Active Detection
via Adaptive Submodularity

Yuxin Chen[†], Hiroaki Shioi[‡], Cesar Antonio Fuentes Montesinos[†]
Lian Pin Koh[†], Serge Wich[¶] and Andreas Krause[†]

| ICML Beijing | June 23, 2014 |

† ETH Zürich
‡ 東京大学 THE UNIVERSITY OF TOKYO
¶ LIVERPOOL JOHN MOORES UNIVERSITY




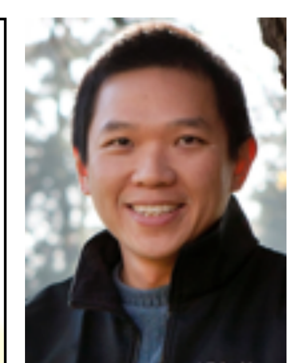

Motivating Example: Biodiversity Monitoring

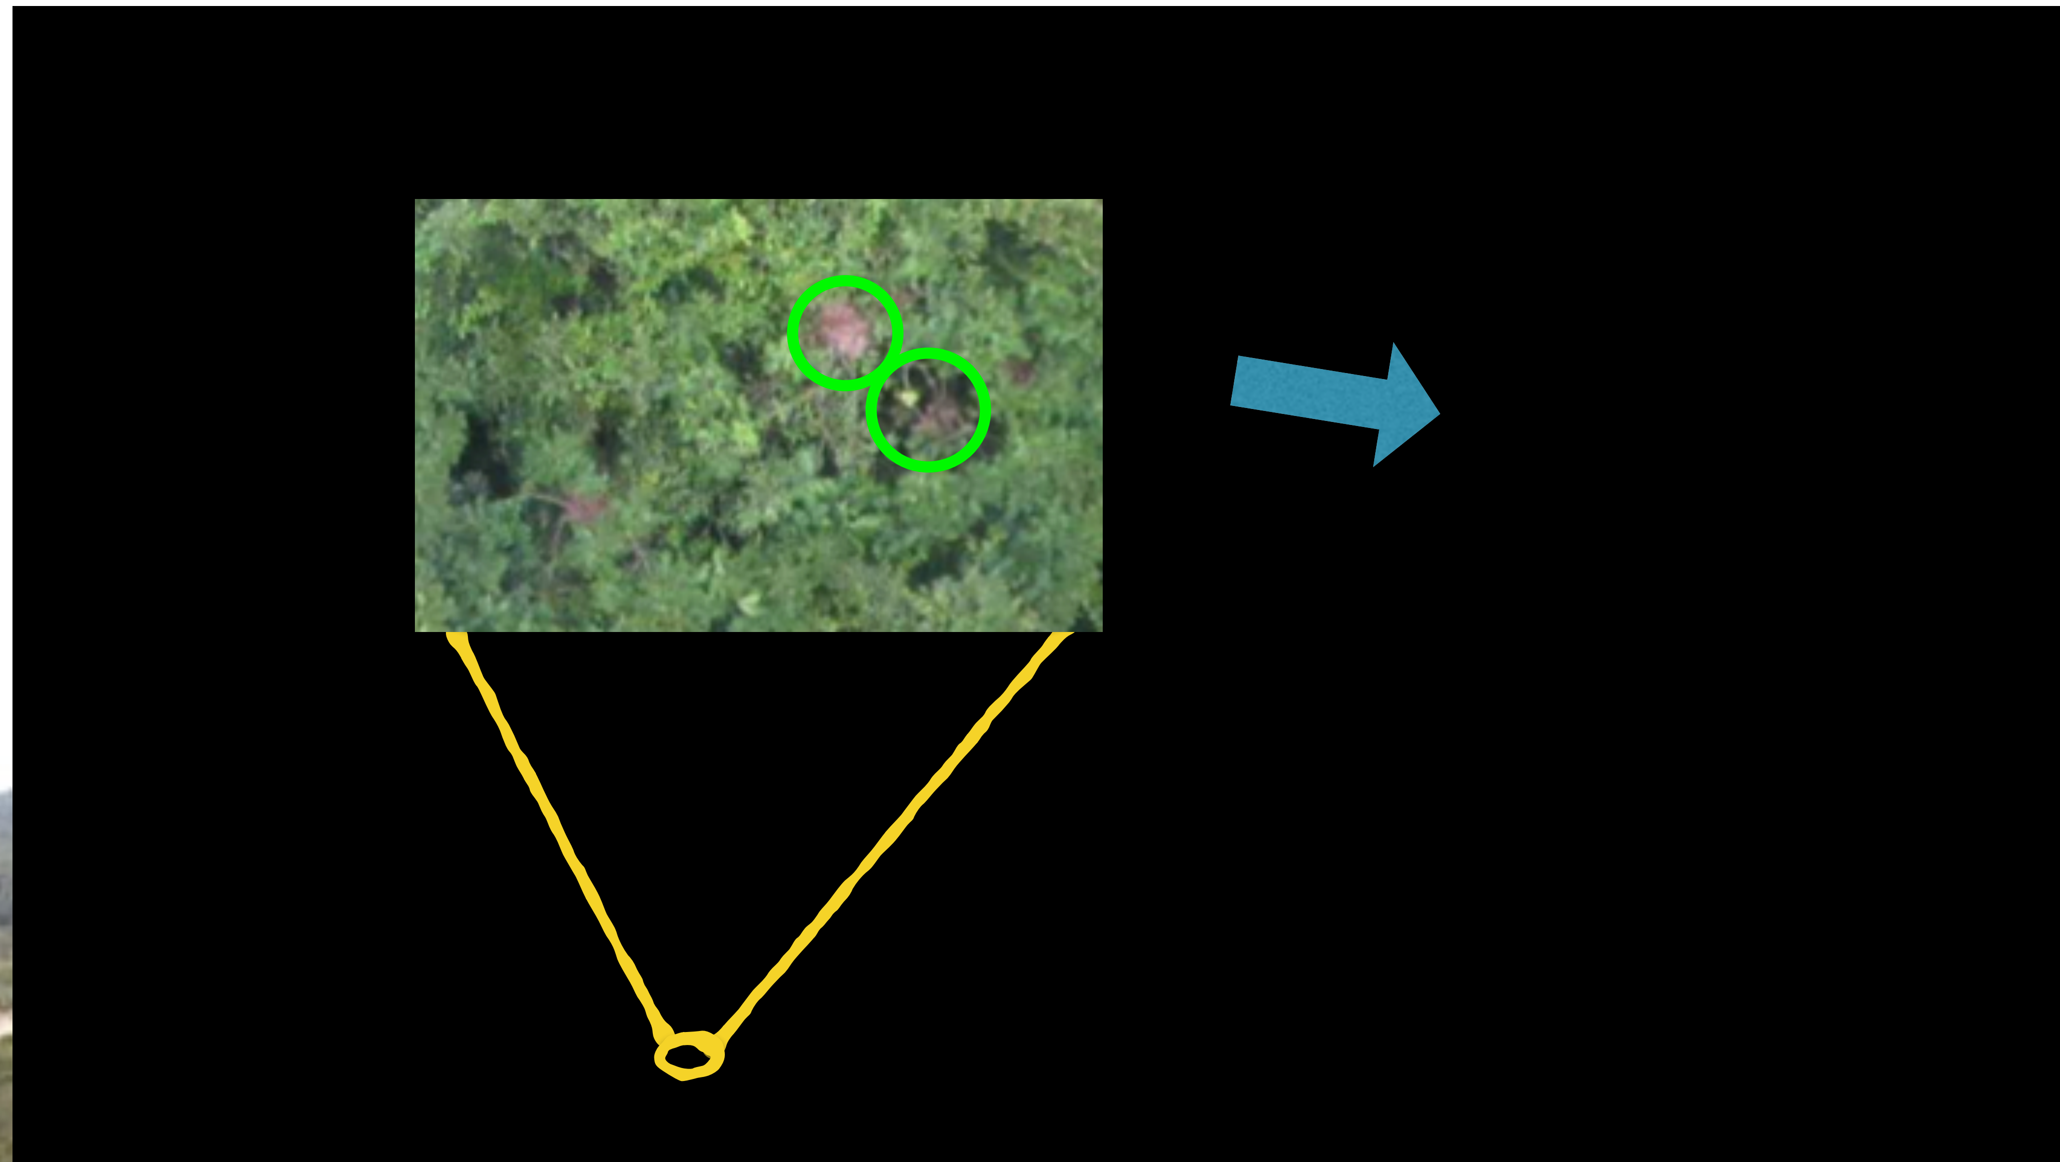

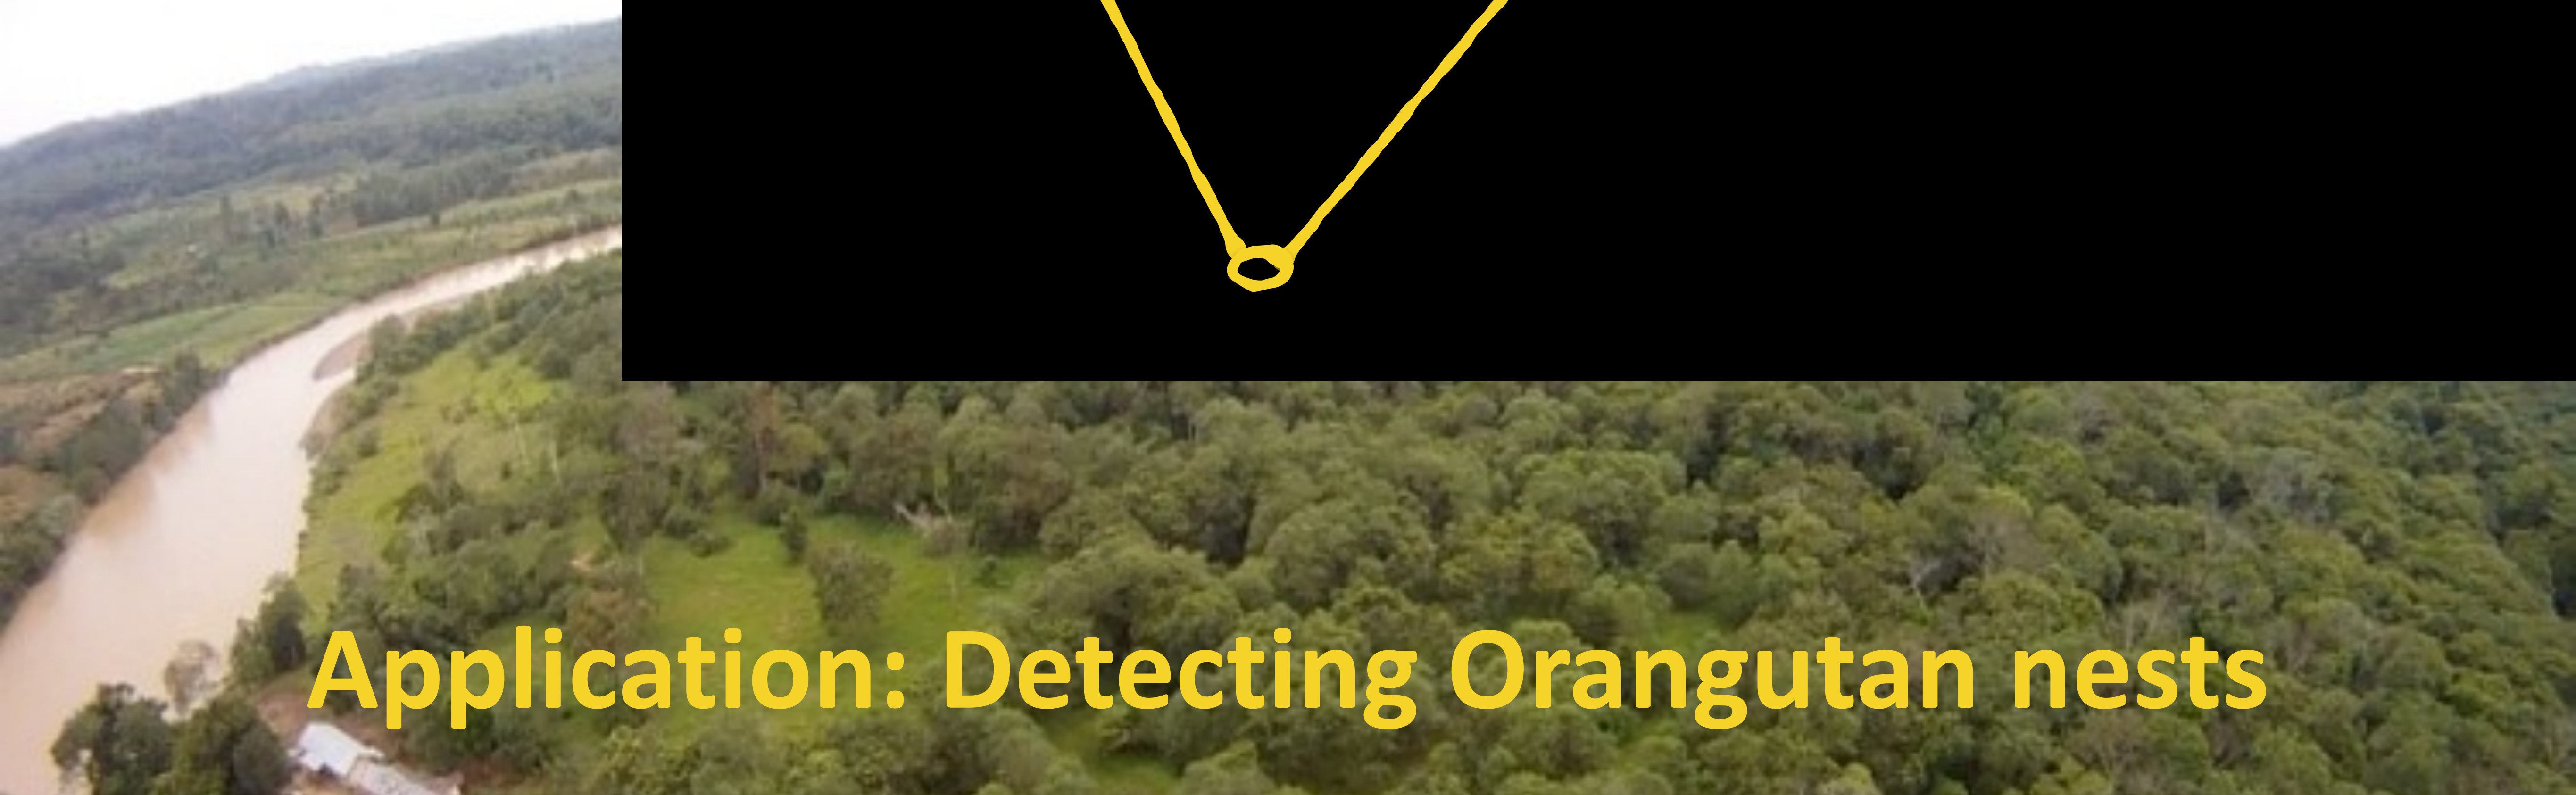

Application: Detecting Orangutan nests

Automatic, Open-loop Computer Vision System

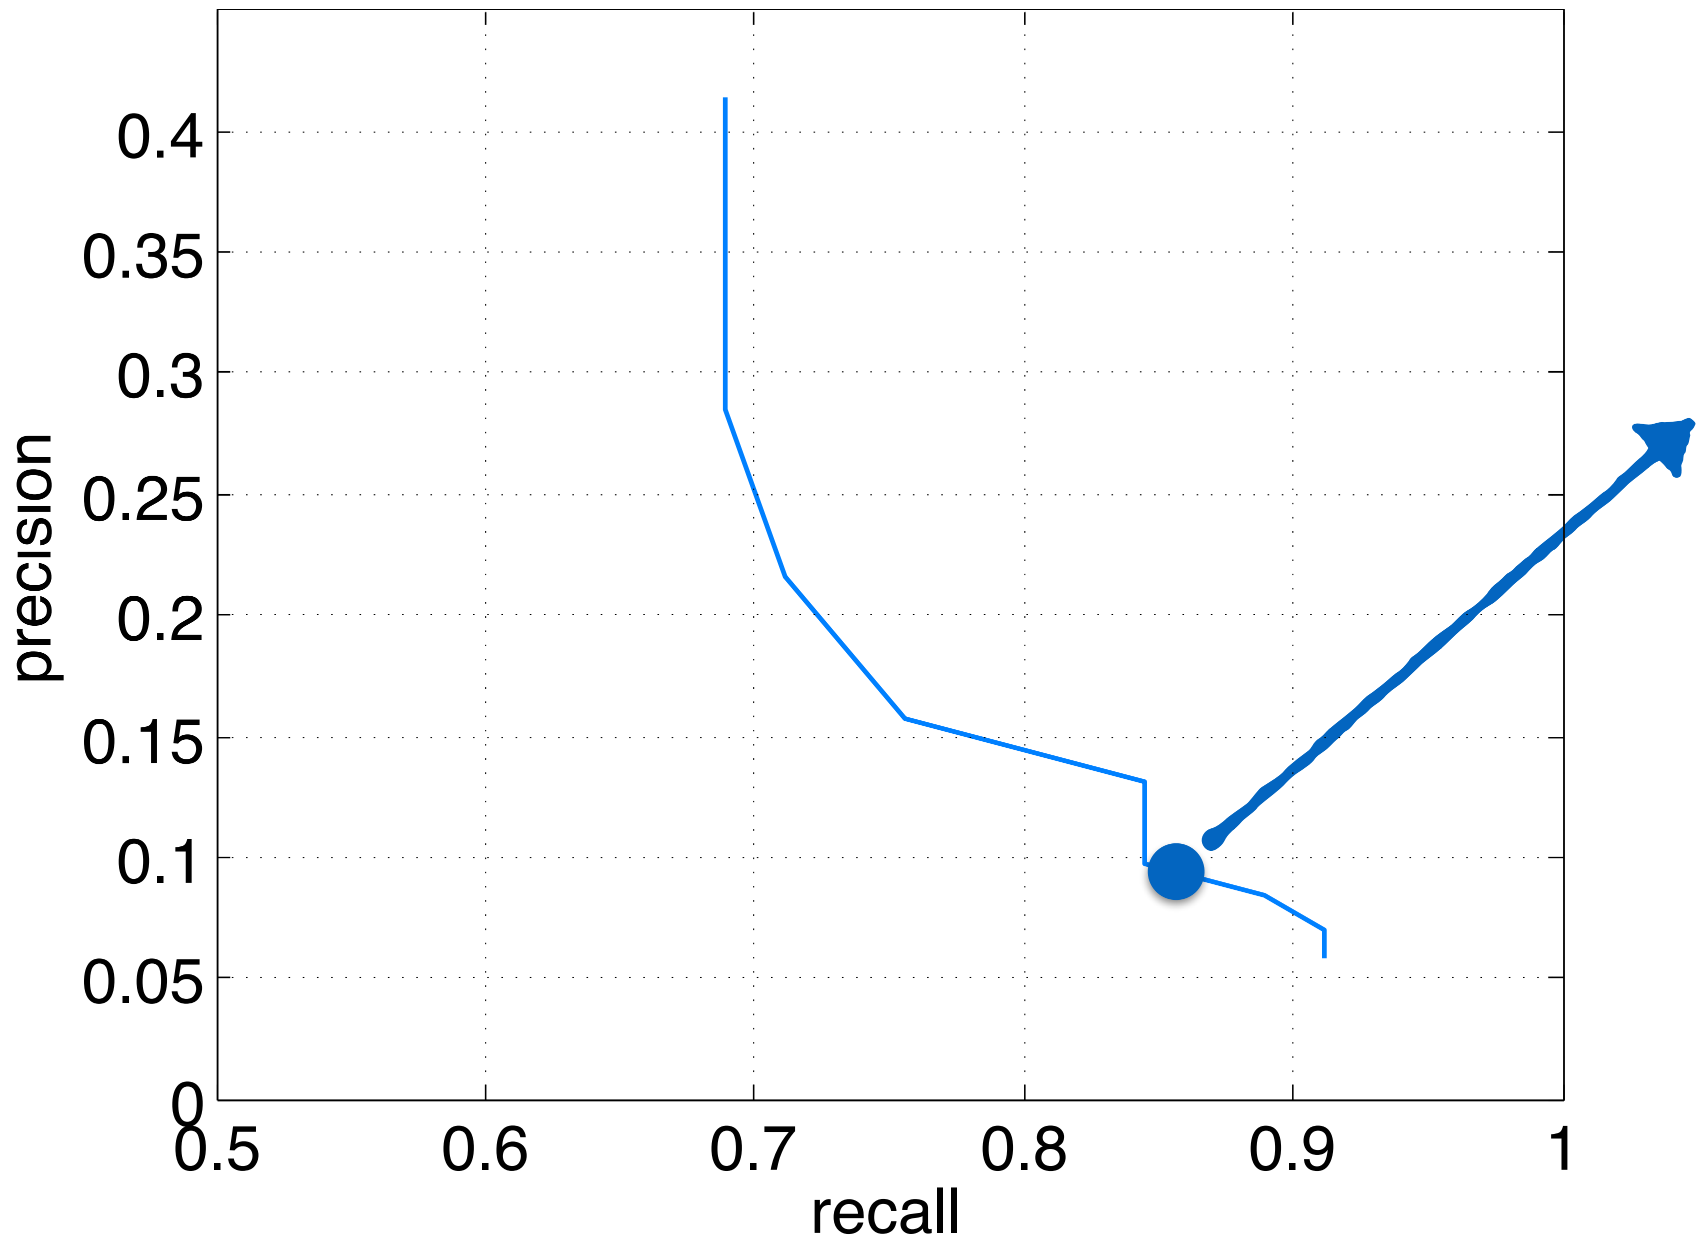

85% recall

10% precision ☹

Interactive Detection

How can human experts best help the detection task?

