

Special Session:
Low-Power Image Recognition Competition
(LPIRC)
www.lpirc.net

Yung-Hsiang Lu, Purdue University

Alex Berg, University of North Carolina



Summary

Two Papers in this Session

1. Overview of LPIRC. Results from 2015 and 2016.
2. 2016 Winner

LPIRC 2017 will be held on July 21 in Hawaii (before CVPR). Travel grants will be available.



IEEE Rebooting Computing

- To **rethink the computer “from soup to nuts”** including all aspects from device to user interface
- Started in 2012 – TAB-funded through 2016
- NIC funding starts 2017
- Participation from 10 S/C + SA



IEEE COMPONENTS, PACKAGING AND
MANUFACTURING TECHNOLOGY SOCIETY



IEEE Council on Electronic Design Automation



Nanotechnology Council®



IEEE CIRCUITS AND SYSTEMS SOCIETY



Rebooting Computing Summit



Rebooting Computing Summit



Rebooting Computing Summit



Low-Power Image Recognition Challenge (LPIRC)

www.lpirc.net



Kent Gauen, Rohit Rangan, Anup Mohan, Yung-Hsiang Lu
Purdue University

Wei Liu, Alexander C. Berg
University of North Carolina

Why LPIRC?

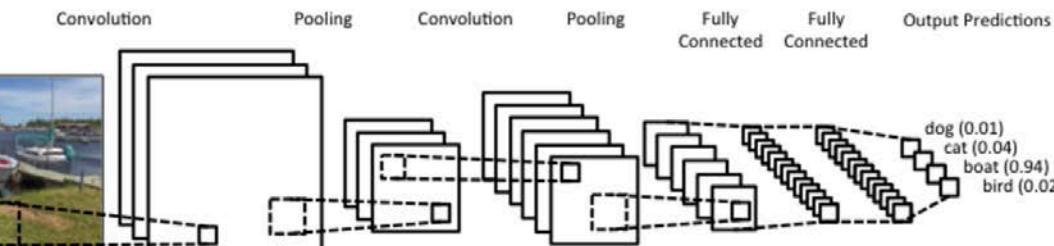
Recent Success in Machine Learning



clarifai



AlphaGo



kaggle™ **Microsoft**



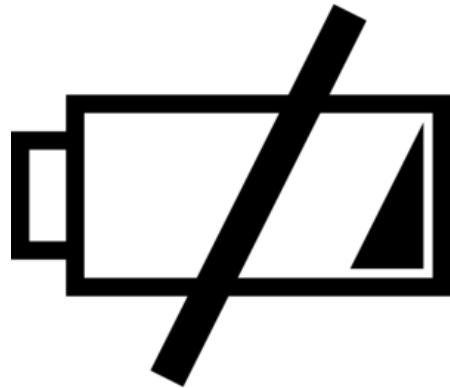
Why LPIRC?

Low-Power World



Many datasets exist to benchmark image processing.

There is no known low-power image processing benchmark.

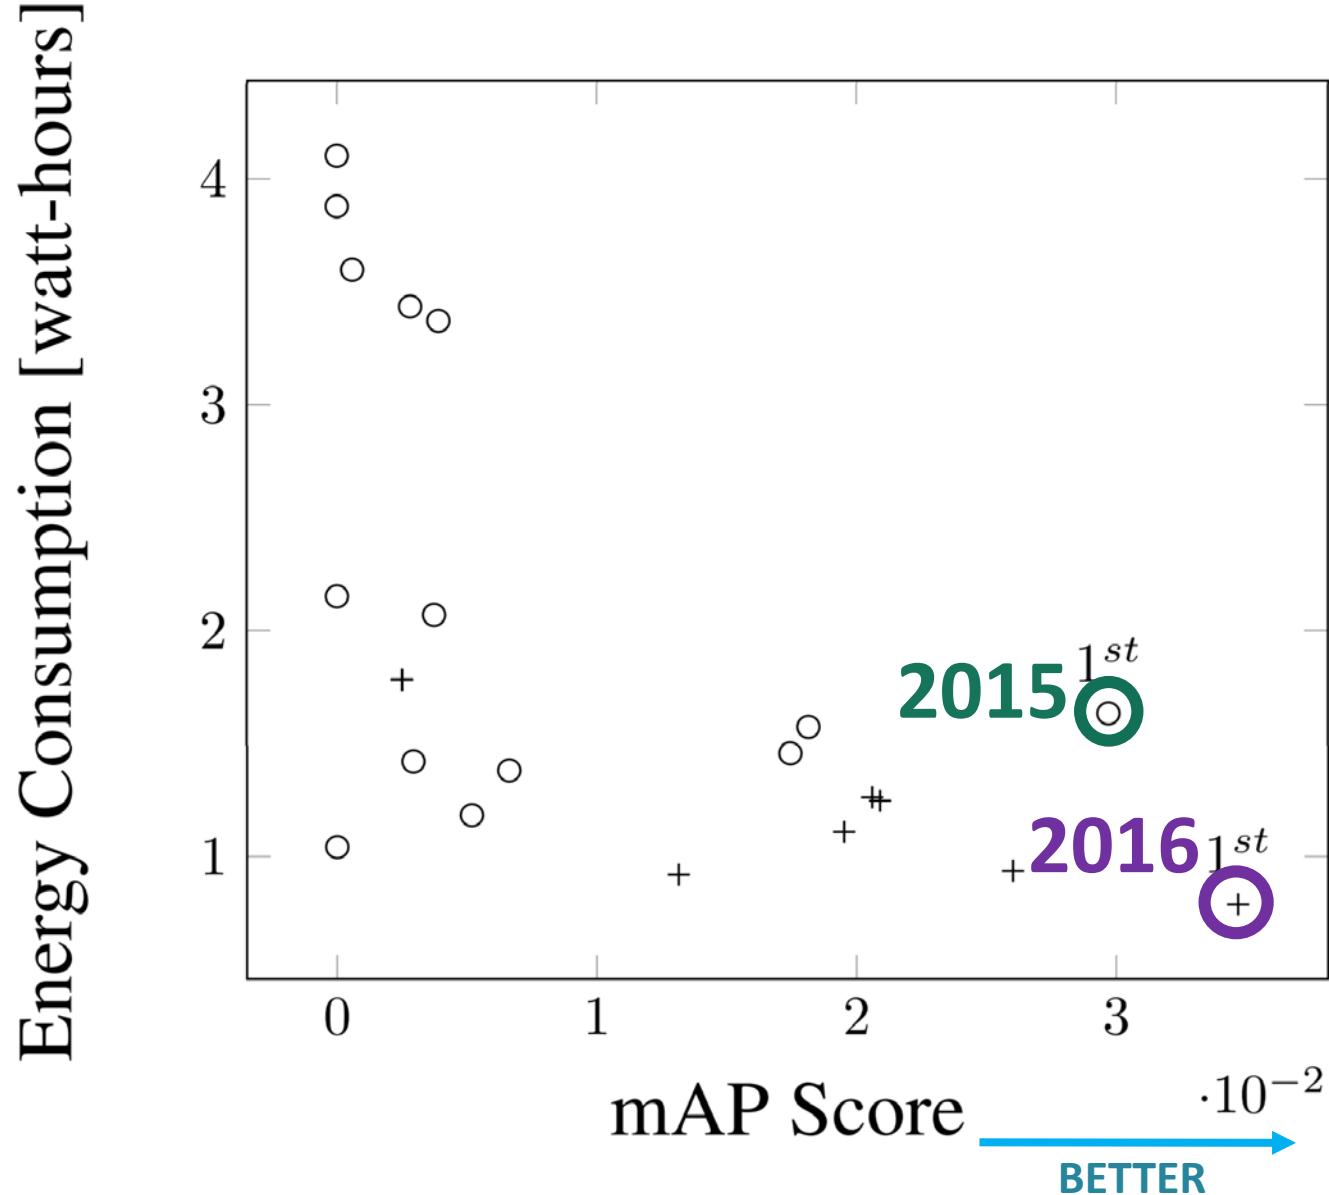


Attracts Attention



Why LPIRC?

