LLM-Driven Exam Autograder

Project Overview

Help design a lightweight platform that uses cutting-edge Large Language Models (LLMs) to automatically grade scanned, handwritten exams and produce individualized study feedback. Instructors will simply upload PDF scans and a rubric; the system will return a downloadable gradebook plus a brief, AI-generated report for each student.

Key Goals

1. Integrate LLM APIs (e.g., OpenAI o4-mini, Gemini 2.5) for automated scoring and natural-language feedback.
2. Build a simple web front-end for file upload and results download.
3. Maintain an efficient and accurate grading process.
4. Pilot the platform in CSCI 128 and MATH 111 and coordinate with course instructors on grading rubrics and feedback-report design

Desired Student Skill Set (students will learn as they go)

- Python 3
- LLM API implementation and development
- LLM prompt engineering
- Basic front-end development
- Version control with Git/GitHub

Project Details

Requirement: Students must have completed CSCI 128 and MATH 111 before this field session project.

Team Size: 4 – 5 students (ideal).

Work Location: Mines campus and virtual.

Client Liaison: Dr. Zibo Wang, zibowang@mines.edu