



Weakly Supervised Object Boundaries

Anna Khoreva¹, Rodrigo Benenson¹, Mohamed Omran¹,
Matthias Hein², Bernt Schiele¹

¹ Max Planck Institute for Informatics, Saarbrücken, Germany

² Saarland University, Saarbrücken, Germany

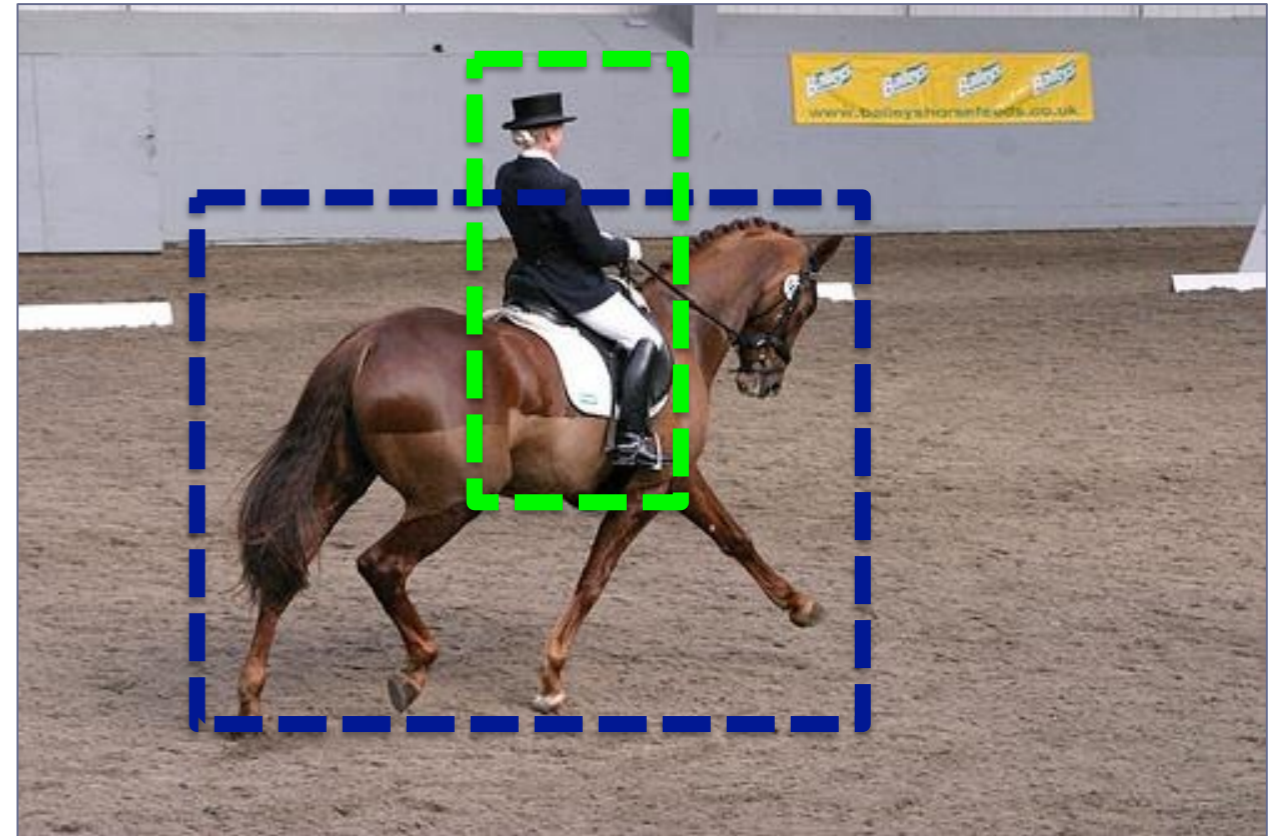


Is it possible to get object boundaries with bounding box supervision?



Full supervision

Time-consuming.




Weak supervision

Only 2 clicks per object.

