


Introduction

• Goal

- Automatically annotate segments in weakly labeled video taken from YouTube



↓ Spatiotemporal segmentation



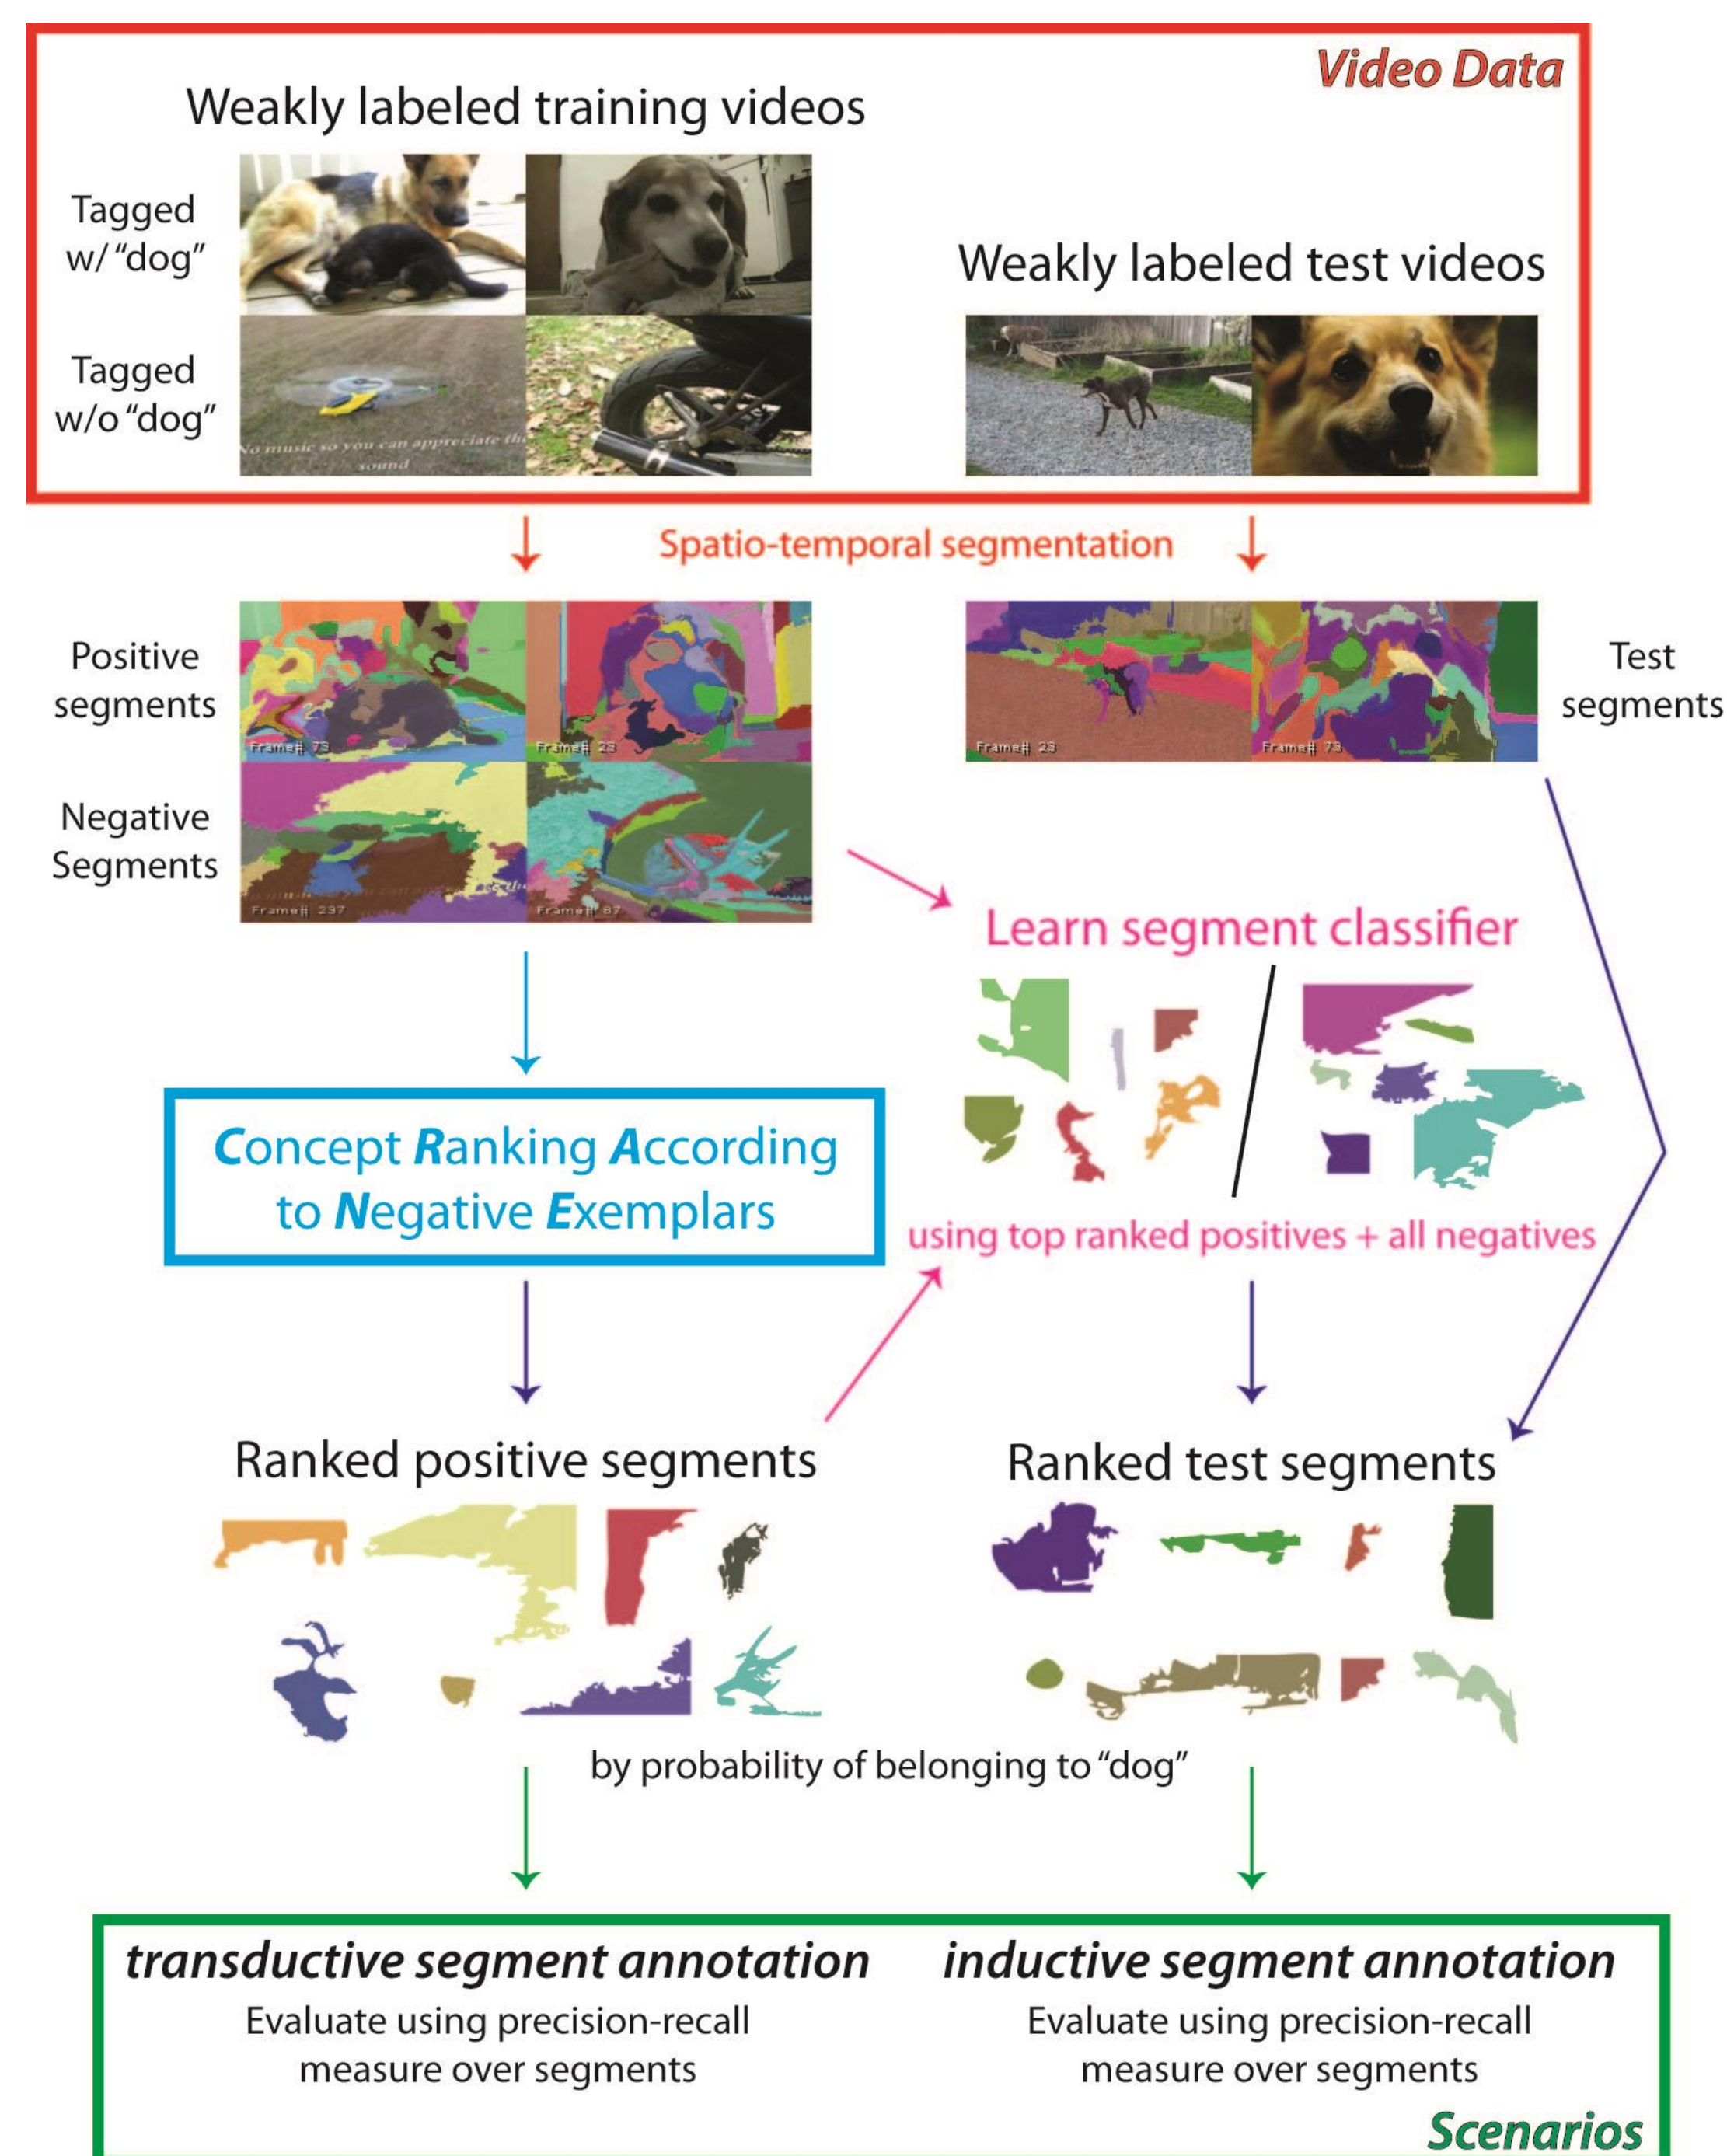
