



Image processing for Earth Observation

3d classification
Best practices
Devis TUIA

Content (6 weeks)


- W1 General concepts of image classification / segmentation
Traditional supervised classification methods (RF)
- W2 Traditional supervised classification methods (SVM)
Best practices
- W3 Elements of neural networks
- W4 Convolutional neural networks
- W5 Convolutional neural networks for semantic segmentation
- W6 Sequence modeling, change detection

Some good practices


- Fighting overfitting
- Crossvalidation
- Spatial splits
- Accuracy measures

The enemy: overfitting


- Sometimes a simple linear prediction is not optimal




- Sometimes a simple linear prediction is not optimal
- We would prefer something nonlinear
- We would avoid to be too close to the training data




- Sometimes a simple linear prediction is not optimal
- We would prefer something nonlinear
- We would avoid to be too close to the training data
- Otherwise, when a new unseen point comes...
- This is called **overfitting**




- Sometimes a simple linear prediction is not optimal
- We would prefer something nonlinear
- We would avoid to be too close to the training data
- Otherwise, when a new unseen point comes...
- This is called **overfitting**



- Sometimes a simple linear prediction is not optimal
- We would prefer something nonlinear
- We would avoid to be too close to the training data
- Otherwise, when a new unseen point comes...
- This is called **overfitting**



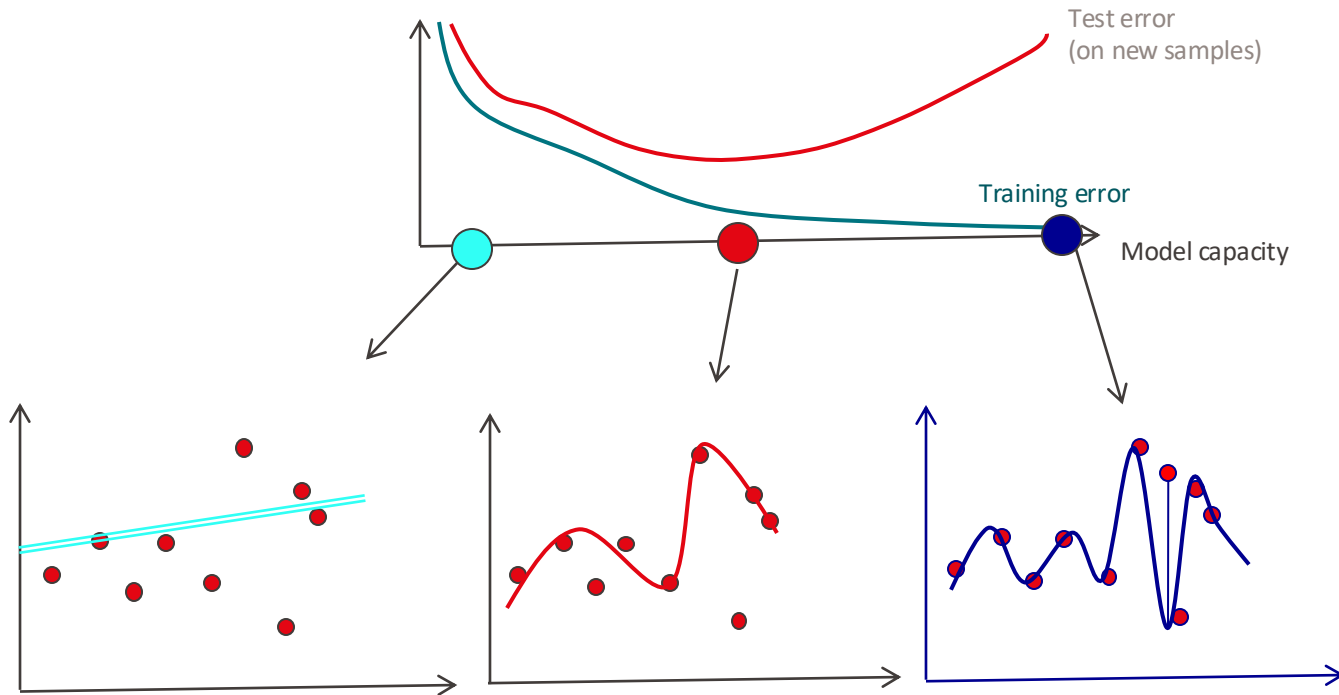
- Sometimes a simple linear prediction is not optimal
- We would prefer something nonlinear



