Boundary Edge Detection (1)

- Boundary edge detection. New dataset: Pascal Boundaries.
- V. Premachandran et al. 2016.
- Current Edge Dataset: BSDS
- Problems:
- (i) performance saturated.
- (ii) small by modern standards
- (iii) different types of edges.
- Pascal Boundaries 10,000 images.
- Deep Net – M-DSBD -variant of HED.

Motivation For New Dataset

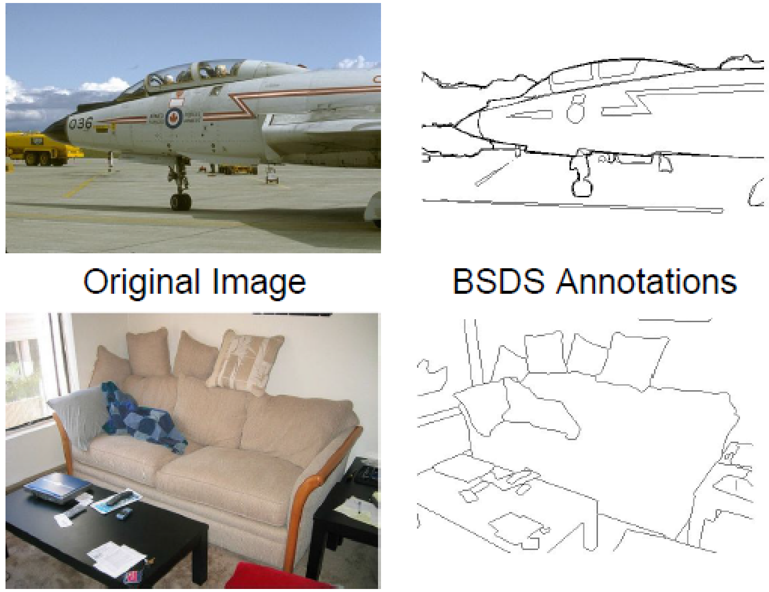

Original Image PB Annotations

- BSDS edge annotations are ambiguous (internal edges)
- BSDS500 is small by today's standards and saturated

Boundary Edge Detection (2)

- HED and SED transfer fairly well to Pascal Boundaries.
- But our M-DSBD outperforms both on Pascal Boundaries.

PASCAL Boundaries Dataset

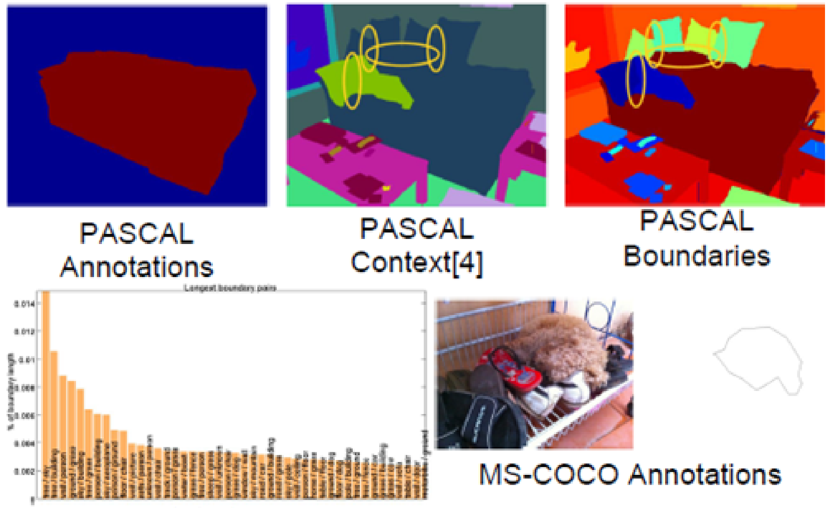

Fig. We annotate twice as many edges as PASCAL 20 masks and MS-COCO foreground masks. A significant portion of the edges are because of the background annotations.

- 20x bigger than BSDS500 with 10k annotated images.
- 2x as many annotations as PASCAL masks
- Groundtruth Boundaries between 459 semantic classes.
- Boundaries between both foreground & background classes.
- Instance-level masks for all foreground objects and many background classes.