↓ Semantic object segmentation



• Challenges

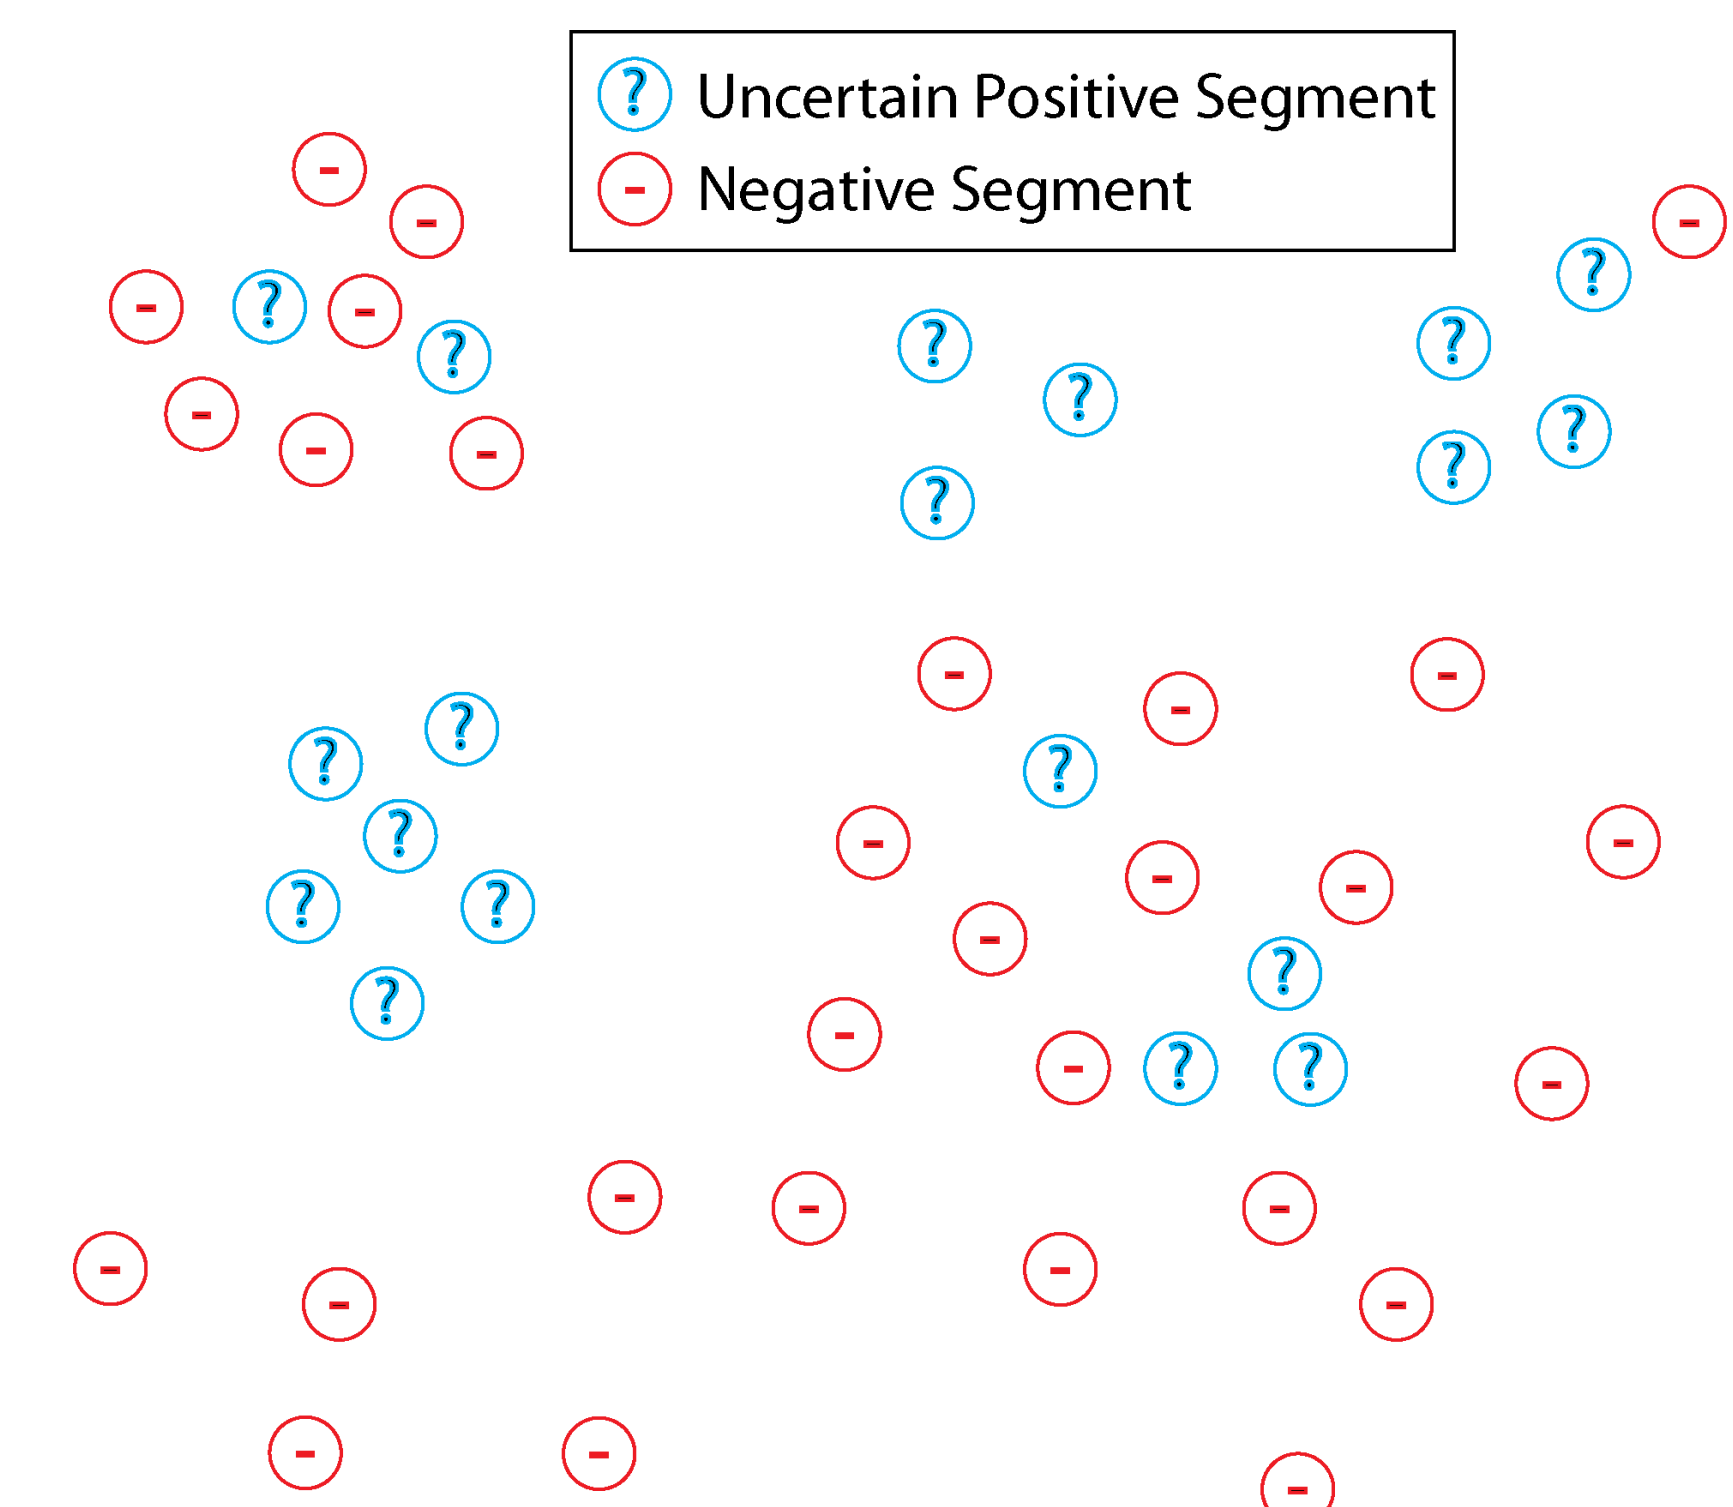
- Learning from weakly labeled data
- Handling label noise in YouTube tags
- Parallelize to deploy over large amounts of YouTube data