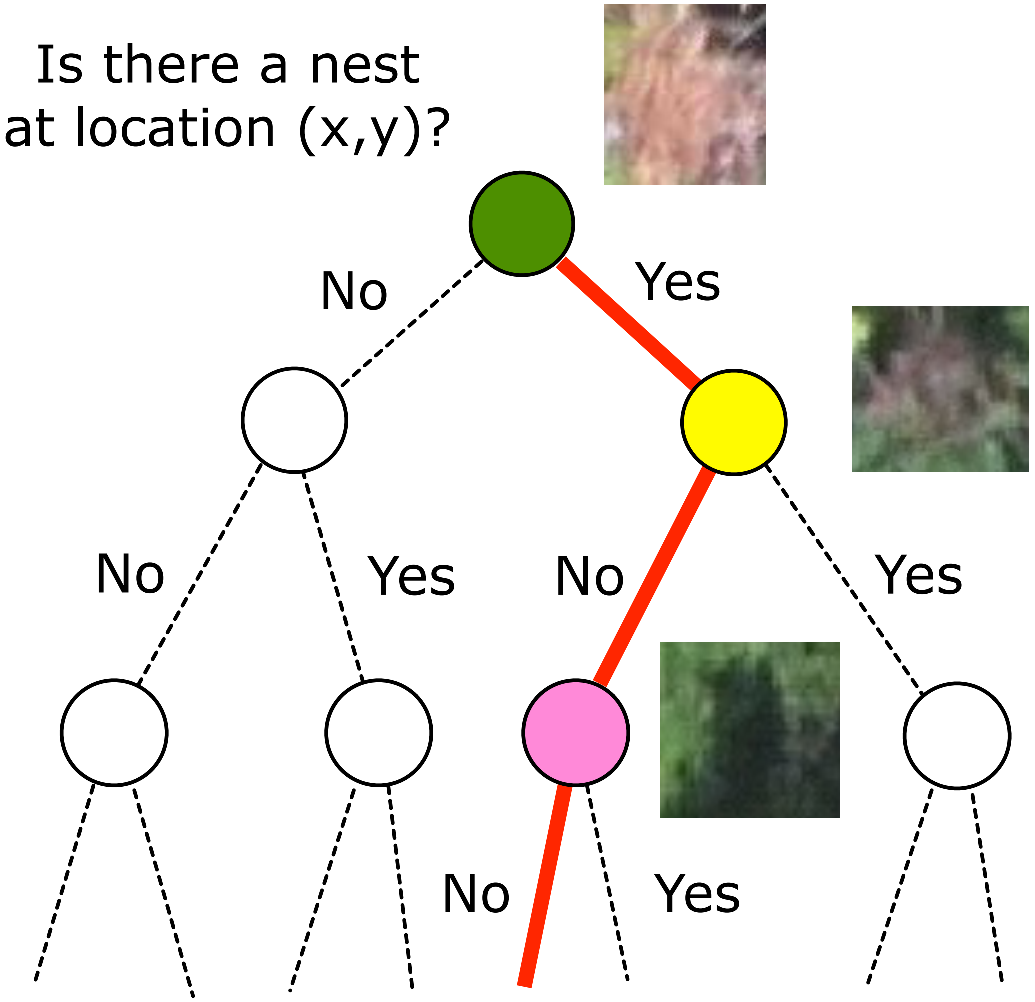

The adaptive policy

Open-loop (Passive) System

✗ ✔ ✗ ✗ ✗ ✔ ✗ ✔

Closed-loop (Active) System

✗ ✔ ✔ ✗ ✔ ✗ ✗ ✗

Evidence for Detection

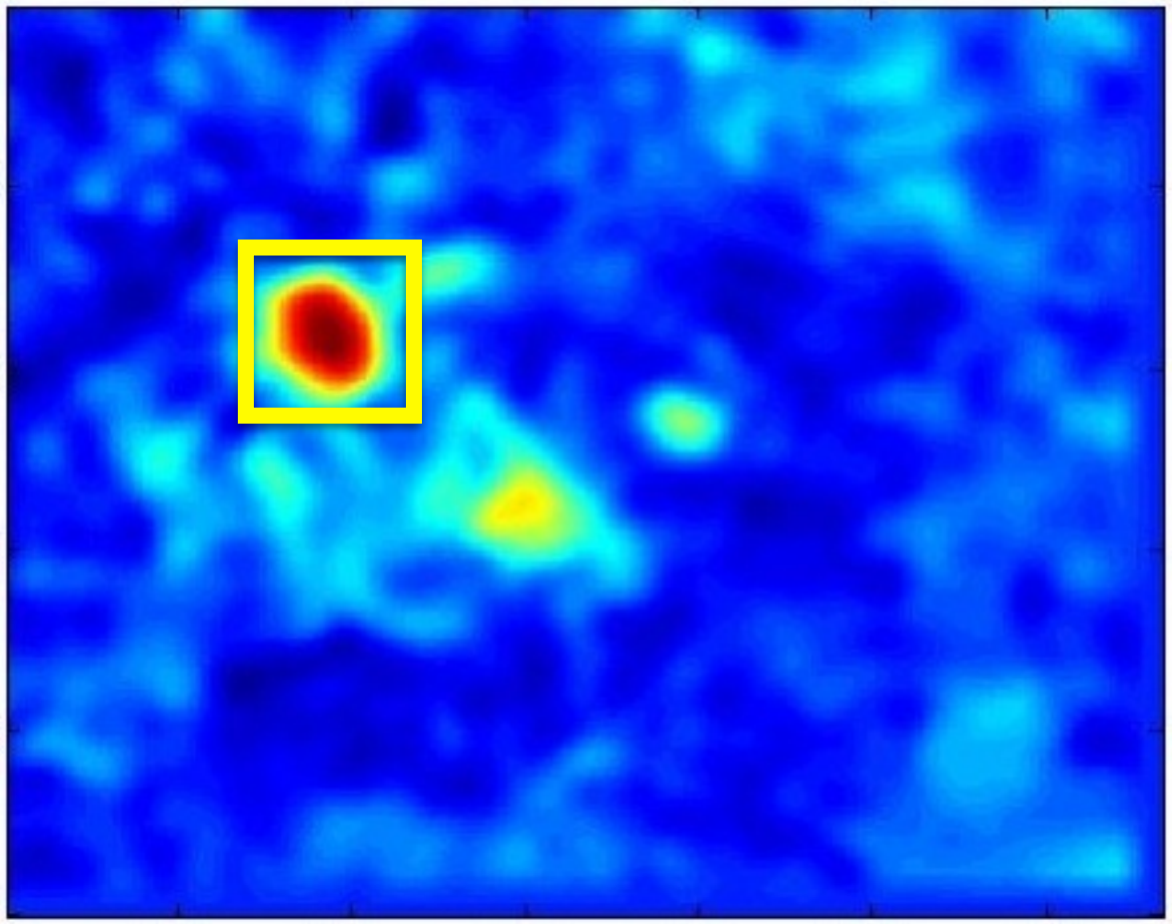

- Train classifier (e.g., SVM; conv. Neural Network, etc.) on **45 positive** and **148 negative** examples
- Use sliding window to produce "**response images**"
- Which detection should be proposed next?

Votes and Hypotheses

Hypotheses

$\mathcal{H} = \{h_1, \ldots, h_4\}$

Voting elements

$\mathcal{V} = \{v_1, \ldots, v_8\}$

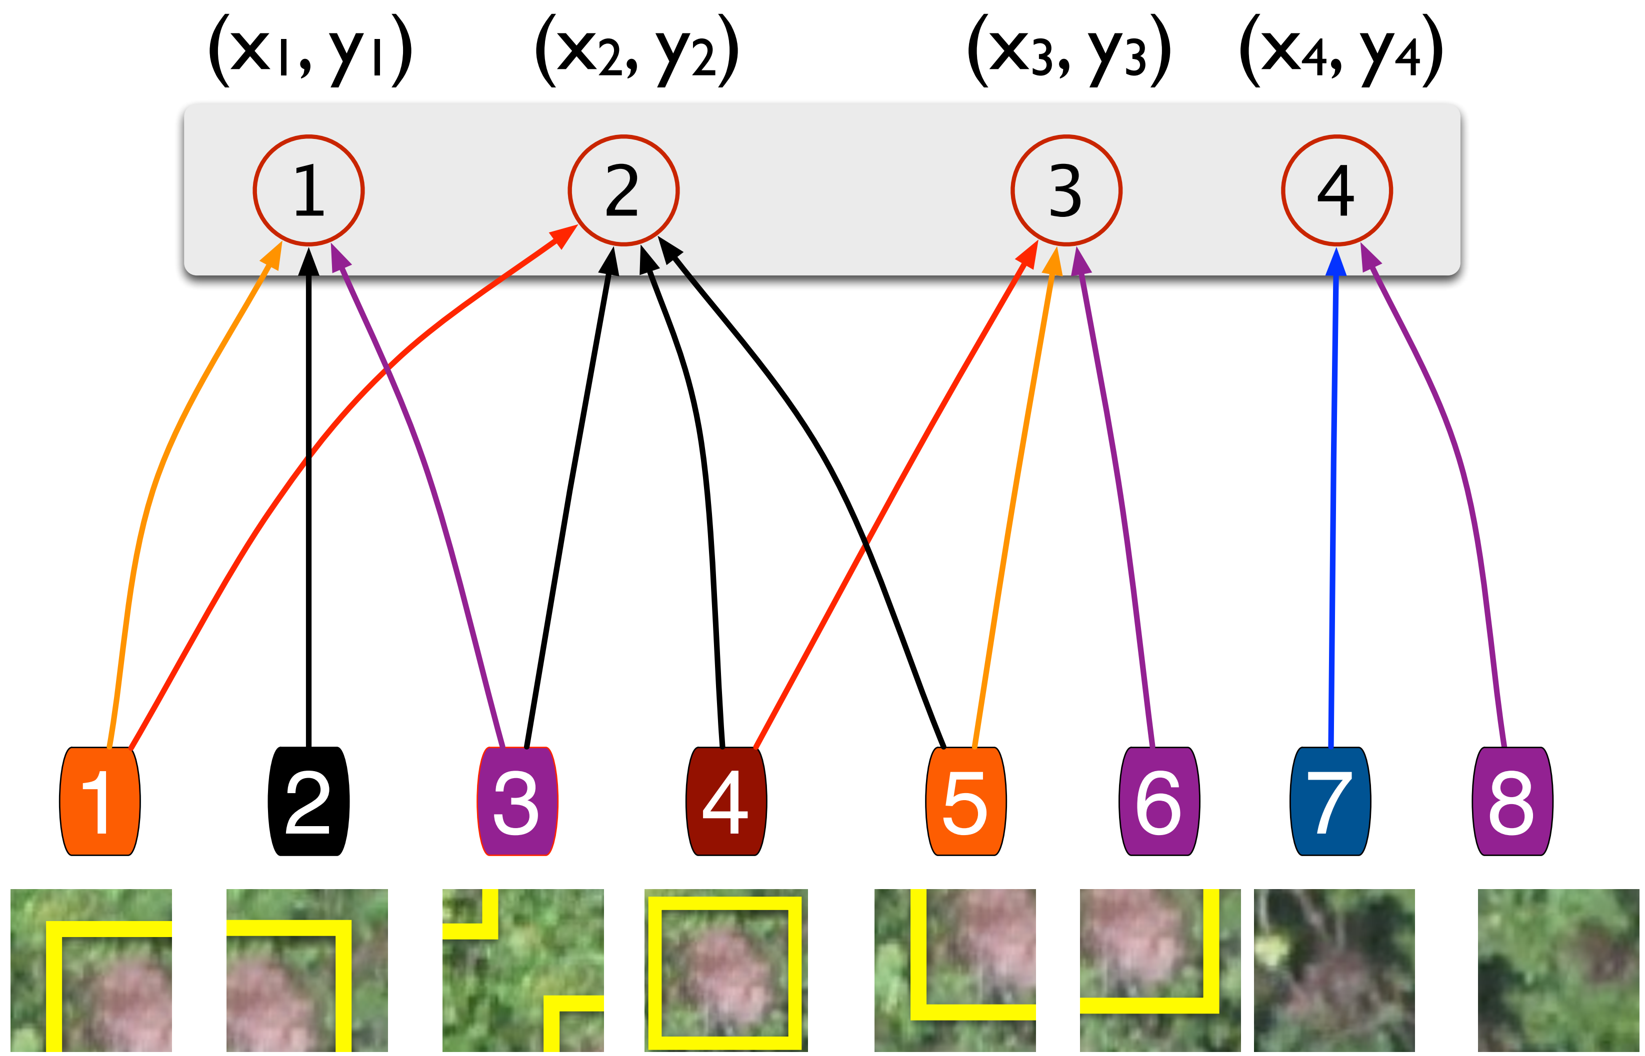

Interactions between voting elements and hypotheses: $\mathcal{G} = (\mathcal{V}, \mathcal{H}, \mathcal{E})$

[Hough '59; Gall et al, '09; Barinova et al, '11]

Active Detection as an Adaptive Optimization Problem

Positive coverage:

Votes can be fully explained /covered by a true hypotheses.

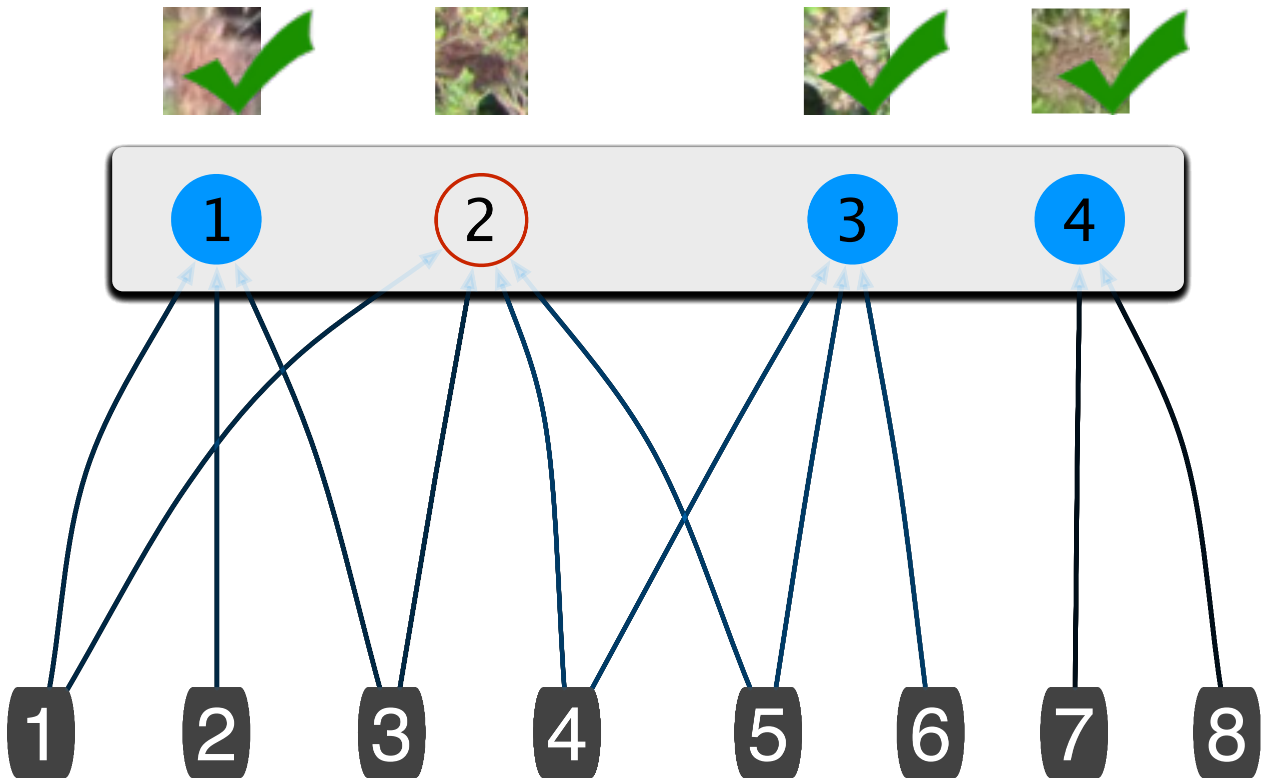

Now we observe that hypotheses 3 is true.