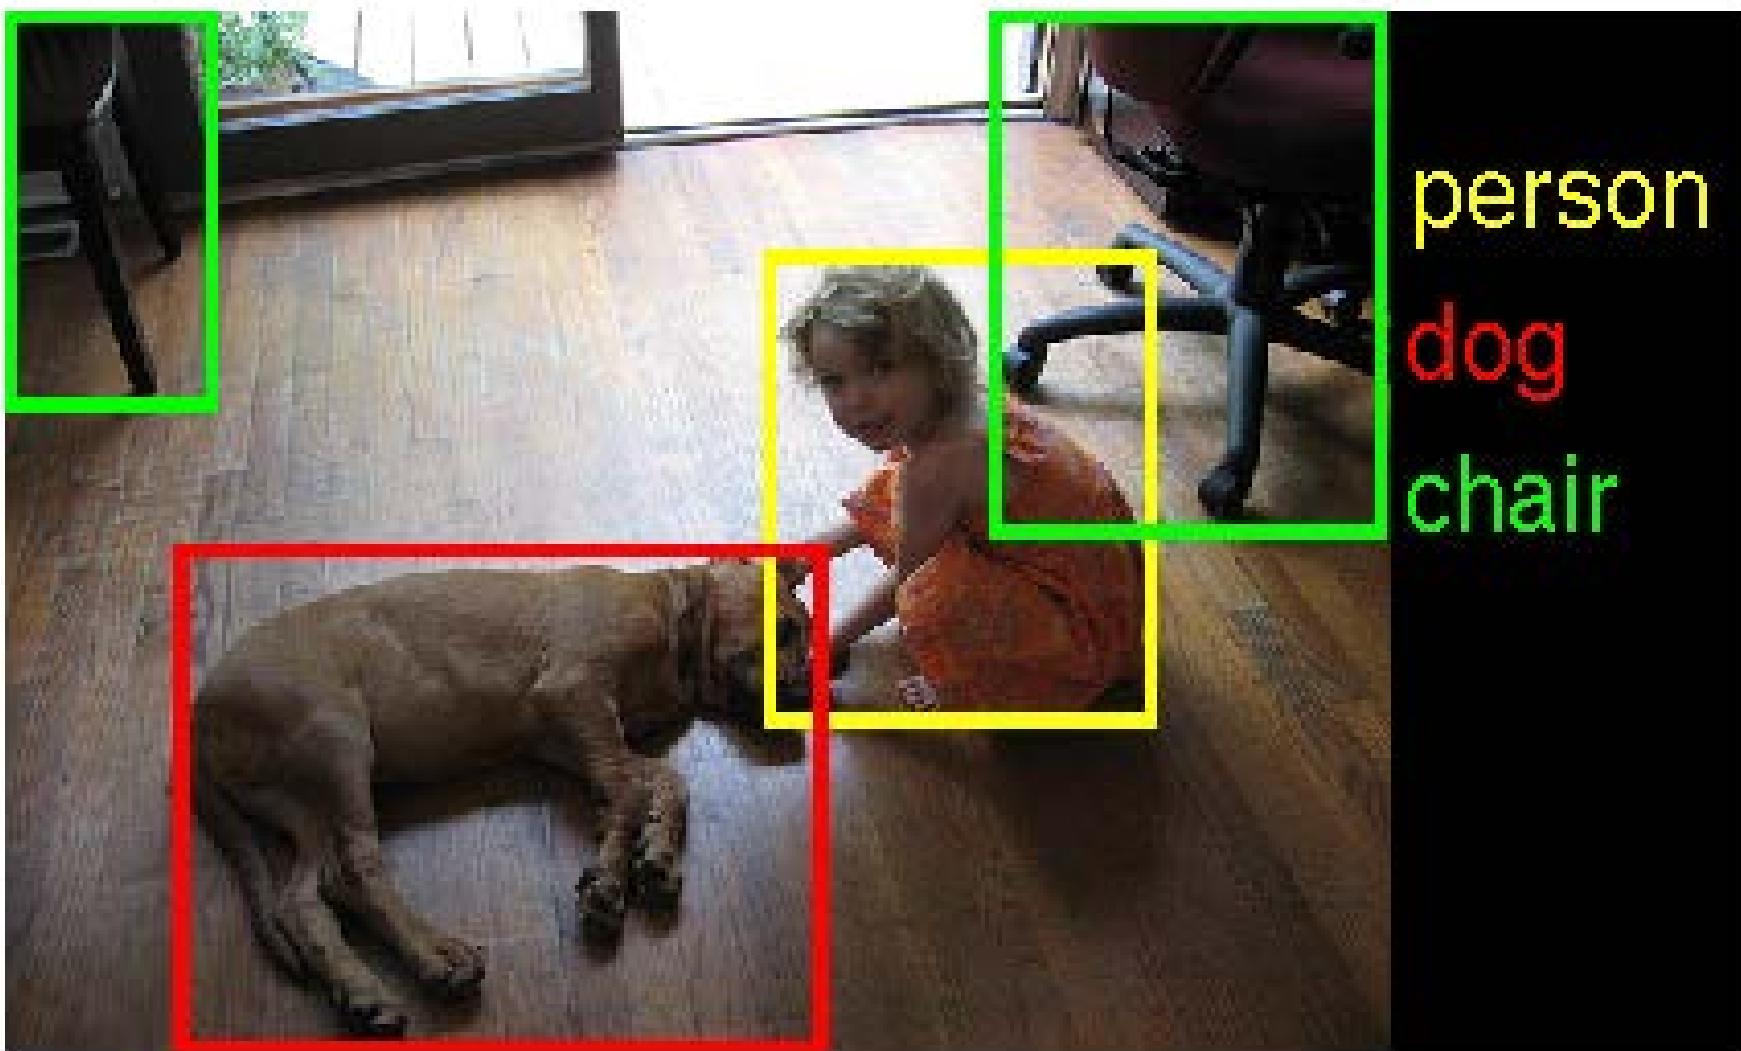
Drives Innovation



x2
Improvement
from 2015 to
2016

What is LPIRC?

Sample Output



$$LPIRC = \frac{\textit{Image Evaluation}}{\textit{Energy}}$$

- LPIRC = Low-Power Image Recognition Challenge
- Bigger is Better

“Detect all relevant objects in as many images as possible of a common test set from an ImageNet like object detection data set within 10 minutes.”

What is LPIRC?

Main Goal

“Detect all relevant objects in as many images as possible of a common test set from an ImageNet like object detection data set within 10 minutes.”

What is LPIRC?

Object Detection

Classification



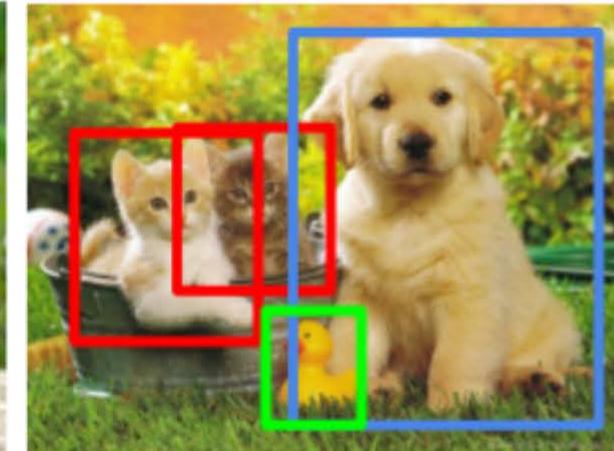
CAT

Classification + Localization



CAT

Object Detection



CAT, DOG, DUCK

Instance Segmentation



CAT, DOG, DUCK

What is LPIRC?

Object Detection

Classification



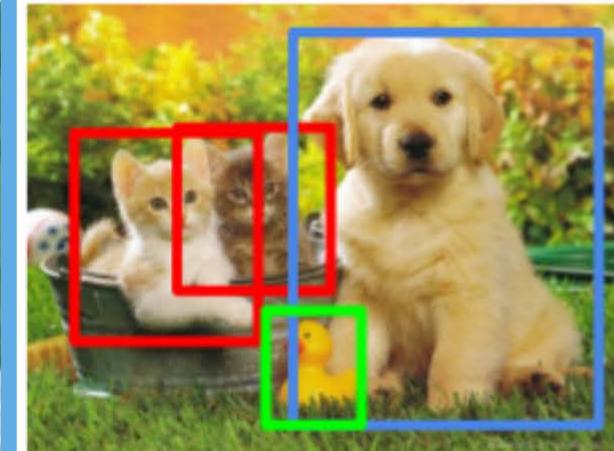
CAT

Classification + Localization



CAT

Object Detection



CAT, DOG, DUCK

Instance Segmentation



CAT, DOG, DUCK

What is LPIRC?

Main Goal

“Detect all relevant objects in as many images as

possible of a common test set from an ImageNet

like object detection data set within 10 minutes.”

What is LPIRC?

ImageNet Dataset

- Clear
- Object of Interest is Large



IMAGENET LSVRC2014 Object Detection Dataset

- Training images collected and fully annotated with all 200 object categories for ILSVRC2014
- Training images annotated with 1-2 object categories from ILSVRC2013
- Validation images fully annotated with all 200 object categories, used in ILSVRC2013 and in ILSVRC2014

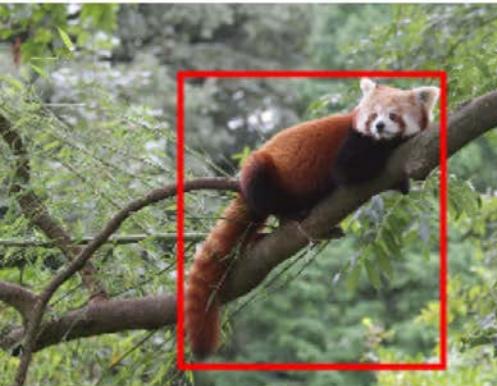
- Images with to object classes
 Images with to object instances

Classes:

- punching bag
- purse
- rabbit
- racket
- ray
- red panda
- refrigerator
- remote control
- rubber eraser
- rugby ball
- ruler

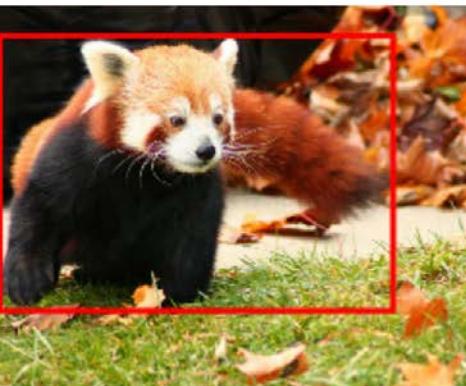
Selection Criteria:

- Training images collected and fully annotated with all 200 object categories for ILSVRC2014
- 1 to 4 object classes
- Images with red panda



• red panda

ILSVRC2014_train_00008818



• red panda

ILSVRC2014_train_00017468



• red panda

What is LPIRC?