Experiments and Results

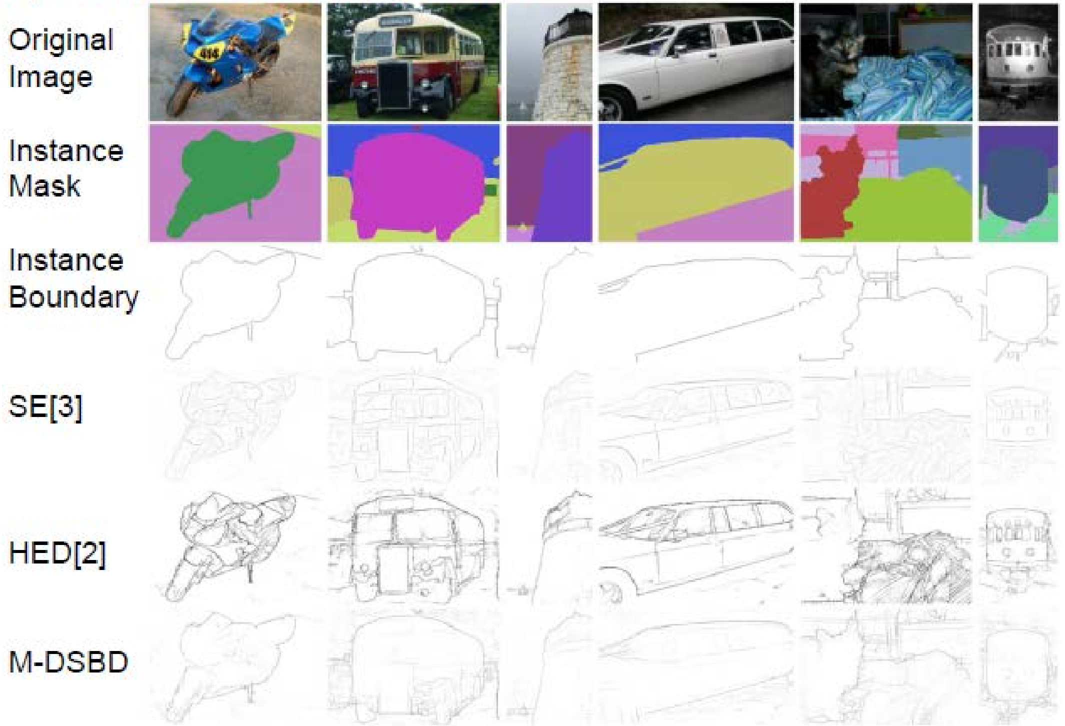

Fig. Notice that HED detects many internal edges while M-DSBD is better at restricting itself to instance-level masks.

Occlusion Detection (1)

- Occlusion boundary detection and border-ownership.
- Motivation – to help detect partially occluded objects.
- New dataset: Pascal object occlusion.10,000 images.
- P. Wang & A. Yuille. 2016.

Fig. 1. Left: Occlusion boundaries represented by orientation θ (the red arrows) which indicates occlusion relationship using the "left" rule where the left side of the arrows is foreground. Right: More examples from our Pascal instance occlusion dataset (PIOD).

Occlusion Detection (2)

- For inference we apply a two stream network (illustrated with the HED network) for predicting the pixel-wise boundaries and the occlusion orientations respectively. Then, we apply non-maximum suppression (NMS) to the boundaries and merge the two predictions, yielding the recovered occlusion boundary.