Overfitting (and generalization)

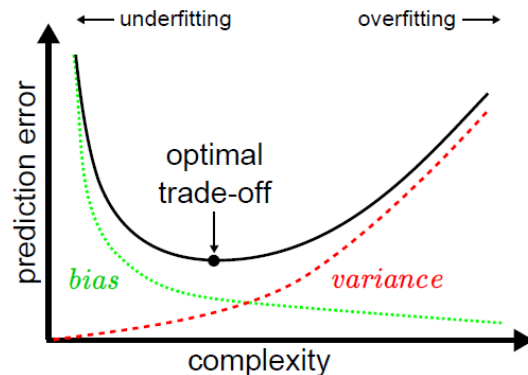
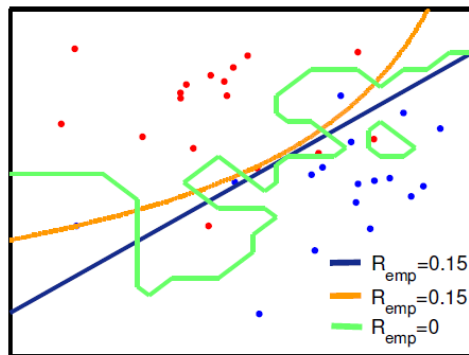
- Overfitting happens when you trust your training data too much.
- Our aim is to model the data structure well, not to represent the training data well
- If we fail miserably on unseen data, the learning was a failure
- So we need to
 - Train well
 - = minimize errors where we can estimate them
 - Limit model overcomplexification
 - = avoid overfitting

Overfitting (and generalization)






More complexity is not always good

- A classification algorithm has free parameters to be tuned
- Typically related to the generalization capability
- A good generalization is a trade-off between :
 - Having a low training error (fitting well enough the training examples)
 - But keeping the model “simple”, with a low complexity



More complexity is not always good

- A classification algorithm has free parameters to be tuned
- Typically related to the generalization capability
- A good generalization is a trade-off between :
 - Having a low training error (fitting well enough the training examples)
 - But keeping the model “simple”, with a low complexity



How do we limit complexity?

- By regularizing solutions

By penalising every time we split into a deeper tree, we avoid taking decisions on smaller and smaller “leaves”

By keeping parameters small, we avoid the model to rely too much on the training data

How do we construct the partitioning?

- Solution
 - A. Stop early (e.g. set a minimum depth)
 - B. Prune the tree using cost:

$$\sum_{m=1}^{|T|} \text{Gini}(m) + \alpha |T|$$


- First grow a very large (deep) tree T_0 . You now have the solution for $\alpha = 0$.
- For each α , there is a subtree $T \subset T_0$
- Increase α and re-run. The regularizer is the price to pay for increasing the number of terminal nodes
- Use a k -fold cross-validation approach to evaluate the error function above. Use the average error as final measure.
- Once minimized, grow an optimal tree using all data.

Tolerance to errors?

- One single missclassification can make the classifier overfit!

$$\min_{\mathbf{w}} \|\mathbf{w}\|^2 + C \sum_i \xi_i$$

- Every time we make a mistake, we penalize by ξ_i (= the distance of the sample from the classification plane)
- If no mistake, $\xi_i = 0$
- Meaning : the bigger the mistake, the more we penalize



How do we limit complexity?

- By regularizing solutions

How do we construct the partitioning?

