Symmetry Axis Detection

Alan Yuile

Symmetry Axes: Examples.

- Symmetry Axes play an important role in early computer vision.
- Many studies were performed on binary images. There were many alternative definitions of symmetry axes, which are roughly equivalent, and algorithms for finding them.
- They relate to representing objects in terms of 3D geons (Biederman)

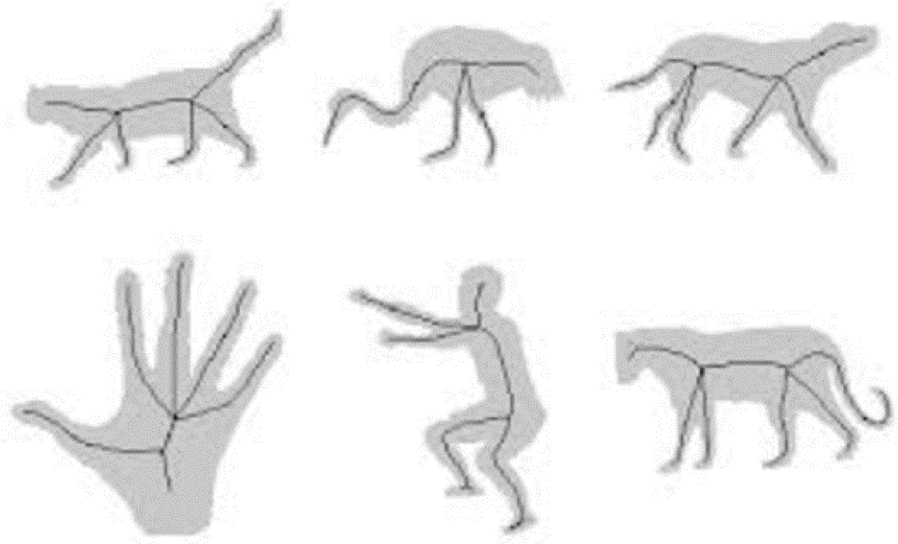

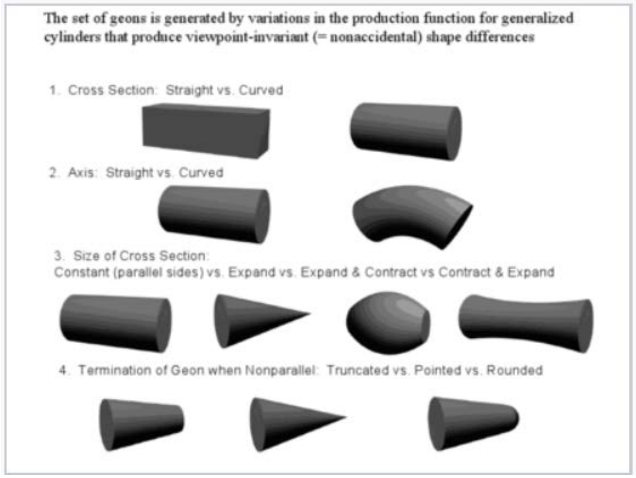

Symmetry Axes: Algorithms

- Key Idea: find the maximal circle which touches both boundaries of the shape and trace the center of this circle as we move along the shape (left)
- Split this central axes into several branches when they bifurcate (right)

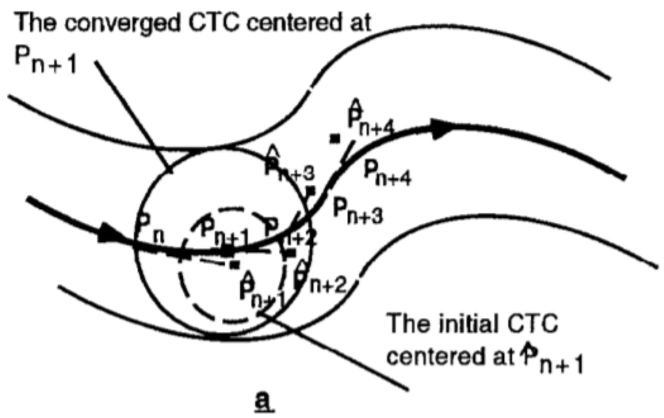

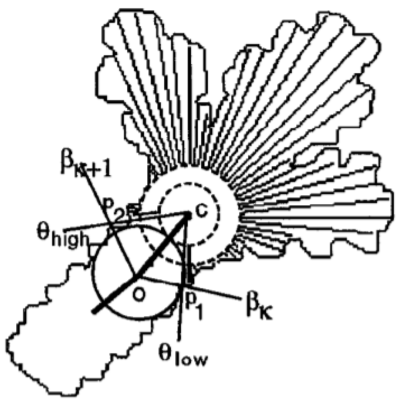

Symmetry Axis and Object Representation

- The symmetry axis can be used to segment a shape into different parts. E.g., decompose a dog into head, torso, legs, and tail (left)
- This yields a representation of objects into parts. It enables systems that can match and recognize objects despite considerable variation in shape (right).