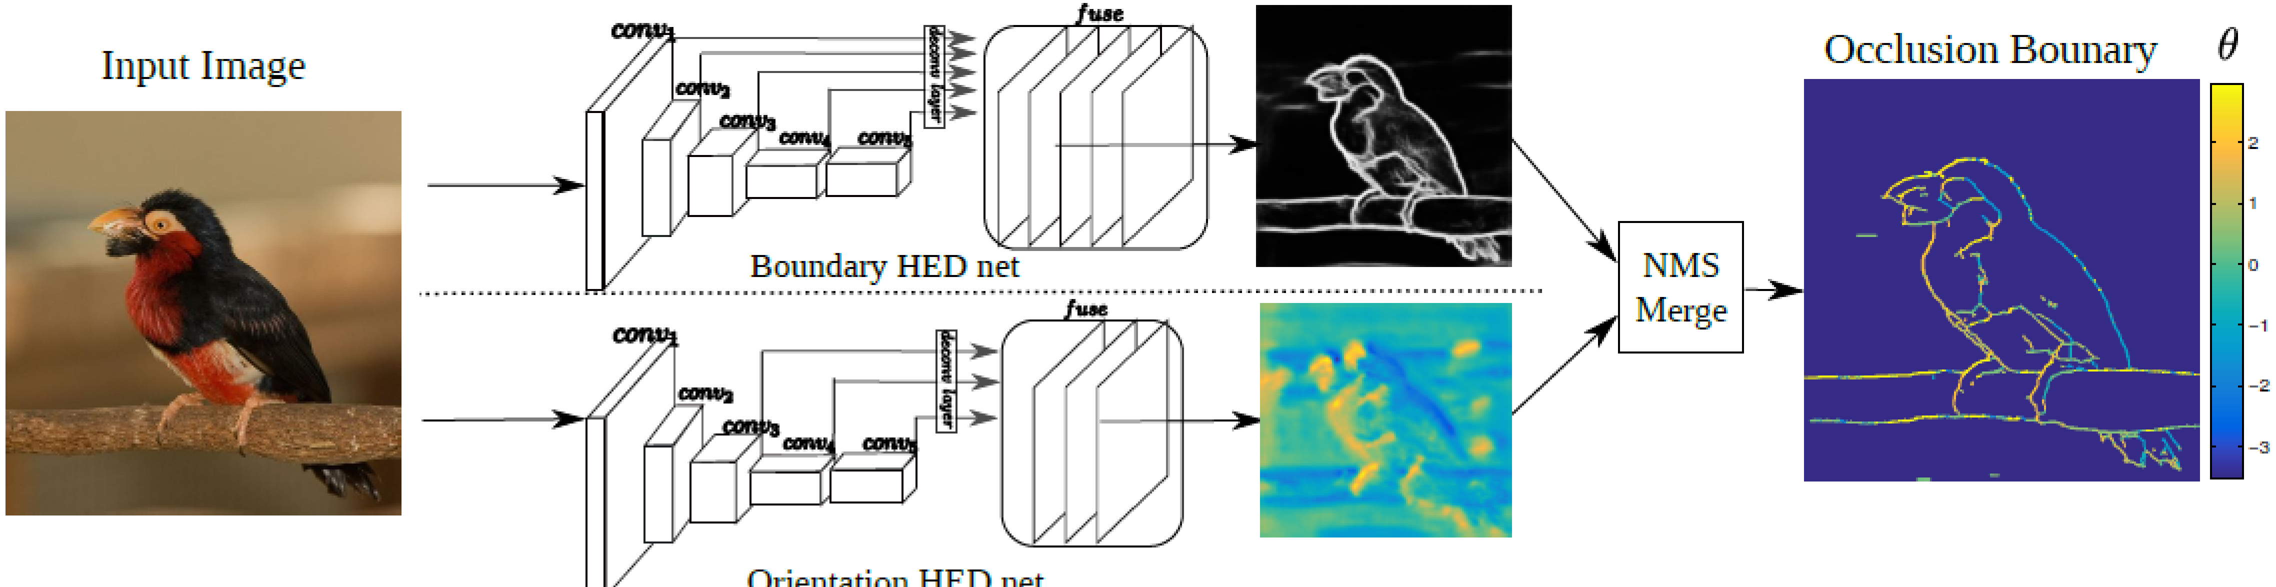

Semantic Segmentation (1)

- DeepLab – (code available) – deep net (“holes” trick from wavelets) Plus fully connected CRF (for long-range context).
- L-C Chen et al. 2015. G. Papandreou et al. 2015.
- Pascal (20 classes), Coco (60) classes.
- Strong supervision. Weak supervision.
- Also new dataset:
- Pascal Context (60 classes, 400 classes).
- Performance improved by local adaption
- to scale: L-C Chen et al. 2016.

Semantic Segmentation (2)

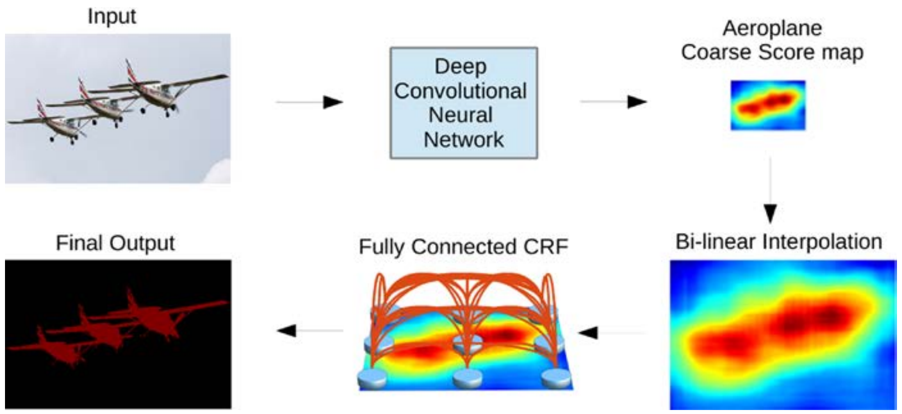

- Task: assign labels per-pixel (fine-grained).
- Method: fully convolutional deep net (DeepLab).
- In addition: conditional random fields (CRF) – sideways context.
- Fully supervised (left) , weakly supervised (right)