

Machine Learning tips and tricks cheatsheet

stanford.edu/~shervine/teaching/cs-229/cheatsheet-machine-learning-tips-and-tricks

Classification metrics

In a context of a binary classification, here are the main metrics that are important to track in order to assess the performance of the model.

Confusion matrix — The confusion matrix is used to have a more complete picture when assessing the performance of a model. It is defined as follows:

		Predicted class	
		+	-
Actual class	+	TP True Positives	FN False Negatives Type II error
	-	FP False Positives Type I error	TN True Negatives

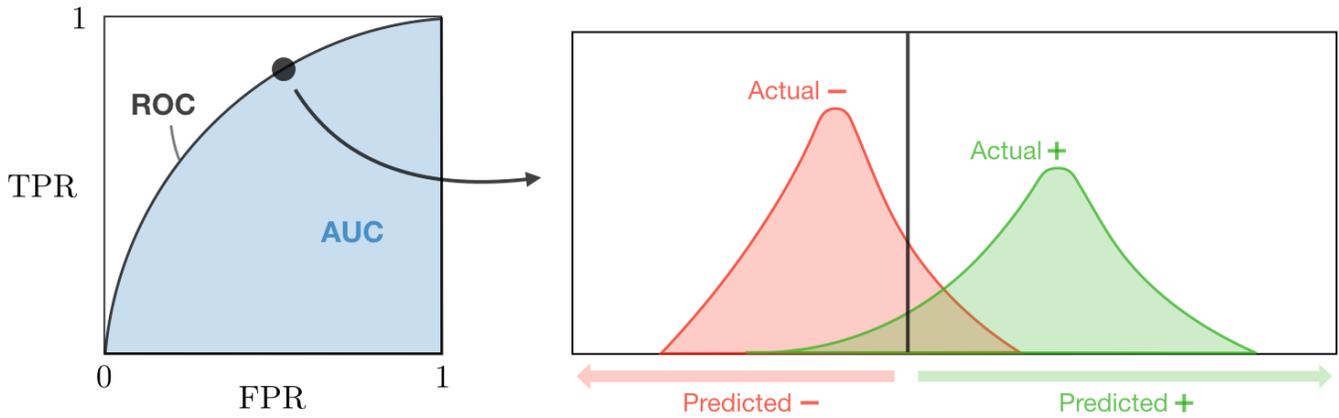
Main metrics — The following metrics are commonly used to assess the performance of classification models:

Metric	Formula	Interpretation
Accuracy	$\frac{TP+TN}{TP+TN+FP+FN}$	Overall performance of model
Precision	$\frac{TP}{TP+FP}$	How accurate the positive predictions are
Recall Sensitivity	$\frac{TP}{TP+FN}$	Coverage of actual positive sample
Specificity	$\frac{TN}{TN+FP}$	Coverage of actual negative sample
F1 score	$\frac{2TP}{2TP+FP+FN}$	Hybrid metric useful for unbalanced classes

ROC — The receiver operating curve, also noted ROC, is the plot of TPR versus FPR by varying the threshold. These metrics are summed up in the table below:

Metric	Formula	Equivalent
True Positive Rate TPR	$\frac{TP}{TP+FN}$	Recall, sensitivity
False Positive Rate FPR	$\frac{FP}{FP+TN}$	1-specificity

AUC — The area under the receiving operating curve, also noted AUC or AUROC, is the area below the ROC as shown in the following figure:



Regression metrics

Basic metrics — Given a regression model f , the following metrics are commonly used to assess the performance of the model:

Total sum of squares	Explained sum of squares	Residual sum of squares
$SS_{\text{tot}} = m \sum_{i=1}^m (y_i - \bar{y})^2$	$SS_{\text{reg}} = m \sum_{i=1}^m (f(x_i) - \bar{y})^2$	$SS_{\text{res}} = m \sum_{i=1}^m (y_i - f(x_i))^2$

Coefficient of determination — The coefficient of determination, often noted R^2 or r^2 , provides a measure of how well the observed outcomes are replicated by the model and is defined as follows:

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}$$

Main metrics — The following metrics are commonly used to assess the performance of regression models, by taking into account the number of variables n that they take into consideration:

Mallow's C_p	AIC	BIC	Adjusted R^2
$SS_{\text{res}} + 2(n+1)\hat{\sigma}^2$	$2[(n+2) - \log(L)]$	$\log(m)(n+2) - 2\log(L)$	$1 - (1 - R^2) \frac{m-1}{m-n-1}$

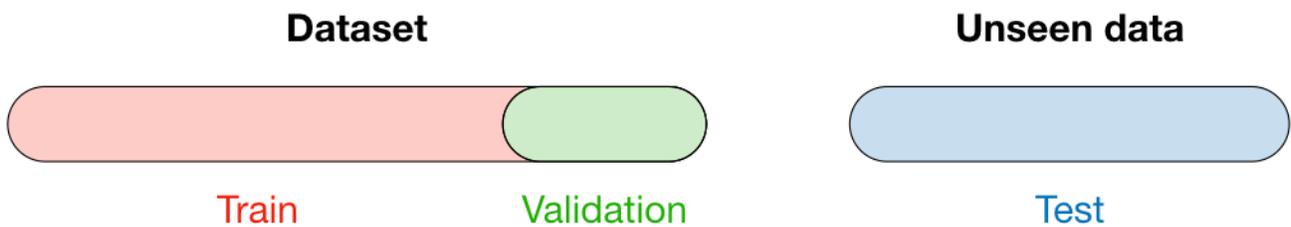
where L is the likelihood and $\hat{\sigma}^2$ is an estimate of the variance associated with each response.

Model selection

Vocabulary — When selecting a model, we distinguish 3 different parts of the data that we have as follows:

Training set	Validation set	Testing set
<ul style="list-style-type: none"> Model is trained Usually 80% of the dataset 	<ul style="list-style-type: none"> Model is assessed Usually 20% of the dataset Also called hold-out or development set 	<ul style="list-style-type: none"> Model gives predictions Unseen data

Once the model has been chosen, it is trained on the entire dataset and tested on the unseen test set. These are represented in the figure below:



Cross-validation — Cross-validation, also noted CV, is a method that is used to select a model that does not rely too much on the initial training set. The different types are summed up in the table below:

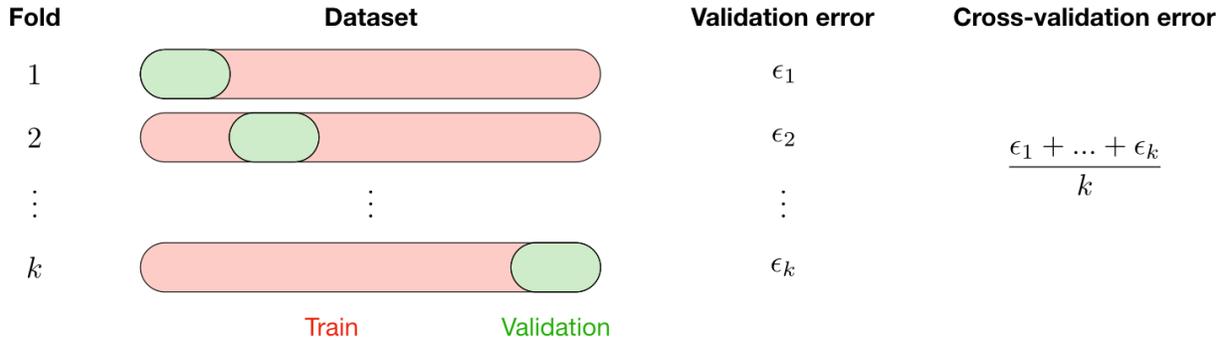
k-fold

Leave-p-out

- Training on $k-1$ folds and assessment on the remaining one
- Generally $k=5$ or 10