- Solution
 - A. Stop early (e.g. set a minimum depth)
 - B. Prune the tree using cost:

$$\sum_{m=1}^{|T|} Gini(m) + \alpha|T|$$


- First grow a very large (deep) tree T_0 . You now have the solution for $\alpha = 0$.
- For each α , there is a subtree $T \subset T_0$
- Increase α and re-run. The regularizer is the price to pay for increasing the number of terminal nodes
- Use a k -fold cross-validation approach to evaluate the error function above. Use the average error as final measure.
- Once minimized, grow an optimal tree using all data.

Tolerance to errors?

- One single missclassification can make the classifier overfit!

$$\min_{\mathbf{w}} \|\mathbf{w}\|^2 + C \sum_i \xi_i$$

- Every time we make a mistake, we penalize by ξ_i (= the distance of the sample from the classification plane)
- If no mistake, $\xi_i = 0$
- Meaning: the bigger the mistake, the more we penalize



Optimization objective: \min (errors on training data + regularizer)

How do we limit complexity?

- By regularizing solutions
- By choosing parameters that do not prone overfitting (e.g. min number of samples in each leaf, which has a similar effect of the regularizer seen in previous slide)

How do we limit complexity?

- By regularizing solutions
- By choosing parameters that do not prone overfitting (e.g. min number of samples in each leaf, which has a similar effect of the regularizer seen in previous slide)
- By early stopping (important in neural networks)

In other words: know when to stop minimizing the training error.

(cross) validation is a possible way to estimate that.

Some good practices


- Overfitting
- (Cross)validation
- Spatial splits
- Accuracy measures

Use a validation set when you have it!

- A classification algorithm has free parameters to be tuned
- Cross-validation can allow to estimate the generalization capability (=accuracy on unseen test data)
- If you have a lot of data, you can take out part of it for validating how well your model is doing on **data never seen during training**

