


Hierarchical Scene Annotation

Michael Maire¹, Stella X. Yu², and Pietro Perona¹

¹California Institute of Technology - Pasadena, CA 91125

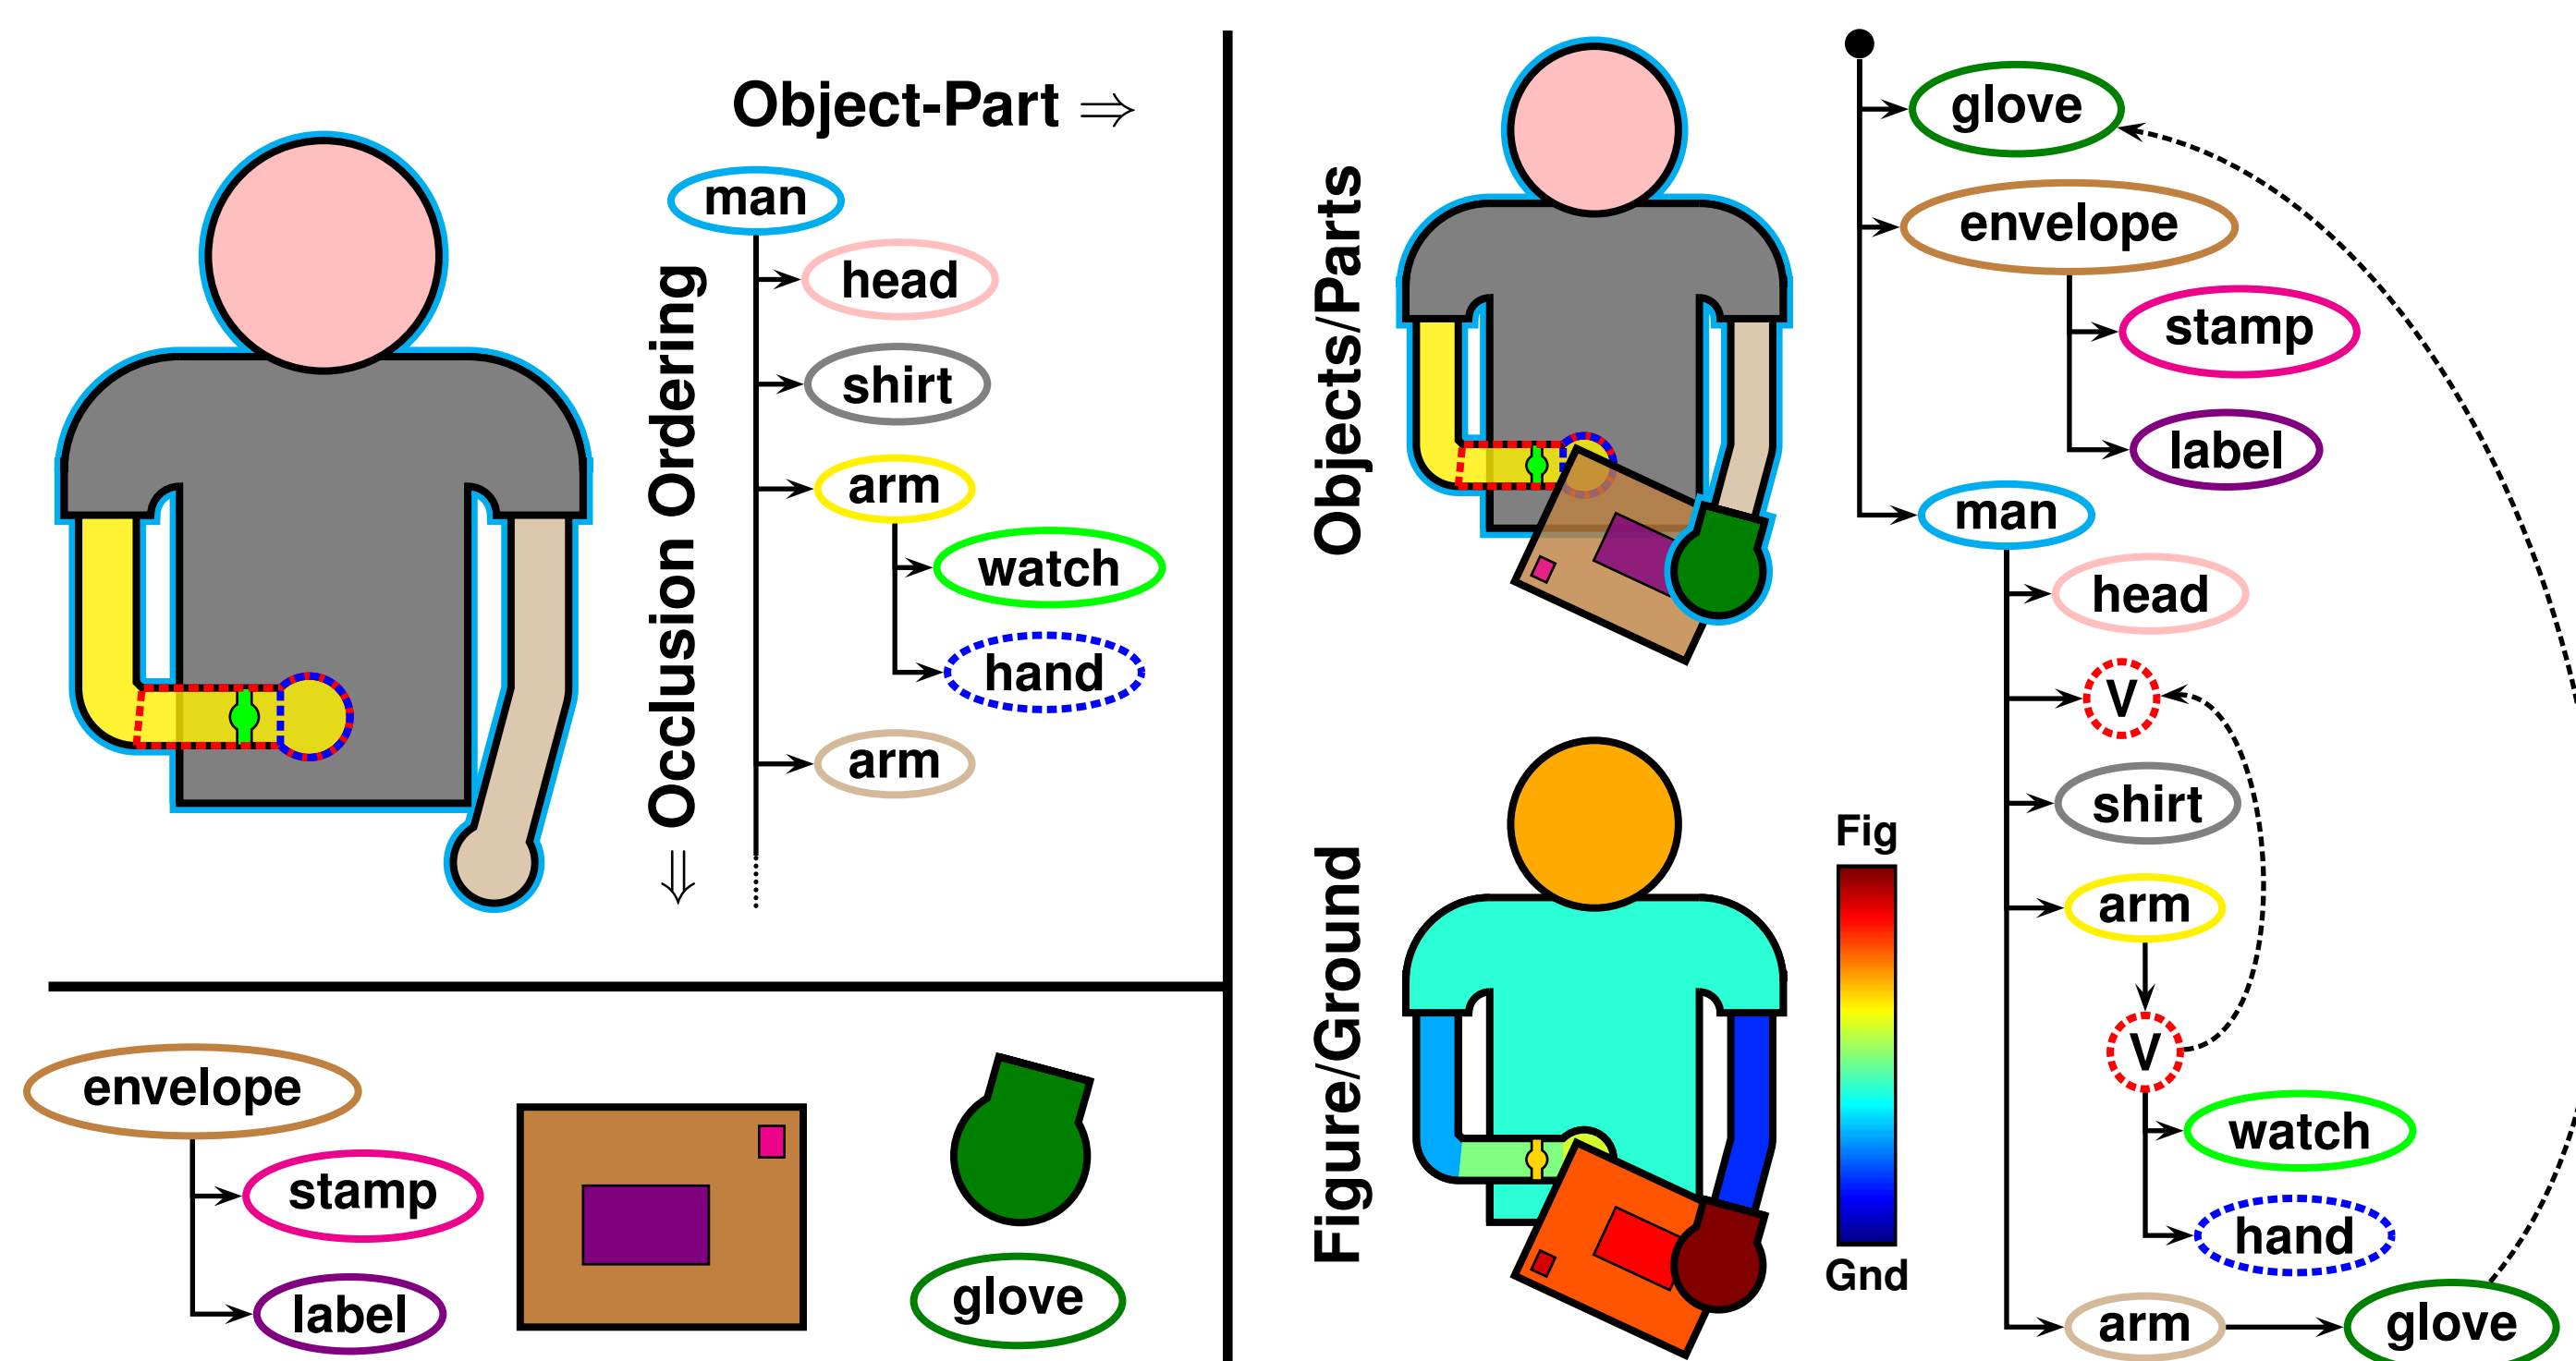
²University of California at Berkeley / ICSI - Berkeley, CA 94704