Main Goal

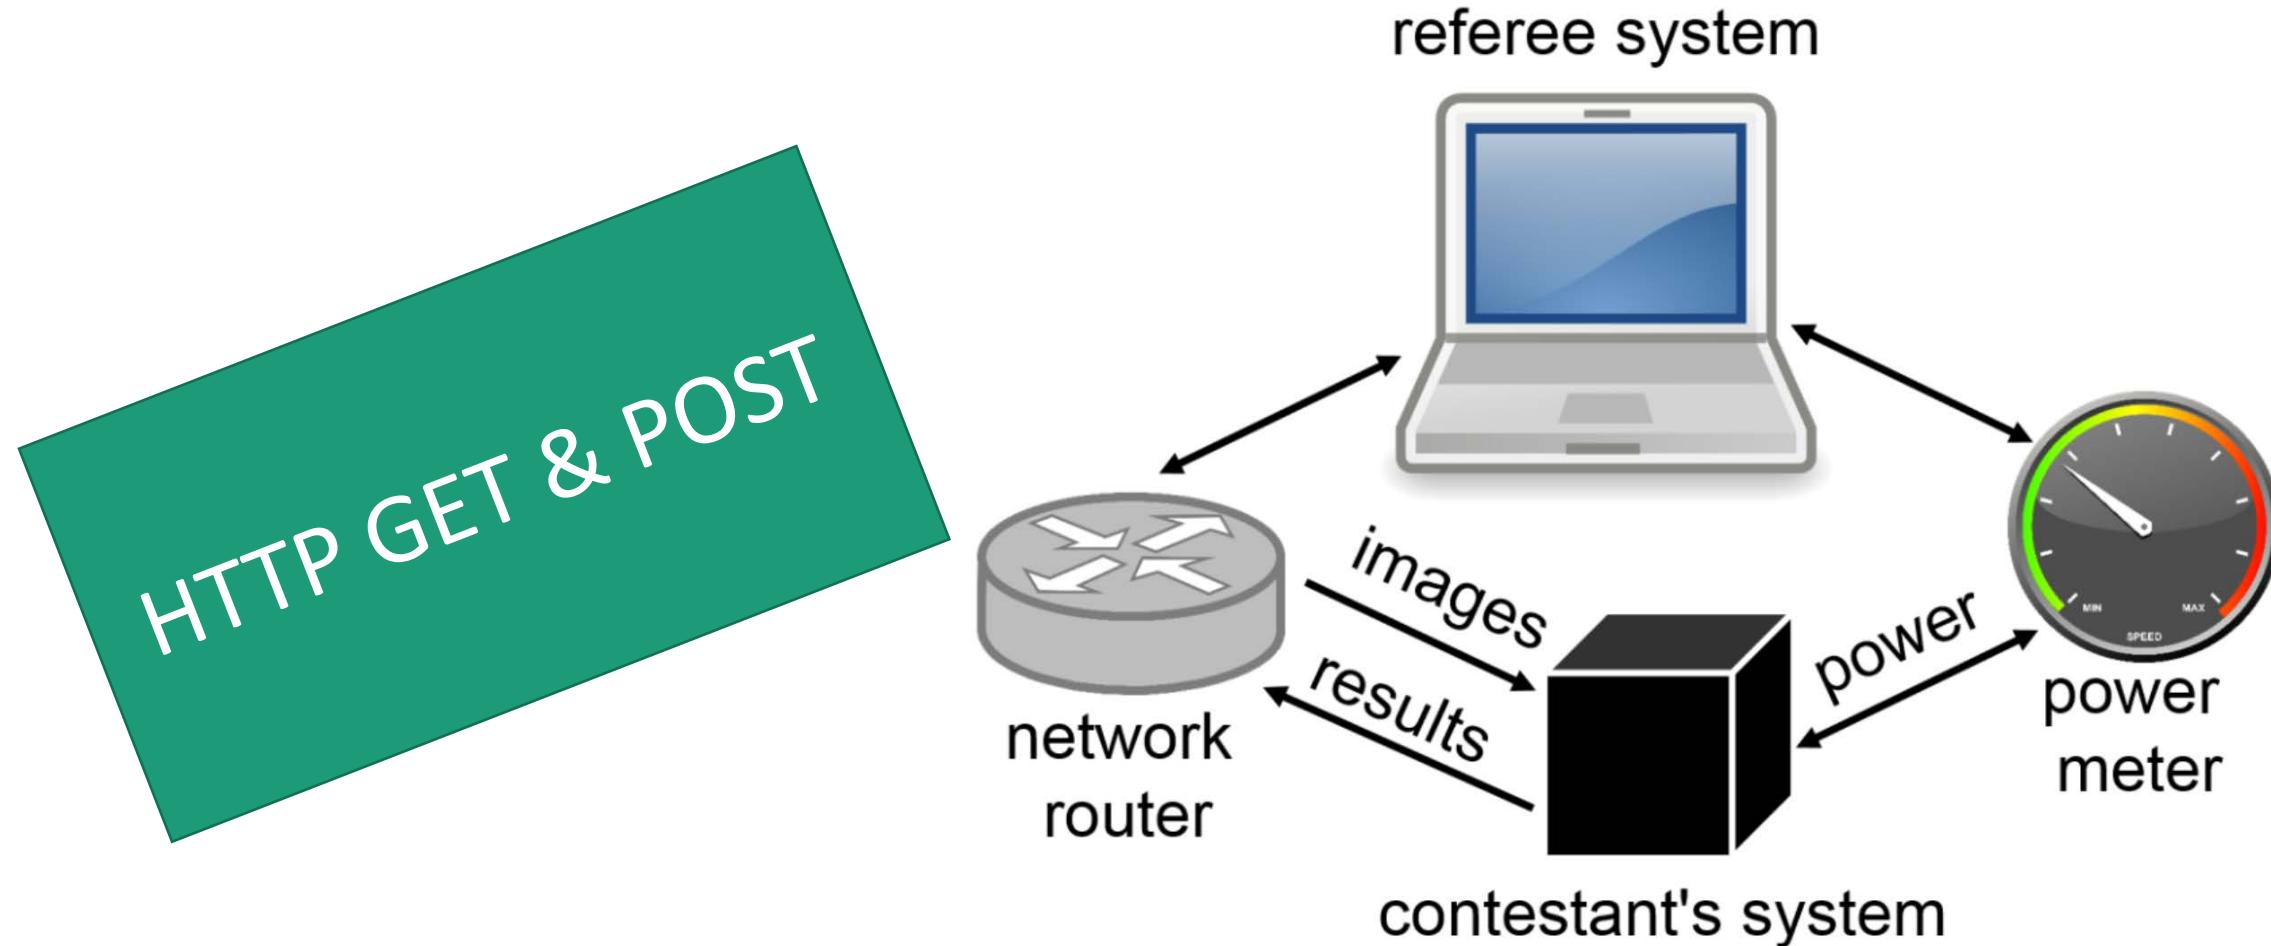
“Detect all relevant objects in as many images as possible of a common test set from an ImageNet like object detection data set **within 10 minutes.**”

What is LPIRC?