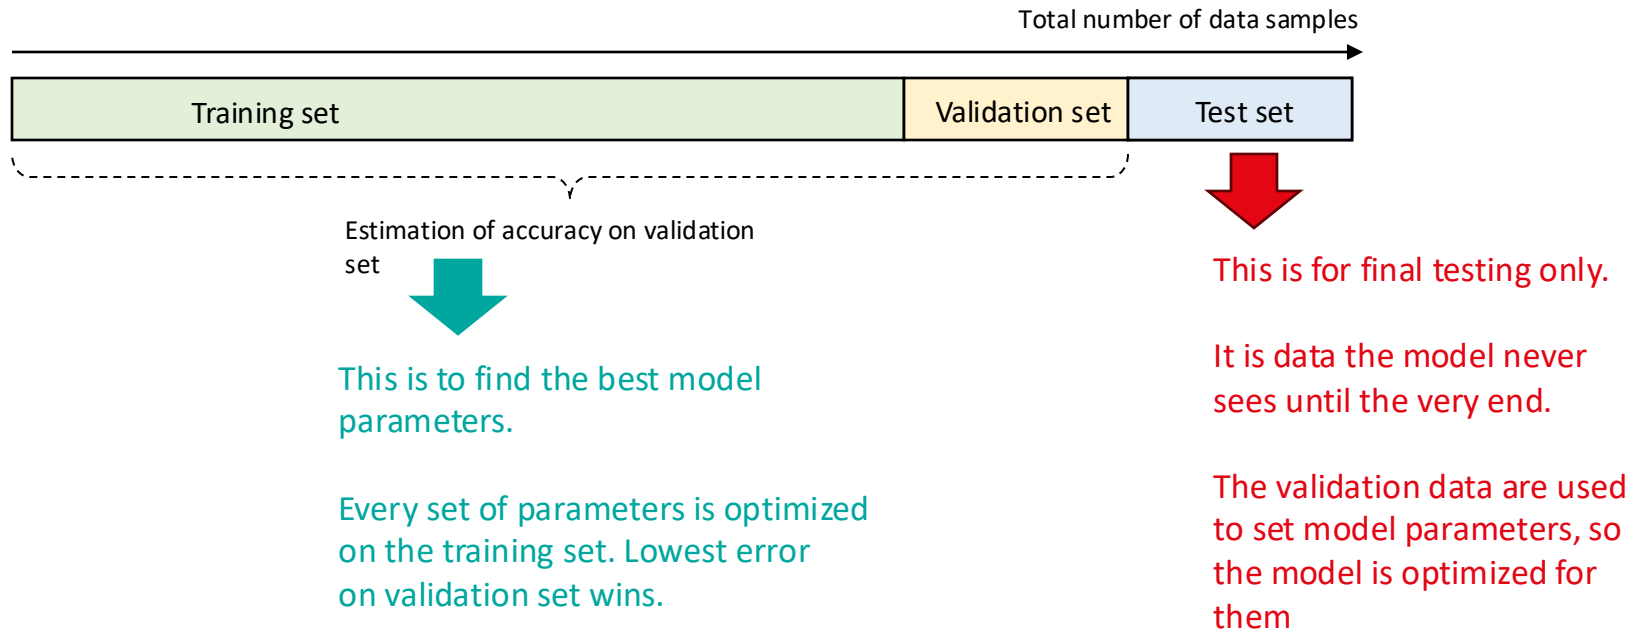
Validation with some left out data

- Validation of model parameters



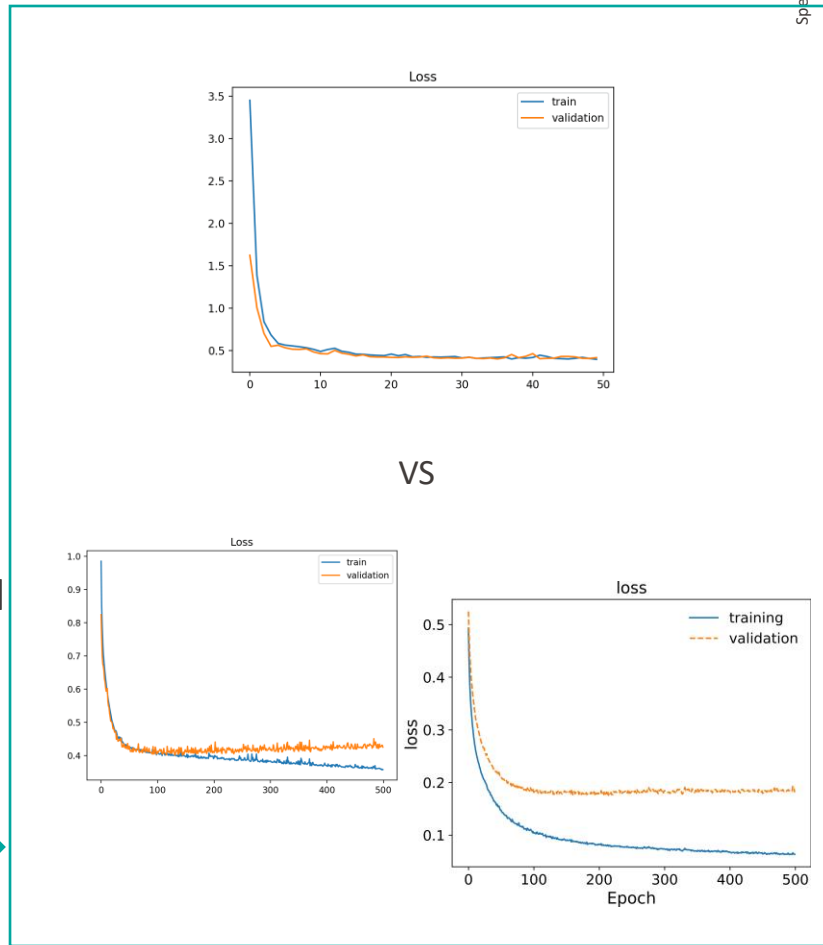
Validation with some left out data

- Validation of model parameters



Use a validation set when you have it!