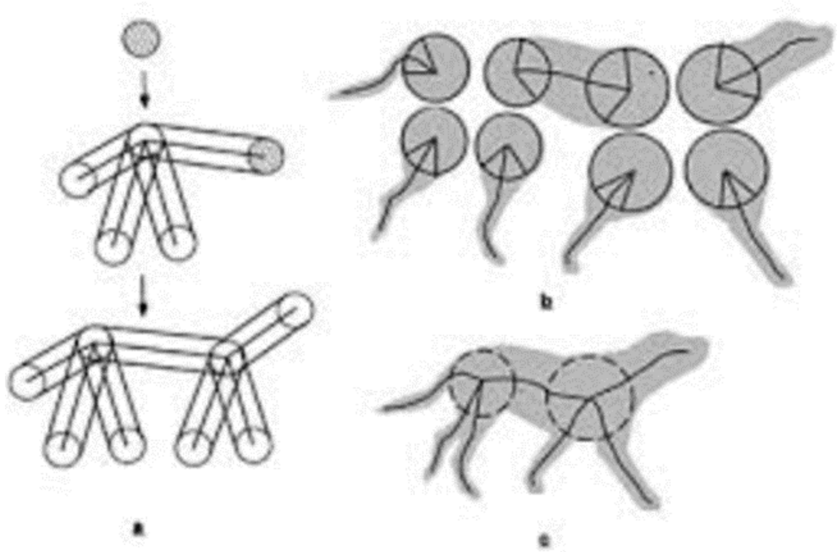

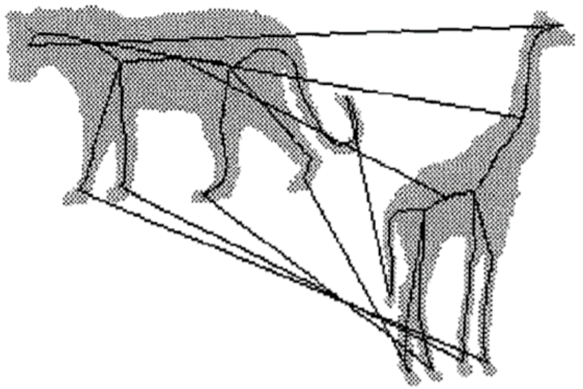

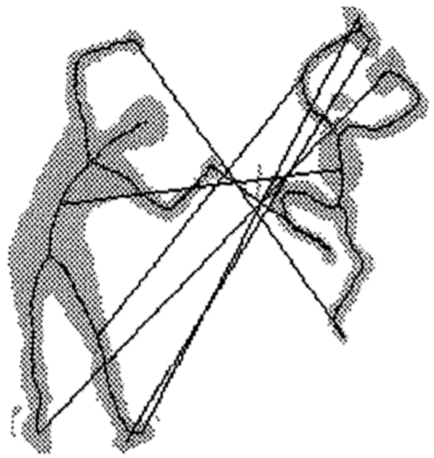

Symmetry Axes and Object Classification

- Enables object recognition despite considerable variations in shape.