Time Limit

10 Minutes





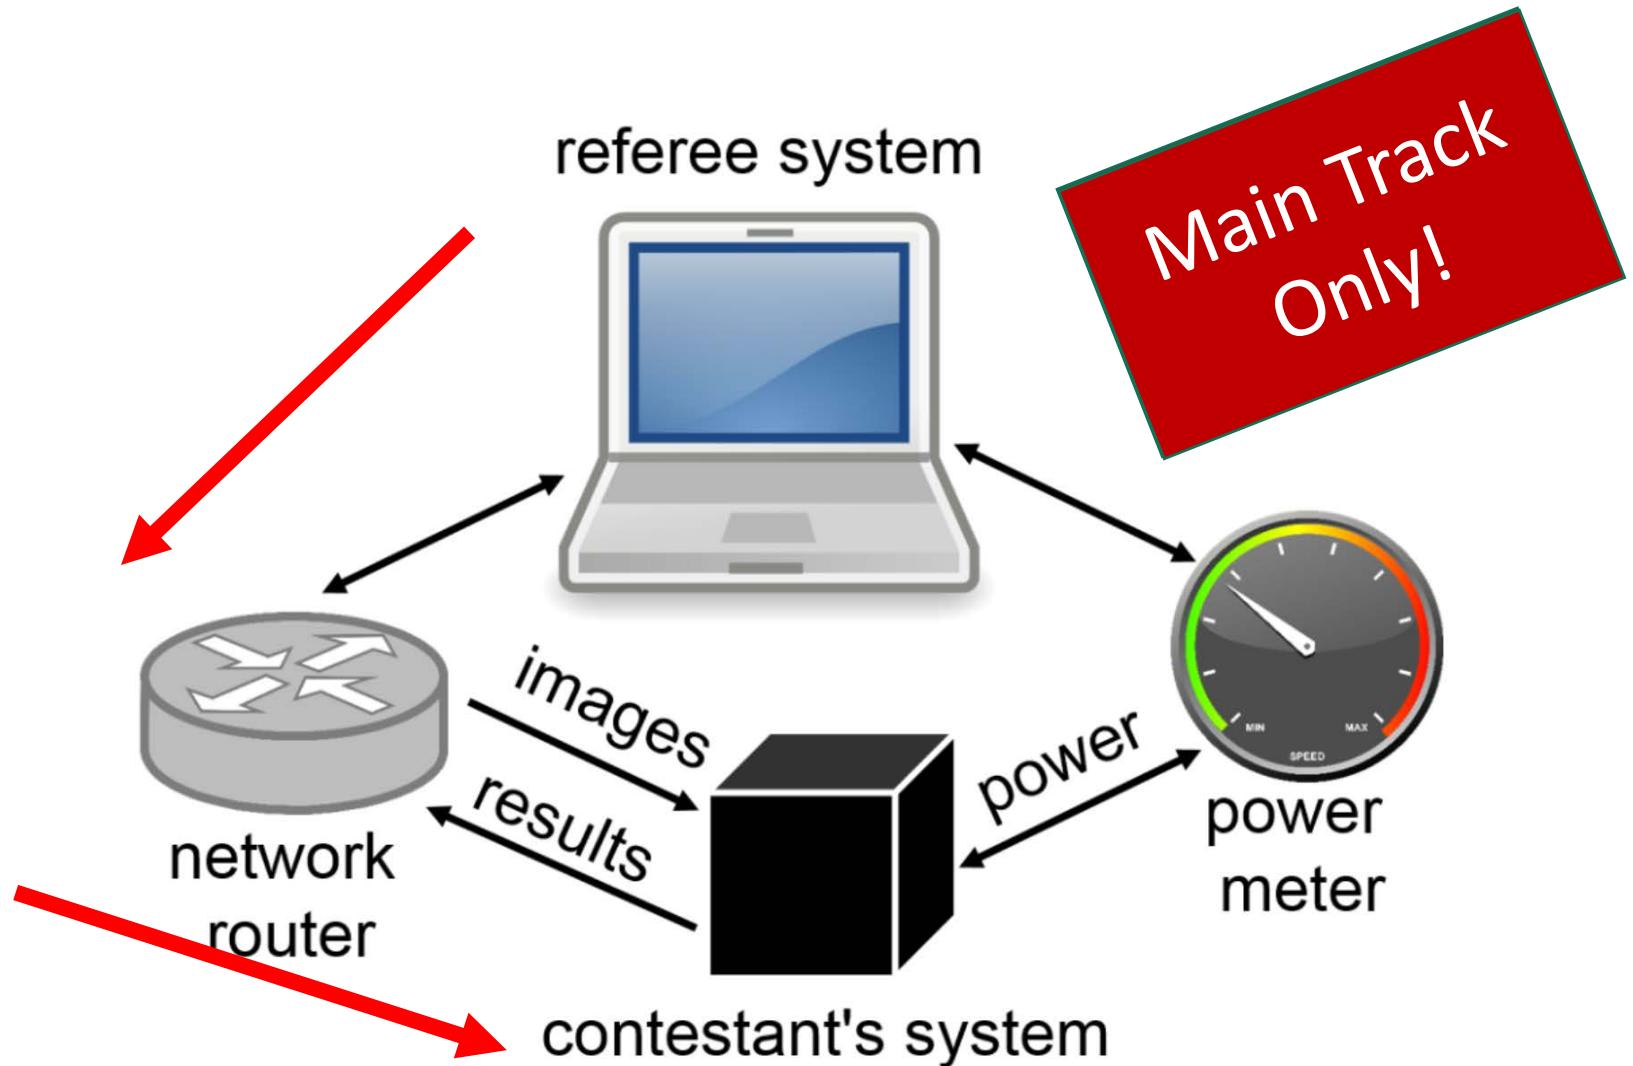
Energy (Power × Time) Measurement

Power Meter: Yokogawa WT 300

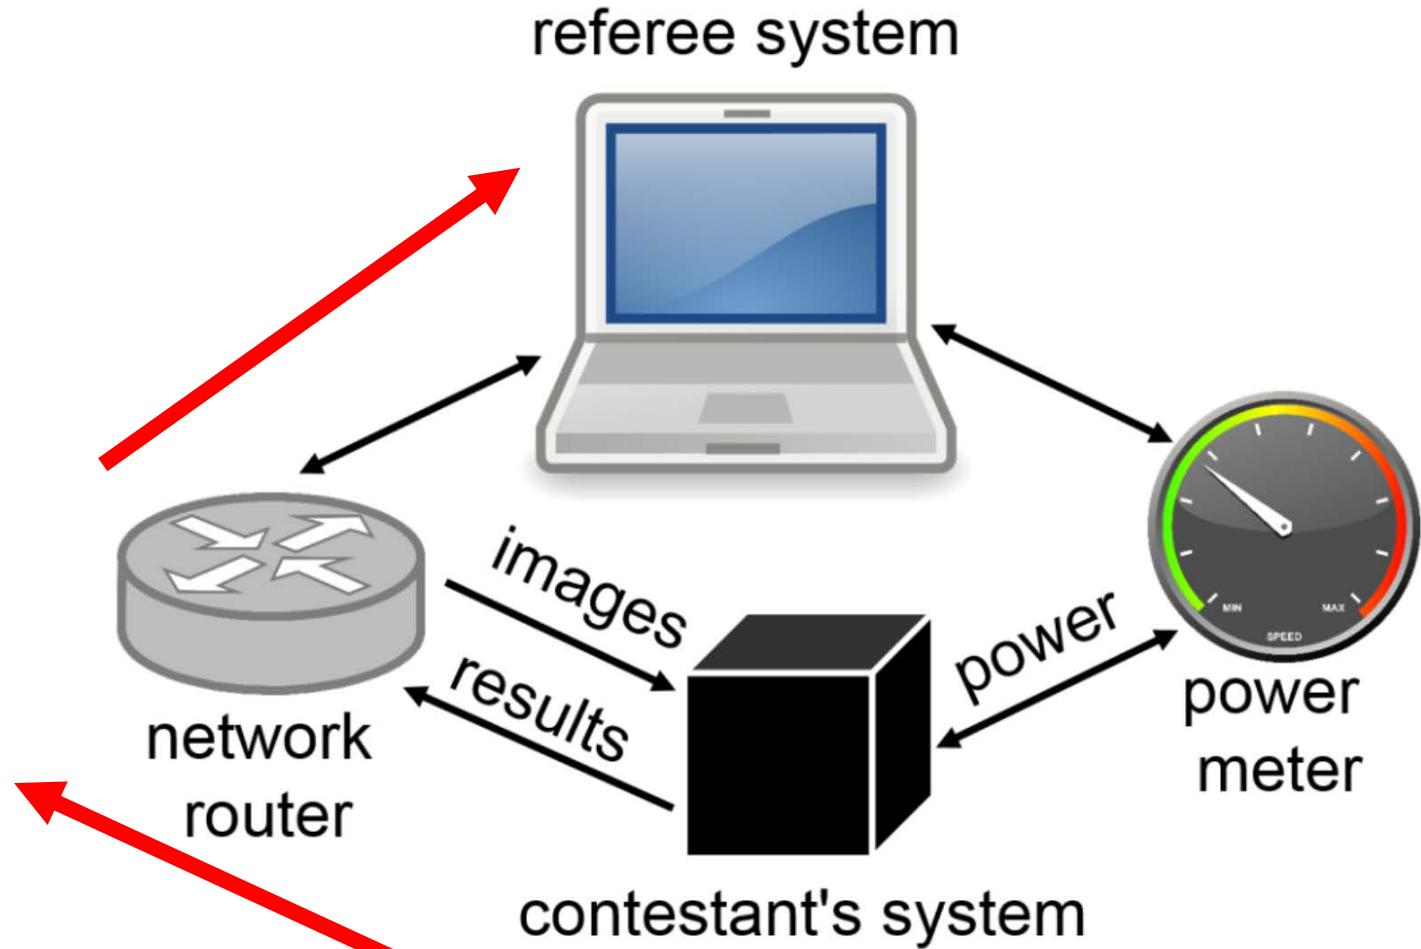


download referee source code: lpirc.net