Task: detection of object boundaries




Image



Ground truth



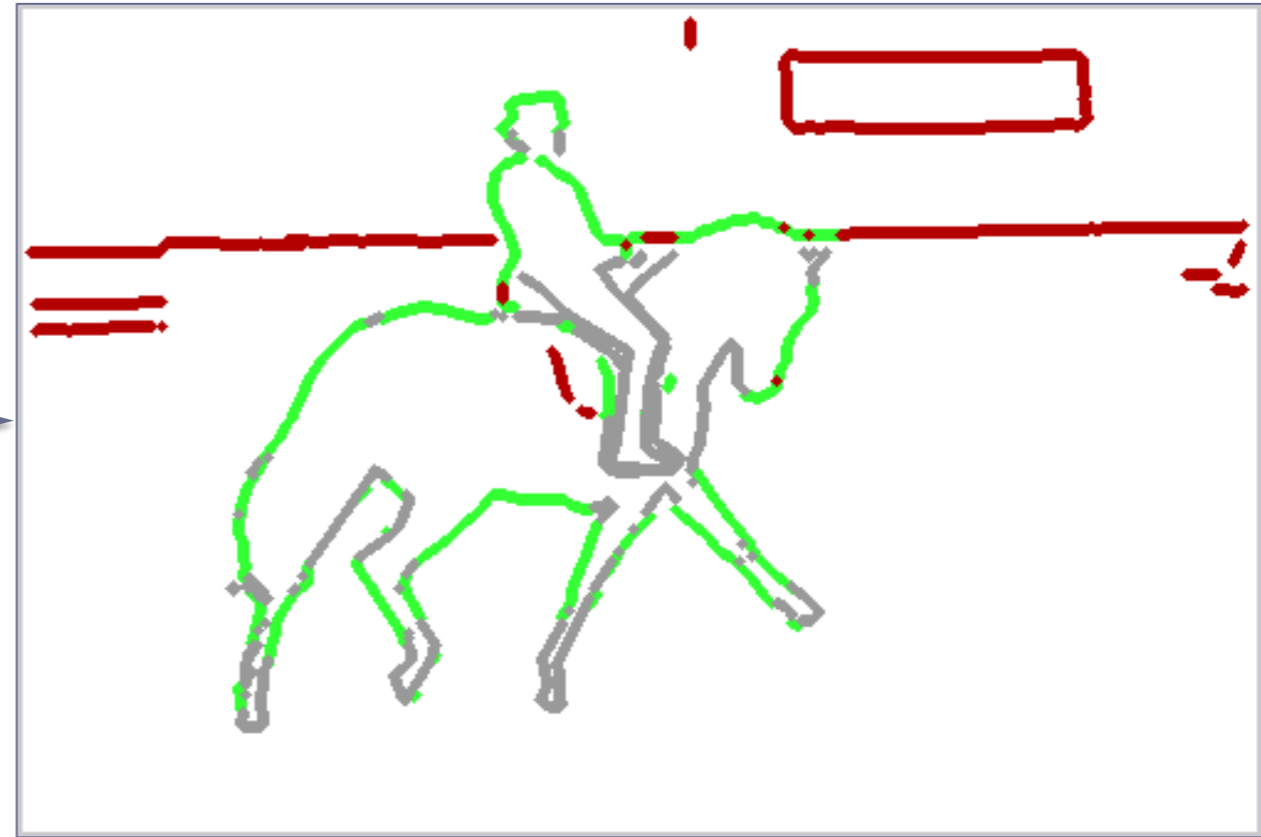
Image






Generic boundary detector output



Image

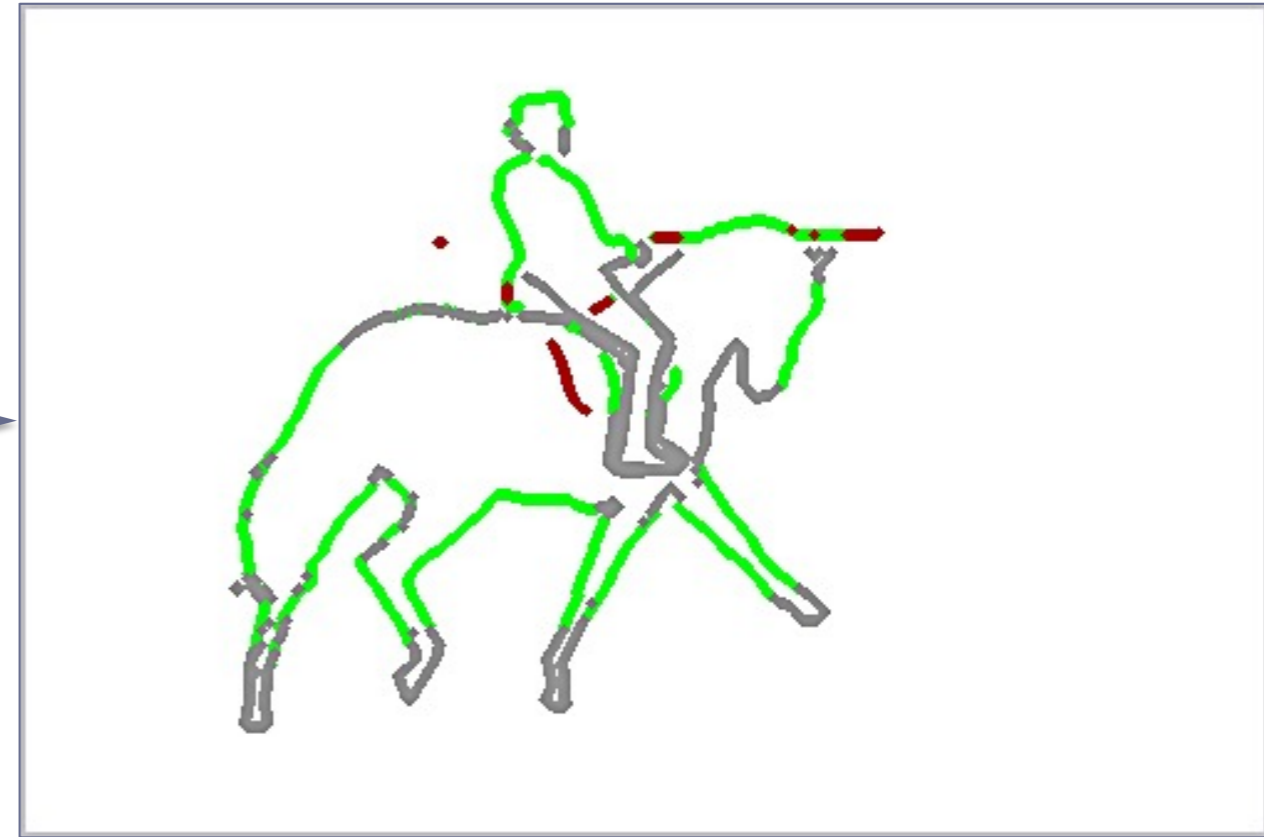


Generic boundary detector output




-  True positive boundaries
-  False positive boundaries
-  False negative boundaries




Full supervision




Object boundary detector output




-  True positive boundaries
-  False positive boundaries
-  False negative boundaries








Weak supervision



Object boundary detector output

-  True positive boundaries
-  False positive boundaries
-  False negative boundaries

It is possible to get object boundaries with bounding box supervision.



Image

Fully supervised

Weakly supervised

 True positives

 False positives

 False negatives

Pipeline:

Input






Image and bounding boxes



Generation of annotations

Regular boundary detector training

Output




Object boundary detections

Task: detection of generic boundaries

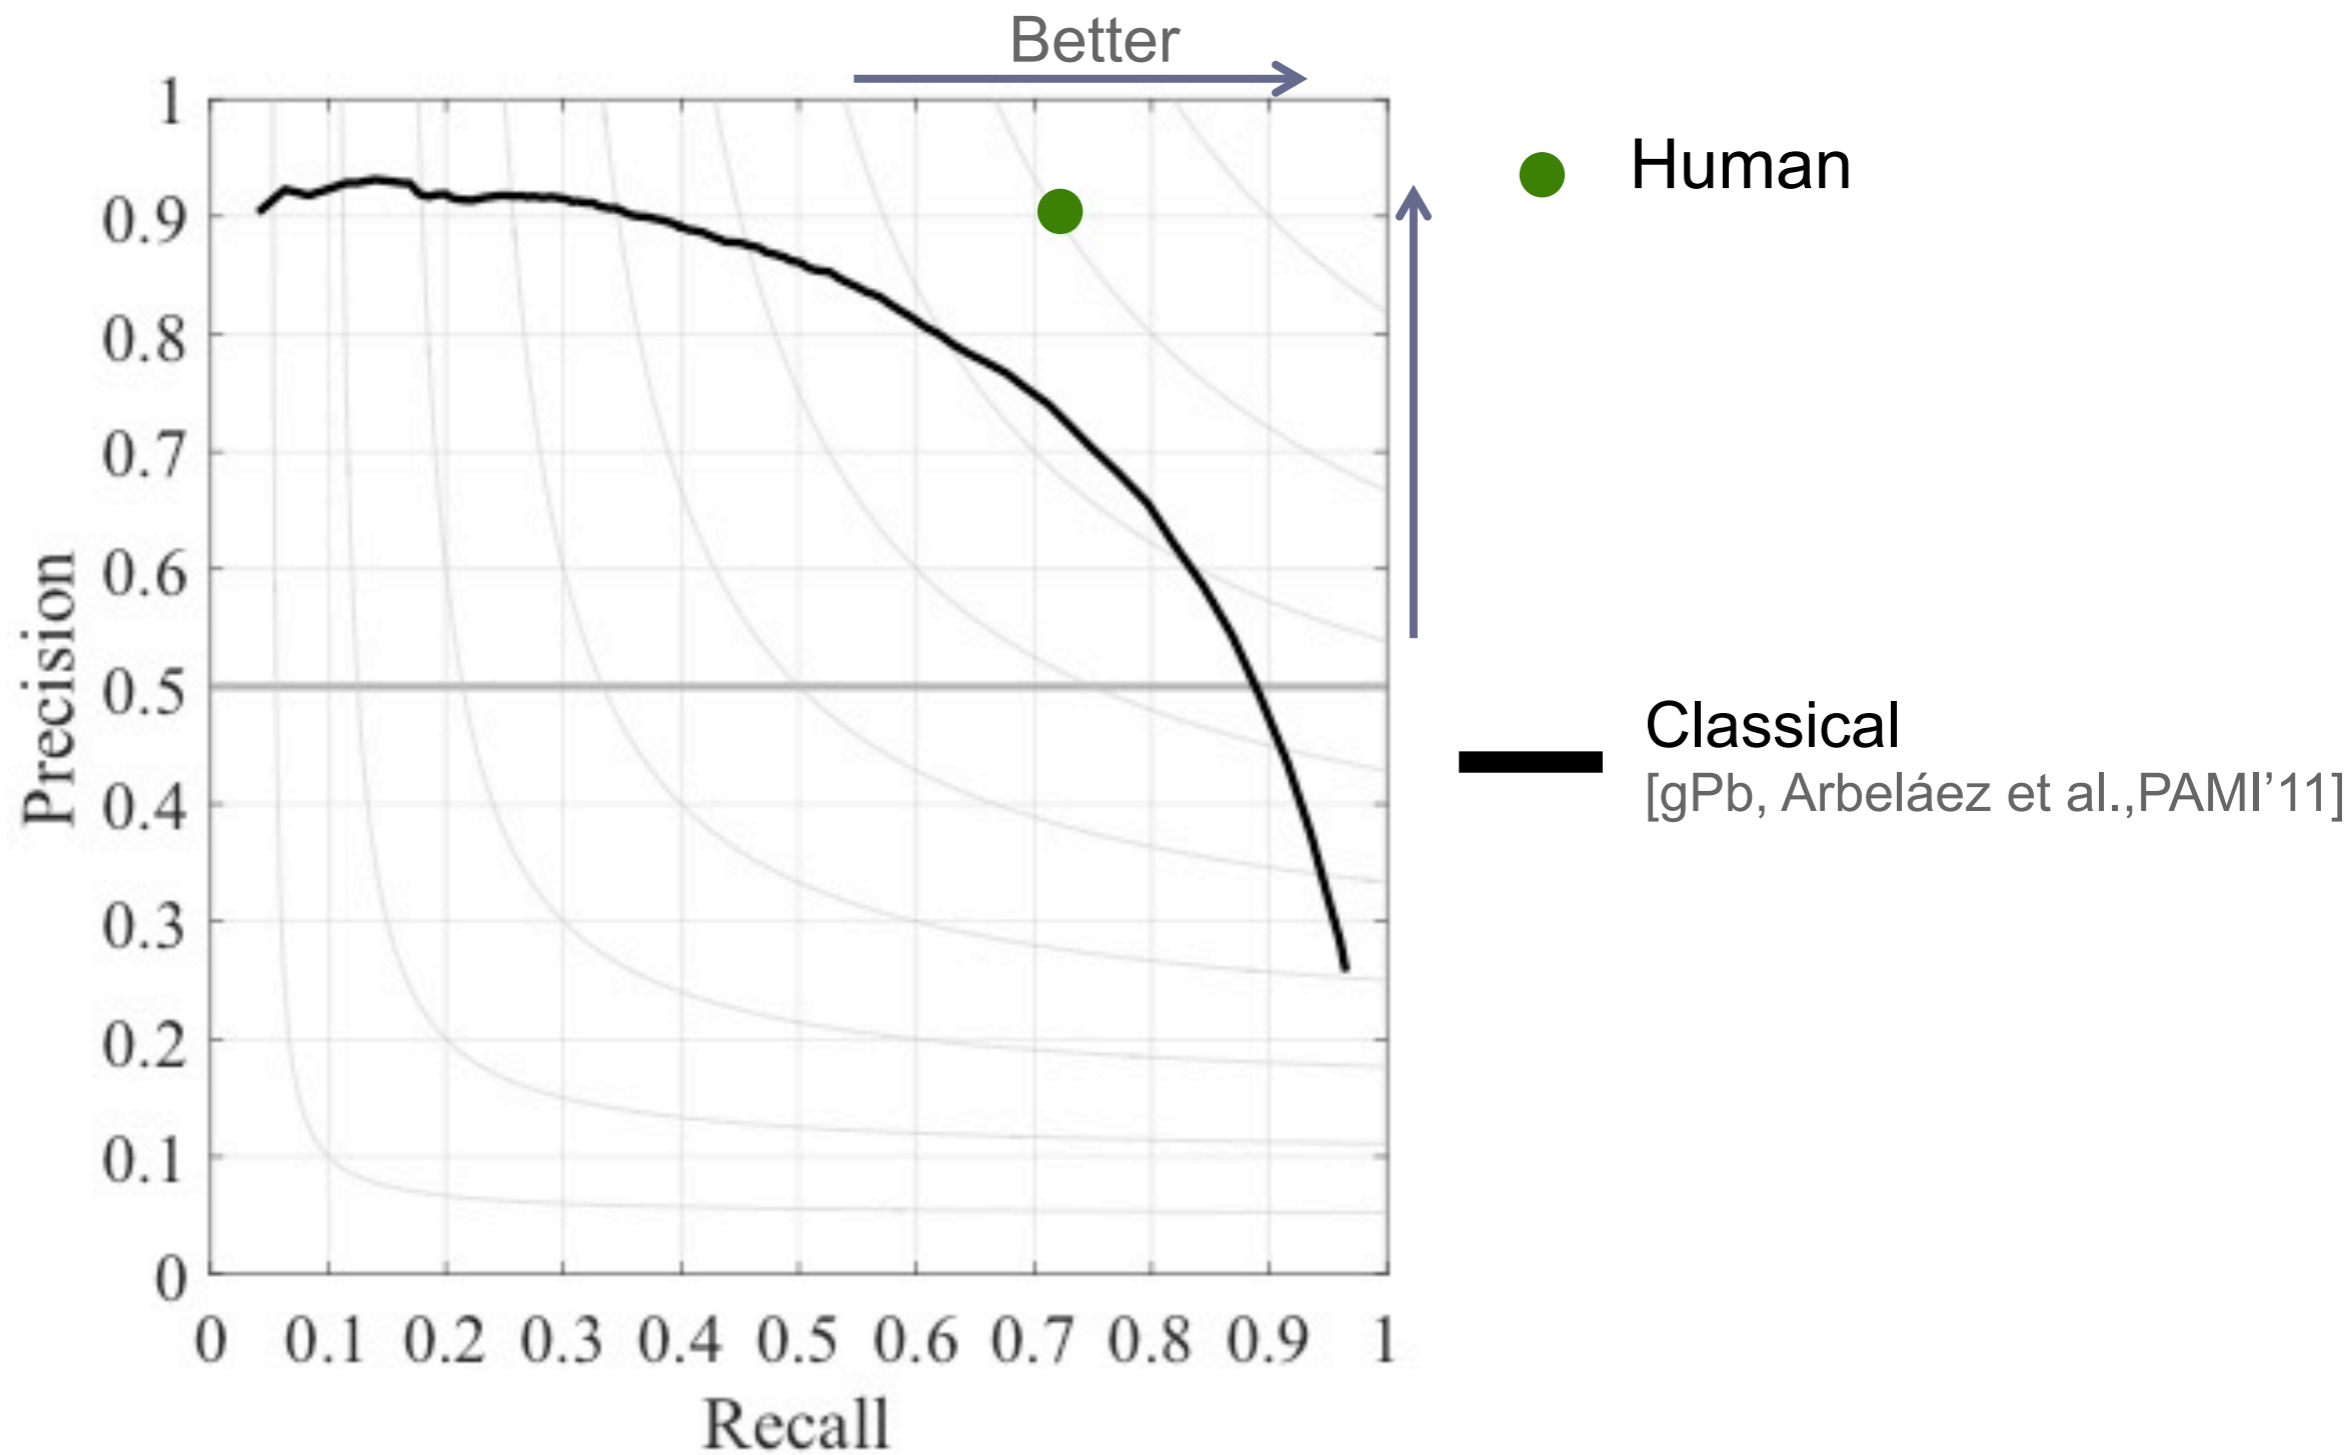


Image




Generic boundaries


Dataset: BSDS [Martin et al., ICCV'01; Arbeláez et al., PAMI'11]



Task: generic boundaries, BSDS dataset.




Task: generic boundaries, BSDS dataset.



Task: generic boundaries, BSDS dataset.

