Multiscale Conditional Random Fields for Semi-supervised Machine Vision

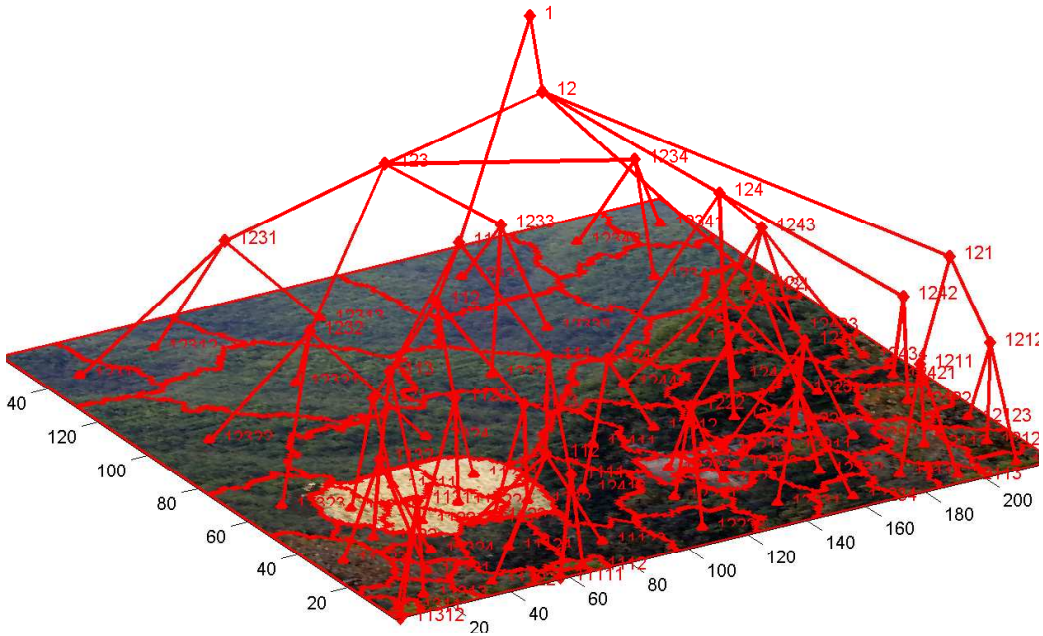

David Duvenaud, Benjamin Marlin and Kevin Murphy
University of British Columbia

Object Detection and Localization

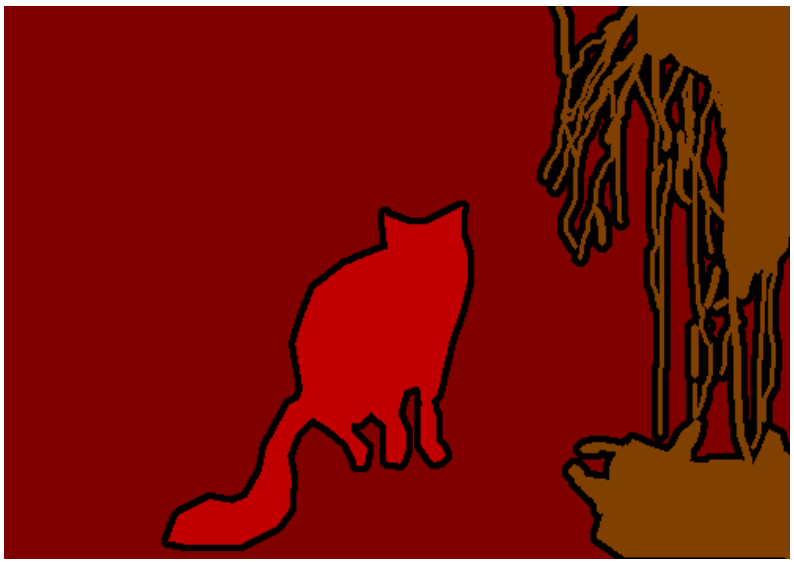

- Solvable using local image models, and supervised learning.
- Big problem: Pixel-labeled data is scarce and expensive.