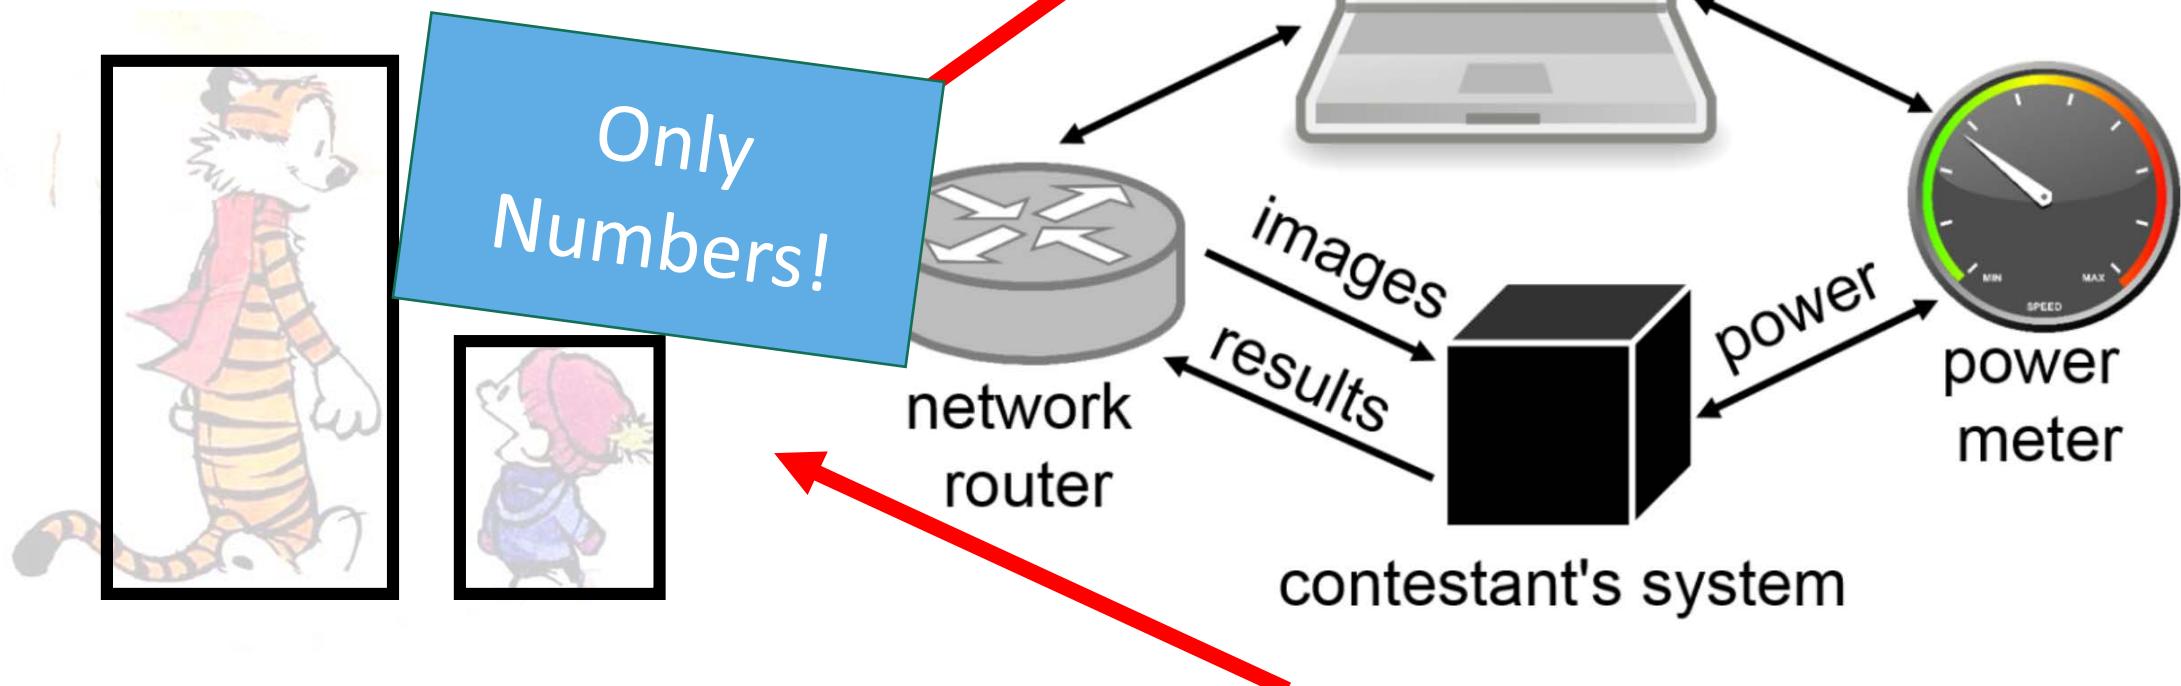
Step 1: Send data to contestant's system



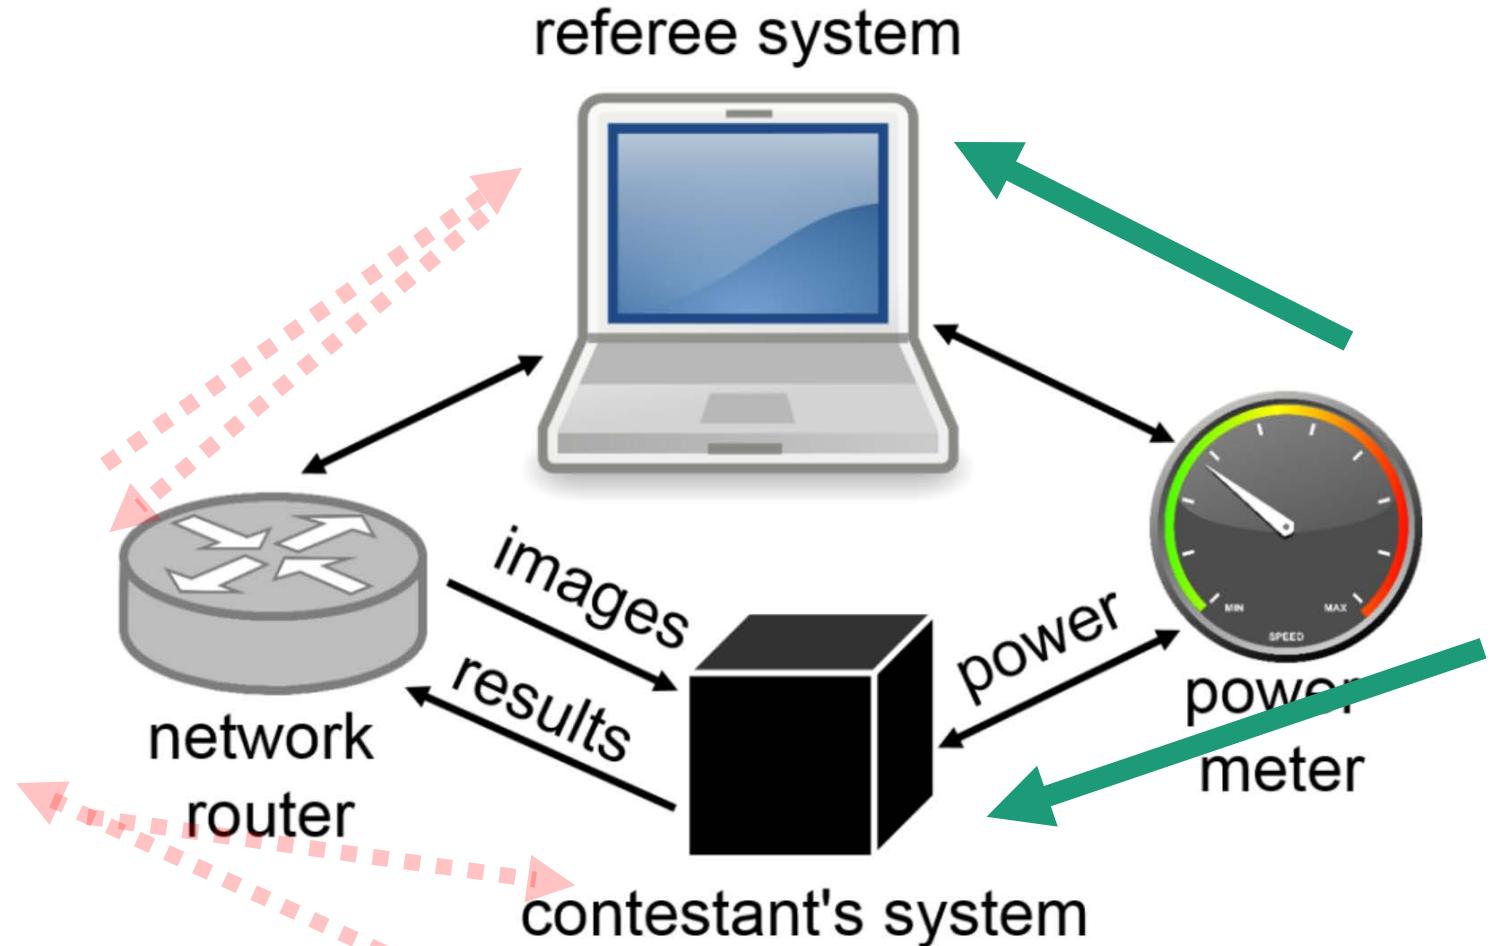
Step 2: Contestant's system responds with image labels