- A classification algorithm has free parameters to be tuned
- Cross-validation can allow to estimate the generalization capability (=accuracy on unseen test data)
- If you have a lot of data, you can take out part of it for validating how well your model is doing on **data never seen during training**
- This can be used for:
 - Comparing different parameters sets
 - Seeing if you are **overfitting**



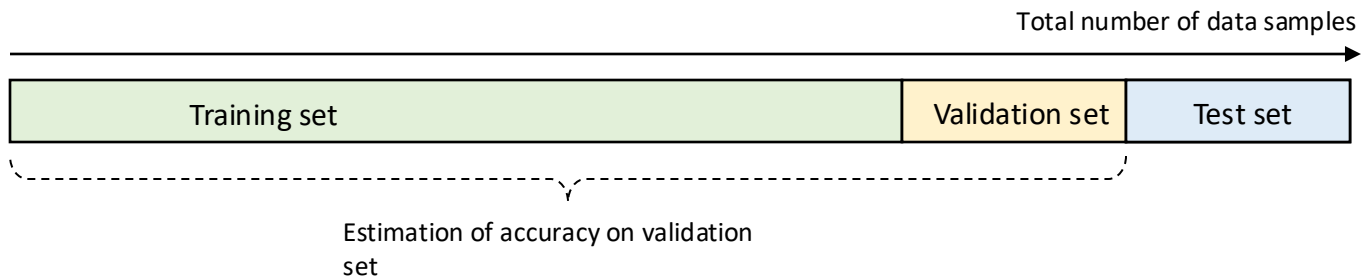
- When you don't have many data, you need to work with what you have...

- Idea behind cross-validation:
 - Split the training examples into different subset or “folds”
 - Leave one fold for testing and estimating accuracy
 - Use the remaining folds for training the model
 - Repeat above steps until all folds have been once tested
 - Take average error as performance metric