Motivation

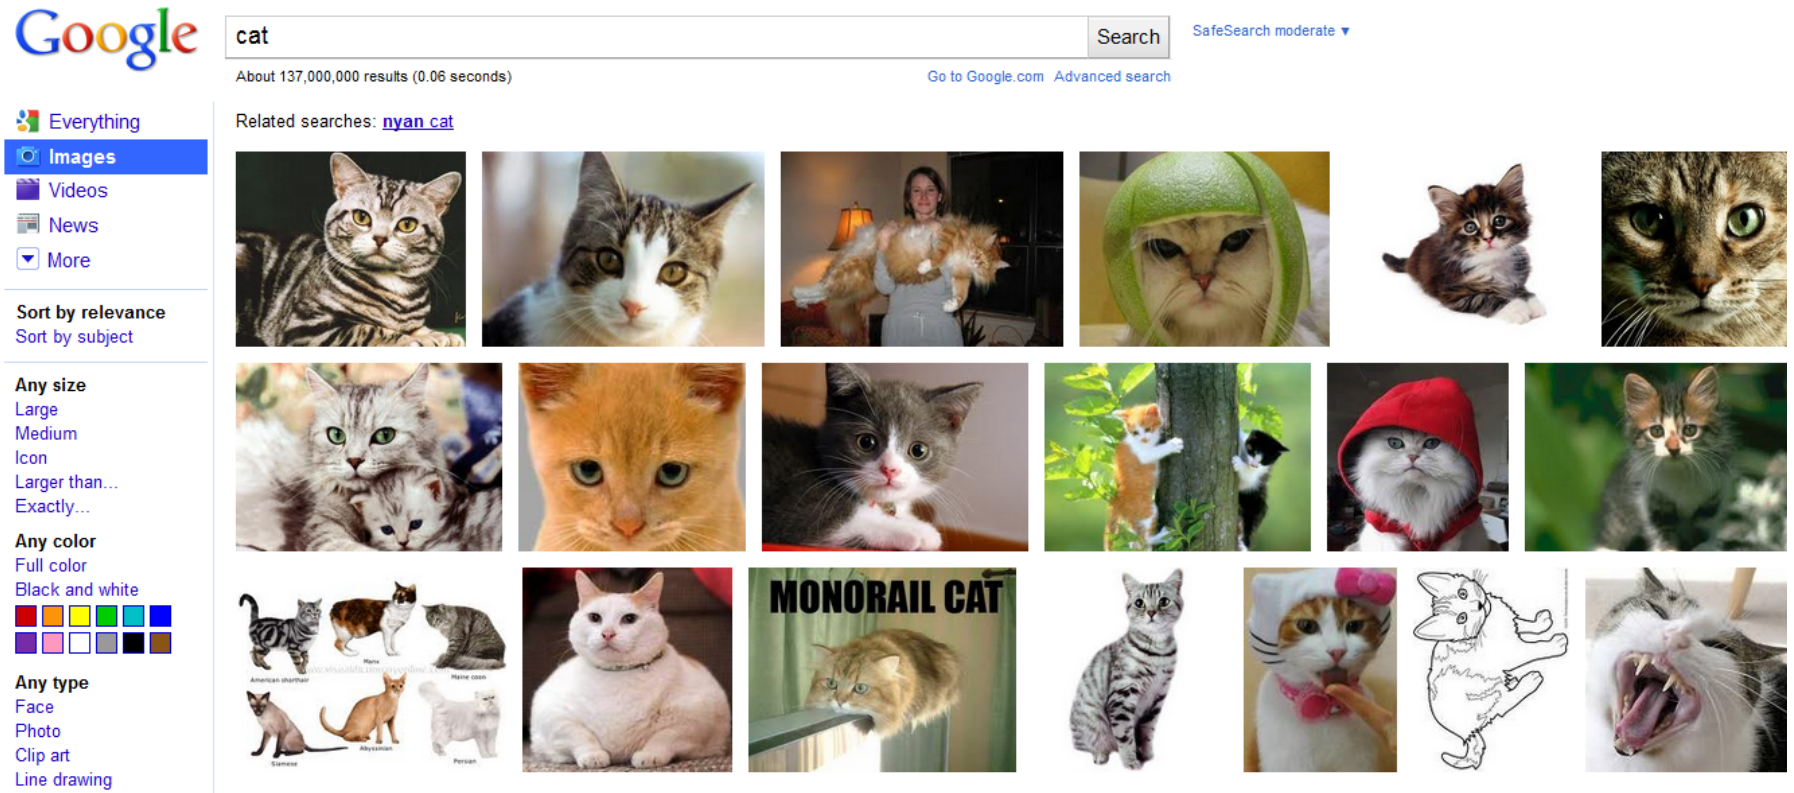

- Much easier to find images known to contain an object, without knowing where.