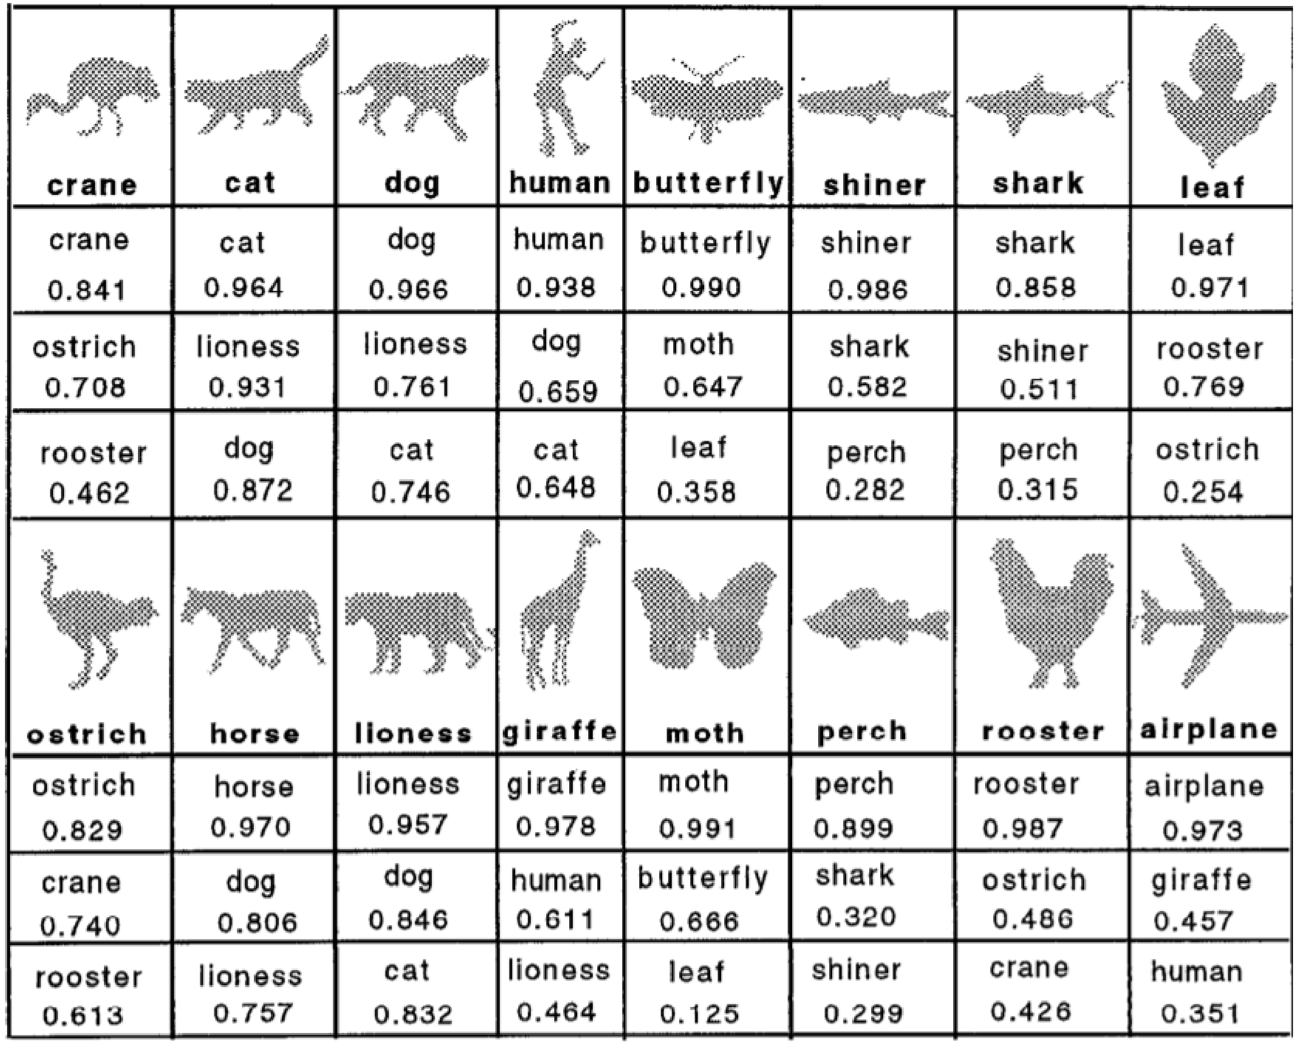

Symmetry Axes Limitations

- Most work on symmetry axes was done on binary images. How to extract symmetry axes from real images? (see next few slides).
- Symmetry axes are only good at decomposing objects into parts from some object configurations. It breaks down is the object is self-occluded (which often happens) or is occluded by other objects.
- Nevertheless symmetry axes relate to geon models of objects – where objects are represented in terms of a combination of (approximately) cylindrical parts.

Segmenting by Seeking the Symmetry Axis

- T-L Liu, D. Geiger, & A. Yuille (ICPR 1998).
- A method for segmenting a shape from an image and simultaneously determining its symmetry axis. The symmetry is used to help the segmentation and in turn the segmentation determines the symmetry. The problem is formulated as one of minimizing a goodness of fitness function and Dijkstra's algorithm is used to find the global minimum of the cost function.