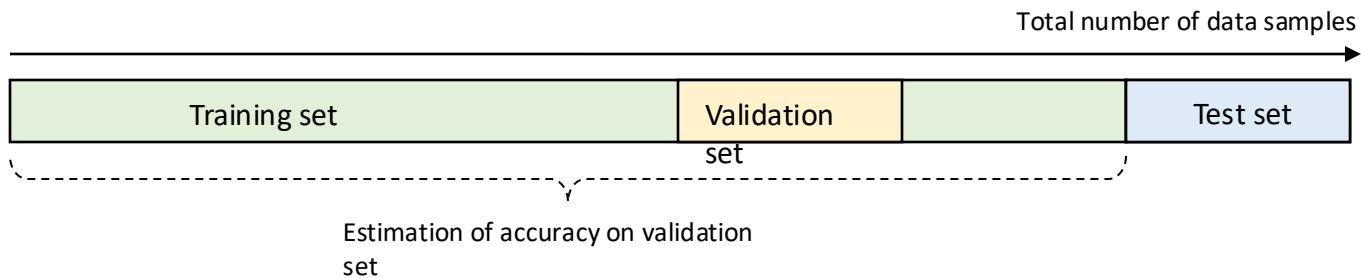
- K-folds cross-validation



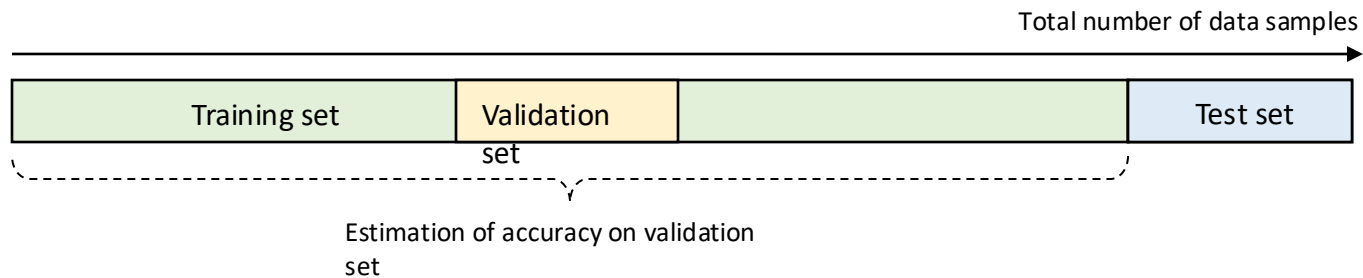
- K-folds cross-validation



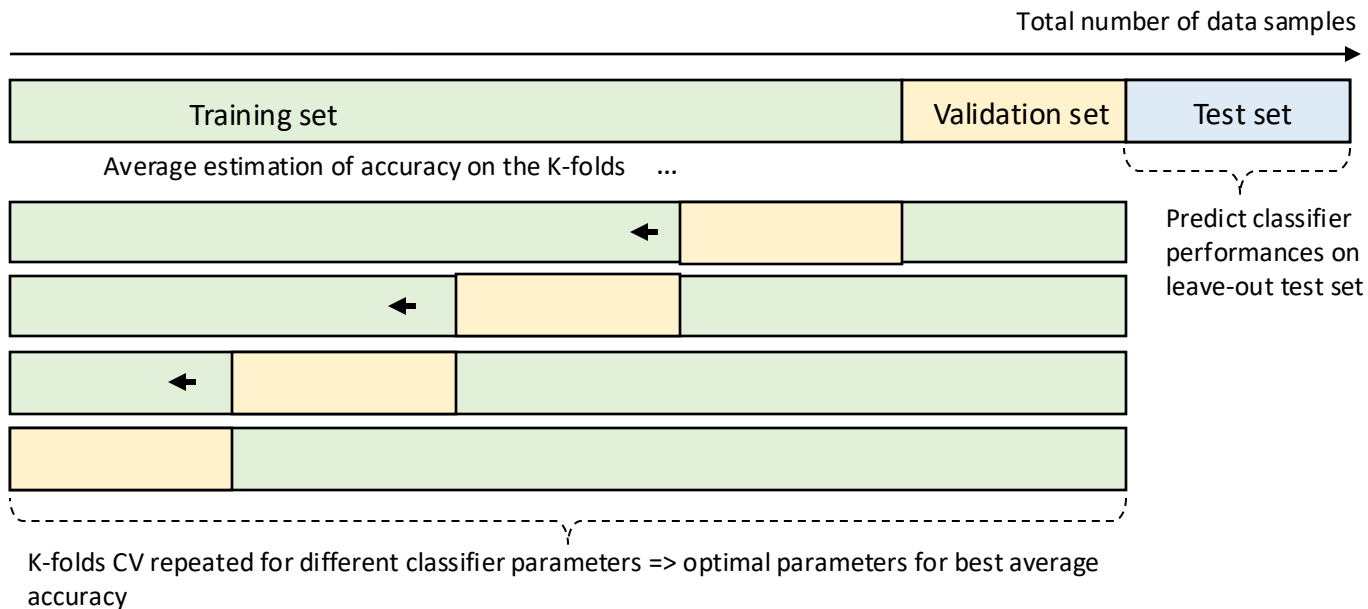
- K-folds cross-validation



- K-folds cross-validation




- K-folds cross-validation



Some good practices

- Overfitting
- Crossvalidation
- Spatial splits
- Accuracy measures

Am I doing a good job?




Predicted classes



True labels

We can compare our predictions to the true labels




Predicted classes




True labels

And we better do it only on areas we didn't use for establishing the model (to avoid positive biases)



Then we compare the true and predicted pixel by pixel



Relates known truth pixels and predictions

Sample from class C wrongly predicted in class A


Sample from class A Correctly predicted in class A

True labels y		Predicted labels y^*				PA
		Class A	Class B	Class C	Class D	
Class A	Class A	n_{AA}				$n_{AA} / n_{A\bullet}$
	Class B		n_{BB}			$n_{BB} / n_{B\bullet}$
	Class C			n_{CC}		$n_{CC} / n_{C\bullet}$
	Class D				n_{DD}	$n_{DD} / n_{D\bullet}$
UA		$n_{AA} / n_{A\bullet}$	$n_{BB} / n_{B\bullet}$	$n_{CC} / n_{C\bullet}$	$n_{DD} / n_{D\bullet}$	OA and AA