Motivation

- How to train a local image model, using globally-labeled data?
- An important problem.
- Seems like it should be possible...

Multiscale model

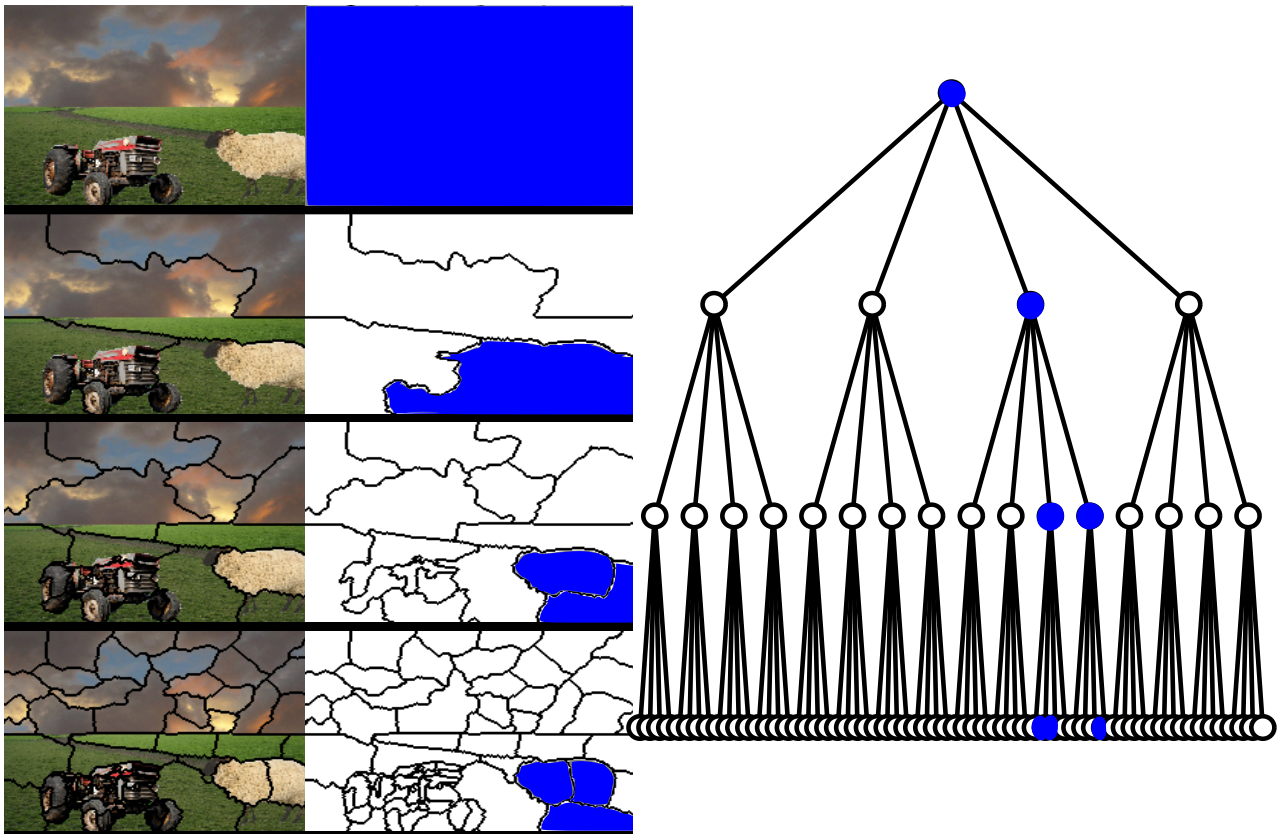

- Perform segmentation at several scales.
- Define a joint object presence model over all scales.
- Evidence from global scale can flow downwards to local scale.
- Enough evidence should resolve ambiguities.

Recursive Segmentation

Recursive Segmentation

How to Build a CRF

$\rightarrow$

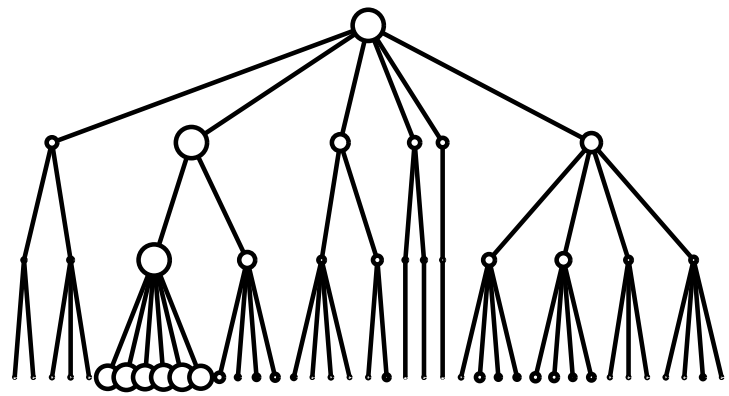

- Need to estimate local object probabilities for each patch, and combine them.