Then we observe that hypotheses I is true.

Then we observe that hypotheses 4 is true.

Assume that each vote carries unit weight

Active Detection as an Adaptive Optimization Problem

Negative coverage:

Votes that are similar with false votes should be discounted.

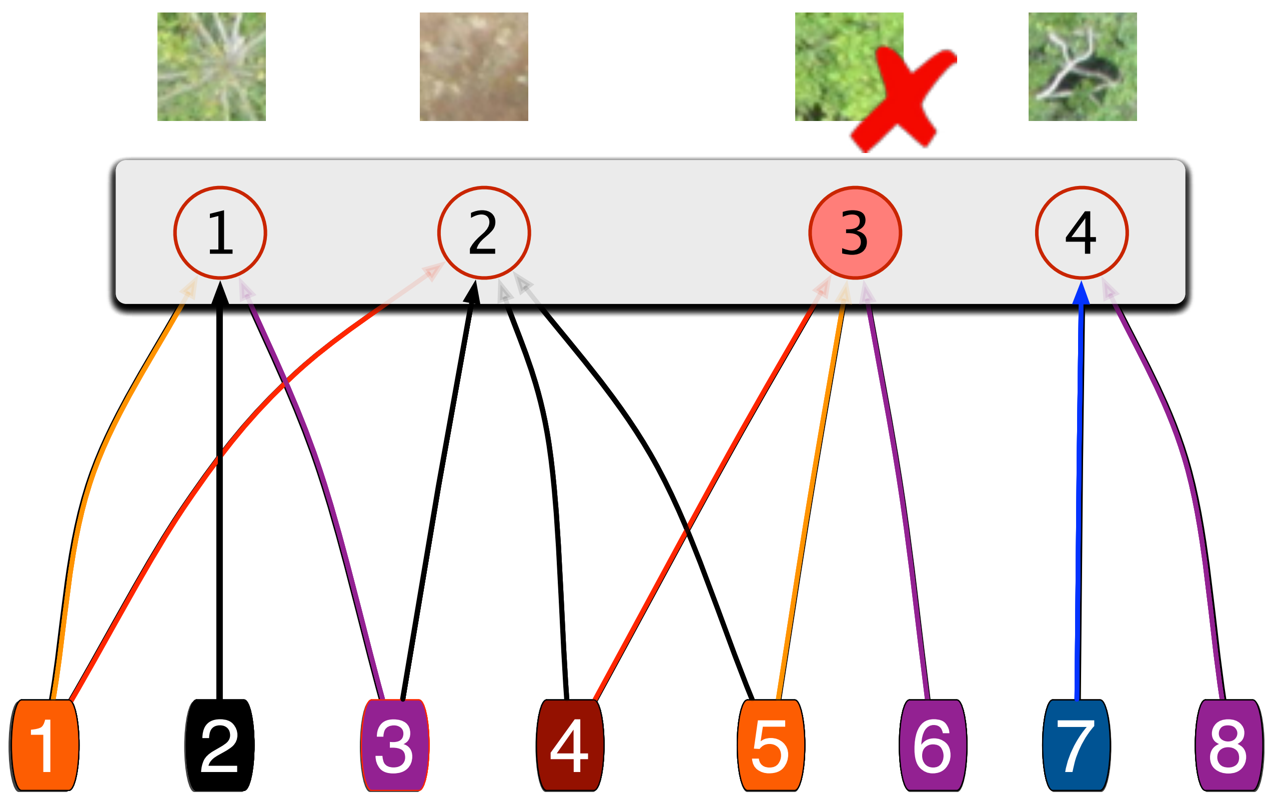

Suppose that we can cluster similar votes.

Now we observe that hypotheses 3 is false.

We also need to discount the votes that are similar with the false votes!!

Assume that each vote carries unit weight

The general case: Real-votes setting

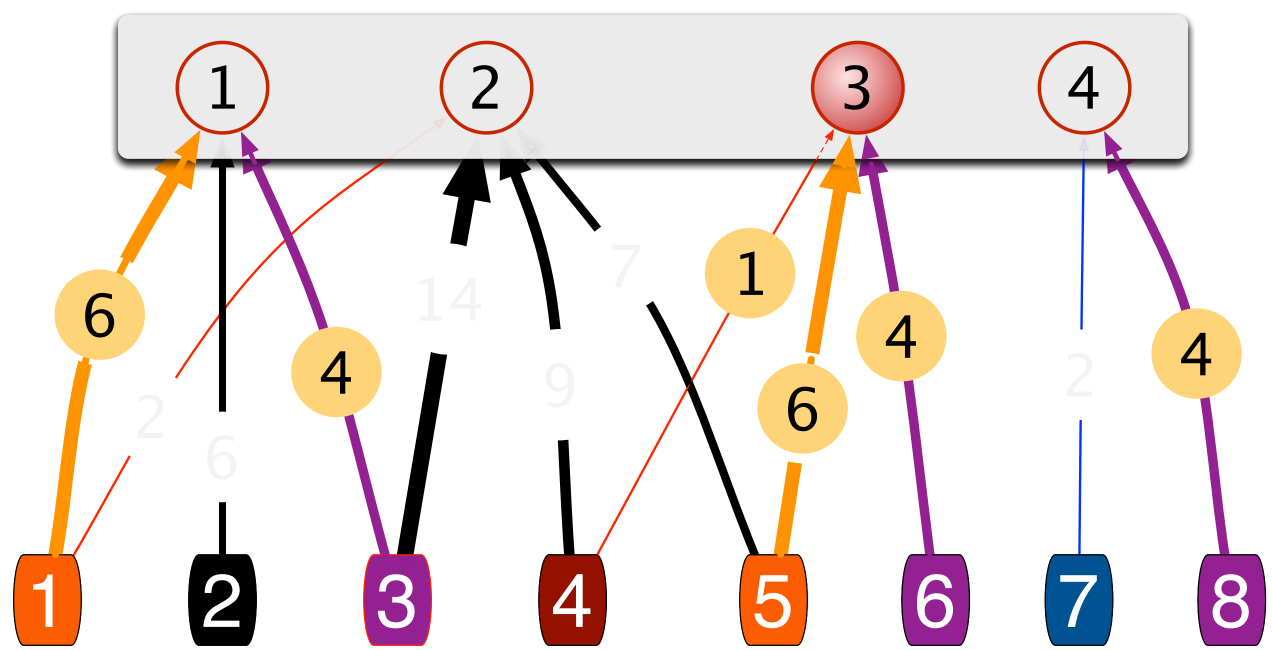

Now observe that hypothesis 3 is false

Voting elements with real-value votes

Active Detection in a Nutshell

Positive observations:

Votes can be fully explained / covered by a true hypotheses.

Negative observations:

Votes that are similar with false votes should be discounted.

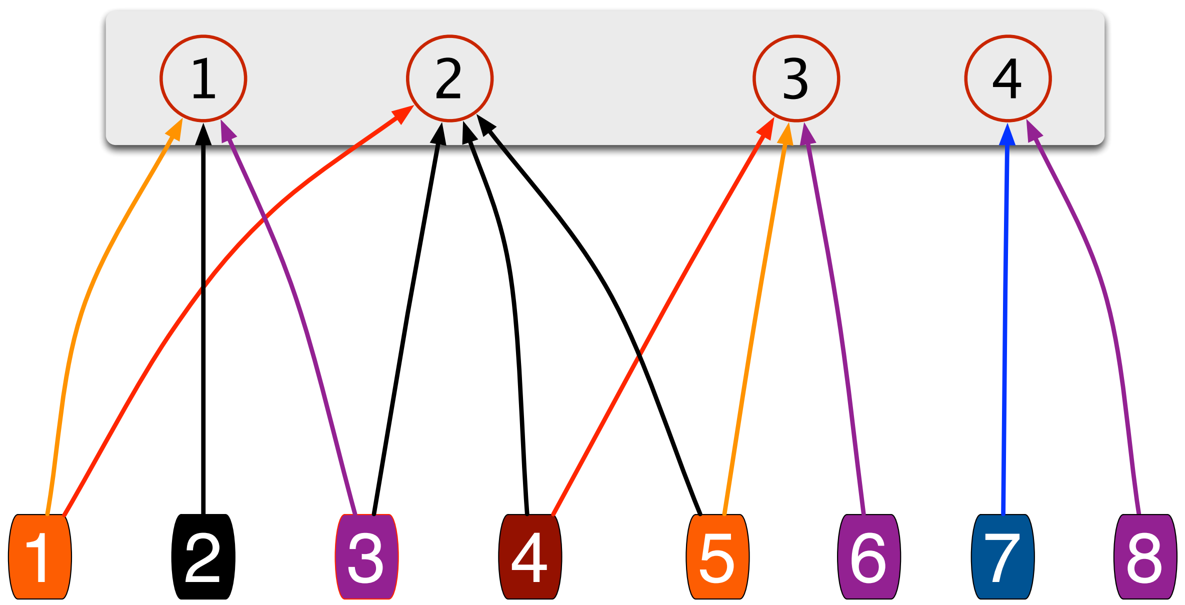

The Objective

Coverage for edge (v,h) = Coverage due to positive observations + Coverage due to negative observations

$$\text{Coverage of } \mathcal{G} = (\mathcal{V}, \mathcal{H}, \mathcal{E}) = \sum_{(v,h)} \text{Coverage for edge (v,h)}$$