- **Overall accuracy:** percentage of correct classifications
= sum of the diagonal of the CM / num of pixels
- **Average accuracy:** average of the per-class accuracies
= the honest one. If you fail on one small class, it is going to show!
- **User's accuracy:** commission errors, % that a pixel classified into a class belongs to that class.
= sum of the column sums of the CM
- **Producer's accuracy:** omission errors, % that a reference sample has been classified correctly.
= sum of the row sums of the CM

User vs producer's

Ground truth



Your map



		Predicted		
		F	U	
True	F	6	0	$6/6 * 100 = 100\%$ Producer Accuracy "how many of true GT has been predicted correctly"
	U	3	1	$1/4 * 100 = 25\%$

$6/9 * 100 = 66\%$
 $1/1 * 100 = 100\%$
 User Accuracy
 "how many of the classifier's predictions were correct"


How to evaluate object detectors

- Object detectors are binary (you detect an object or not)



How to evaluate object detectors

- Object detectors are binary (you detect an object or not)




How to evaluate object detectors

- Object detectors are binary (you detect an object or not)
- In this case, we compare detections (in red) against reference (in blue)
- We first define a detection threshold (e.g. in terms of % overlap between reference and detection)



How to evaluate object detectors

- We then compute the confusion matrix



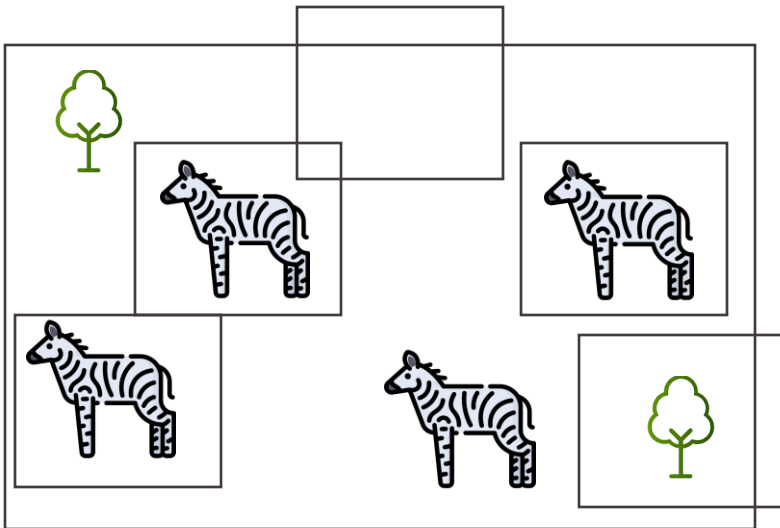
Classification

		Reference		
		+	-	Σ
+	+	TP	FP	TP + FP
	-	FN	TN	FN + TN
Σ		TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives.
- **N** is the total sample size.

How to evaluate object detectors

- We then compute the confusion matrix
- Let's say the model has detected those:

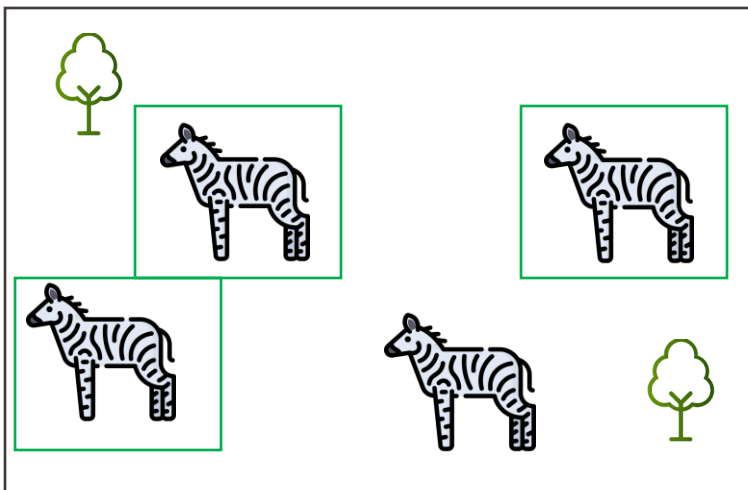


Classification

		Reference		Σ
		+	-	
+	+	TP	FP	TP + FP
	-	FN	TN	FN + TN
Σ		TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives.
- **N** is the total sample size.