Image Segments to Class Probabilities

$\rightarrow$

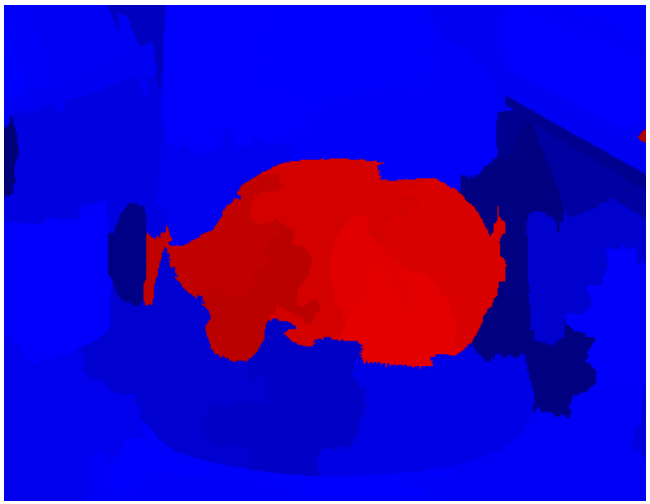

- Patch Features: Colour histograms, HoG, textons, location.
- Each feature vector has 826 dimensions.
- Could use GIST at top level.

Class probabilities to CRF

$\rightarrow$

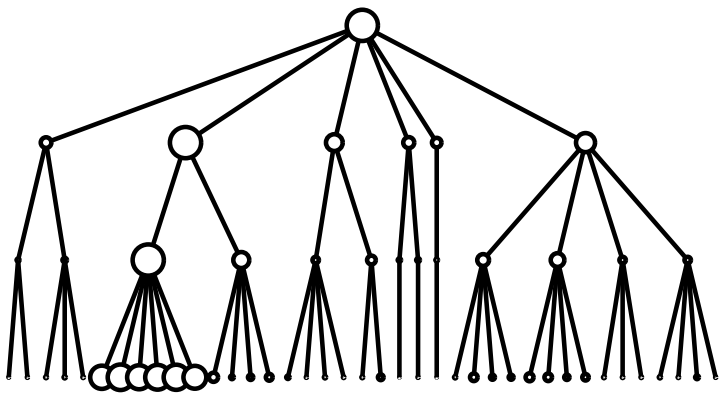

- Local class probabilities become nodes on a Conditional Random Field
- CRF has factors enforcing consistency

