

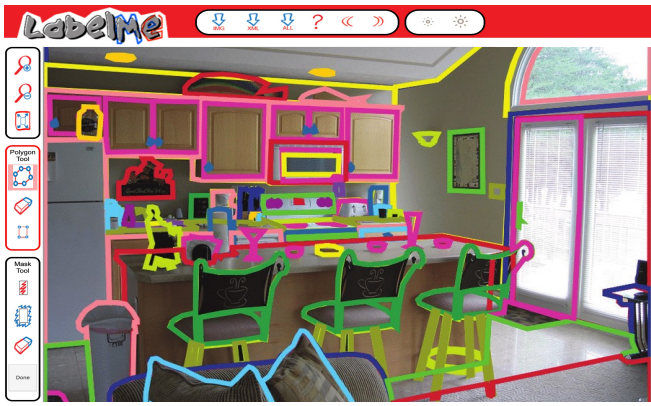
Scaling up instance annotation via label propagation

Dim P. Papadopoulos*, Ethan Weber*, Antonio Torralba



Instance segmentation

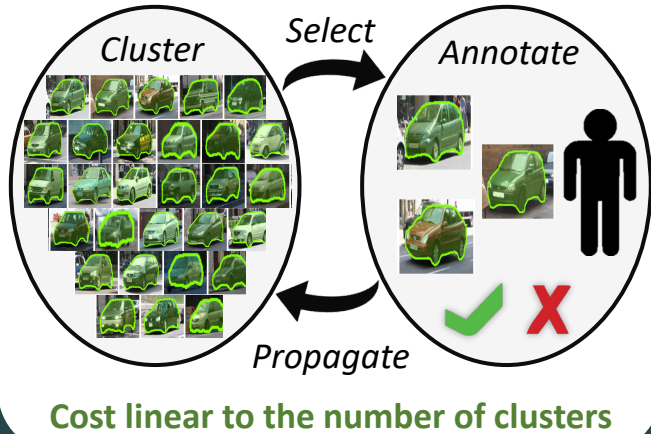
Obtaining training data



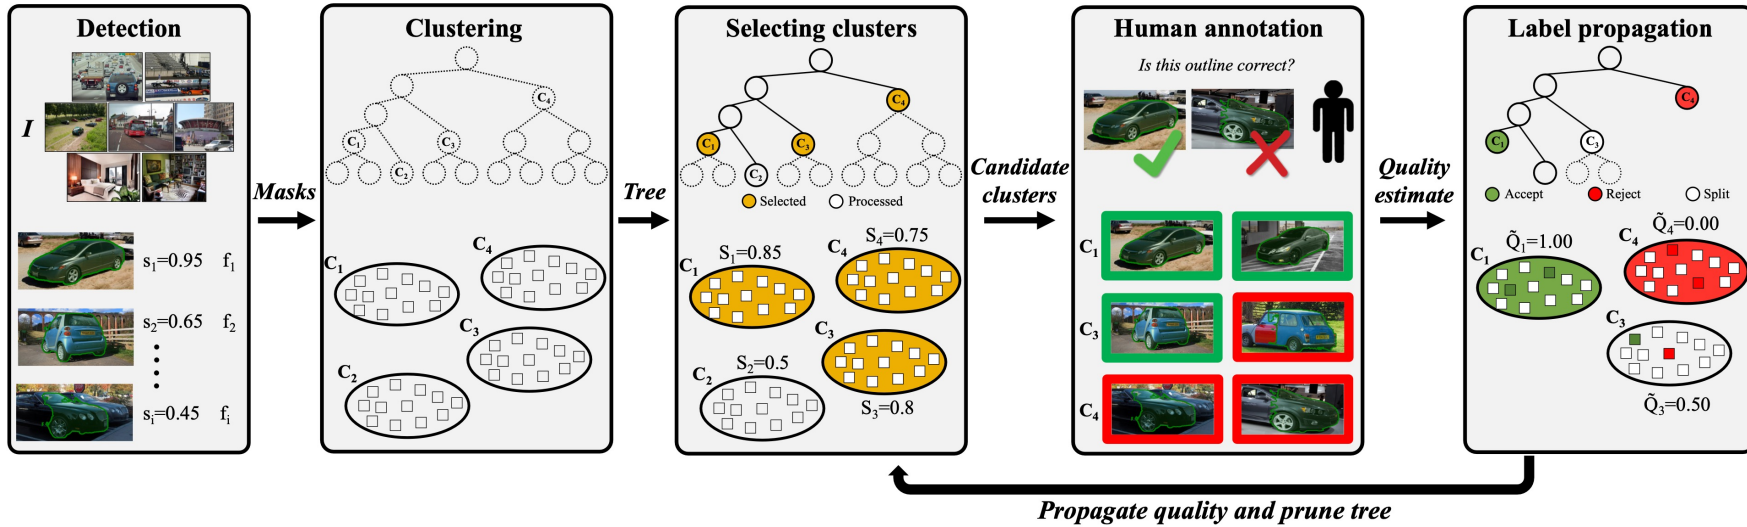
- **Manual annotation:** expensive (80s per object)
- **Interactive segmentation:** faster human interaction

Cost linear to the number of instances

Our method



Method

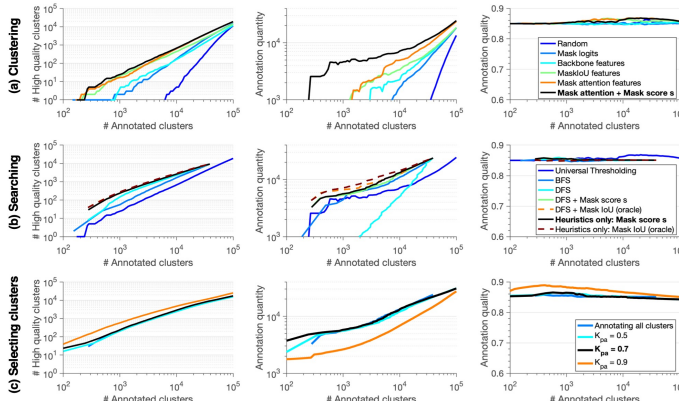


Experiments

Simulation (ADE20K)


- Train initial model on 10K images
- Unlabeled pool of 10K images
- PointRender with ResNet50
- 80 object classes

- Feature representation
- Searching the tree
- Selecting clusters
- Initial training set
- Number of samples
- Quality of clusters
- IoU quality of masks




Real large-scale (Places)

- Quantity (1M masks)
- Usefulness (12.7% mAP)



train	AP	Time (h)
ADE 1K	8.1	217h
ADE 1.4K	9.6	290h
Ours	12.7	290h

- Quality (81.4% mIoU)



<http://scaling-anno.csail.mit.edu/>