Our Problem Setup



Our Algorithm: CRANE


- Input: uncertain positive segments, large set of negative segments
- Output: ranked positive segments by probability of belonging to our concept



Intuition: Positive segments are less likely to belong to our concept if they are near many negative segments.

$$S_{\text{CRANE}}(s_i) = - \sum_{z \in \mathcal{N}} \mathbf{1} \left[s_i = \arg \min_{t \in \mathcal{P}} (\text{dist}(t, z)) \right] \cdot f_{\text{cut}}(\text{dist}(s_i, z))$$

Sample Object Segmentations

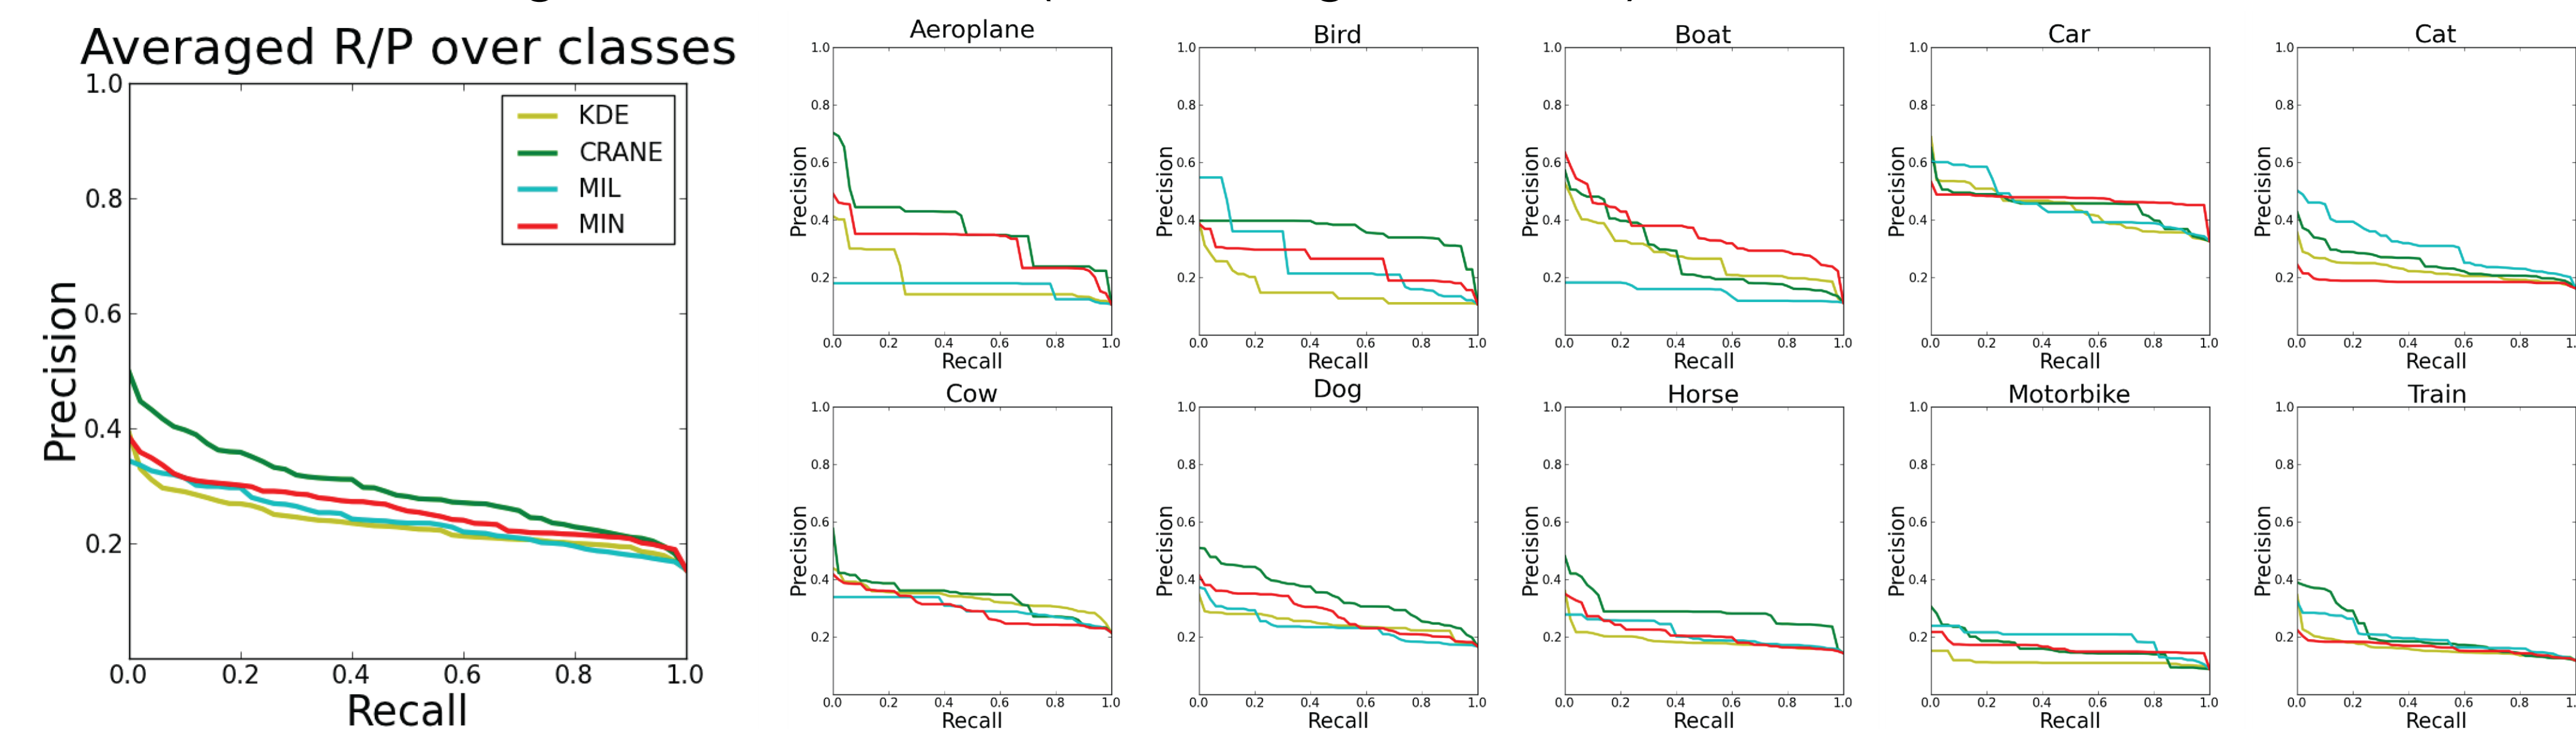


Inductive Segment Annotation [top two rows]
Transductive Segment Annotation [bottom two rows]