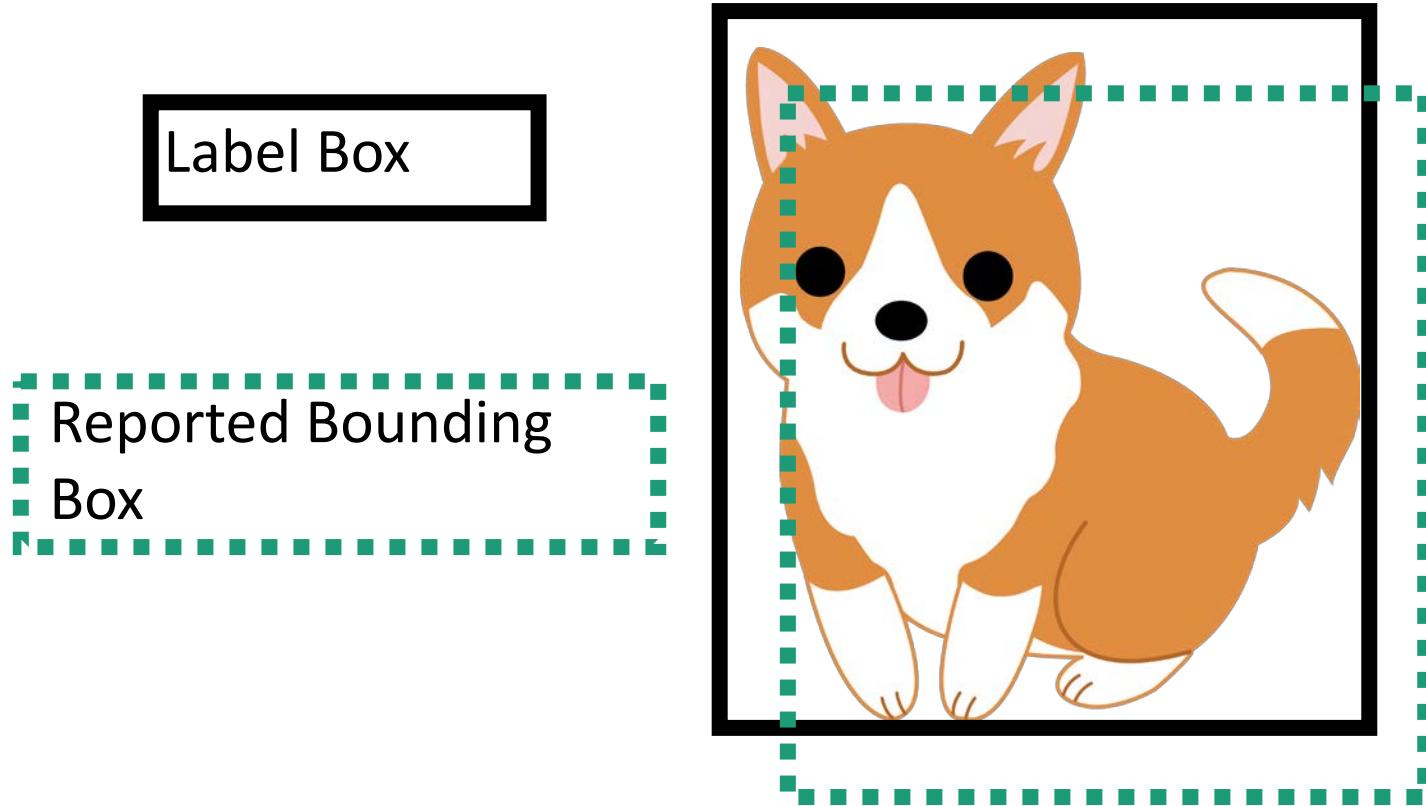


Step 2: Contestant's system responds with image labels



Step 1.5: Period sampling of the power meter





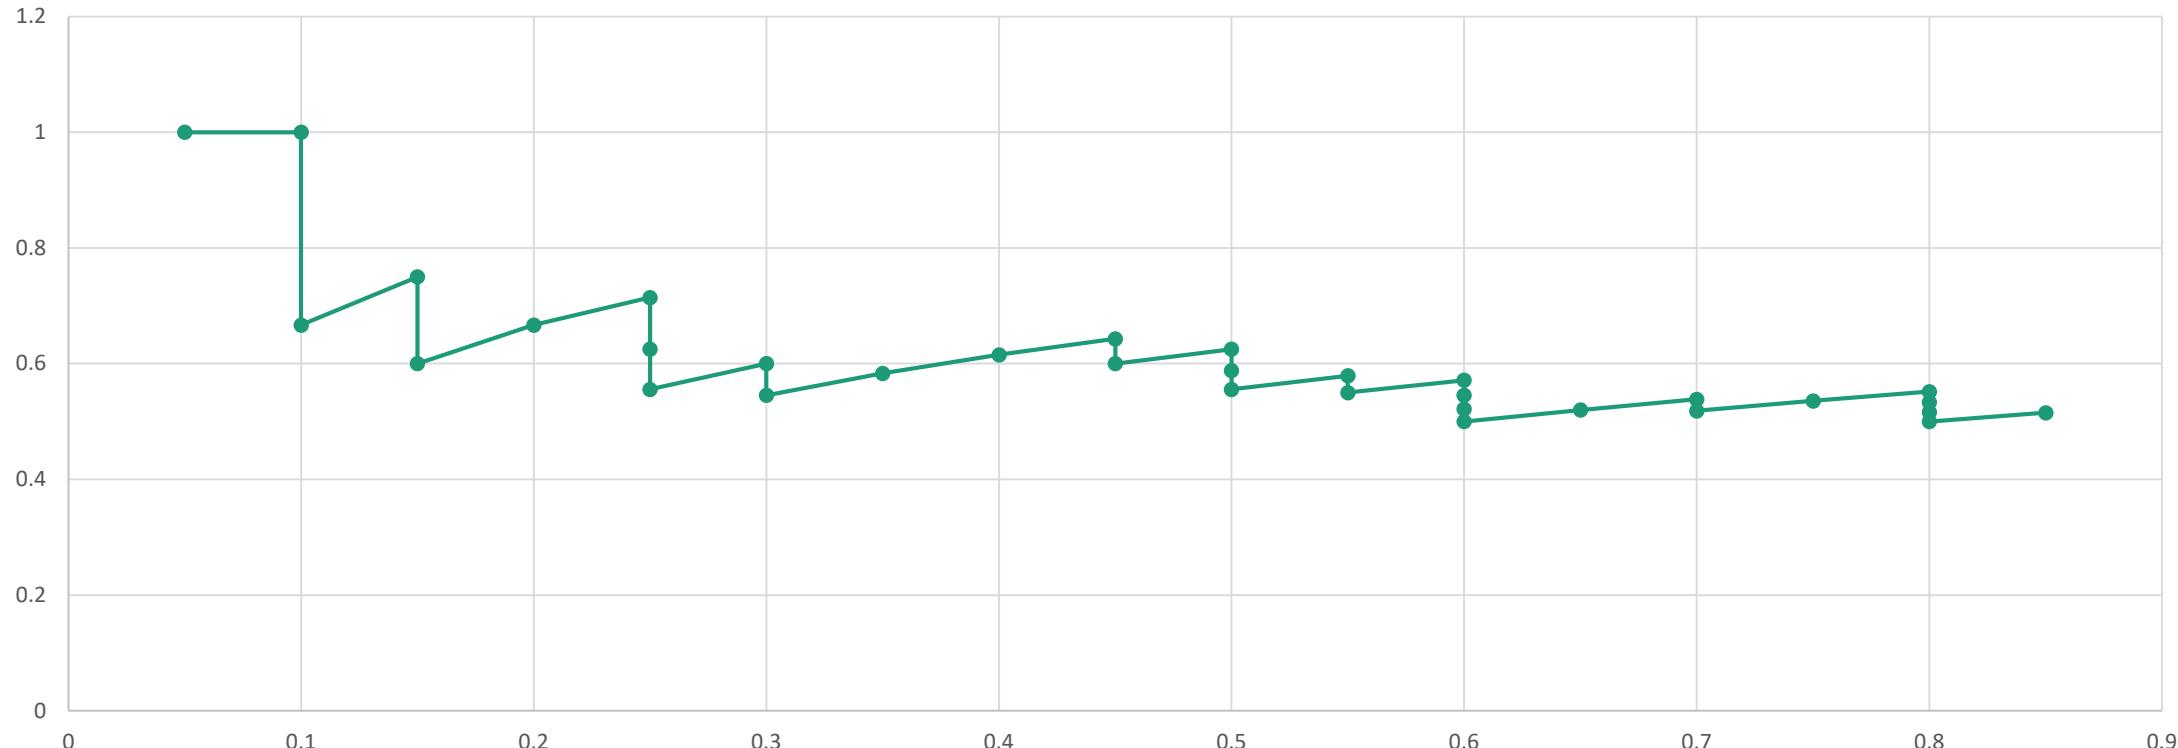
Correct if

$$\frac{Label \cap Report}{Label \cup Report} > 0.5$$

Referred to as...
“Intersection over Union”

- mAP = mean average precision
- Bigger is better

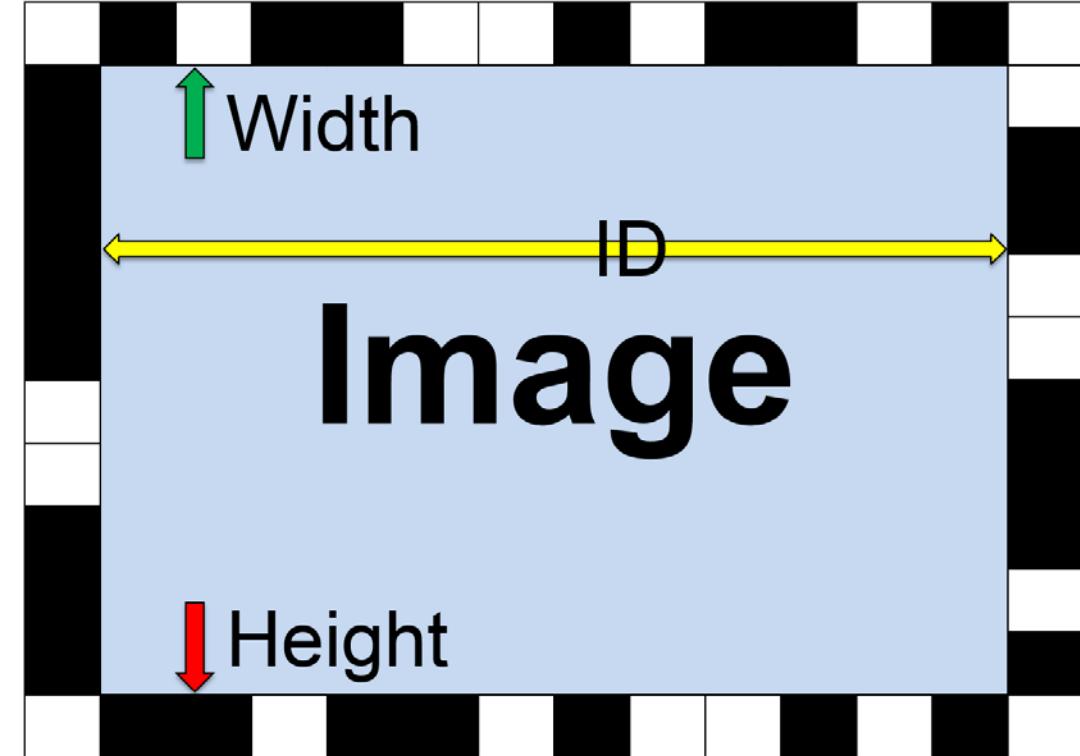
Precision v.s. Recall



* Olga Russakovsky*, Jia Deng*, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg and Li Fei-Fei. (* = equal contribution) **ImageNet Large Scale Visual Recognition Challenge**. IJCV, 2015