Evidence Flow

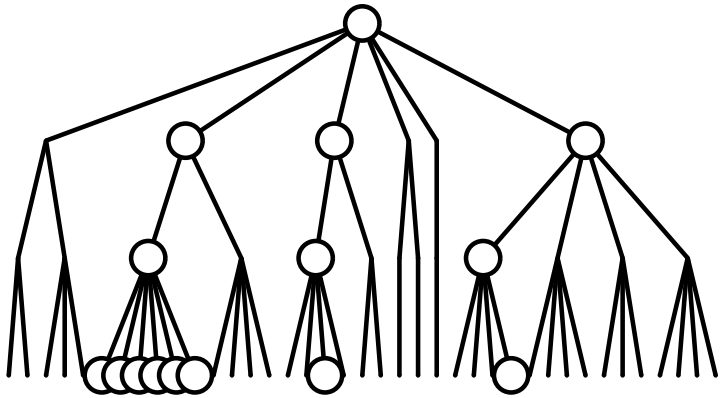

True Labels

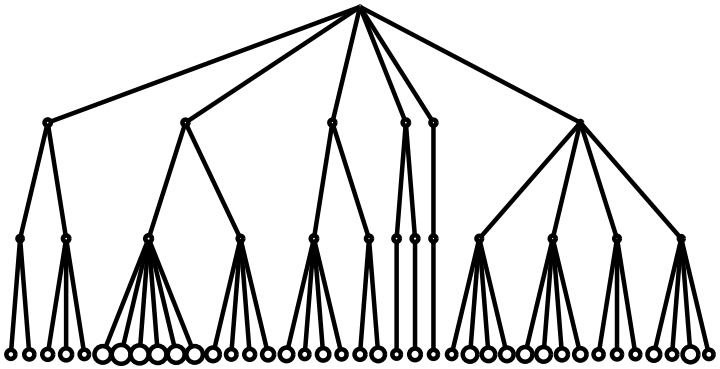

Local Evidence

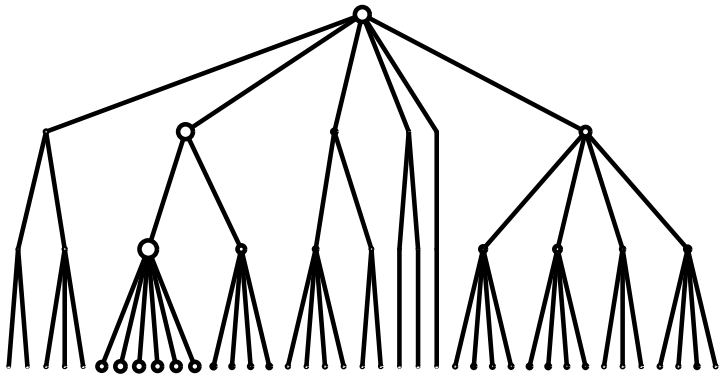

Tree Marginals

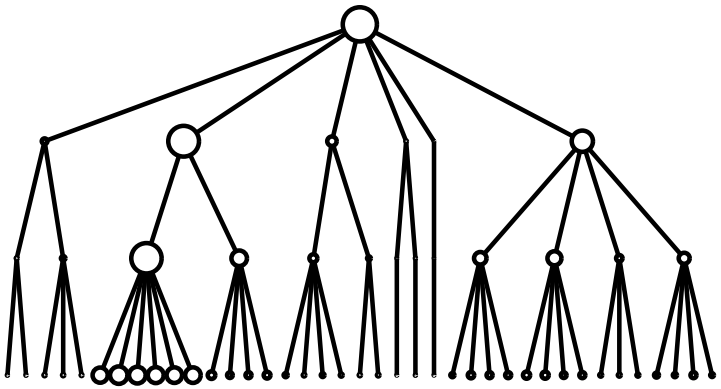

Clamped Marginals

Semi-supervised Learning

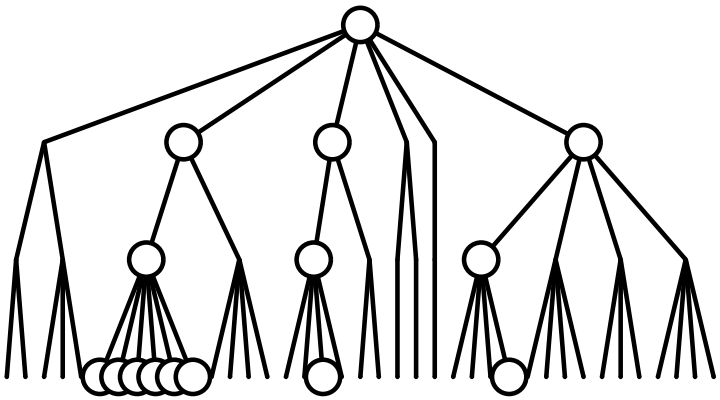

Fully Labeled

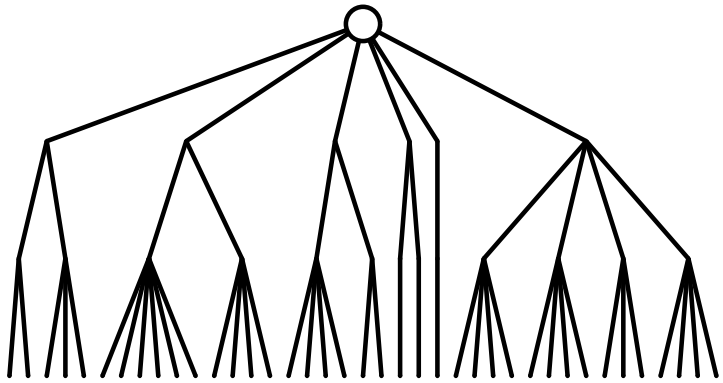

Globally Labeled

- Supervised setting: Trained on 400 fully-labeled examples.
- Semi-Supervised: Trained on an additional 400 globally-labeled examples.
- Used E-M for learning, Belief Propagation for inference.
- Learning the local image models, and the CRF joint factors.

