

Machine learning · supported focus

International Baccalaureate · IB Diploma · Computer Science · HL

Scope and prerequisites

Supported HL focus. First assessment 2027 target; official PDF returns 403; older acquired brief is final assessment 2026. Remaining guide, assessment and practical requirements retain their recorded holds.

Prerequisites: read the stated quantities and units, use arithmetic and the model conditions below. Each lesson develops its own method before independent transfer.

These are original or explicitly fictional teaching examples, not actual measurements or completed assessed learner investigations.


Data models and responsible machine learning

■ What would explain this observation?

- A model can score well by learning a shortcut in the training data. Evaluation must test the intended task on appropriate unseen examples.
 - Start with a prediction. State the quantities or features you would compare, then decide what evidence could distinguish two explanations.
-

■ Build the model

- A relational database uses tables, keys and relationships to organize data. Machine learning estimates patterns from data. Training and evaluation data serve different roles.
- **primary key** : An attribute set uniquely identifying a record; **data leakage** : Use of information unavailable at the intended prediction time.



■ Choose evidence that can test it

- A primary key uniquely identifies a record; a foreign key relates records. In prediction, leakage can expose information unavailable at the real decision time. Accuracy alone may hide an unbalanced target distribution.
 - Use fictional booking records to identify entities, attributes and relationships. For a learning exercise, use a public non-sensitive dataset, separate training and test data, and describe who may be affected by errors. The current 2027 CS objective scope awaits the acquired guide.
-

■ Work from known quantities

- State the known values and their units. Choose the relation because its assumptions fit this case, then rearrange before substitution.
- Known: a classifier makes 90 correct predictions among 100 cases. $\text{Accuracy} = \frac{\text{correct}}{\text{total}} \times 100 = 90\%$. If 90 cases belong to one class, always predicting that class gives the same accuracy and can still fail every minority-class case.

Example:

A model makes 72 correct predictions out of 80. Find accuracy percentage. Use the same sequence: known quantities $\rightarrow$ model $\rightarrow$ relation $\rightarrow$ substitution $\rightarrow$ unit and interpretation.

■ Check the conclusion and its limits

- High accuracy is not evidence of fairness or causal understanding. A database primary key is not simply whichever field looks important.
- Return to the original observation. Explain what the result supports, which conditions it assumes, and one way to test a competing explanation.

Warn:

A classifier with 90% accuracy must treat every class equally well. This claim is false: High accuracy is not evidence of fairness or causal understanding. A database primary key is not simply whichever field looks important.

Key:

Data models and responsible machine learning: A primary key uniquely identifies a record; a foreign key relates records. In prediction, leakage can expose information unavailable at the real decision time. Accuracy alone may hide an unbalanced target distribution.

■ Runnable trace and boundary

```
tp, fn, fp, tn = 4, 1, 10, 85
total = tp + fn + fp + tn
accuracy = (tp + tn) / total
recall = tp / (tp + fn) if tp + fn else None
print(accuracy, recall)
positives = 0
print(0 / positives if positives else "UNDEFINED")
```

Expected output:

```
0.89 0.8
UNDEFINED
```

Accuracy divides correct predictions by all cases; positive recall divides true positives by actual positives. With no positive cases, recall is undefined rather than zero. The dataset is fictional, with no trained model or personal data.