final submission deliverables are the same for everyone**. All projects, regardless of the track, must submit the following four items:

1. **Project Report (PDF + LaTeX Source):** A 4-6 page report using the [official ICLR paper style](#). **The paper should be between 4 and 6 pages, neither more or less than these bounds (e.g., 3.5 pages is less than 4). The references section and appendices does not count towards the page length, as is common in AI/ML conferences. Also, figures/tables should neither be excessively small or large (see ICLR papers for reasonable sizes).** The report should be a self-contained document that clearly communicates your project’s motivation, methods, results, and conclusions. Track-specific instructions are below.
2. **Code (ZIP File):** A zip file containing all code. A README.md file must clearly explain the code structure, dependencies, instructions to run the project, and explicitly state which parts of the code were written by you, adapted from prior code, or copied from external repositories. If you are editing prior code, your README should contain exact line numbers that were edited. The code should automatically download any necessary datasets or models if feasible (or if not publicly accessible, give an explanation of how to download the datasets or models or the reason for an exception to this policy such as company proprietary data).
3. **Demo Video (5 minutes):** A presentation summarizing your project. Track-specific content guidelines are provided below.
4. **Poster (PDF):** A poster summarizing the project in the BetterPoster format ([Generation 1](#), [Generation 2](#)), suitable for a virtual poster session. The key is that the key takeaway should be highlighted and very large so it is easy to read from a distance and at a glance. **The poster should have the aspect ratio of 48” x 36” (landscape orientation) and be designed to be viewed on a computer screen.**

Project grades will be determined through a multi-faceted evaluation process. This will involve a combination of AI-assisted tools to assess specific criteria, manual evaluation by the teaching staff, and possibly a peer-review component where you will provide and receive constructive feedback from your classmates.

Important Submission Requirement: To ensure fair and unbiased grading, all project submissions must be **anonymous (except for the 5-min video, which will have at least your voice)**. Your names, student IDs, or any other personally identifiable information should not appear in your report, code, or poster. Submissions will be handled in a manner similar to double-blind research conference reviews.

1 Track 1: TinyReproductions - The Art of Scientific Distillation

This track challenges students to distill and verify **one or more of the core insights** of a significant research paper. The goal is not to replicate state-of-the-art results but to create a minimal, elegant, and fast set of experiments that convincingly demonstrate *why* the paper’s method works. Students should aim to **reproduce the trends and relative performance differences shown in the paper’s key result tables and figures**, using at least one essential baseline for comparison. It’s about building a deep, intuitive understanding by recreating a paper’s experimental story in miniature.

- **Core Challenge:** Clever experimental design and scientific simplification.
- **Focus:** Reproducibility, critical analysis, and pedagogical demonstration.
- **Success Looks Like:** A clear, compelling set of experiments that anyone with a single GPU can run to validate the paper’s central claims and experimental narrative.

1.1 Required Components

- **Paper Selection:** Identify a high-impact paper with a clear, testable claim and accessible baselines.
- **Hypothesis Distillation:** Isolate the core claims and identify the key tables or figures that support them in the original paper.
- **Minimalist Experimental Design:** Thoughtfully select a smaller dataset, a simpler model architecture, and a shorter training schedule that preserves the essence of the original problem.
- **Faithful Implementation:** Implement the paper’s proposed method and at least one primary baseline for comparison.
- **Analysis of Results:** Generate plots and tables that are analogous to those in the original paper and analyze the extent to which your small-scale results match the original conclusions.

1.2 Example Projects

- **Vision (GANs):** Reproduce the core qualitative claims of the Wasserstein GAN (WGAN) paper. The goal would be to show, using a simple dataset like MNIST or Fashion-MNIST, that the WGAN loss function leads to more stable training and avoids mode collapse compared to the original GAN formulation.

- **NLP (Transformers):** Recreate a key ablation study from the “Attention Is All You Need” paper on a small-scale translation task (e.g., IWSLT’14 German-English). The project would demonstrate the performance drop when removing components like multi-head attention or positional encodings, matching the trend in the original paper’s table.
- **RL (Policy Gradients):** Validate the performance improvements of Proximal Policy Optimization (PPO) over a standard A2C baseline on 2-3 classic control environments from the OpenAI Gym (e.g., `CartPole-v1`, `Acrobot-v1`). The goal is to reproduce the learning curves showing PPO’s superior sample efficiency and stability.

1.3 Deliverables

Your deliverables must follow the shared requirements outlined in the “Deliverables & Evaluation” section. The content should be tailored as follows:

- **Project Report:** Structure the report as a reproduction study, detailing the original paper’s claims, the simplified experimental setup, comparative results, and an analysis of the findings.
- **Demo Video:** The presentation should briefly introduce the original paper’s claim, explain your tiny-reproduction setup, and walk through the results, showing the replicated figures/tables and comparing them to the original.

1.4 Evaluation Criteria

- **Scientific Rigor & Experimental Design:** Quality and justification of the simplification. Fairness of the comparison between the method and baseline.
- **Implementation & Reproducibility:** Correctness of the code. Clarity of the `README.md` and ease of running the experiments.
- **Analysis & Written Report:** Depth of the analysis comparing reproduced results to the original. Clarity and quality of the written report.
- **Poster & Video Presentation:** Effectiveness of the poster and video in communicating the project’s goals, methods, and outcomes.

2 Track 2: ProductPrototype - From Problem to Prototype

This track is for students who want to build something real. Students will identify a concrete, real-world problem, define a target user or client, and engineer an AI-powered prototype to solve it. The emphasis is on the end-to-end application, from understanding user needs to deploying a functional system, rather than inventing a new algorithm.

