

The PASCAL 2006 Visual Object Classes Challenge

Mark Everingham, Luc van Gool, Chris Williams, Andrew Zisserman

September 2006


The Challenge

- Ten object classes
 - bicycle, bus, car, cat, cow, dog, horse, motorbike, person, sheep
- Classification
 - predict whether at least one object of a given class is present
- Detection
 - predict bounding boxes of the objects in a given class
- Competitions
 - train on supplied data, or any non-test data

Dataset

- Images from 3 sources
 - personal photos, MSR Cambridge, flickr
- Annotation
 - bounding box
 - viewpoint (front, rear, left, right, unspecified),
 - “truncated” flag (bounding box \neq object extent)
 - “difficult” flag
- Complete annotation of all specified classes (as per guidelines)

Examples



- PAspersonTrunc
- PASHorseTrunc
- PASHorseTrunc

Examples



- PASpersonRear
- PASpersonFrontalTrunc
- PASbusRearTrunc
- PASbusRearTruncDifficult

Annotation Procedure

- All annotations performed by a group of annotators in one location over a few days
- Detailed guidelines developed on
 - What to label
 - Viewpoint
 - Bounding box (don't extend greatly for a few pixels)
 - Truncation
- Data checked by a single annotator afterwards

Dataset Statistics

	train+val		test	
	img	obj	img	obj
bicycle	270	323	268	326
bus	174	235	180	233
car	553	854	544	854
cat	386	429	388	429
cow	206	313	197	315
dog	365	422	370	423
horse	247	326	254	324
motorbike	235	275	234	274
person	666	1156	675	1153
sheep	251	421	238	422
Total	2618	4754	2686	4753

Participation

- 22 teams submitted results, from 14 different institutions
- 28 different methods
 - 19 for classification only
 - 4 for detection only
 - 5 for classification and detection

Key Issues

- Evaluation of classification and detection
- Detection requires class vs non-class decisions
- Evaluation of classification without a “forced choice” paradigm
- More training examples per class than Caltech 101
- LabelMe data is incompletely labelled and has inconsistencies
⇒ unsuitable for *testing*

Planning for VOC2007

- Mark Everingham, Luc van Gool, Chris Williams, John Winn, Andrew Zisserman
- Tasks: classification, detection (bounding boxes), taster on pixelwise labelling, parts spatial layout (?)
- More classes
- Tentative timeframe:
 - Week beginning Jan 15 2007: annotation party
 - early Feb: release of training data
 - early April: release of test data
 - mid May: deadline for submission of results
 - subsequently: Challenge workshop

Improvements of VOC2006 over VOC2005

Problems from VOC2005

- Errors in annotation
- Limited amount of training data
- Difficulty of test2 sets compared to training data
- Release of annotations on test data led to poor practice by some participants

These issues were addressed in the VOC2006 challenge to the satisfaction of the participants