Diminishing Evidence in Detection

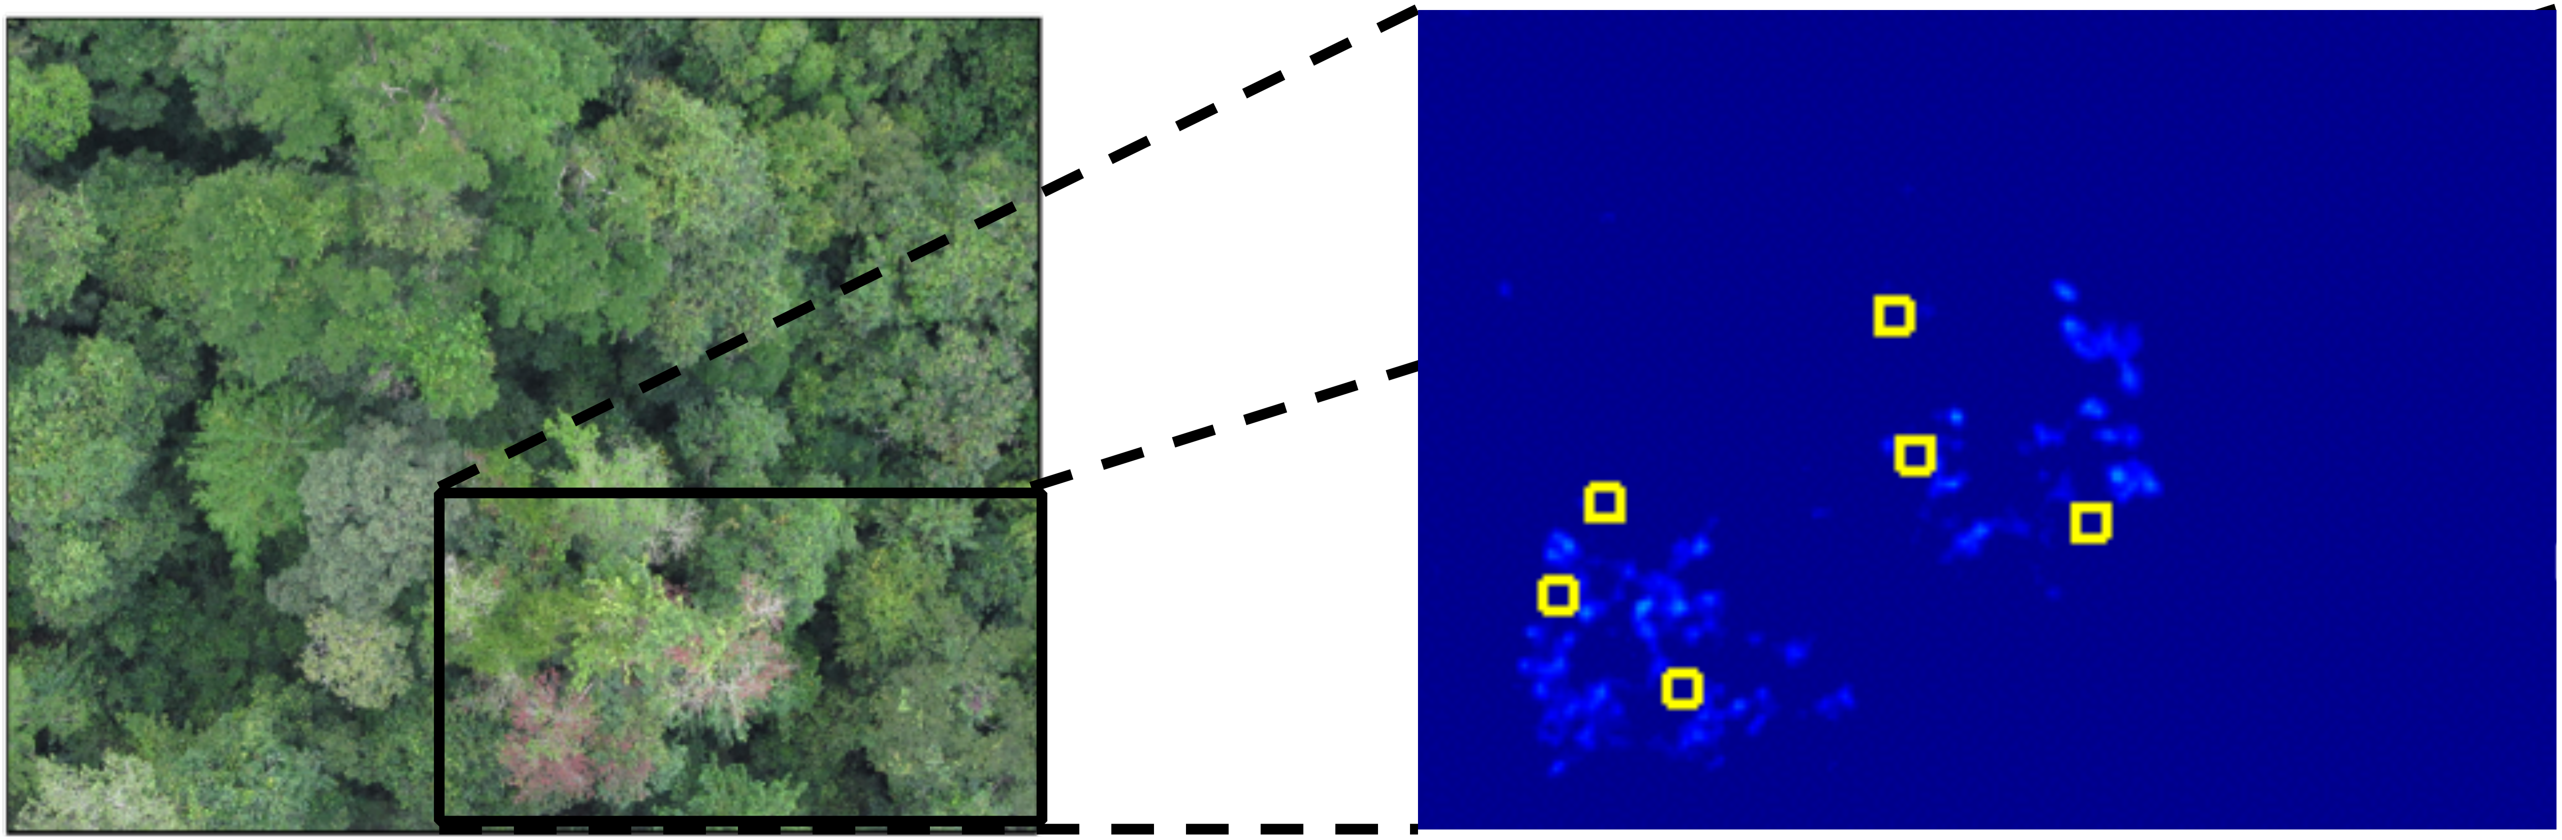

- Positive observations explain "response" in local areas
- Negative observations explain "response" in similar areas

Adaptive submodular objective can capture this diminishing returns effect

Adaptive Submodularity [Golovin & Krause, 2011]

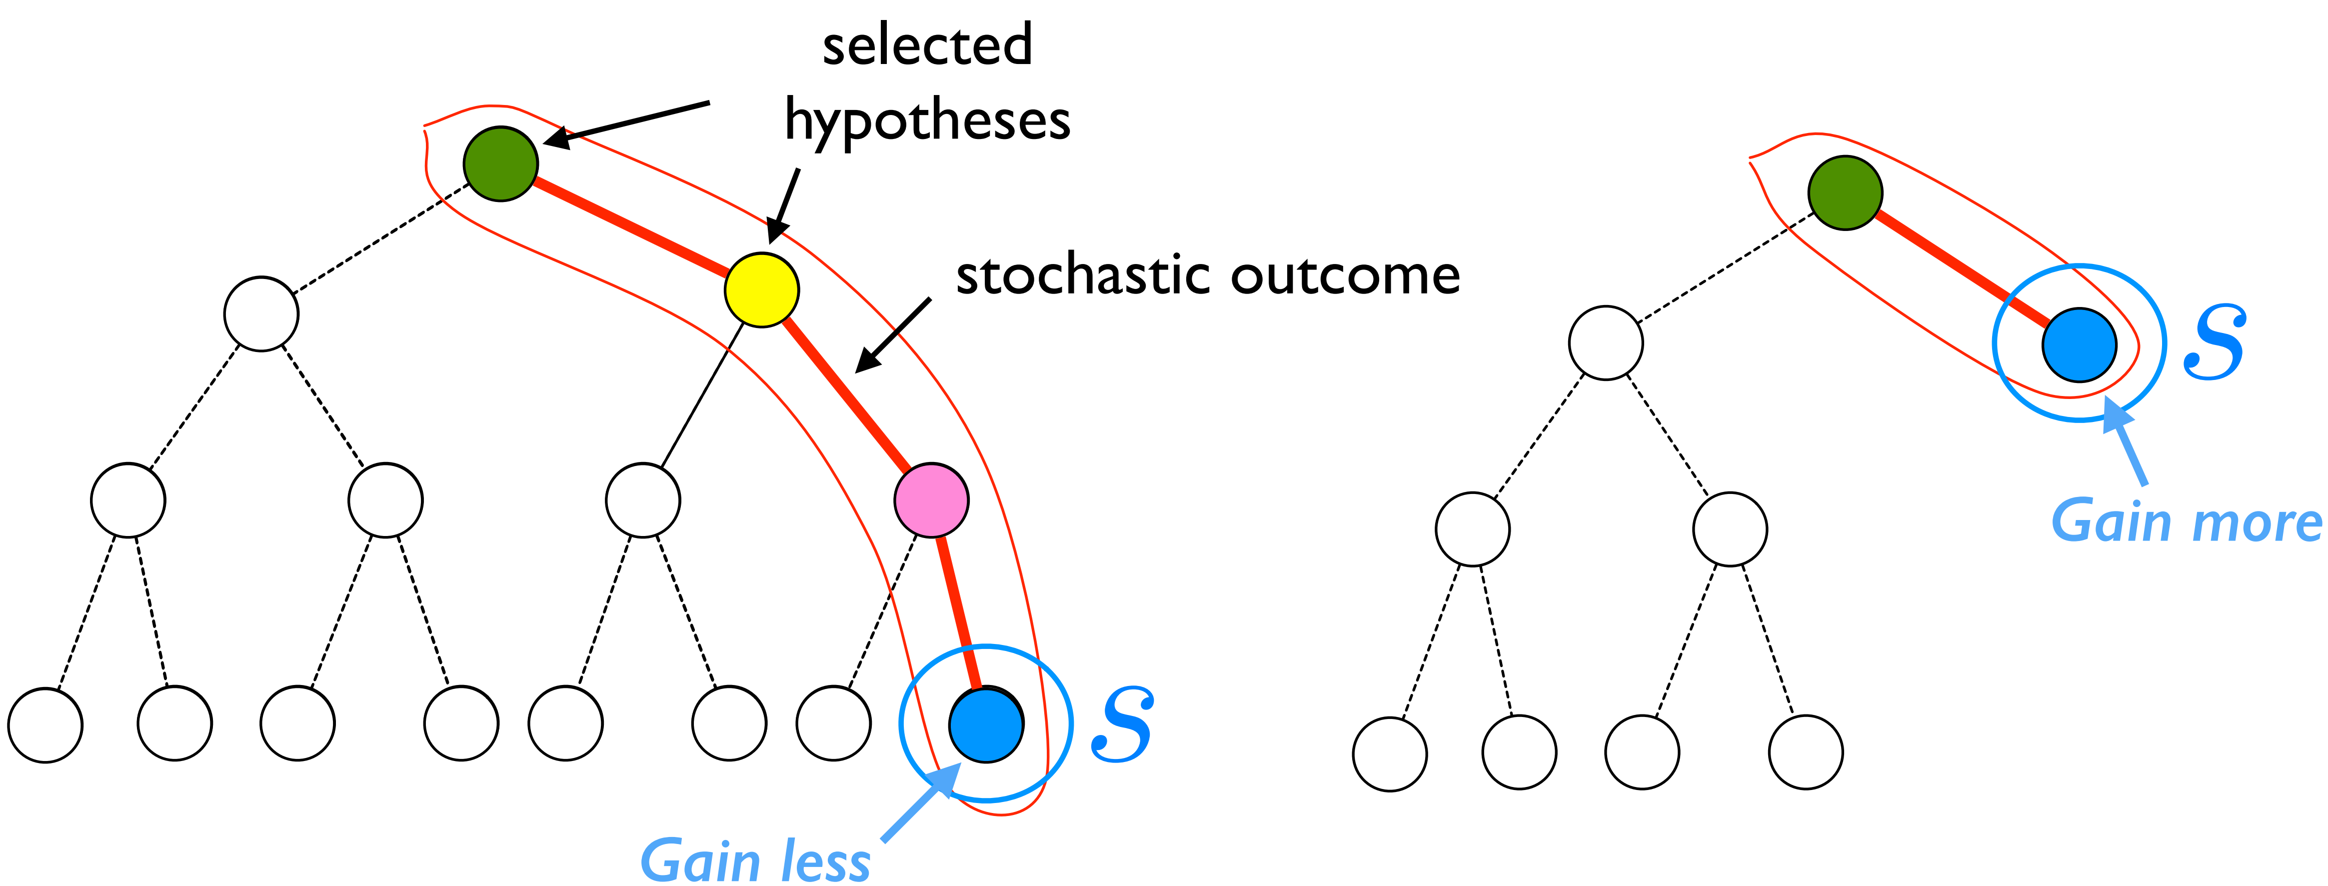

Receiving observation earlier (i.e., at an ancestor) only increases its ***expected*** marginal benefit.

taken over its outcome

Greedy vs. Optimal

Assume that:

- The optimal policy achieves a maximum coverage of Q
- The greedy policy achieves a maximum coverage of Q-β

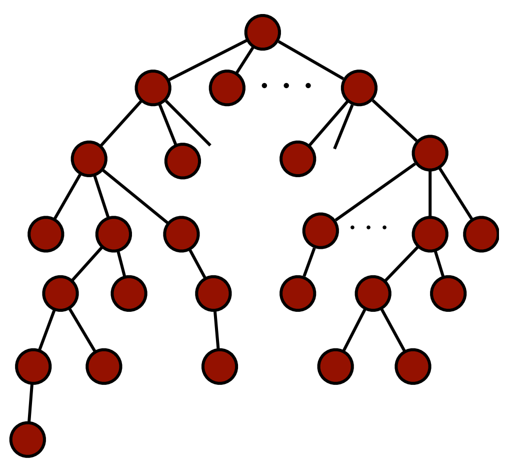

$$\left| \; \cdot \; \right| \leq \left(\ln \frac{Q}{\beta} + 1\right) \cdot \left| \; \cdot \; \right|$$