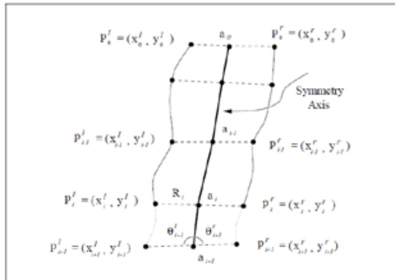

Figure 1. The representation of the shape and the symmetry axis. The left and right boundary points, $\{\vec{p}_i^{\,l}\}$ and $\{\vec{p}_i^{\,r}\}$ are shown. So are the symmetry axis $\{\vec{a}_i\}$ and the rib angles $\{\theta_i^l\}$ and $\{\theta_i^r\}$.

(a) (b) (c)

Symmetry Detection with Deep Networks

- W Shen, K Zhao, Y Jiang, Y Wang, X Bai, A Yuille.
- DeepSkeleton: Learning multi-task scale-associated deep side outputs for object skeleton extraction in natural images. IEEE Transactions on Image Processing 26 (11), 5298-5311. 2017.

Symmetry Axis

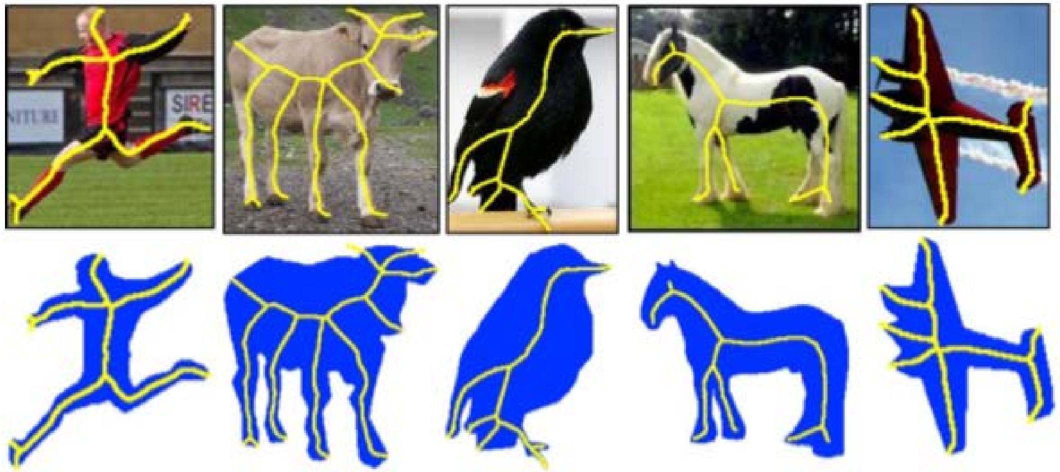

- Examples of symmetry axes.
- Human perception.
- Neuroscience.
- Object skeleton extraction in natural images. The skeletons are in yellow. Top: Skeleton localization. Bottom: Scale prediction which enables object segmentation (blue regions are the segments reconstructed from skeletons according to the scales).

Symmetry is detected at different scales

- Different Scales:

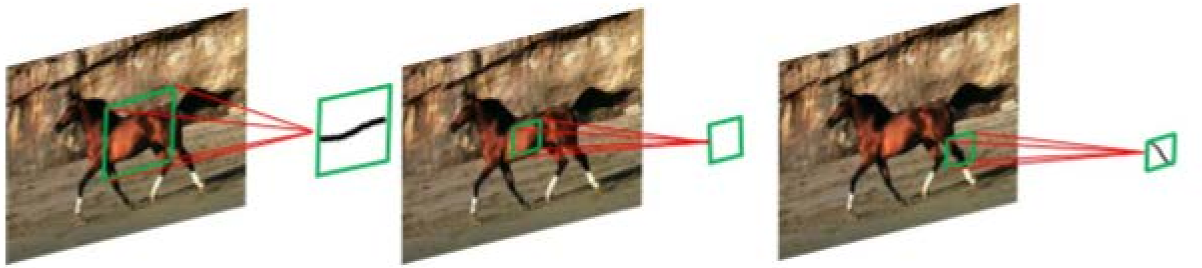

- Using filters (the green squares on images) of multiple sizes for skeleton extraction. Only when the size of the filter is larger than the scale of current skeleton part can the filter capture enough context feature to detect it.