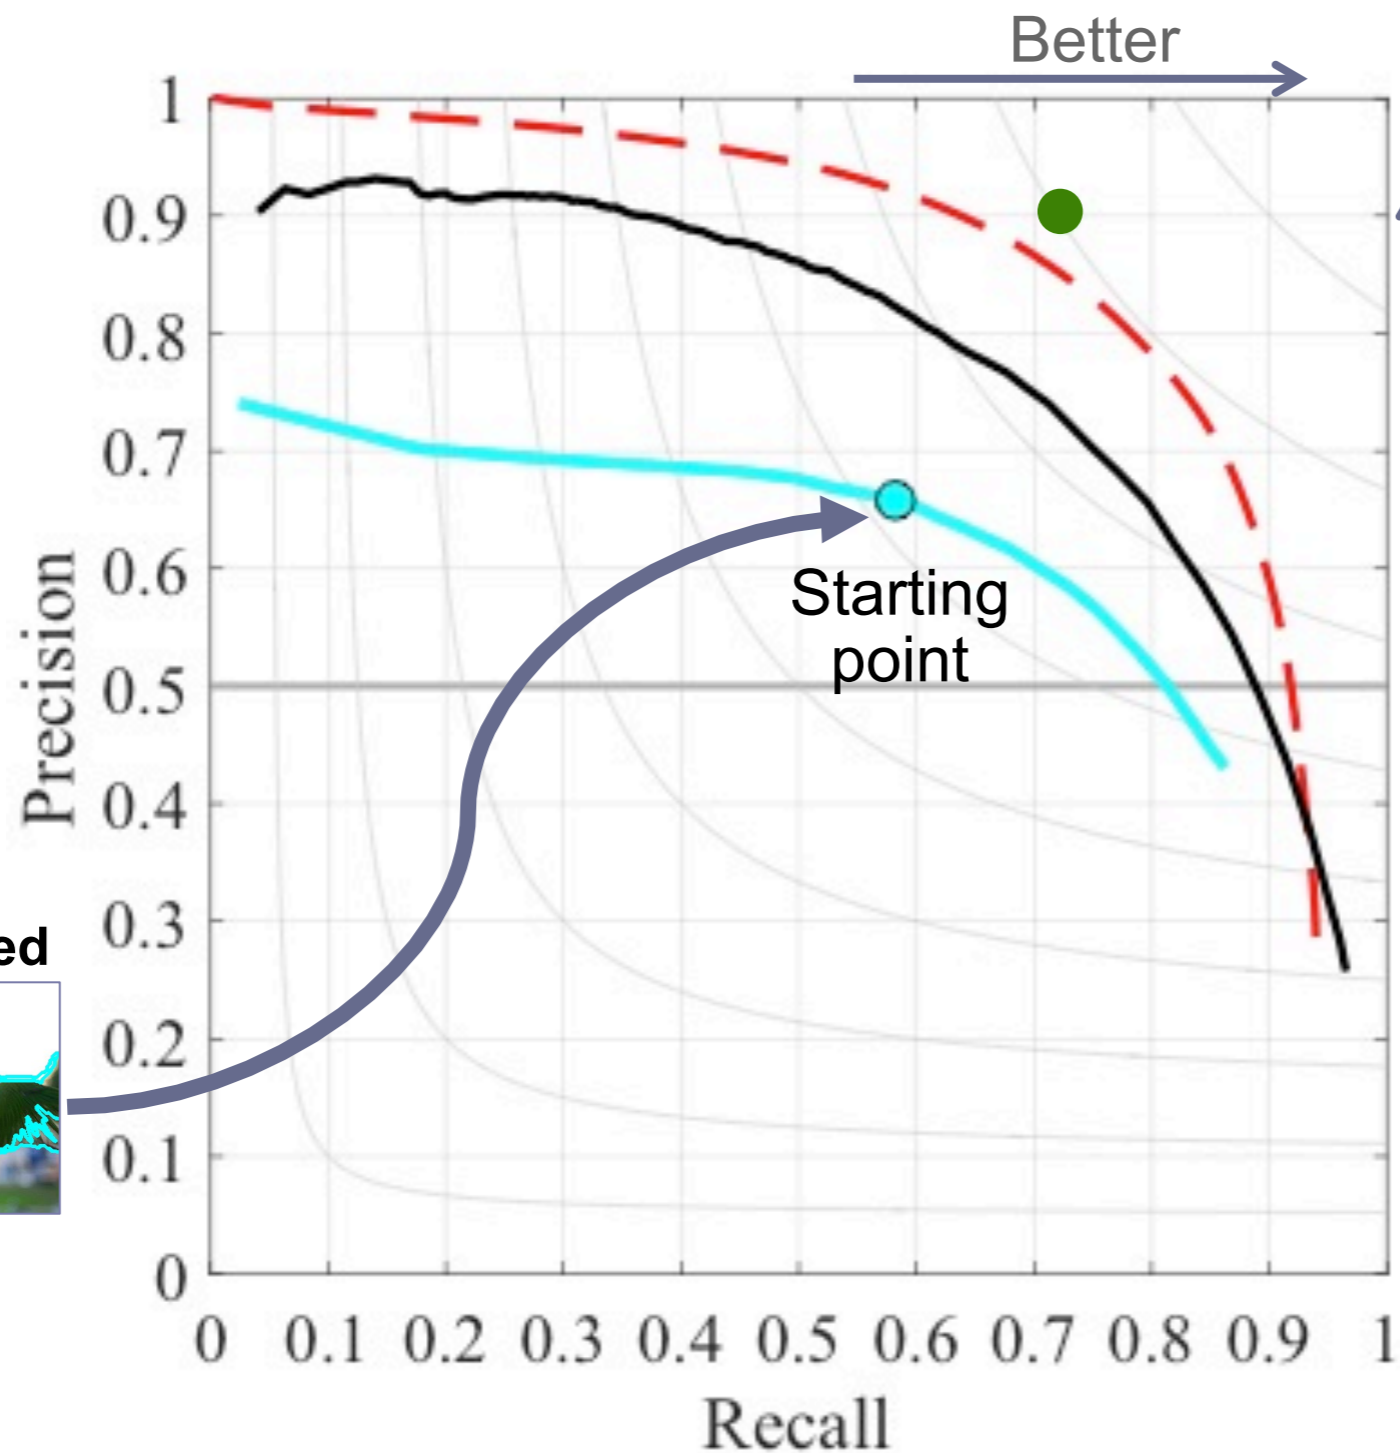
Full supervision



Unsupervised

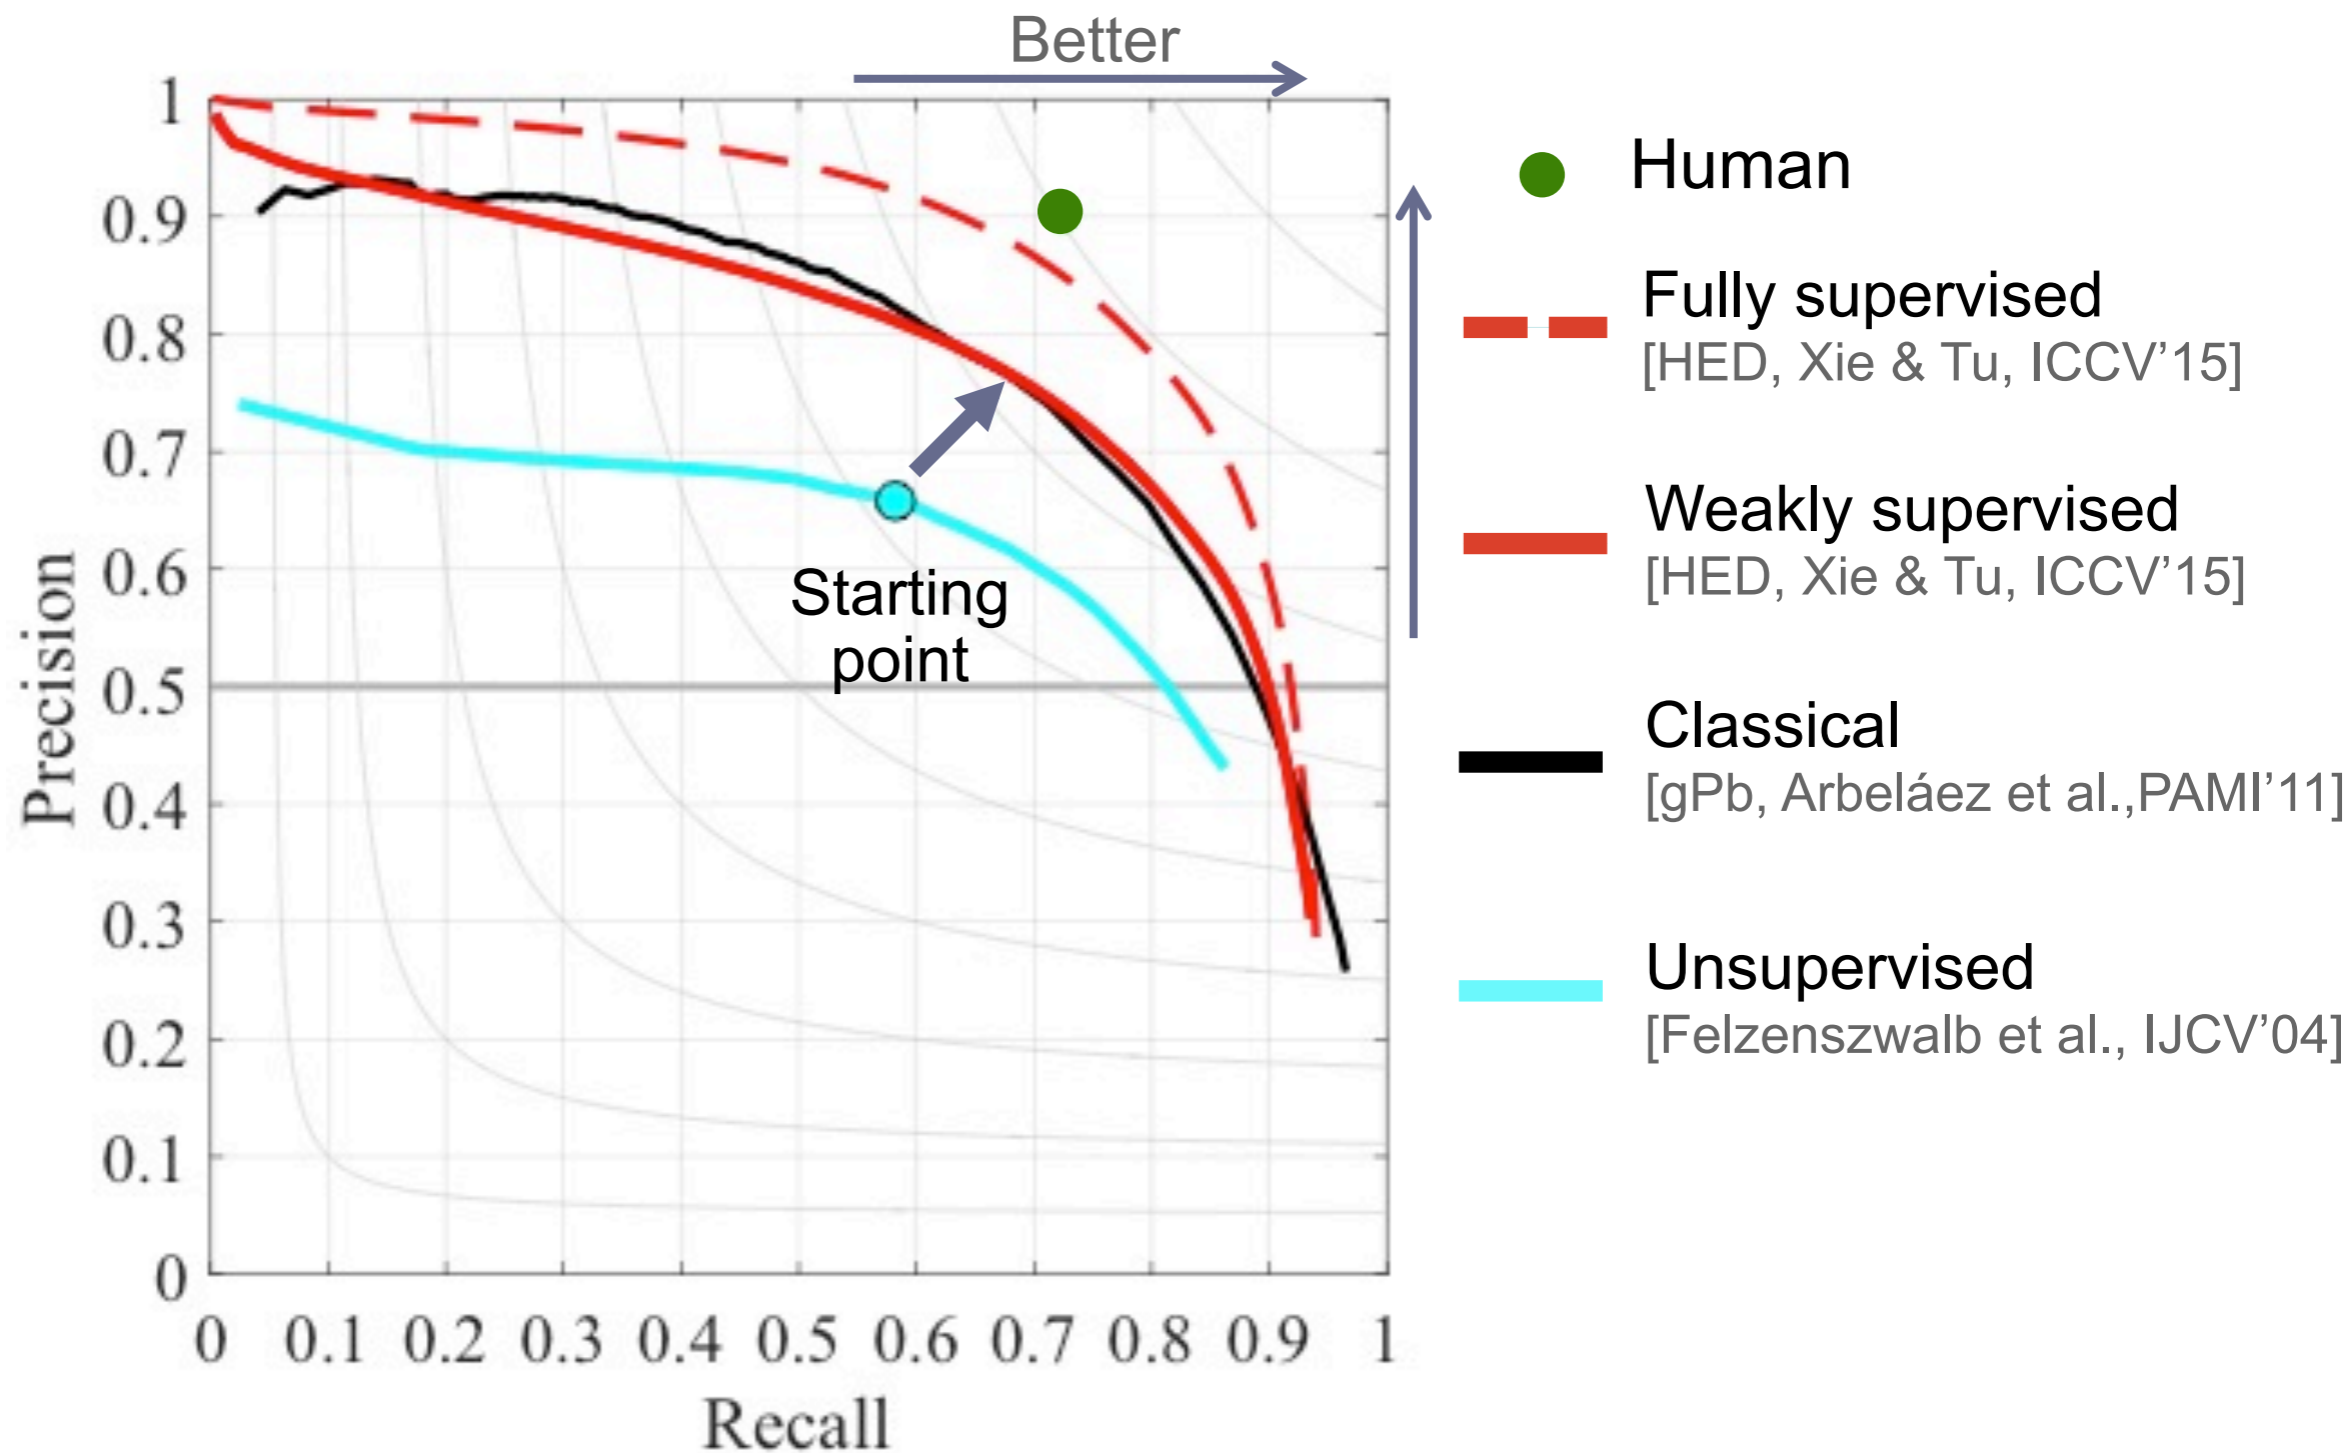


[Felzenszwalb et al., IJCV'04]




- Human
- Fully supervised
[HED, Xie & Tu, ICCV'15]
- Classical
[gPb, Arbeláez et al., PAMI'11]
- Unsupervised
[Felzenszwalb et al., IJCV'04]

Task: generic boundaries, BSDS dataset.



Task: generic boundaries, BSDS dataset.

Boundary detectors are robust to annotation noise.




Pipeline:

Input



Image and bounding boxes



Generation of annotations

Regular boundary detector training

Output




Object boundary detections

Task: detection of object boundaries

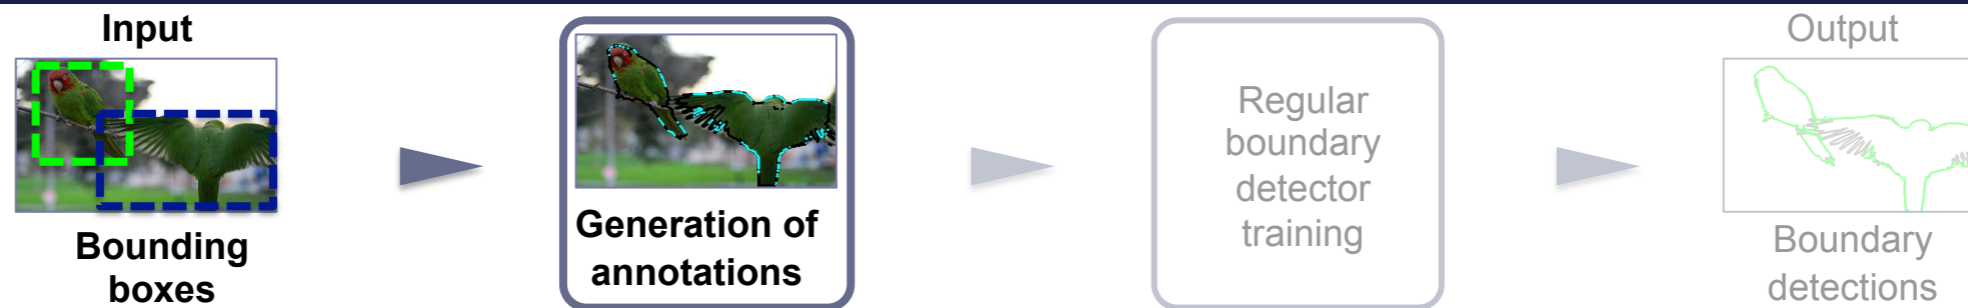


Image




Object boundaries

Dataset: VOC [Everingham et al., IJCV'15]



Generation of annotations:

Detection bounding boxes





[Fast-RCNN, Girshick, ICCV'15]

Object proposals



[SeSe, Uijlings et al., IJCV'13]



	Positive boundaries
	Ignore region
	Negative boundaries



Combining weak annotations from:

Object proposals



GrabCut




Graph-based segmentation



[SeSe, Uijlings et al., IJCV'13]

[Rother et al., SIGGRAPH'04]

[Felzenszwalb et al., IJCV'04]



Consensus between different methods

	Positive boundaries
	Ignore boundaries
	Negative boundaries


Ground truth



Generated annotations




 Positive boundaries

 Ignore boundaries

 Negative boundaries

Ground truth



Generated annotations



 Positive boundaries

 Ignore boundaries


 Negative boundaries

Pipeline:

Input




Image and bounding boxes




Generation of annotations

Regular boundary detector training

Output

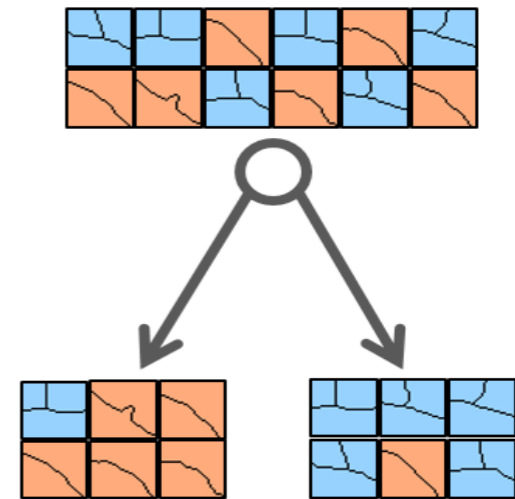



Object boundary detections

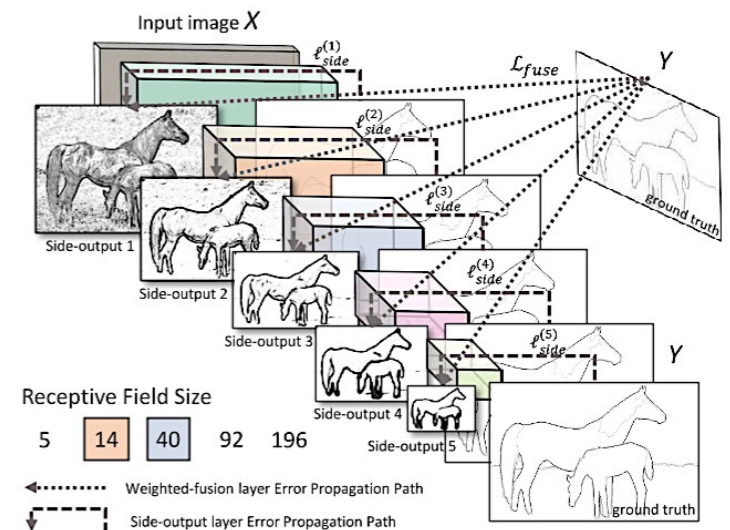



Boundary detectors:

- Structured Edge Forests [SE, Dollar et al., PAMI'15]




- Holistically-nested Edge Detection [HED, Xie & Tu, ICCV'15]



Task: detection of object boundaries



Image



Object boundaries


Datasets: VOC [Everingham et al., IJCV'15]

SBD [Hariharan et al., ICCV'11]




VOC [Everingham et al., IJCV'15]

SE results





Generic boundaries





Task: object boundaries, VOC dataset.

SE results





Task: object boundaries, VOC dataset.

SE results



Task: object boundaries, VOC dataset.


HED results





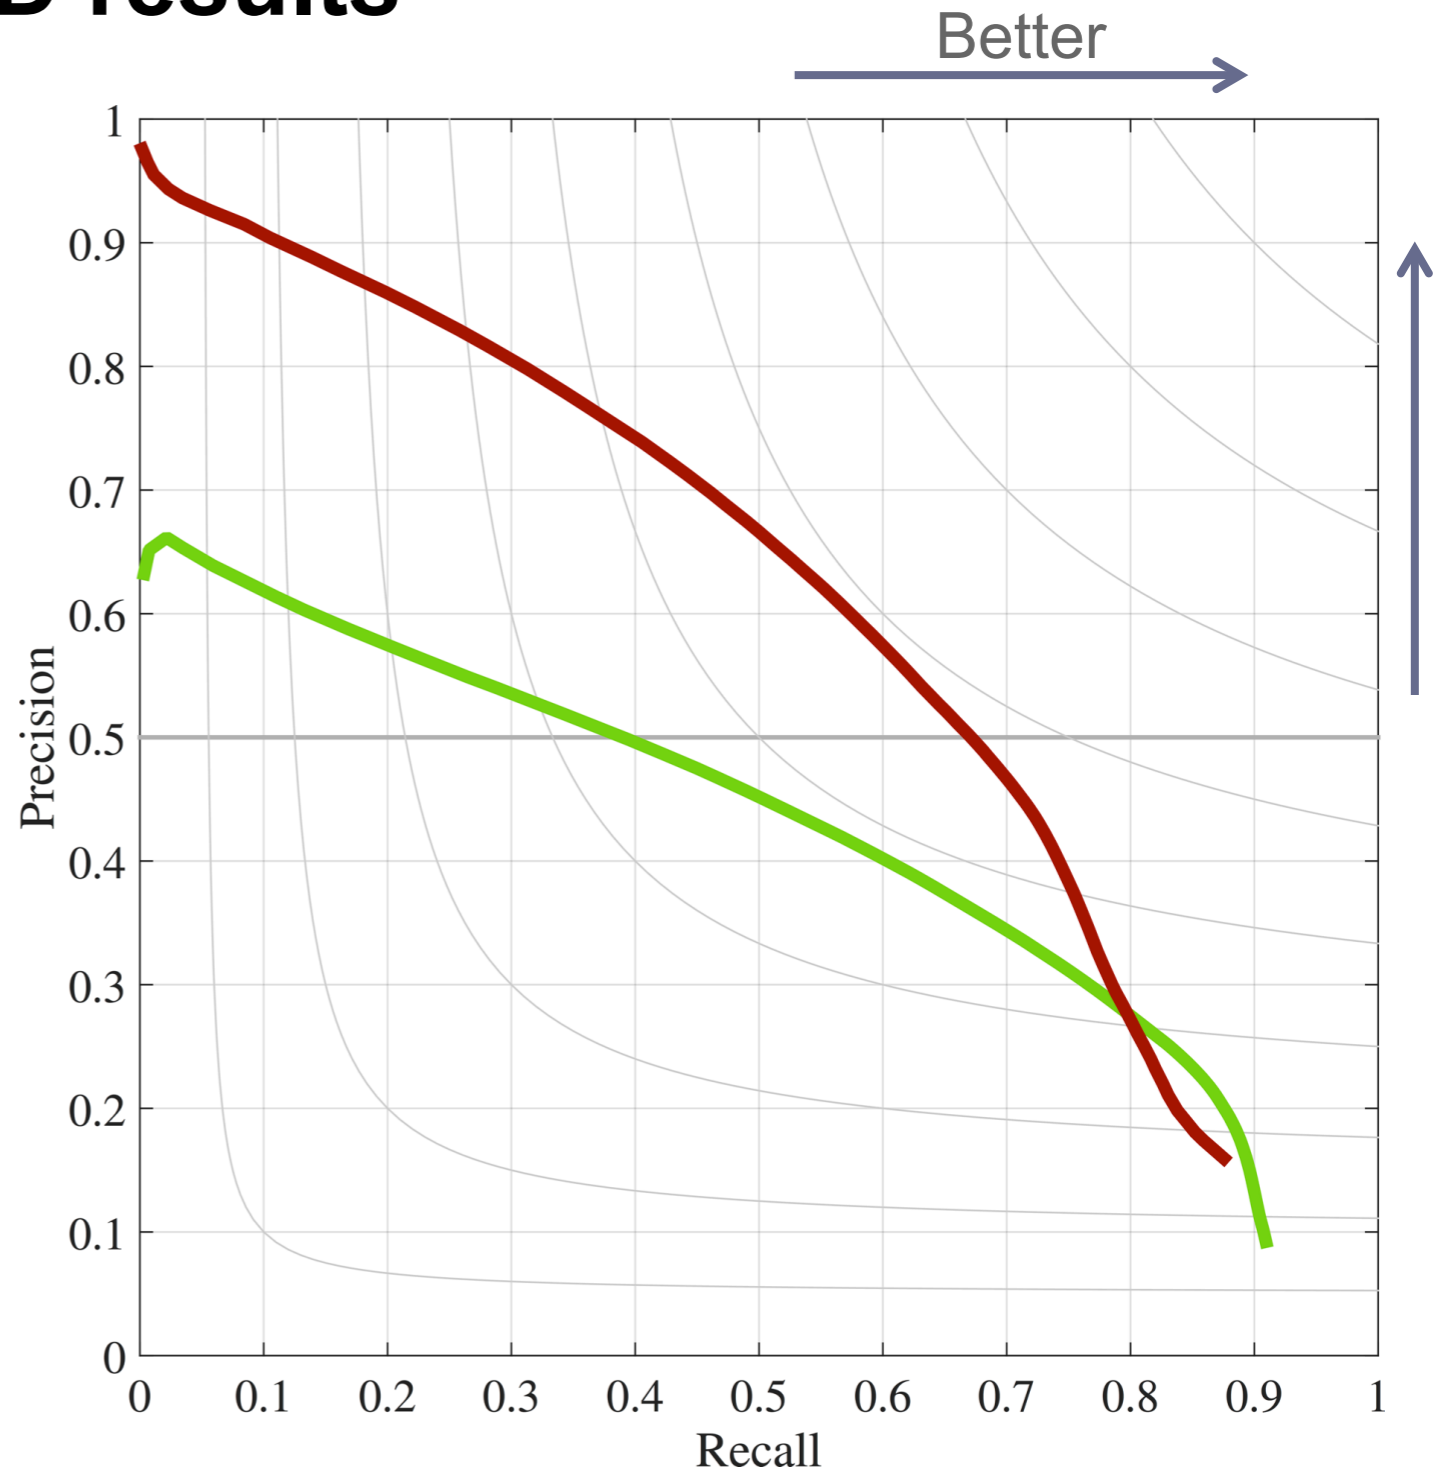
Generic boundaries

Task: object boundaries, VOC dataset.

