Logistic regression and a neural-network classifier — review sheet

Use this after attempting the worked examples and practice. Explain each claim in words before calculating.

Key ideas

- Define labels, score thresholds and the equality case.
- Fit preprocessing and weights with training data only.
- A nonlinear hidden map changes representational capacity.
- Trace shapes before treating a matrix output as a prediction.
- Confusion counts reveal errors hidden by one accuracy number.
- Compare with simple and scientifically meaningful baselines.
- A fixed synthetic test result is limited evidence, not a general guarantee.

Common errors and repairs

Calling a bounded score a calibrated probability. Calibration requires separate evidence; the present model comparison does not establish it.

Transposing confusion counts without relabeling axes. State actual rows and predicted columns before computing metrics.

Returning zero for every undefined metric. Report absent denominators explicitly; this helper uses None.

Choosing a model from final test scores. Use separate validation for choices and preserve untouched final evidence.

Confusing a classification target with a measured mechanism. The fitted classifier describes the supplied labels; its predictions do not establish a causal mechanism.

Retrieve without your notes

- How many trainable parameters are in the 2-to-8-to-1 network?
- Which denominator defines precision?
- What does removing hidden nonlinearity do?
- Why does a radius comparator belong in the discussion?

Readiness check

- State the meaning of a binary label and a declared score threshold.
- Calculate confusion counts, accuracy, precision and recall with explicit denominators.
- Trace shapes through a two-input, eight-hidden-unit neural network.
- Compare frozen models using the same held-out rows and a transparent baseline.
- Distinguish training improvement, representation choices and evidence for generalization.