- **Core Challenge:** System design, engineering execution, and mapping user needs to technical solutions.
- **Focus:** Applied AI, user-centric design, and software engineering.
- **Success Looks Like:** A convincing demonstration of a functional prototype that clearly solves a well-defined problem for a specific user, backed by solid engineering principles.

2.1 Required Components

- **Problem Definition:** A clear description of a real-world problem and the target user/client who faces it.
- **User Requirements:** A list of needs and success criteria from the user’s perspective. What would make this tool useful for them?
- **Literature Review & Technology Selection:** A survey of existing AI/ML techniques that could solve the problem, culminating in a justified choice of model(s) and tools for the prototype.
- **System Design:** An architecture plan for the end-to-end system, including data acquisition/processing, model inference, and the user interface (even if it’s a command-line tool).
- **Prototype Implementation:** Building a functional, demonstrable version of the system.

2.2 Example Projects

- **ECE (Wireless Communications):** An “Automated RF Signal Classifier” for a wireless engineer. The prototype would be a Python tool that takes a file with raw I/Q signal data, passes it through a trained CNN, and outputs the predicted modulation scheme (e.g., BPSK, QAM16, GFSK).
- **ECE (Power Systems):** A “Smart Grid Load Forecaster” for a university facilities operator. The system would use a time-series model (like an LSTM) trained on historical energy usage and weather data to predict the next day’s campus energy demand, presented in a simple dashboard.
- **General Engineering:** A “GitHub Issue Triage Bot” for a software development team. The tool would use a fine-tuned BERT-based model to automatically classify new GitHub issues with relevant labels (e.g., `bug`, `feature-request`, `UI/UX`) to streamline the development workflow.

2.3 Deliverables

Your deliverables must follow the shared requirements outlined in the “Deliverables & Evaluation” section. The content should be tailored as follows:

- **Project Report:** Focus on the problem definition, user requirements, system design, technical choices (justified by literature), and an evaluation of the prototype’s performance against the user’s needs.
- **Demo Video:** The video **must** be a screencast that shows the prototype in action. It should state the problem, demonstrate how a user would interact with the tool to solve it, and finally show the result.
- **Poster:** The poster should place a strong emphasis on the “Problem” and “Solution” sections.

2.4 Evaluation Criteria

- **Problem Definition & User Focus:** Clarity and real-world relevance of the problem. Depth of thought regarding the target user’s needs and how the solution addresses them.
- **System Design & Engineering Execution:** Quality, functionality, and robustness of the prototype. Soundness of the technical choices.
- **Written Report & Justification:** Clarity of the report in justifying the design based on user needs and research.
- **Poster & Demo Video:** Effectiveness of the live demo. How well do the poster and video communicate the value proposition of the prototype?

3 Track 3: NovelResearch - The First Step Towards Discovery

This track mirrors the process of conducting original research and writing a workshop-level conference paper. Students will identify a research gap, formulate a novel hypothesis, propose a new method or analysis, and design experiments to provide initial supporting evidence. The final report should be a self-contained paper that tells a complete research story.

- **Core Challenge:** Identifying a meaningful research gap and rigorously testing a novel idea.
- **Focus:** Research methodology, critical thinking, and scientific communication.
- **Success Looks Like:** A self-contained paper that makes a clear, albeit small, novel contribution, with convincing evidence that would be suitable for submission to a reputable AI/ML workshop.

3.1 Required Components

- **Literature Review:** A targeted review of a specific research area to clearly identify a gap, limitation, or unexplored question in prior work.

- **Hypothesis and Contribution:** A precise research question or hypothesis and a clear, bulleted list of the 1-2 novel contributions of the project.
- **Proposed Method:** A formal description of the novel algorithm, model modification, or analysis technique being proposed.
- **Experimental Design:** A rigorous plan to validate the hypothesis, including appropriate datasets, strong baselines, and meaningful evaluation metrics.
- **Implementation and Experimentation:** Implementation of the proposed method and execution of the designed experiments to gather evidence.

3.2 Example Projects

- **Methodology:** Propose a simple, novel data augmentation technique for a specific data modality (e.g., for time-series sensor data from wearables) and demonstrate its effectiveness for improving downstream classification accuracy compared to standard augmentation methods.
- **Analysis:** Conduct a targeted empirical study analyzing a failure mode of a popular model. For instance, design a suite of logic puzzles to systematically test and categorize the reasoning failures of Large Language Models (LLMs), and hypothesize about the architectural reasons for these failures.
- **Theory:** Propose a modification to a standard loss function, such as adding a new regularization term to the cross-entropy loss that encourages robustness to a specific type of noise. Provide theoretical motivation and empirical evidence of its benefit on a relevant dataset (e.g., image classification with label noise).

3.3 Deliverables

Your deliverables must follow the shared requirements outlined in the “Deliverables & Evaluation” section. The content should be tailored as follows:

- **Project Report:** The report must be structured as a formal research paper and include standard sections: Abstract, Introduction, Related Work, Proposed Method, Experiments, and Conclusion.
- **Demo Video:** The video should be a concise, “conference-style” oral presentation of your research paper, walking through the motivation, proposed method, key results, and conclusion.

3.4 Evaluation Criteria

- **Novelty & Problem Formulation:** Significance, originality, and clarity of the identified research gap and proposed contribution.

- **Methodology & Experimental Rigor:** Soundness of the proposed method. Rigor and fairness of the experimental setup, baselines, and evaluation.
- **Quality of Written Paper:** Clarity, structure, and persuasiveness of the research paper. Does it tell a complete and convincing story?
- **Poster & Video Presentation:** Professionalism and clarity of the research presentation.