Symmetry

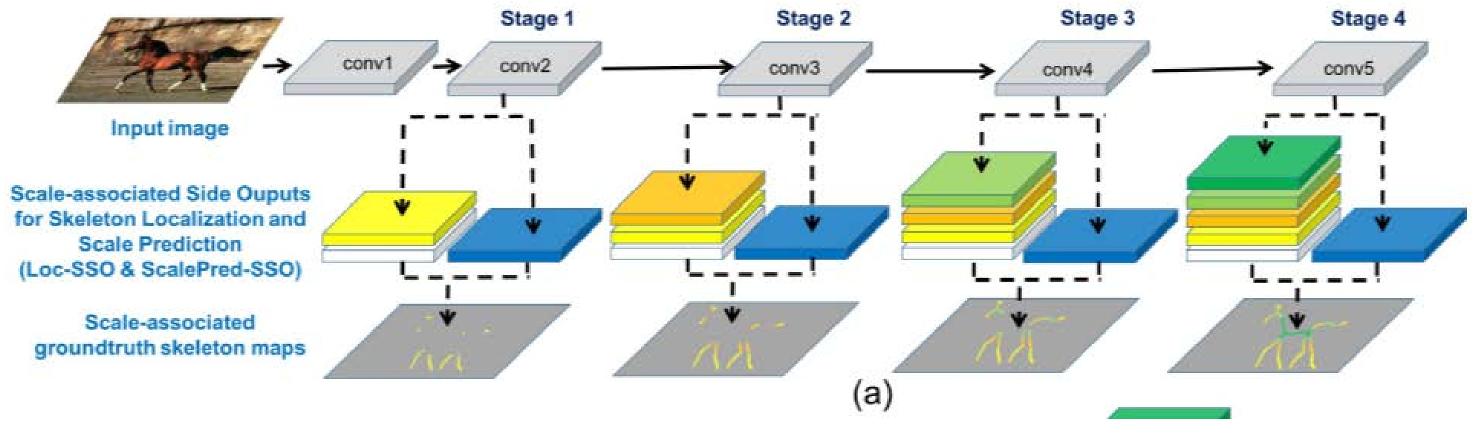

- Multi-task Scale-associated side outputs (SSOs) learning. The network has 4 stages with SSO layers connected to the convolutional layers. Each stage branches into two sibling SSO layers, one for skeleton localization and the other for scale prediction, denoted by Loc-SSO (the left multi-color blocks) and ScalePred-SSO (the right blue block), respectively. The SSOs in each stage are guided by a scale-associated groundtruth skeleton map.
- The skeleton pixels with different quantized scales are in different colors. Each block in a Loc-SSO is the activation map for one quantized scale, marked by the corresponding color.

Symmetry and Scale

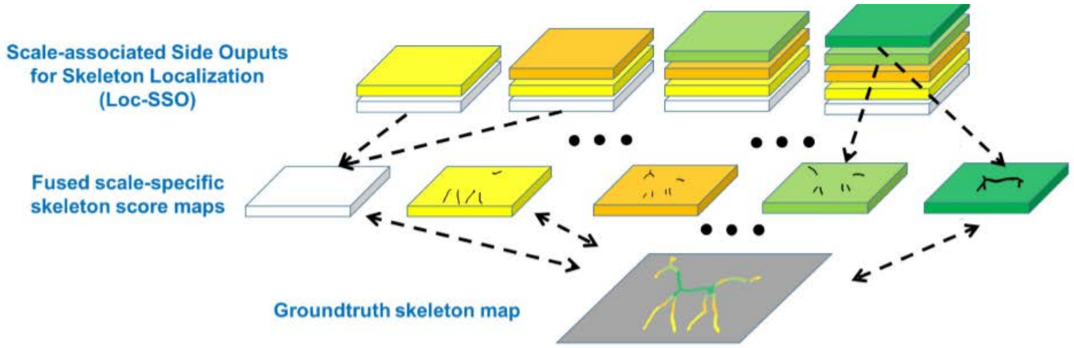

(b)

- Scale-specific fusion. Each Loc-SSO provides a certain number of scale-specific skeleton score maps (identified by stage number-quantized scale value pairs). The score maps of the same scales from different stages will be sliced and concatenated. Five scale-specific weighted-fusion layers are added to automatically fuse outputs from multiple stages.

Symmetry And HED

- Comparison with HED edge detector.