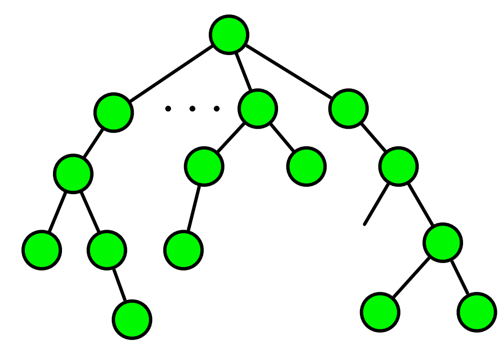

Cost of the **Greedy algorithm** w.r.t. F

Cost of **optimal** policy

Detection Results

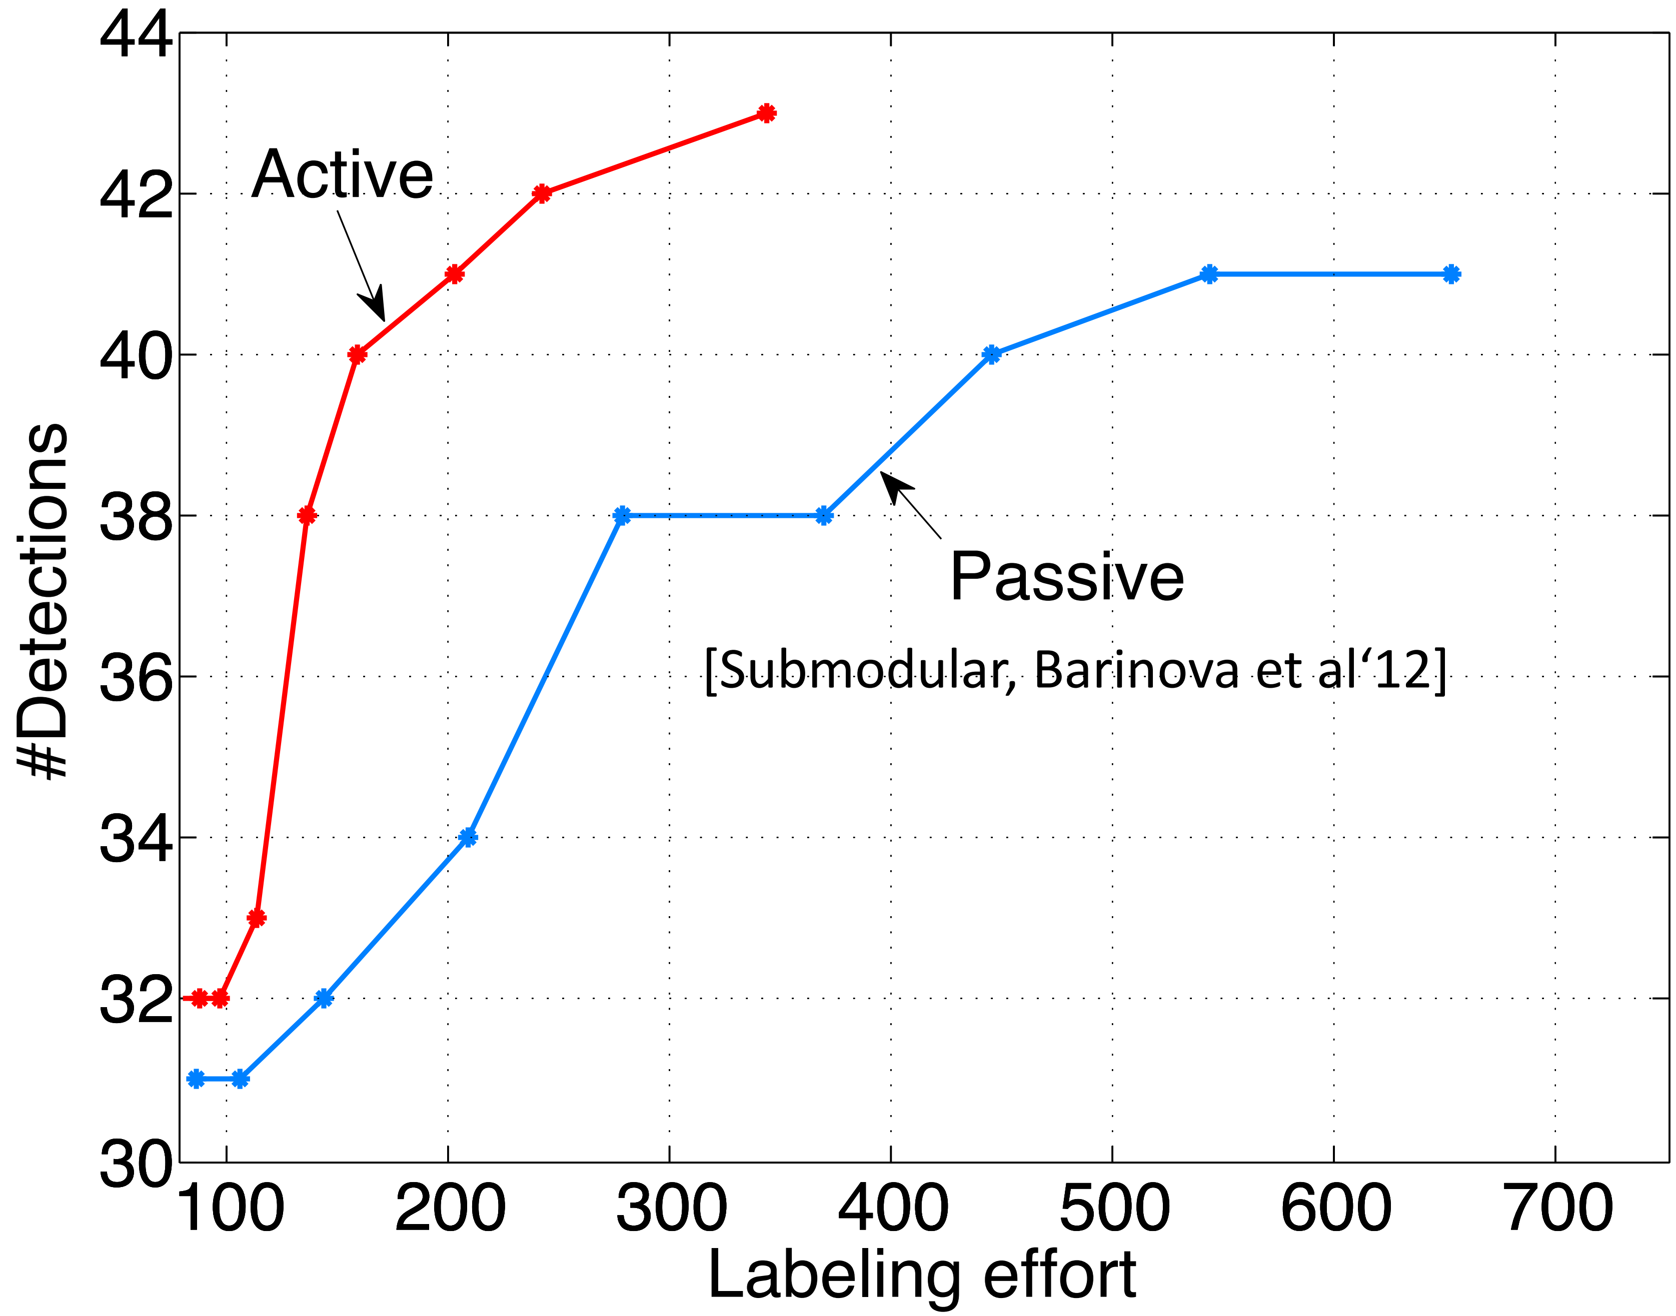

Active detection improves precision and recall

TUD-pedestrian: Pedestrian Detection

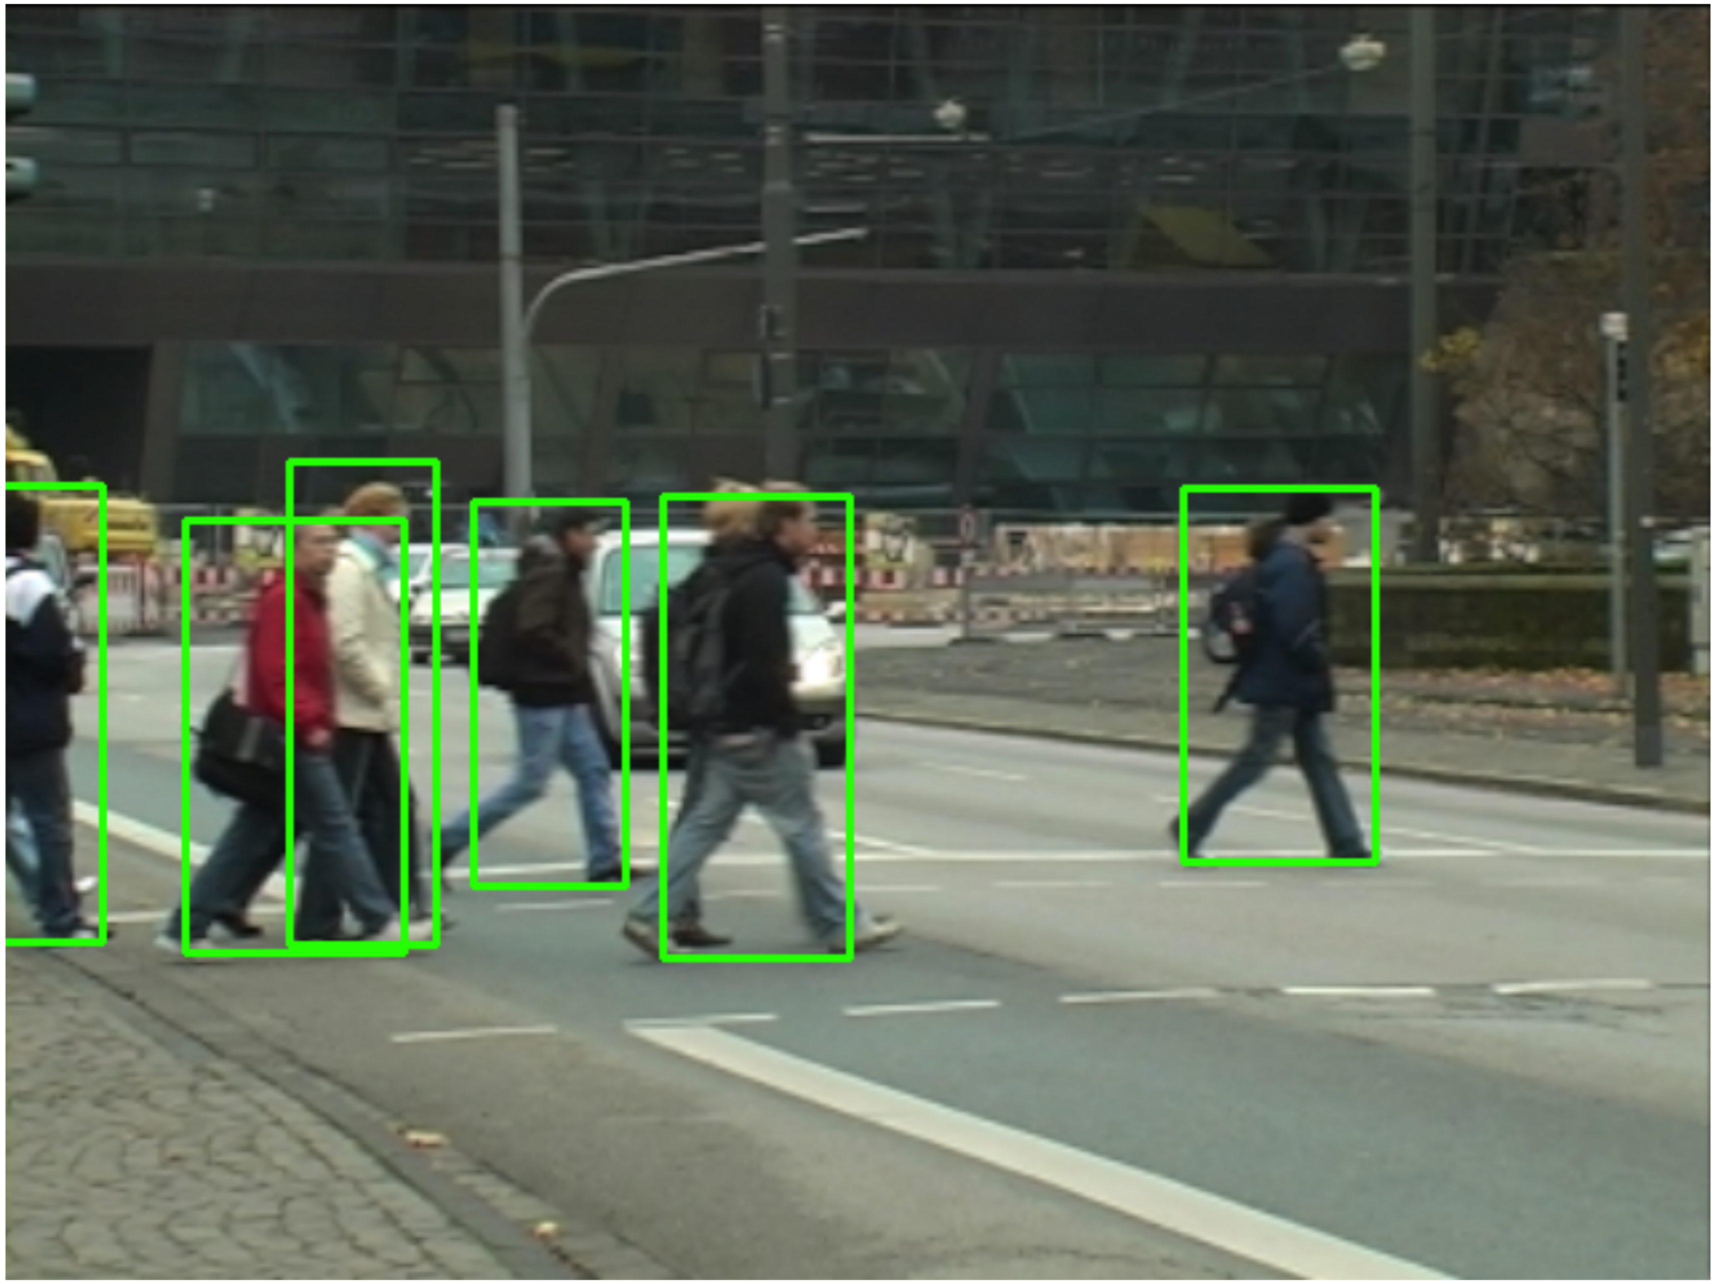

Votes and Hypotheses
Hough-forest Based Detector

Response Image
Original Image
[Hough-forest, Gall et al, CVPR'09]