Models

Image

Segmentation

True labels

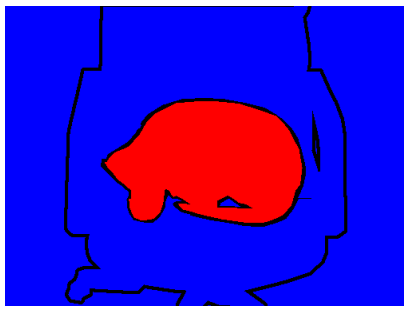

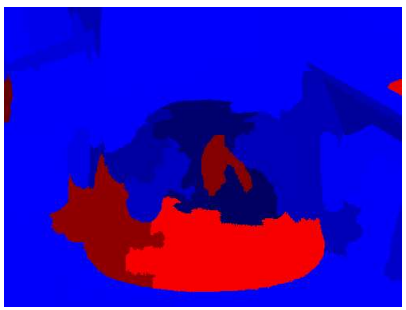

Independent

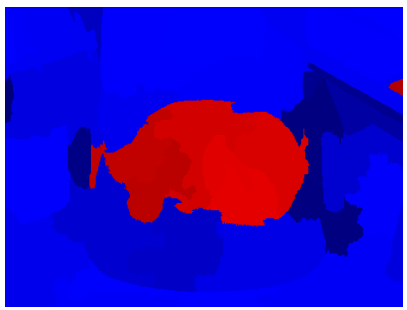

Pairwise Trees

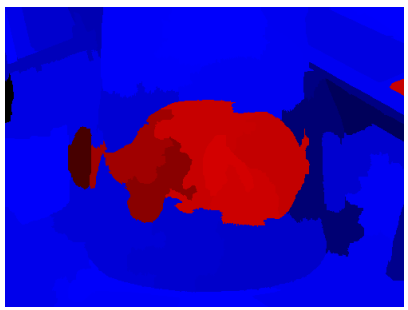

Noisy-or Trees