Overview

- Annotate multiple modalities:
 - Objects, parts, subparts
 - Object-part containment
 - Segmentation
 - Occlusion (figure/ground)
- Unifying abstraction: *region trees*
 - Web-based annotation tool:
 - Computer-assisted segmentation
 - Object segmentation dataset & benchmark
 - Motivation: joint detection & segmentation [2]

Scene Model




- Image \Rightarrow regions $\{R_1, R_2, \dots, R_n\}$ (R_i, R_j possibly overlapping)
- Map each region R_i to a node in doubly-ordered tree:
 - Parent-child node relationships correspond to object-part containment
 - Relative ordering of sibling nodes resolves occlusion ambiguities
- Tree traversal:
 - Recovers object-part hierarchy
 - Converts local occlusion relationships into global figure/ground ordering
- Virtual nodes (dotted ovals): parts without visible boundaries
- Virtual links (dotted arrows): remap occlusion ordering (handle self-occlusion)

Interactive Labeling



- Browse object-part hierarchy by expanding/collapsing visible subtrees (A)
- Drag and drop to rearrange region hierarchy (B)
- Fade occlusion layers by depth to visualize figure/ground ordering (C)
- Interactively edit regions (D)
 - Editor enforces parent-child region containment invariant
 - Superpixel selection brush speeds region definition:




Dataset

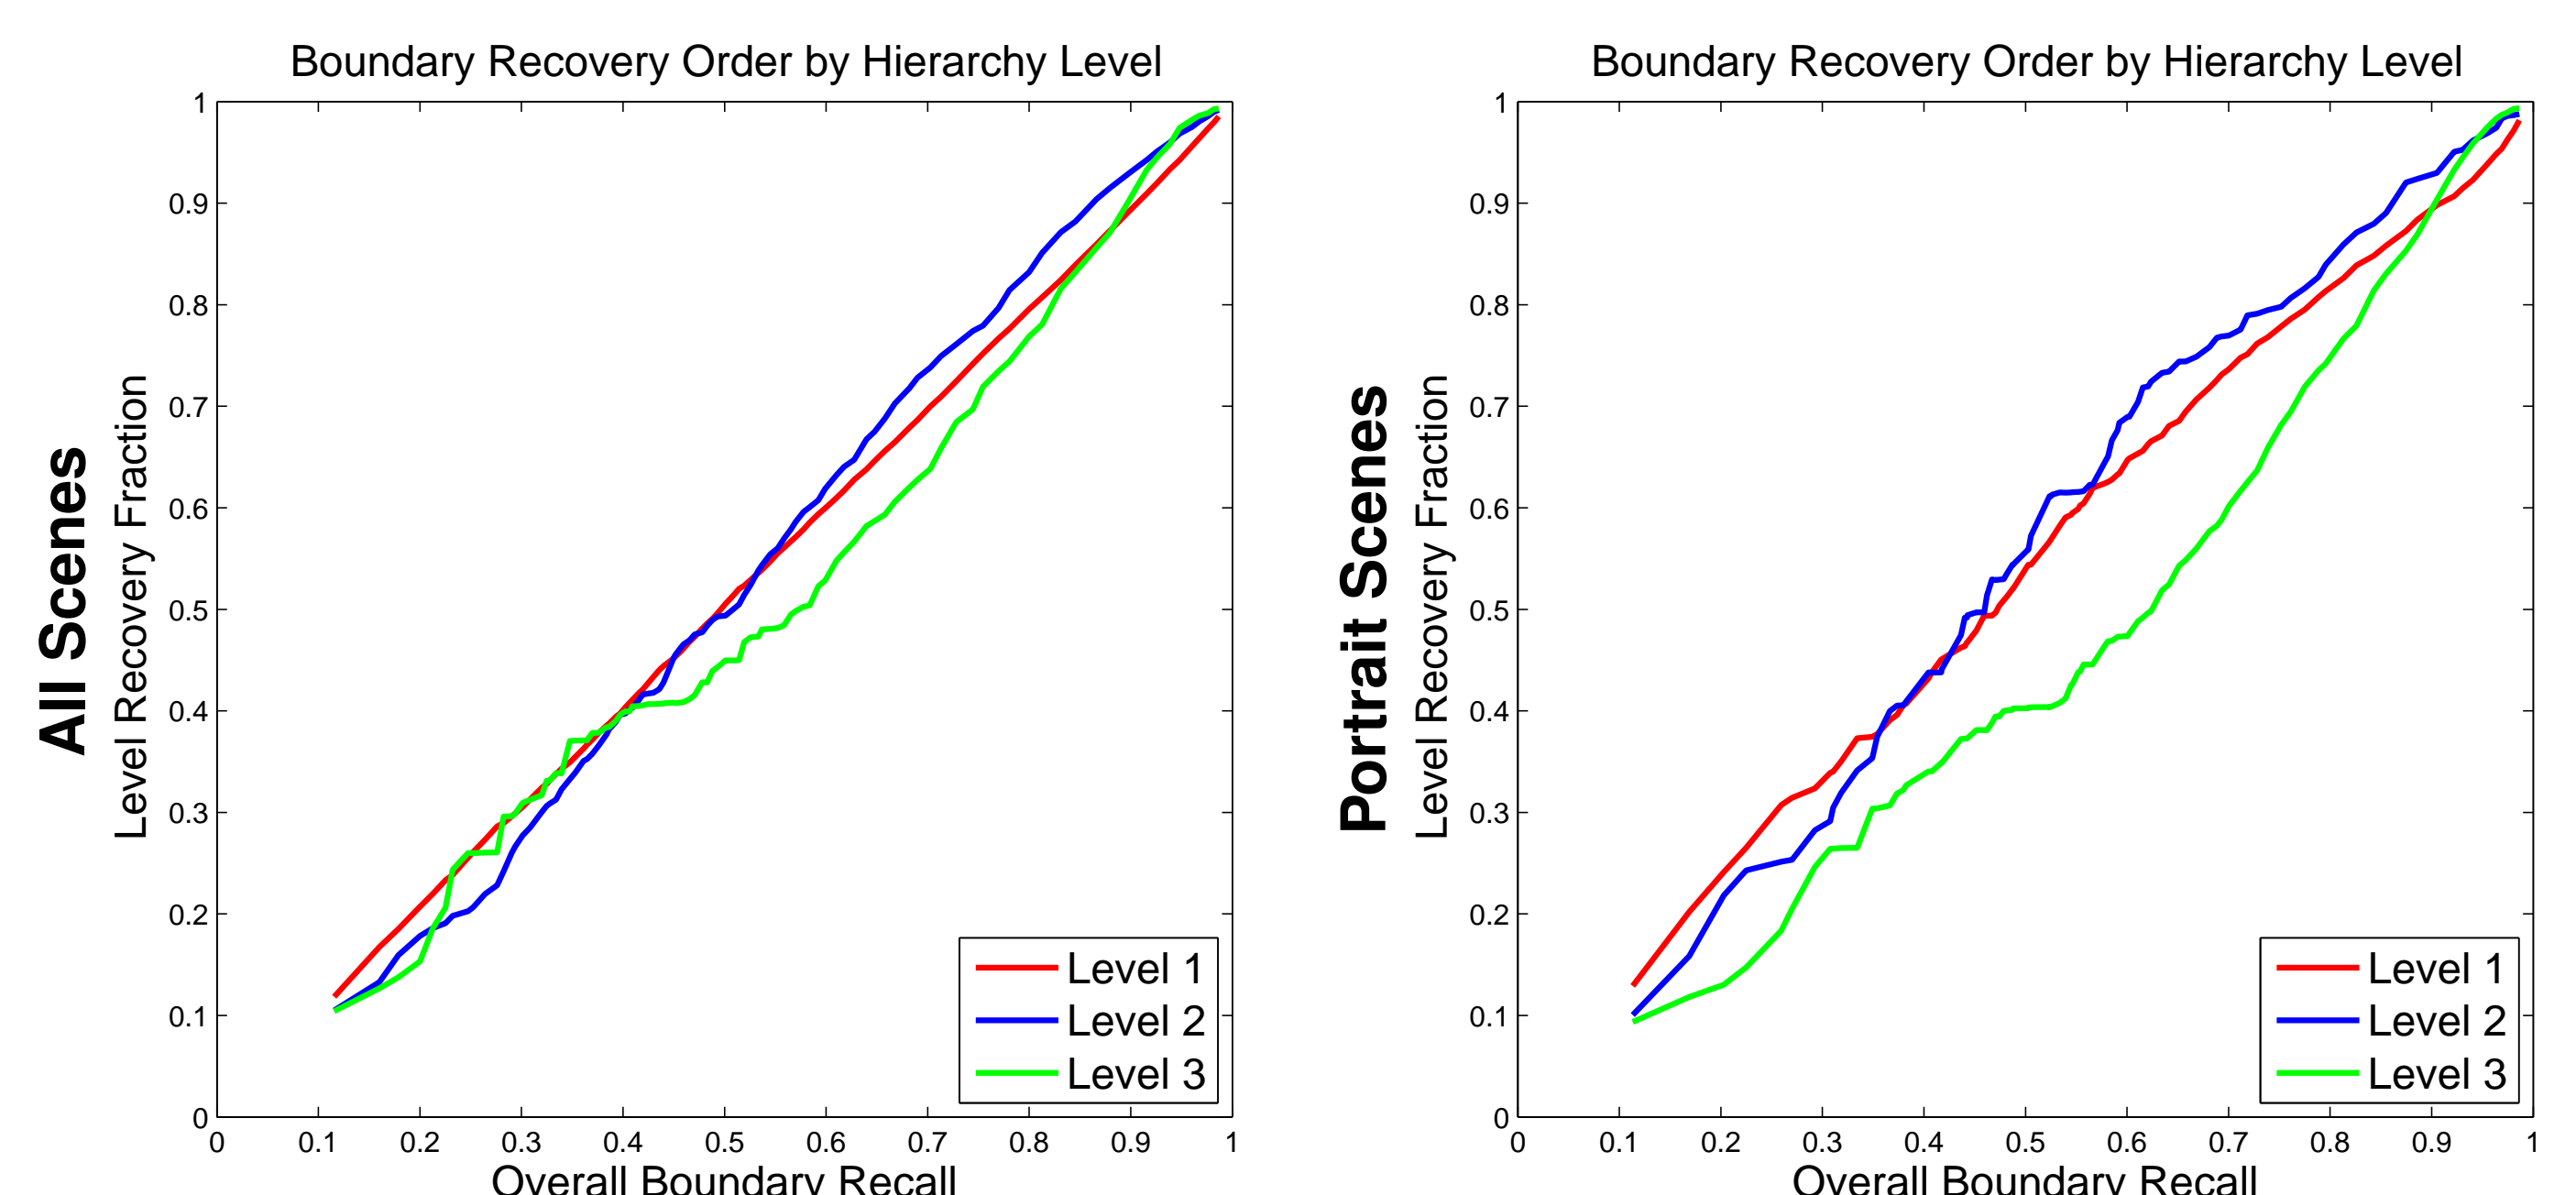


- Complete labeling of 97 complex scenes photographed by artist Stephen Shore [4]
- Compared to Berkeley Segmentation Dataset [3]: more objects, greater scale range

Benchmarks



- Construct groundtruth Ultrametric Contour Map (UCM):
 - Weight each boundary by level at which it appears in object-part hierarchy
 - Determine boundary visibility by region tree traversal
- Does machine hierarchical segmentation [1] respect object-part containment?
 - Residuals (above) show differences by:
 - Type: *false positive* / *false negative* / *incorrect level*
 - Severity: color intensity reflects magnitude of error
 - Plots (below) measure recovery order of groundtruth boundaries:
 - Ideally recover top-level objects, then parts, then subparts (levels 1, 2, 3)
 - gPb-UCM only recovers groundtruth hierarchy on simpler scenes (portraits)



- P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik. Contour Detection and Hierarchical Image Segmentation. *PAMI*, 2011.
- M. Maire, S. X. Yu, and P. Perona. Object Detection and Segmentation from Joint Embedding of Parts and Pixels. *ICCV*, 2011.
- D. Martin, C. Fowlkes, D. Tal, and J. Malik. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. *ICCV*, 2001.
- M. Spain and P. Perona. Measuring and Predicting Object Importance. *IJCV*, 2010.