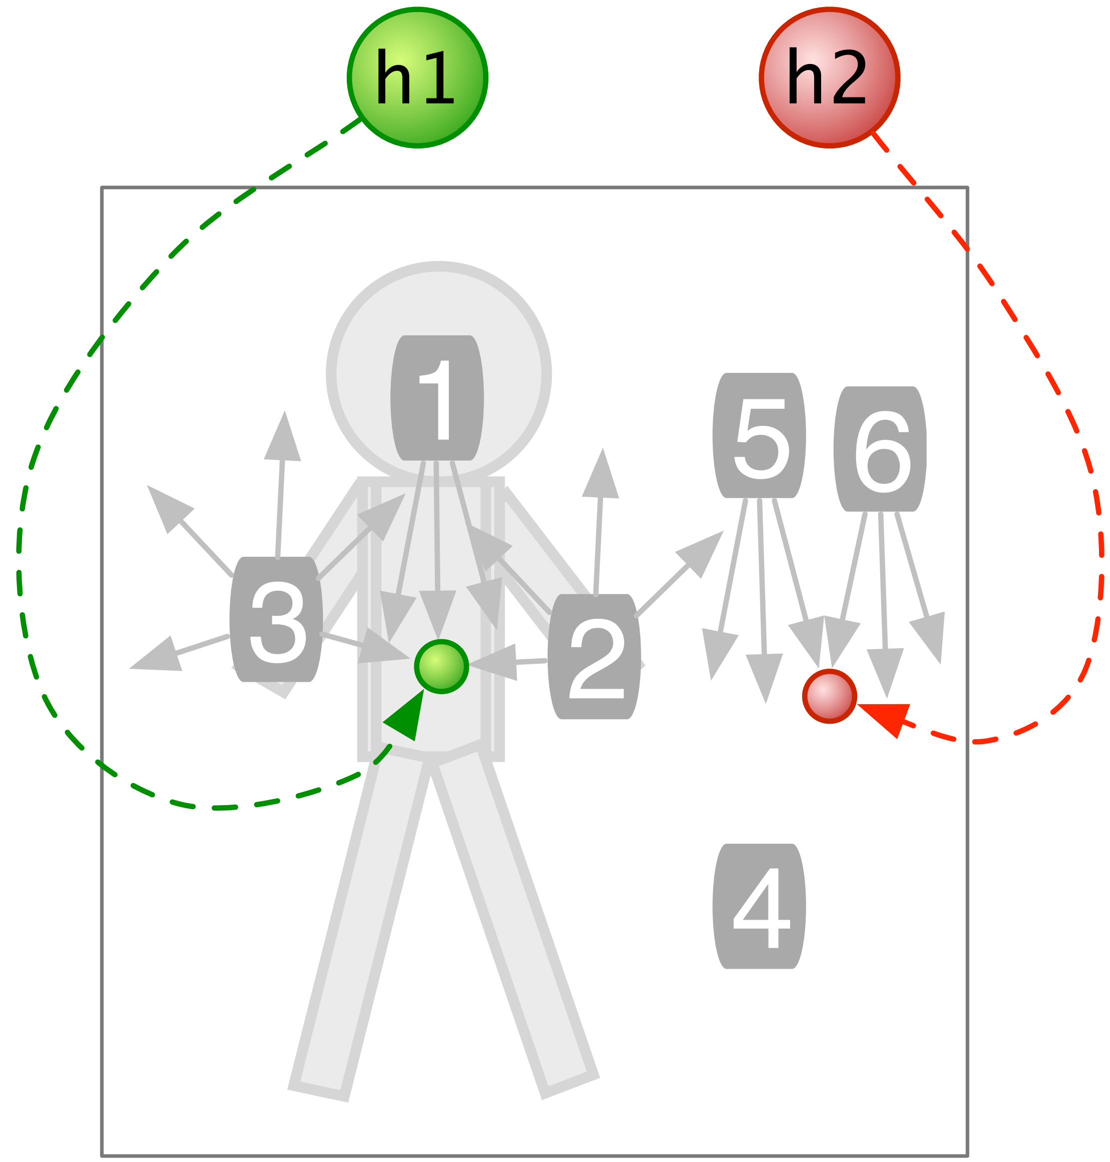

TUD-pedestrian: Detection Results

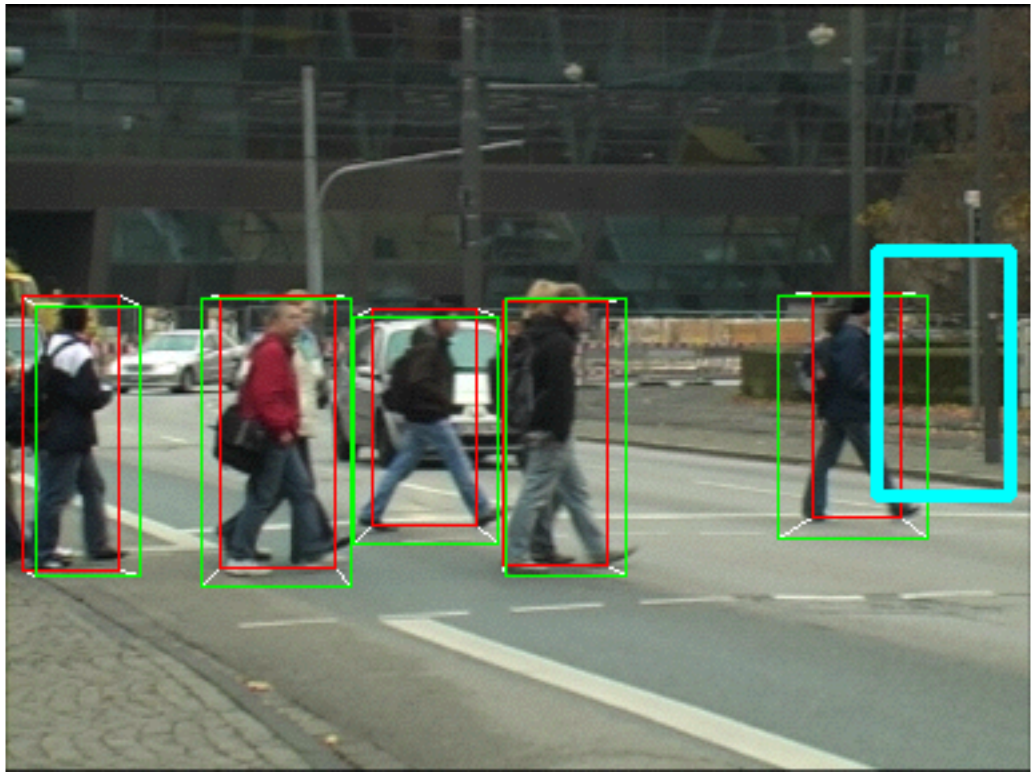

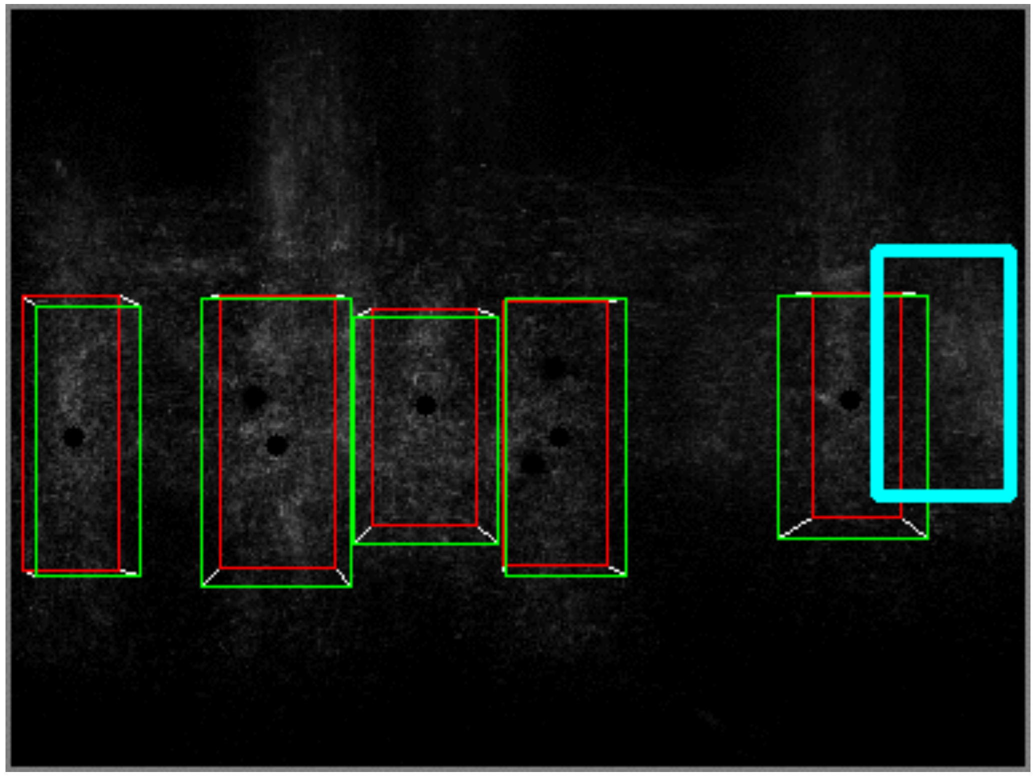

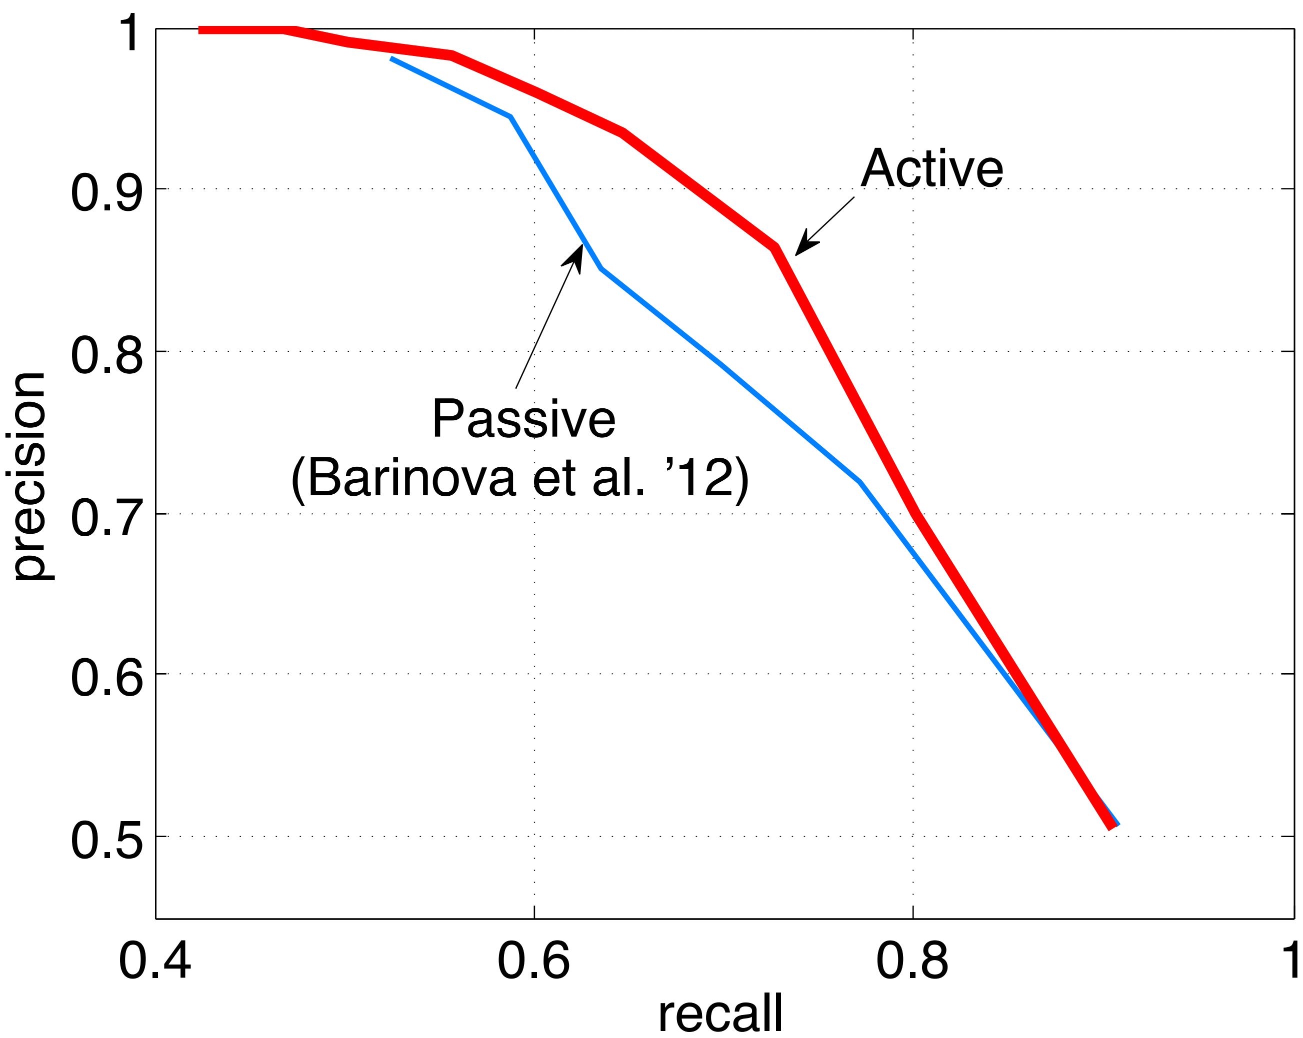

Cyan box: current detection.
Red boxes: ground-truth labels of pedestrians.
Green boxes: detections made by the active detector.

PASCAL 2008 - Person Category

Deformable Parts Model (DPM)

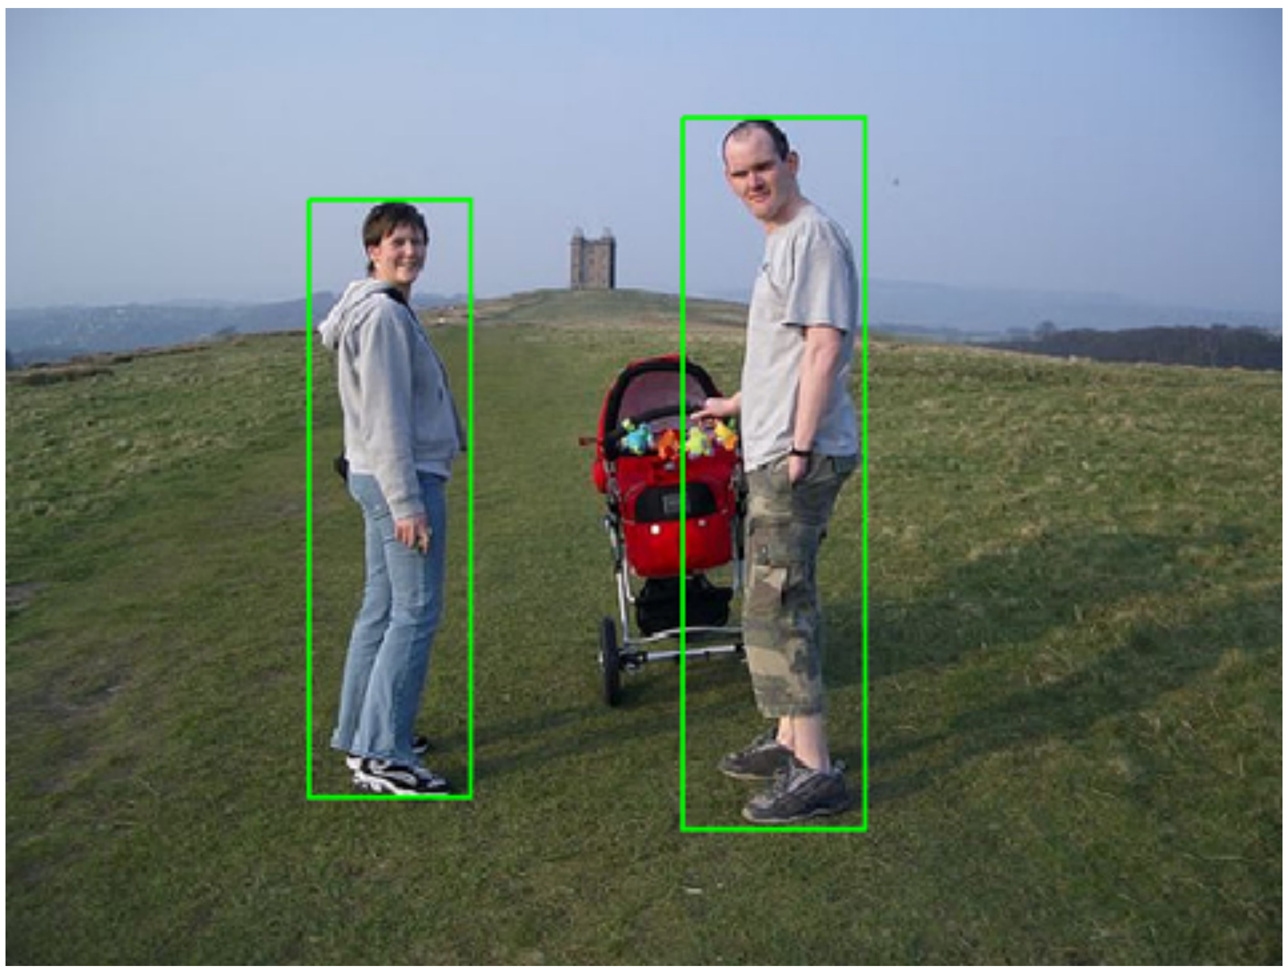

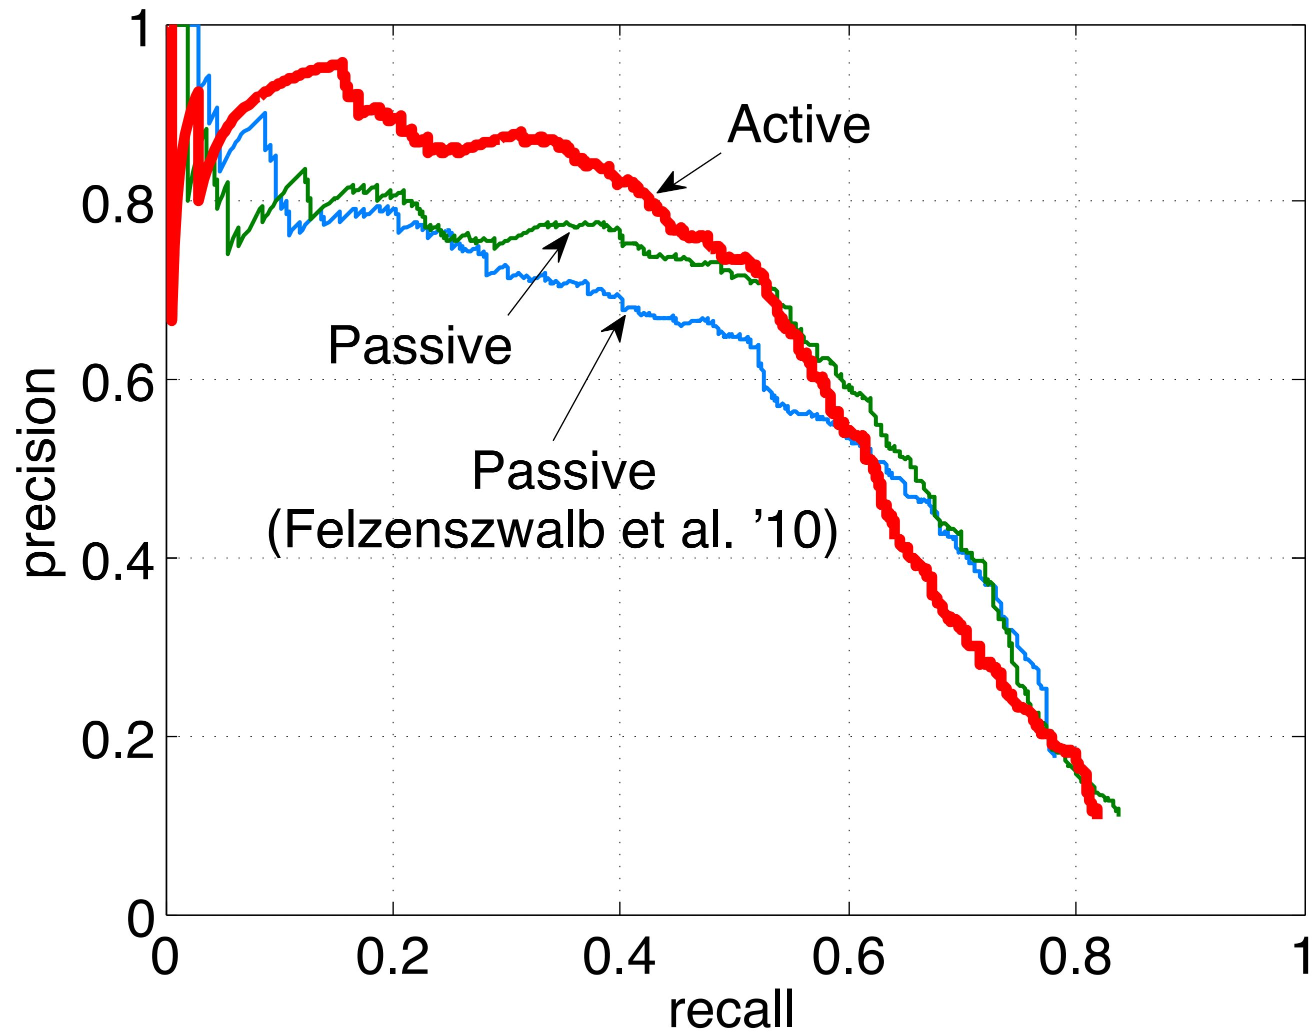

Conclusion

- An active detection framework that enables turning existing base detectors into systems that intelligently interact with users.
- We show that the objective function satisfies adaptive submodularity, allowing us to use efficient greedy algorithms, with strong theoretical guarantees.
- We demonstrate the effectiveness of the active detection algorithm on three different real-world object detection tasks.

Come to our poster on Tuesday for more details !