Results

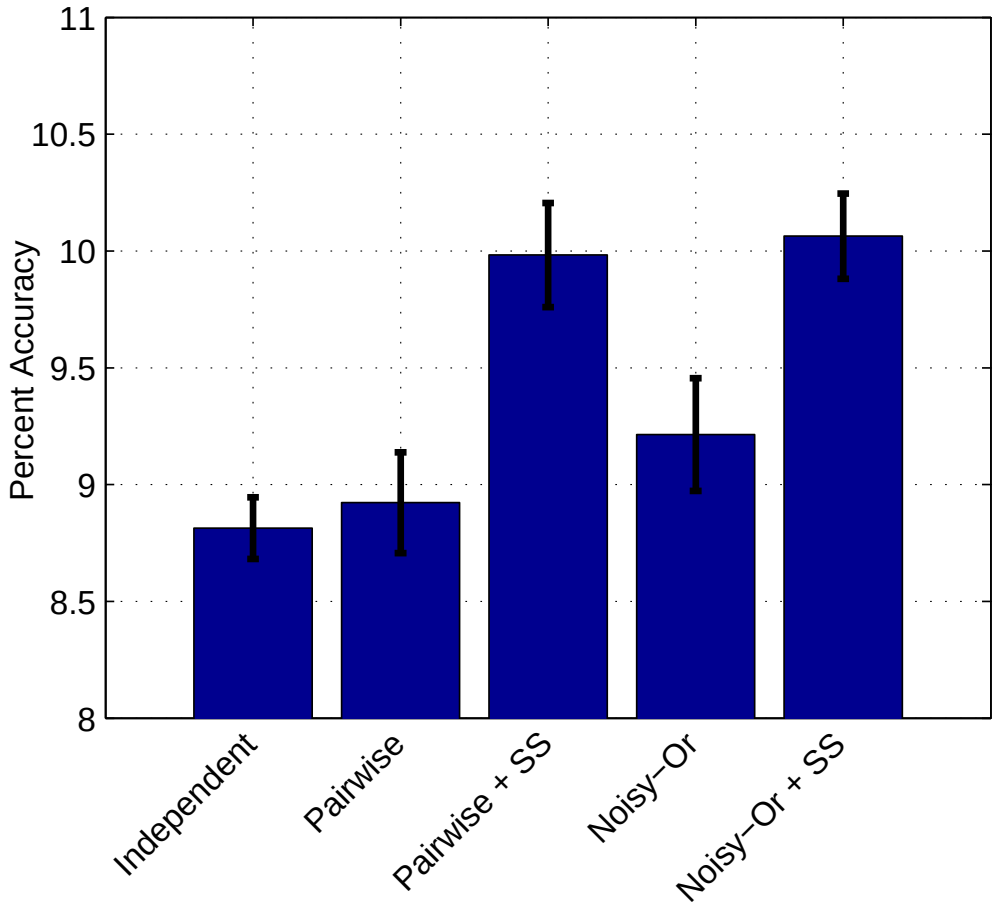

Globally-labeled data improves pixel-level labels!

Pixel-level Improvement from Semi-Supervised

Which classes benefitted?

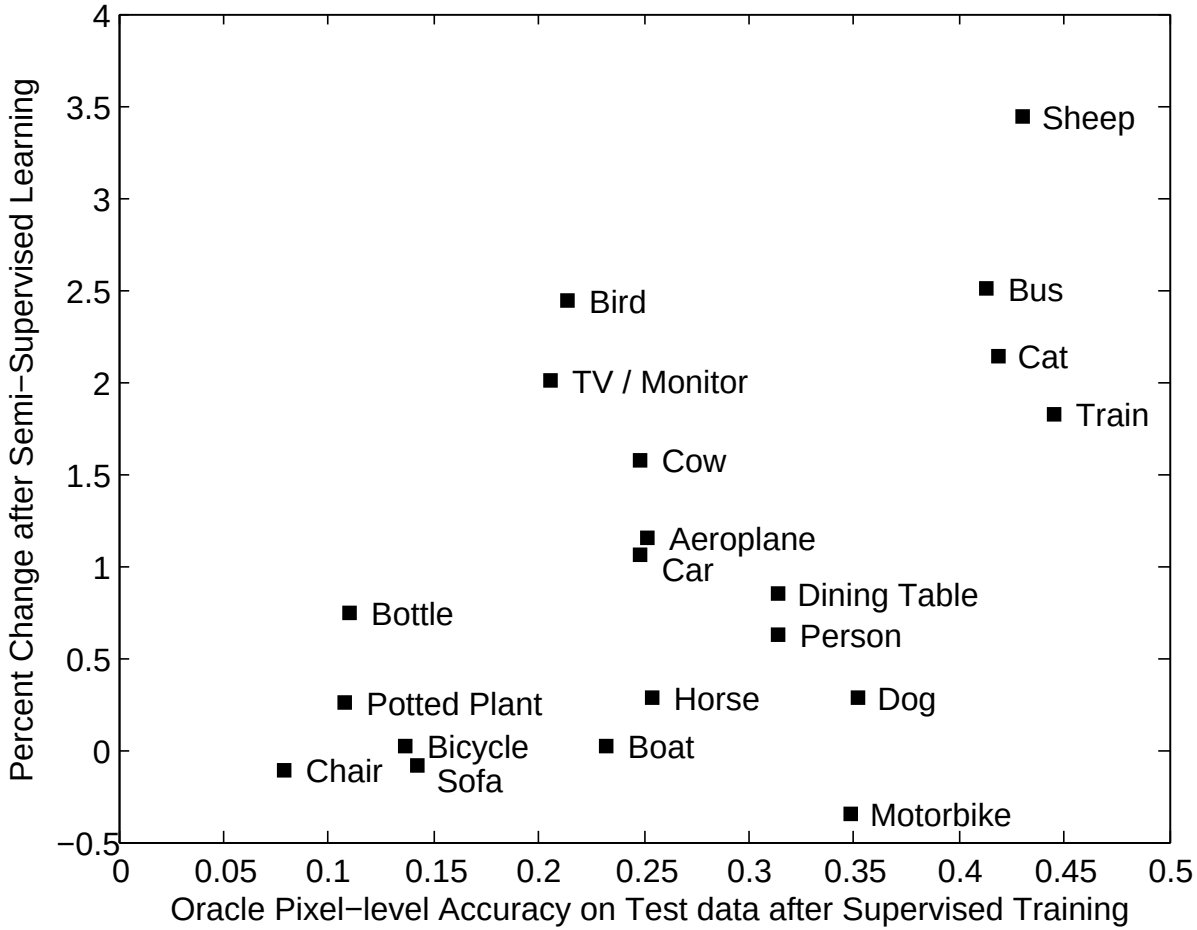

Controlled for better global-level model.