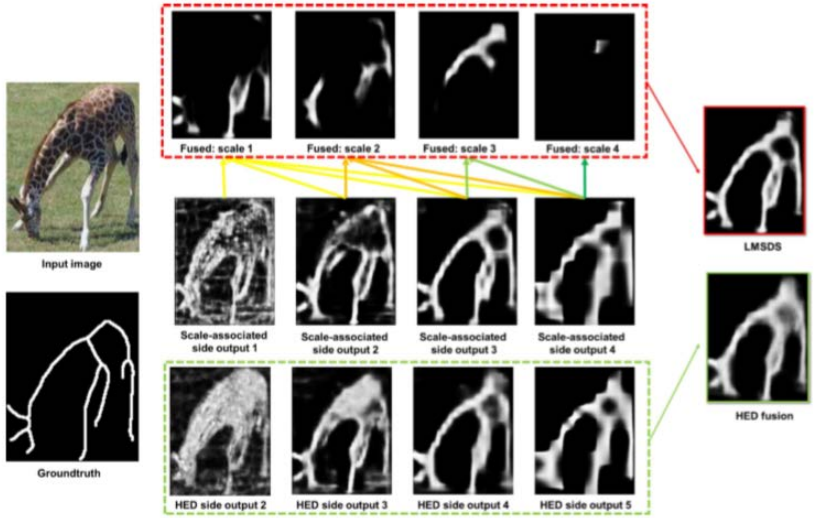

- The comparison between the intermediate results of LMSDS and HED. We can observe that the former are able to differentiate skeleton pixels with different scales, while the latter cannot.

Skeleton and Reconstruction

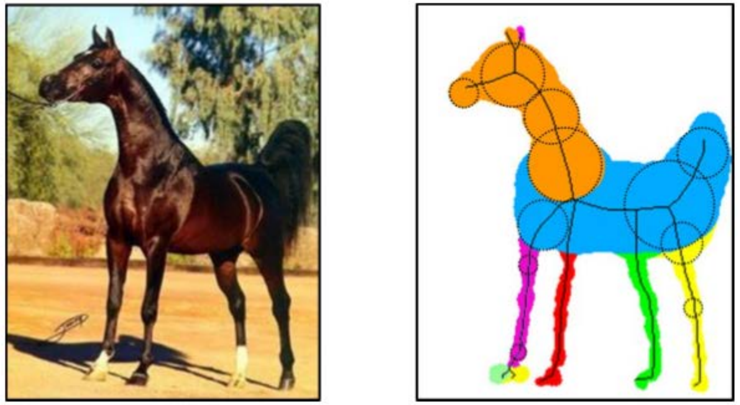

-

- Skeleton based object segmentation. Left: The original image. Right: The object segments reconstructed from the skeleton with scales. Different object segments are marked in different colors. The dashed circles are sampled maximal disks.

Symmetry Comparisons

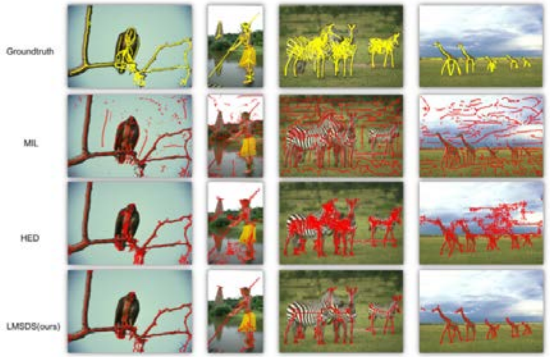

- Illustration of skeleton extraction results on the SYMMAX300 dataset [15] for several selected images. The groundtruth skeletons are in yellow and the thresholded extraction results are in red. Thresholds were optimized over the whole dataset.

More Skeleton Examples

- Illustration of skeleton extraction results on SK-LARGE for several selected images. The groundtruth skeletons are in yellow and the thresholded extraction results are in red. Thresholds were optimized over the whole dataset.

Results: State of the Art (for several months)

- One out of ten tables and five graphs.

TABLE III
SKELETON LOCALIZATION PERFORMANCE COMPARISON BETWEEN DIFFERENT METHODS ON SK-SMALL. †GPU TIME.

Method	F-measure	Avg Runtime (sec)
Lindeberg [13]	0.277	4.03
Levinshtein [16]	0.218	144.77
Lee [17]	0.252	606.30
Particle Filter [18]	0.226	322.25†
MIL [15]	0.392	42.38
HED [19]	0.542	0.05†
FSDS (ours)	0.623	0.05†
LMSDS (ours)	0.621	0.05†