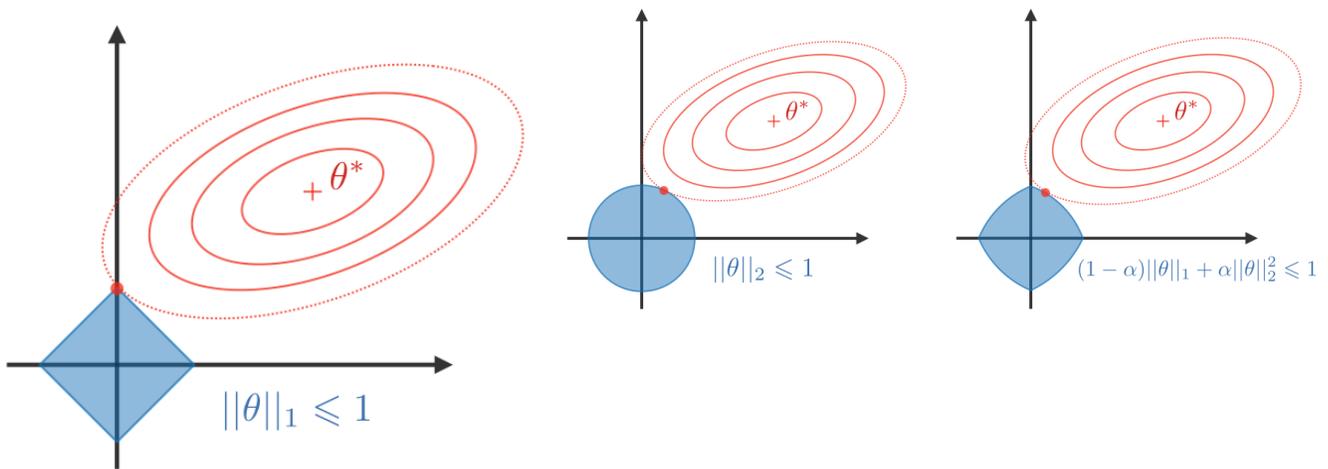
- Training on $n-p$ observations and assessment on the p remaining ones
- Case $p=1$ is called leave-one-out

The most commonly used method is called k -fold cross-validation and splits the training data into k folds to validate the model on one fold while training the model on the $k-1$ other folds, all of this k times. The error is then averaged over the k folds and is named cross-validation error.



Regularization — The regularization procedure aims at avoiding the model to overfit the data and thus deals with high variance issues. The following table sums up the different types of commonly used regularization techniques:

LASSO	Ridge	Elastic Net
<ul style="list-style-type: none"> • Shrinks coefficients to 0 • Good for variable selection 	Makes coefficients smaller	Tradeoff between variable selection and small coefficients



$$\dots + \lambda[(1-\alpha) \|\theta\|_1 + \alpha \|\theta\|_2] \dots + \lambda[(1-\alpha) \|\theta\|_1 + \alpha \|\theta\|_2]$$

$\lambda \in \mathbb{R}, \alpha \in [0, 1]$

Diagnostics

Bias — The bias of a model is the difference between the expected prediction and the correct model that we try to predict for given data points.

Variance — The variance of a model is the variability of the model prediction for given data points.

Bias/variance tradeoff — The simpler the model, the higher the bias, and the more complex the model, the higher the variance.

	Underfitting	Just right	Overfitting
Symptoms	<ul style="list-style-type: none"> • High training error • Training error close to test error • High bias 	<ul style="list-style-type: none"> • Training error slightly lower than test error 	<ul style="list-style-type: none"> • Very low training error • Training error much lower than test error • High variance
Regression illustration			
Classification illustration			
Deep learning illustration			
Possible remedies	<ul style="list-style-type: none"> • Complexify model • Add more features • Train longer 		<ul style="list-style-type: none"> • Perform regularization • Get more data

Error analysis — Error analysis is analyzing the root cause of the difference in performance between the current and the perfect models.

Ablative analysis — Ablative analysis is analyzing the root cause of the difference in performance between the current and the baseline models.