$$LPIRC = \frac{ILSVRC\ Score}{Energy}$$

- LPIRC = Low-Power Image Recognition Challenge
- ILSVRC = ImageNet

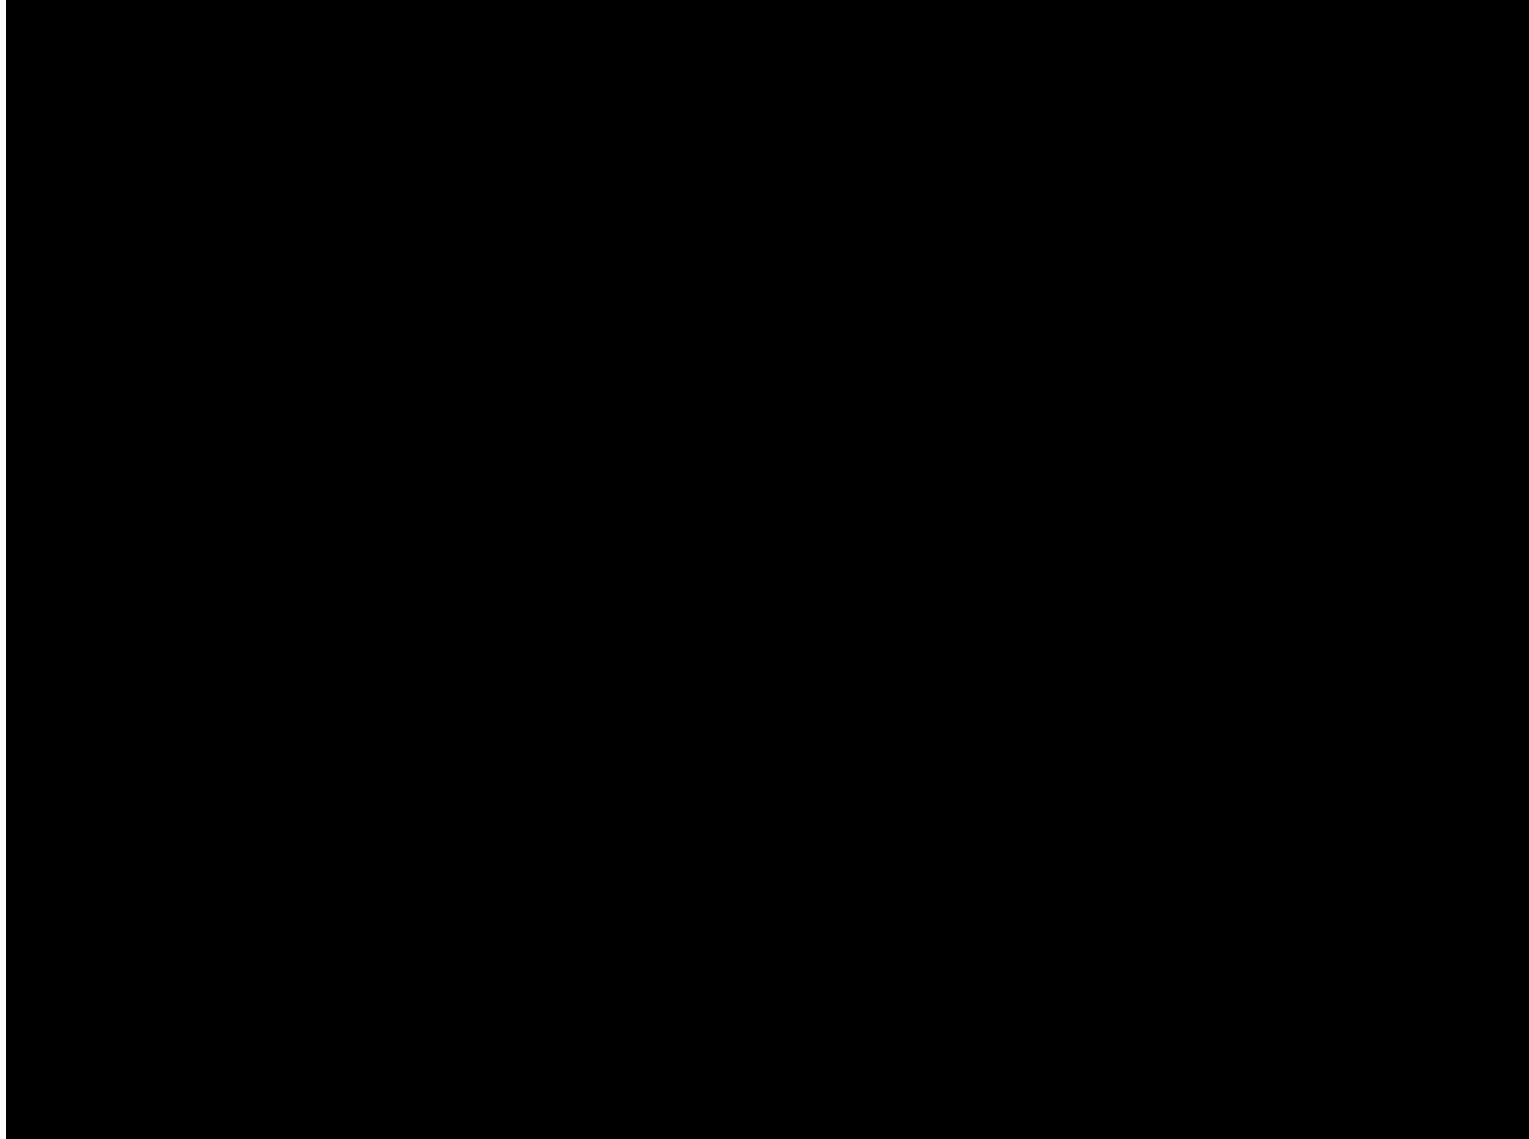