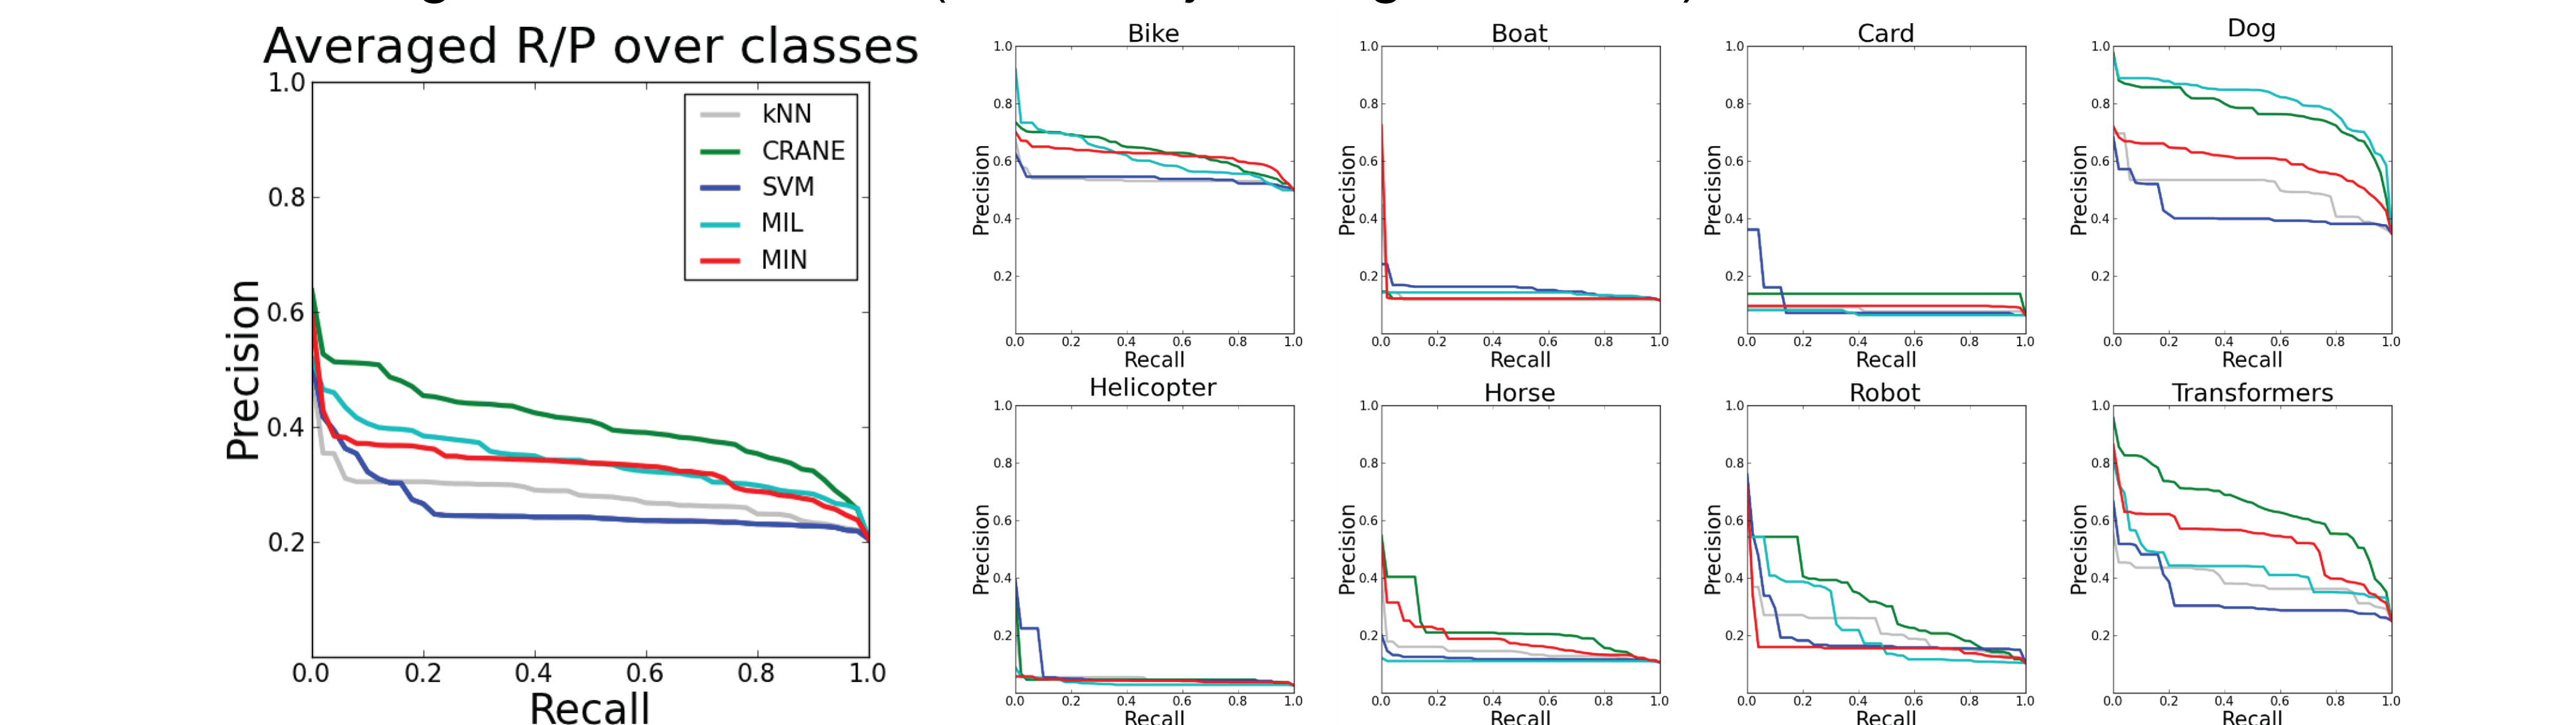
Common Failure Cases

Quantitative Results

- Transductive Segment Annotation (annotating a dataset)



- Inductive Segment Annotation (novel object segmentation)



[1] P. Siva, C. Russell, and T. Xiang. In defence of negative mining for annotating weakly labelled data. ECCV 2012.
 [2] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph-based video segmentation. CVPR 2010.
 [3] G. Hartmann et al. Weakly supervised learning of object segmentations from web-scale video. ECCV 2012 Workshop.
 [4] A. Prest et al. Learning object class detectors from weakly annotated video. CVPR 2012.