

AI for Managers Final Project

Human review package for Team 9. The trained model, saved experiment and candidate quotation prototype are included. The report draft follows the nine-page course structure.

Links

[Open the Colab experiment notebook](#)

[Open the Hugging Face trained adapter](#)

[Open the official California 2024 voter guide](#)

Results

15 cases, five difficult, three conditions, 45 frozen outputs. AI judge strict passes: fine-tuned 7/15; naive 2/15; source-only untuned 3/15. All human second ratings are pending. This is a research demo with known factual failures.

Review first

Read the report. Independently complete all 45 rows in Human Review.xlsx using the blind response viewer and official references. Review candidate speaker context and labels. Add team NetIDs. Then recompute disagreements and approve the report before Canvas submission.

Files

- 01 Report: filled supplied cover, four body pages, Appendices A to C.
- 02 Human Review: blank ratings, rubric and quote review.
- 03 Responses: question, visible answer, raw output and official evidence.
- 04 Quotes: standalone viewer for Garcia and Whitesides, California 27, 2024.
- 05 Data and Results: complete test/evaluation records and split data.
- 06 Prototype and Colab: portable code, notebook and reproducibility records.
- 07 Adapter: trained weights and tokenizer.
- 08 Links and 09 Manifest: URLs and file checksums.

Hosting

Paid GPU inference is inactive. Reading the materials and using the quote viewer needs no GPU. The professor would need only a browser for a later hosted test window. Fresh adapter inference requires the pinned base model and CUDA runtime.

Course submission

One PDF on Canvas, maximum nine pages, due October 15 at 11:59 pm. The final human results and mismatch counts must replace the current pending entries. The assignment does not require a continuously running public app.