

Qualification Round

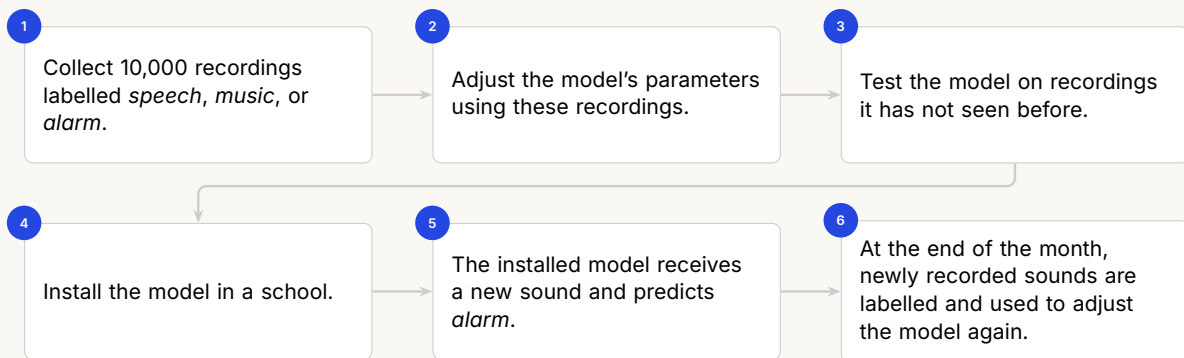
EDITION 2026



Problem A: Machine Learning Lifecycle

(5 Points)

A sound-recognition system is developed through the six steps shown below.

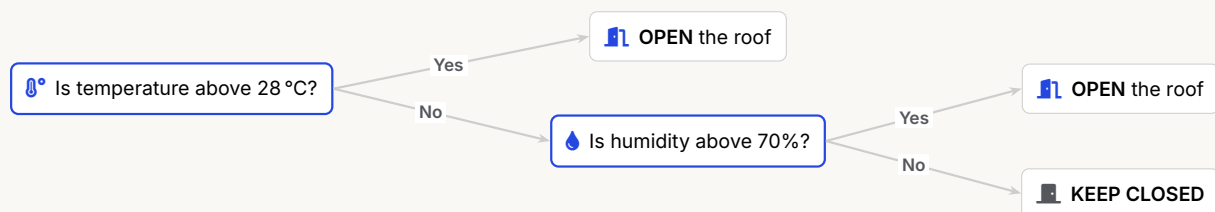


- Label each step as data collection, training, evaluation, deployment, or inference. A step may have more than one label.
- During which steps is the system actually learning? Explain what learning means here.

Problem B: The Greenhouse Tree

(5 Points)

A greenhouse controls its roof with this decision tree:



- What action does the tree predict when the temperature is 26 °C and the humidity is 68%?

- (b) A new CO₂ sensor is installed, and the manager begins logging the conditions together with what the roof should have done:

Temperature	Humidity	CO ₂	Roof
31 °C	60%	700 ppm	OPEN
26 °C	75%	650 ppm	OPEN
25 °C	60%	800 ppm	KEEP CLOSED
27 °C	68%	1,100 ppm	KEEP CLOSED
24 °C	55%	1,400 ppm	OPEN
22 °C	50%	1,550 ppm	OPEN

Adjust the decision tree so that it correctly models every observation in the log.

Problem C: To Fit or Not to Fit

(5 Points)

A model predicts y from x . The training data are:

x	0	1	2	3
y	1	3.2	4.8	7

Two models are proposed:

$$M_1(x) = 0.2x^3 - 0.9x^2 + 2.9x + 1, \quad M_2(x) = 2x + 1.$$

Models are scored using

$$J = \sum (y - \hat{y})^2 + \lambda \sum_i a_i^2,$$

where the second sum runs over the model's coefficients of x , x^2 , and x^3 .

- (a) For $\lambda = 1$, calculate the score J for each model. Which model does the scoring rule select?
- (b) Which model fits the training data better? Explain why the score J nevertheless selects the other model, and what the term $\lambda \sum_i a_i^2$ is for.

Problem D: The Price of Drift

(5 Points)

A language model is trained to maximise a fixed reward so that its outputs better match desired qualities, such as being helpful and aligned with human preferences. Without a constraint, however, the model may drift far from its trusted starting point, the reference model, and learn to game the reward rather than produce genuinely good answers.¹

Let $t \geq 0$ measure the drift from the reference model. Suppose the reward is rt and the drift penalty is βt^2 , with $r, \beta > 0$. The resulting loss is

$$L(t) = -rt + \beta t^2.$$

and training chooses the value of t that minimizes it. In real systems, drift is typically measured using KL divergence; here we use t^2 for simplicity.

¹This is a simplified account of RLHF (reinforcement learning from human feedback). For an introduction, see e.g. <https://huggingface.co/blog/rlhf>.

- (a) Find the shift t^* that minimizes the loss, and the loss value it reaches.
- (b) What happens to t^* as $\beta \rightarrow 0$, and as $\beta \rightarrow \infty$? Explain in one or two sentences what each case means for the language model.
- (c) The reward estimate r is noisy: the engineers only know that $0 < r \leq r_{\max}$. They require the shift to stay below a safety limit T no matter what r turns out to be. Show that this is guaranteed exactly when $\beta \geq \frac{r_{\max}}{2T}$.
- (d) What does the term $\lambda \sum_i a_i^2$ from Problem C have in common with the penalty term βt^2 here?

Problem E: Farming Advice

(5 Points)

An organisation wants to introduce a machine-learning system for farmers in rural Nepal. A farmer provides information about a crop, and the system recommends whether to irrigate, apply treatment, wait, or contact an expert. The organisation wants to deploy an off-the-shelf large language model, such as ChatGPT, Claude, or Gemini. Discuss three opportunities and three risks of introducing this system.



Participation Instructions

- ✓ Write your solutions by hand on sheets of paper or type them on a computer.
- ✓ **Submit your solutions online by the deadline shown on the website.**
Website: <https://imlco.org/submission>
- ✓ You do not need to include the problem statements in your solution document.
- ✓ Show your work to receive full marks. Maximum points: **25**.
- ✓ Clearly label each problem and highlight your final answers.
- ✓ You need to score at least **15/17/20 points** as Junior/Youth/Senior to qualify for the Semi-Final Round. See <https://imlco.org/age-groups> for details.
- ✓ If you have questions, reach out to us at: info@imlco.org