

Answer Key: Evaluating a Loan-Default Classifier

Default is the positive class.

1. Derive the four confusion-matrix values

- **TP = 36** - 36 customers actually defaulted and were correctly predicted as Default.
- **FN = 4** - 40 actual defaults - 36 correctly detected = 4 missed defaults.
- **FP = 9** - 9 customers who did not default were incorrectly predicted as Default.
- **TN = 51** - 60 actual non-defaults - 9 false positives = 51 correctly identified non-defaults.

2. Confusion matrix

Actual / Predicted	Default	No Default
Default	TP = 36	FN = 4
No Default	FP = 9	TN = 51

3. Calculate the metrics

Accuracy	$(TP + TN) / \text{Total} = (36 + 51) / 100 = 87 / 100 = \mathbf{87\%}$
Precision	$TP / (TP + FP) = 36 / (36 + 9) = 36 / 45 = \mathbf{80\%}$
Recall	$TP / (TP + FN) = 36 / (36 + 4) = 36 / 40 = \mathbf{90\%}$
F1-score	$2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$ $= 2 \times (0.80 \times 0.90) / (0.80 + 0.90)$ $= 1.44 / 1.70 = \mathbf{0.847 \approx 84.7\%}$

4. Interpretation

- **Precision = 80%**: Of all customers the model predicted would default, 80% actually defaulted. The remaining 20% of predicted defaults were false alarms.
- **Recall = 90%**: The model identified 90% of all customers who actually defaulted. It missed 10% of the actual defaults.

Teaching point: Accuracy alone does not describe the types of mistakes a classifier makes. Precision and recall reveal different error trade-offs.