Combining evidence

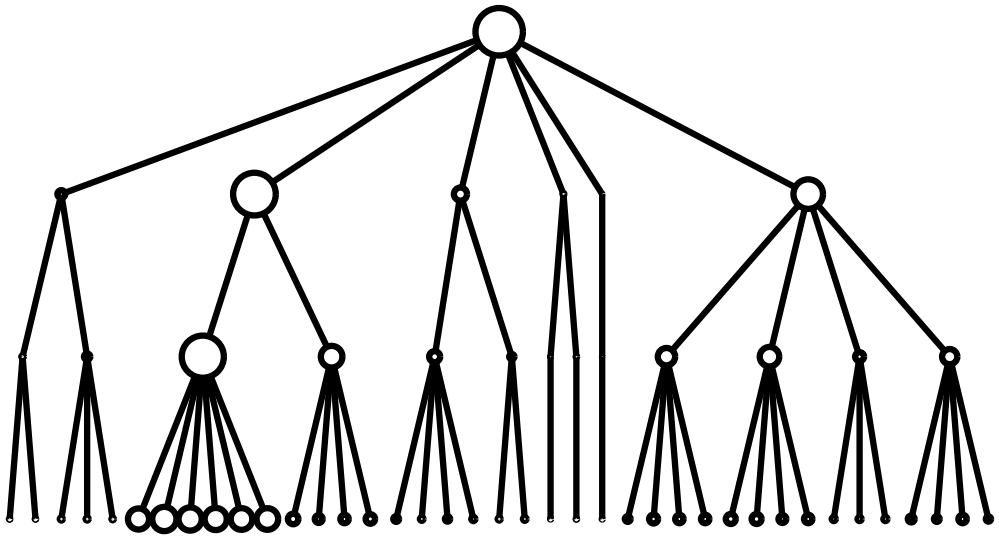

- Could incorporate evidence from other classifiers
- Could handle noisy labels
- Could learn from loose bounding boxes

Conclusion

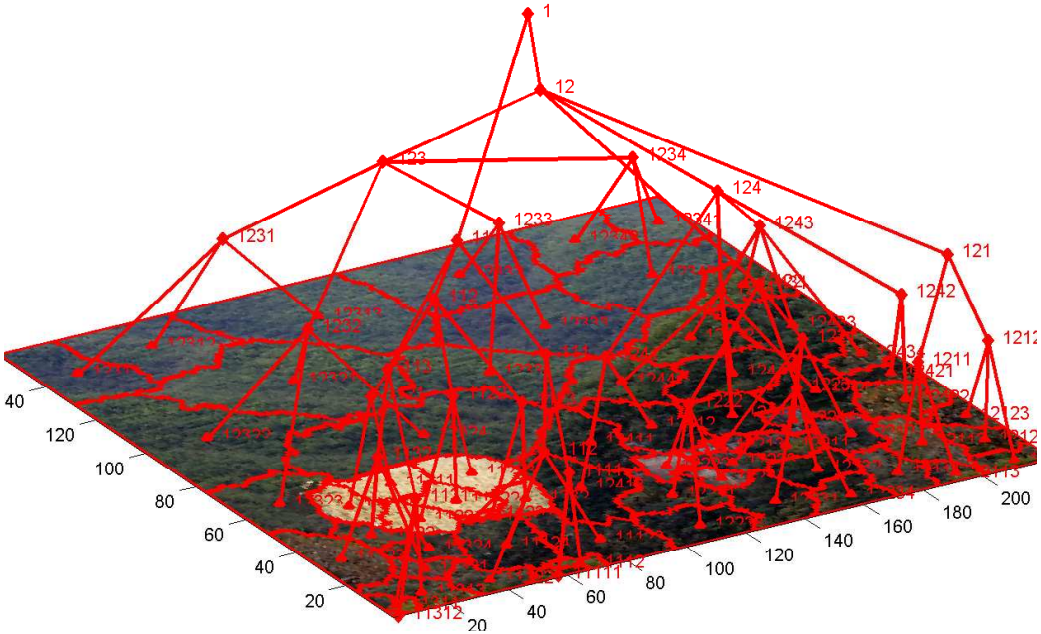

- Learning a local image model from globally-labeled data is possible.
- All you need is a joint probabilistic model over all scales.
- Questions?