How to evaluate object detectors




Classification

		Reference		
		+	-	Σ
+		TP	FP	TP + FP
-		FN	TN	FN + TN
Σ		TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives.
- **N** is the total sample size.

How to evaluate object detectors




Classification

		Reference		
		+	-	Σ
+	+	TP	FP	TP + FP
	-	FN	TN	FN + TN
Σ		TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives.
- **N** is the total sample size.

How to evaluate object detectors

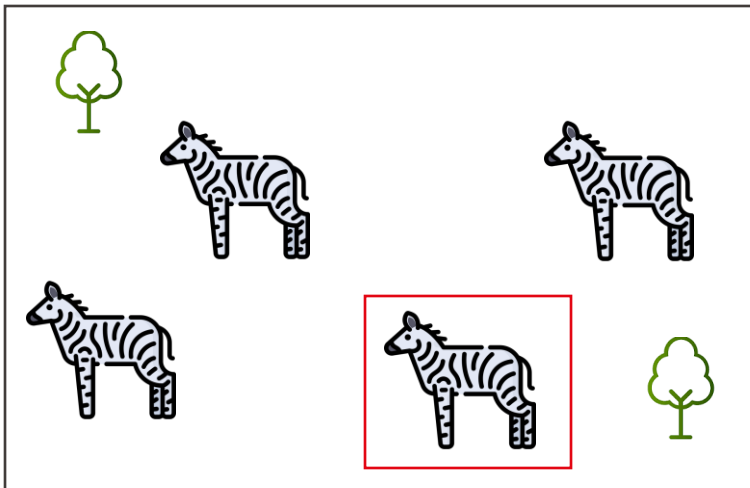


Classification

		Reference		
		+	-	Σ
+	+	TP	FP	TP + FP
	-	FN	TN	FN + TN
Σ		TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives.
- **N** is the total sample size.

How to evaluate object detectors



Classification

		Reference		
		+	-	Σ
+	+	TP	FP	TP + FP
	-	FN	TN	FN + TN
Σ		TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives
- **N** is the total sample size.

How to evaluate object detectors

- We then compute metrics of interest:
 - Precision

$$P = \frac{TP}{TP + FP}$$

- The accuracy of the true positives
How many of the predictions were correct
- Between 0 and 1

		Reference		
		+	-	Σ
Classification	+	TP	FP	TP + FP
	-	FN	TN	FN + TN
	Σ	TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives.
- **N** is the total sample size.

How to evaluate object detectors

- We then compute metrics of interest:
 - Precision
 - **Recall**

$$R = \frac{TP}{TP + FN}$$

- How many of the ground truth objects were detected
- Between 0 and 1

		Reference		
		+	-	Σ
Classification	+	TP	FP	TP + FP
	-	FN	TN	FN + TN
	Σ	TP + FN	FP + TN	N

- **TP** is the number of true positives.
- **TN** is the number of true negatives.
- **FP** is the number of false positives.
- **FN** is the number of false negatives.
- **N** is the total sample size.

How to evaluate object detectors

- We then compute metrics of interest:

- Precision
- Recall
- **F1 score**

$$F1 = \frac{R * P}{R + P} = \frac{2TP}{2TP + FN + FP}$$

- A summary between precision and recall
- Harmonic mean between precision and recall
- Between 0 and 1

In summary

- Machine learning models are very flexible and non-linear, so they can fit your training data very well
- Minimizing training error does not mean good global performance!
- Always (cross)validate and avoid overfitting. You know how now.
- And validate your algorithms so that you don't get cheated by the model. You know how now.