HED results




-  Fully supervised
-  Generic boundaries




Task: object boundaries, VOC dataset.







HED results



- █ Fully supervised
- █ Weakly supervised
- █ Generic boundaries



Task: object boundaries, VOC dataset.



Image

Generic boundaries









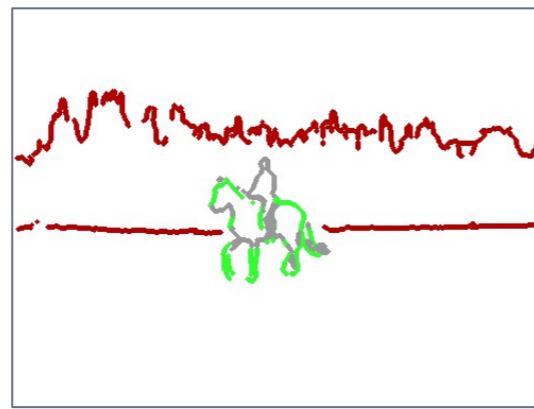
Fully supervised

Weakly supervised

True positives

False positives

False negatives



Image

Generic boundaries

Fully supervised


Weakly supervised

True positives

False positives


False negatives

While training an object detector one can also get a high quality object boundary detector for free.









SBD [Hariharan et al., ICCV'11]



Task: semantic object boundaries, SBD dataset.



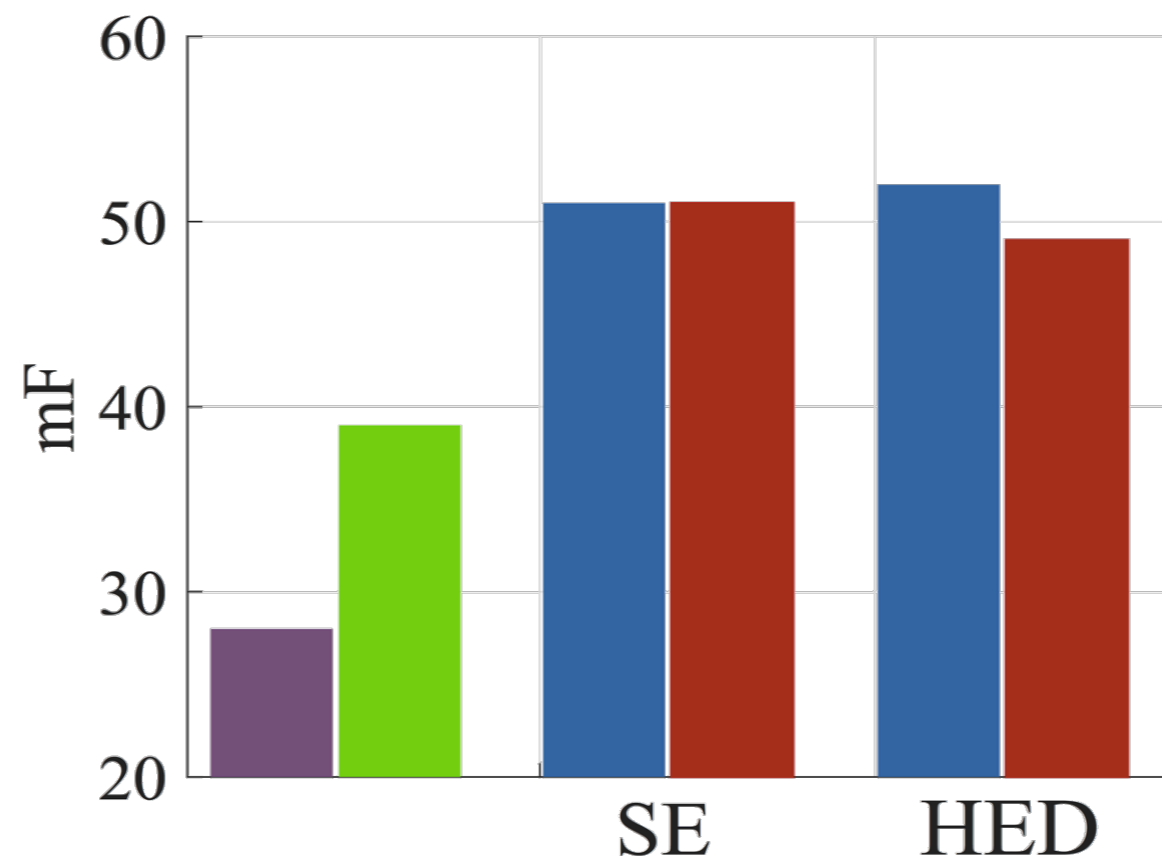
Image

Ground truth

Fully supervised


Weakly supervised

Weakly supervised object boundaries can reach the full supervision quality.

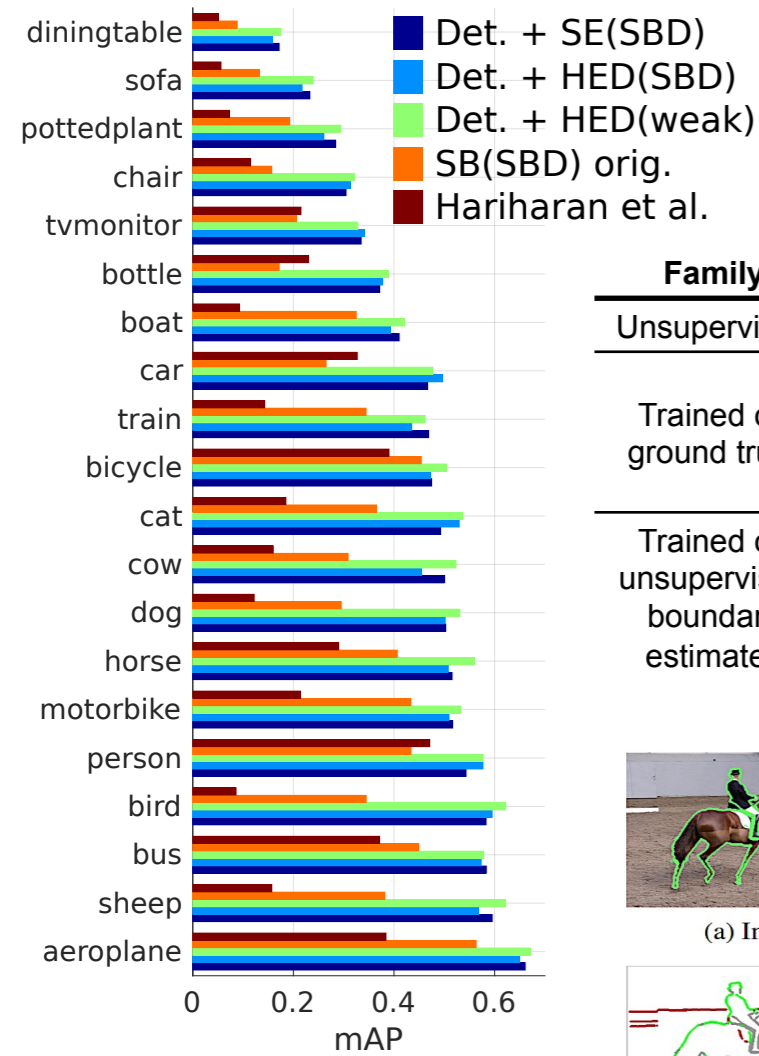


Contributions:



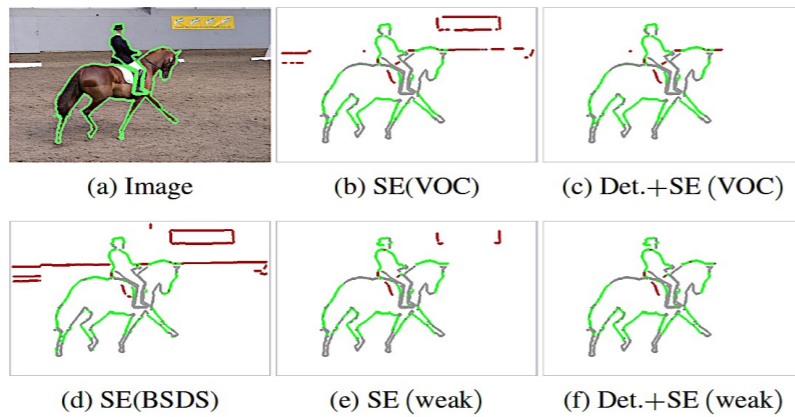
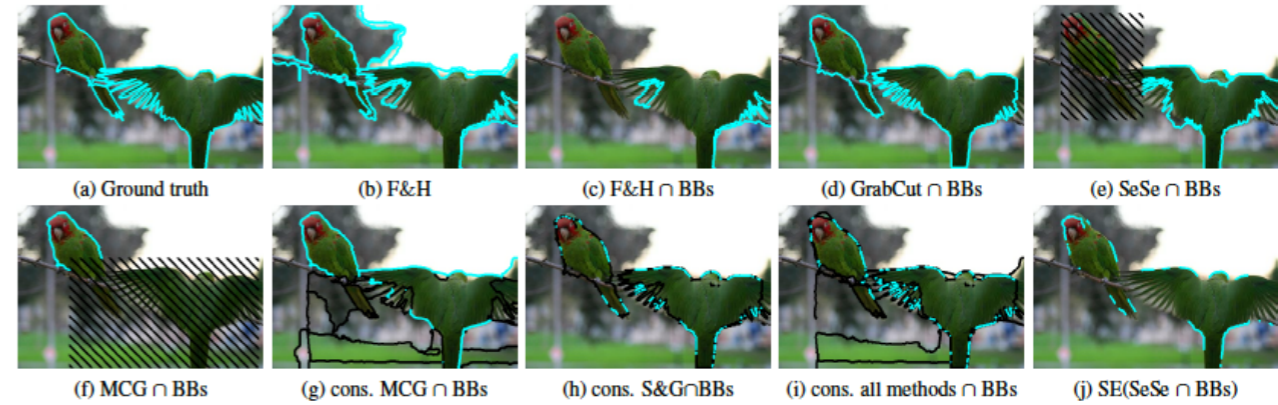
- Introduce the problem of weakly supervised object boundaries.
- Boundary detectors are robust to annotation noise.
- High quality object boundaries can be obtained using box annotations alone.




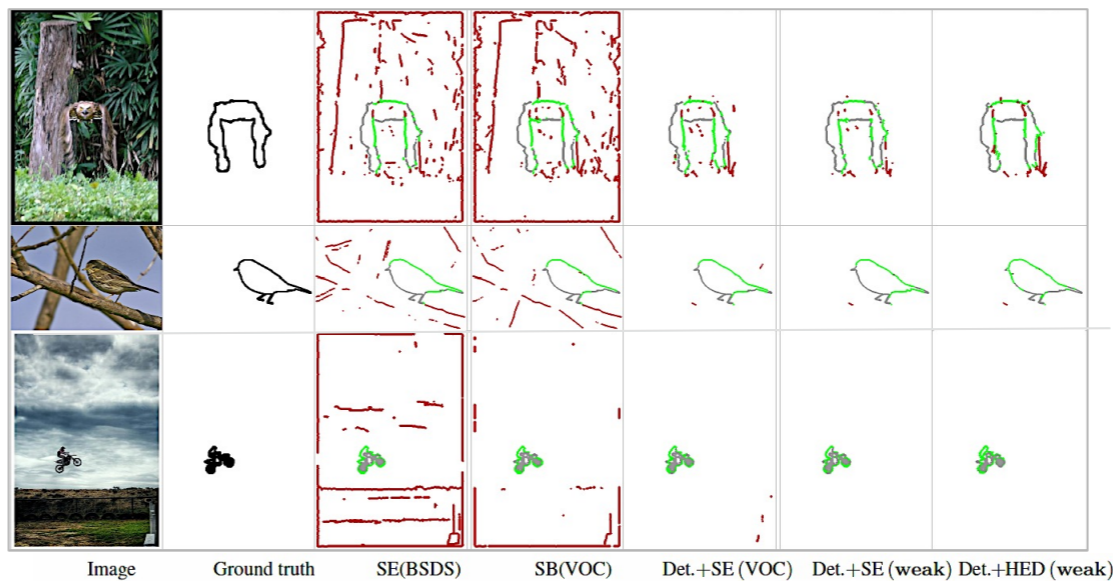
Even more results in the paper!



Family	Method	ODS	OIS	AP
Unsupervised	F&H	64	67	64
	gPb-owt-ucm	73	76	73
Trained on ground truth	SE(BSDS)	74	76	79
	HED(BSDS)	79	81	84
Trained on unsupervised boundary estimates	SE(F&H)	71	74	76
	SE(SE(F&H))	72	74	76
	HED(SE(F&H))	73	76	75



Family	Method	mAP	mAP
Other	GT Hariharan et al.	28	21
	SB(SBD) orig.	39	32
GT	SB(SBD)	43	37
	Det. + SE(SBD)	51	45
SE	Weakly supervised		
	SB(SeSe ∩ BBs)	40	34
	SB(MCG ∩ BBs)	42	35
	Det. + SE(SeSe ∩ BBs)	48	42
	Det. + SE(MCG ∩ BBs)	51	45
GT	HED(SBD)	44	41
	Det. + HED(SBD)	49	45
HED	Weakly supervised		
	HED(cons. MCG ∩ BBs)	41	37
	HED(cons. S&G ∩ BBs)	44	39
	Det. + HED(cons. MCG ∩ BBs)	48	44
	Det. + HED(cons. S&G ∩ BBs)	52	47



Details and more results at poster 20 tomorrow

Project page at <https://www.mpi-inf.mpg.de/wsob>

Weakly Supervised Object Boundaries

Anna Khoreva¹ Rodrigo Benenson¹ Mohamed Omran¹ Matthias Hein² Bernt Schiele¹
¹ Max Planck Institute for Informatics, Saarbrücken, Germany ² Saarland University, Saarbrücken, Germany

Motivation
Goal: High quality object boundaries from bounding box annotations.
Full supervision: Significant annotation effort is required.
Weak supervision: Bounding box annotation requires 2 clicks per object.

Framework
Boundary detection tasks:
Generic boundaries (BDS) [Martin et al., ICCV'01], Object-specific boundaries (VOC) [Everingham et al., UCV'13], Class-specific object boundaries (COCO) [Lin et al., ECCV'14], SBD [Harizan et al., ICCV'11].
Levels of supervision: Fully unsupervised, Detection annotations, Detection and generic boundary annotations.
Baselines:
• Structured Edge Forests [SE, Dollar et al., PAMI'15]
• Holistically-nested Edge Detection [HED, Xie & Tu, ICCV'15]

Robustness to Annotation Noise
BDS: generic boundaries. Family Method ODS OS AP. Trained on: Unsupervised F&H 64 67 64. Trained on ground truth: SE(BSD) 73 76 73, HED(BSD) 74 76 79, SE 79 81 84, HED 71 74 76, HEDSE(F&H) 72 74 76, HEDSE(F&H) 73 76 75. SE and HED are robust to annotation noise during training.

Weakly Supervised Boundary Annotations
Approach: Generate weakly supervised annotations to train boundary detector.
Combination of sources: Detection bounding boxes (B&B, Fast R-CNN, Girshick, ICCV'15), Graph-based segmentation (F&H, Fitzgibbon et al., ICCV'04), Box driven segmentations (S&G, Rother et al., ICCV'04), Object proposals (S&O, Ullings et al., UCV'13; JACO, Post-Tussell et al., arXiv'15).
Generated annotations: F&H A Bbs, SeSe A Bbs, SE(SeSe A Bbs), cons. S&G A Bbs, cons. all methods A Bbs.

Experimental Results
VOC: object-specific boundaries. SE models, HED models. Detection boxes at test time to improve boundaries. Weakly supervised object boundaries can reach the full supervision quality.
SBD: class-specific object boundaries. Family Method mF mAP. Other GT Harizan et al. 28 21. SE(BSD) org 39 32. SE(BSD) Det. + SE(BSD) 43 37. SE(BSD) super 40 34. Weakly supervised Det. + SE(SeSe A Bbs) 42 35. HED Det. + HED(SeSe A Bbs) 48 42. HED Det. + SE(MCO A Bbs) 51 45. HED Det. + HED(SD) 44 41. HED Det. + HED(cons. MCO A Bbs) 49 45. HED Det. + HED(cons. S&G A Bbs) 41 37. HED Det. + HED(cons. MCO A Bbs) 44 39. HED Det. + HED(cons. S&G A Bbs) 48 44. HED Det. + HED(cons. S&G A Bbs) 52 47.

While training an object detector one can also get high quality object boundary detector for free!

Project page at
<https://www.mpi-inf.mpg.de/wsob>

Trained models, generated annotations, and results
are available online.

Thank you for your attention!