

Post-Module Homework

Deadline: Tuesday, August 4, 2026 *before* 11pm.

Reflection (2%)

Please write a response (minimum 250 words) to one of the following prompts:

1. Consider a concrete example like the Amazon hiring case or the COMPAS case. Do you think that it is ever ethically justified to apply different thresholds for different demographic groups in this case, assuming that the classifiers treat those groups differently even after you have improved the training data as much as possible? Defend your answer. If your answer is yes, include any additional conditions that you think are necessary for differential thresholds to be ethically justified.
2. Is LLM sycophancy a form of manipulation? To help you argue for your conclusion, whether yes or no, use the two theories of manipulation (Noggle's norms theory and SRN's hiddenness theory) discussed at the end of the second ethics module.

Your assignment will be graded in the following way:

Pass (1/1): The assignment (i) meets the required word count in a non-repetitive way, (ii) uses some content from the module, and (iii) gives a coherent argument for a position.

Fail (0/1): The assignment misses one or more of the conditions above.

Deliverables: Submit a PDF file named `reflection.pdf` to [MarkUs](#).

Post-Module Survey (1.5%)

Fill out the post-module survey at [this link](#) (also available on our [website](#)).