Input

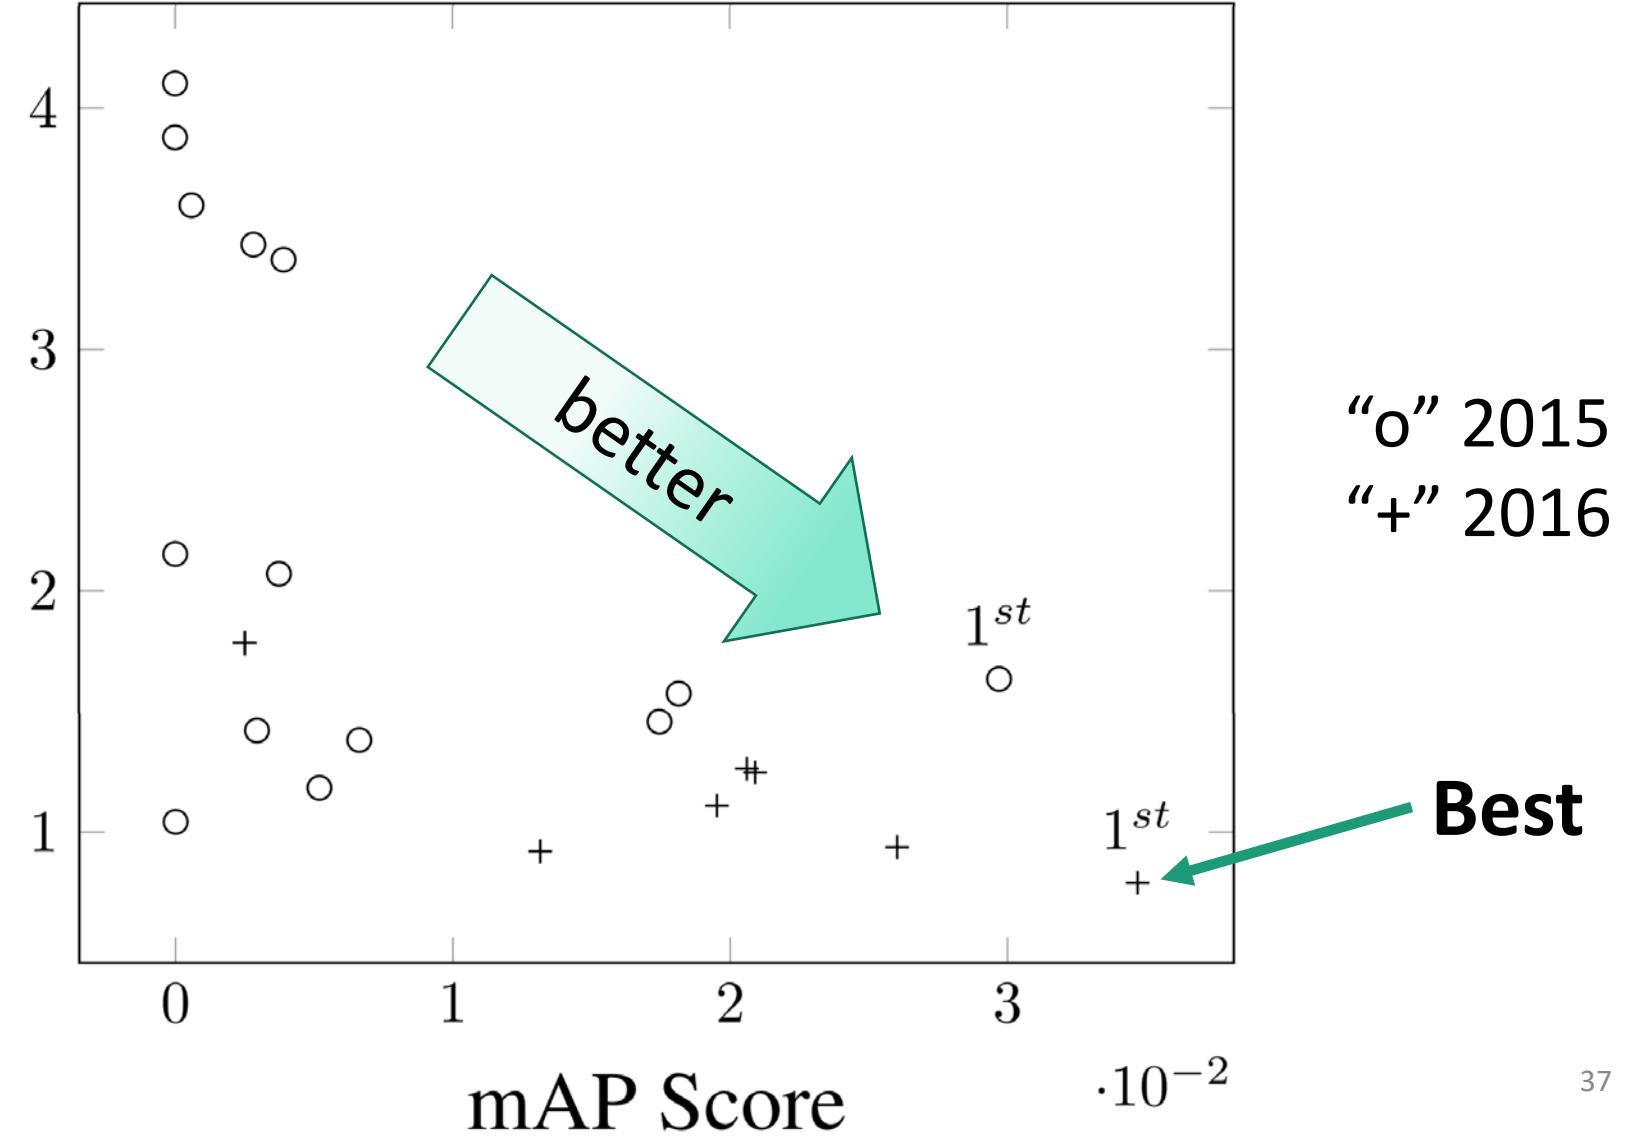


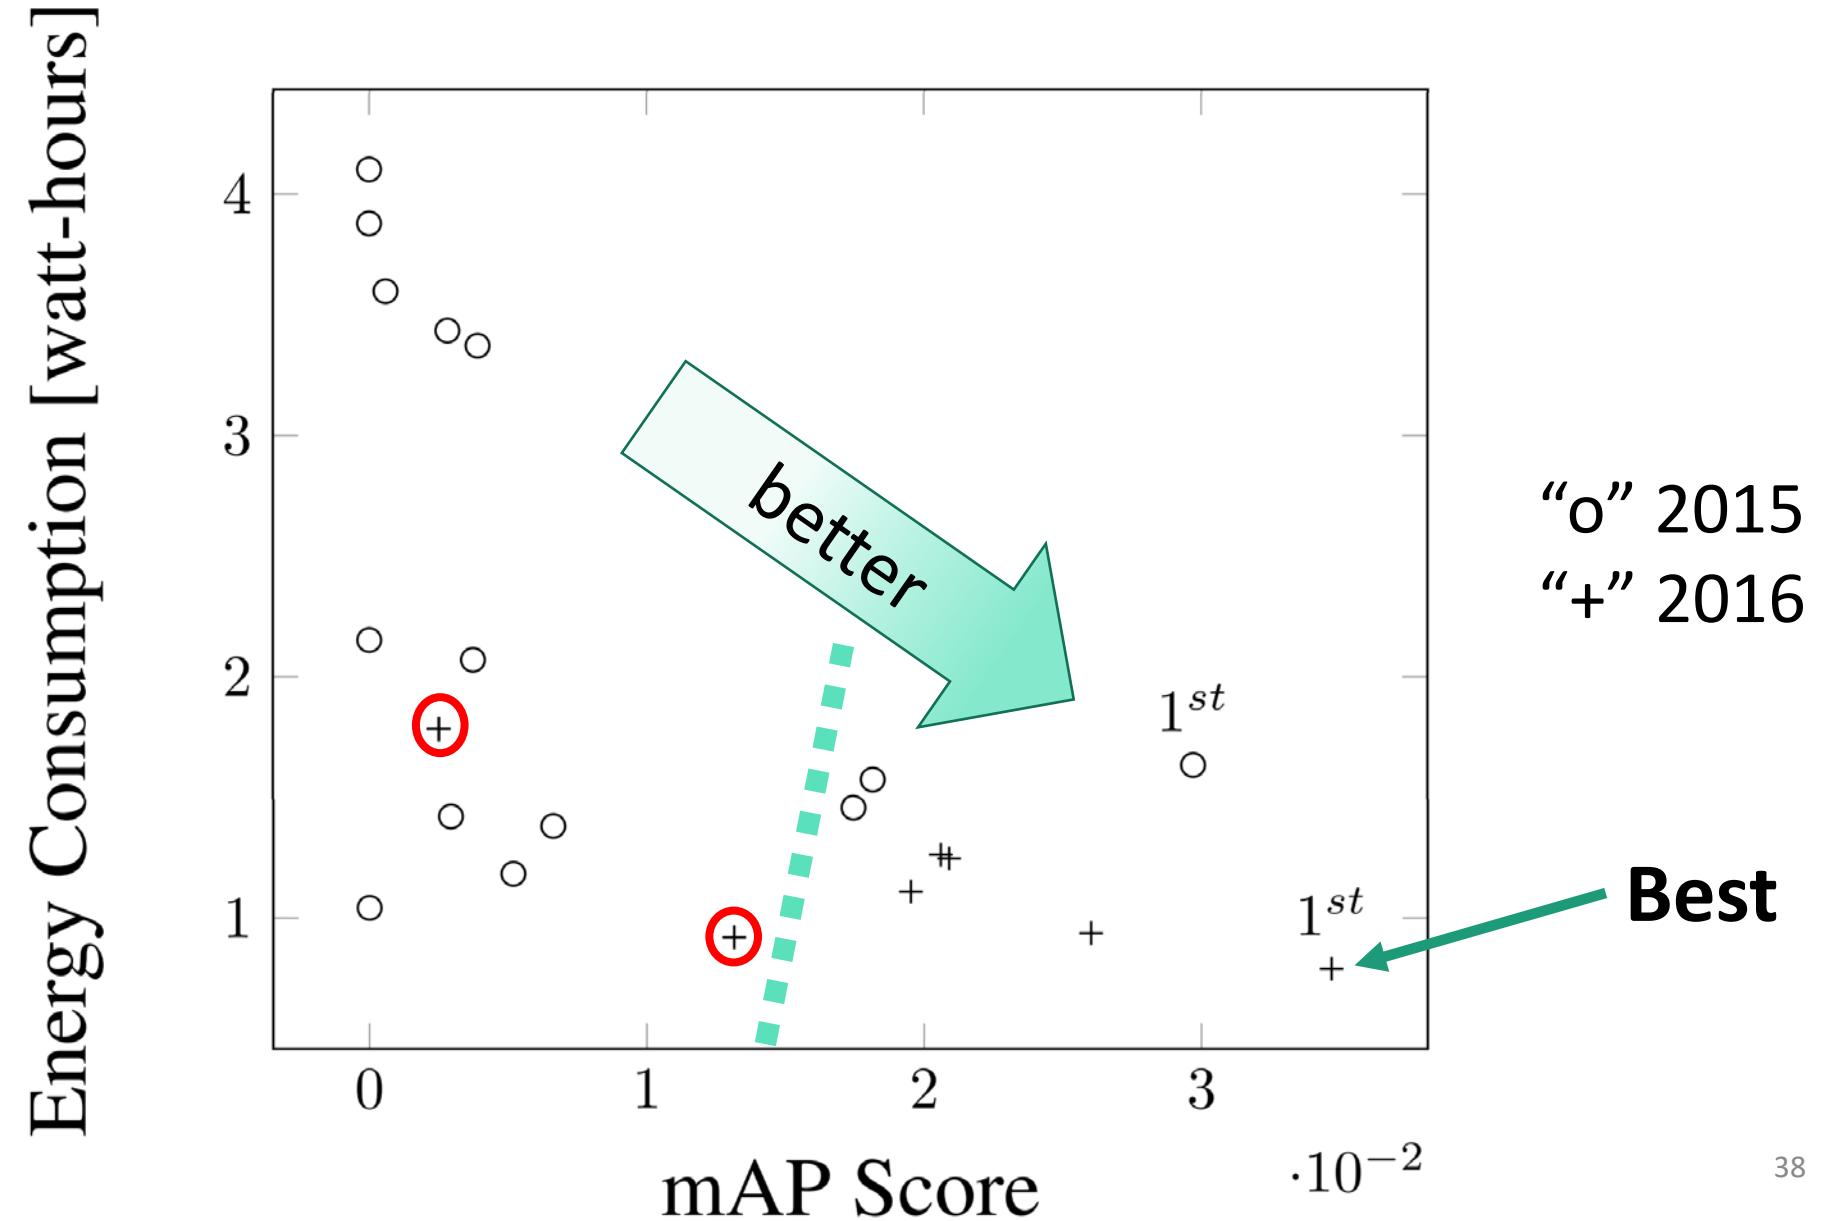
Output



mAP =
Mean
Average
Precision

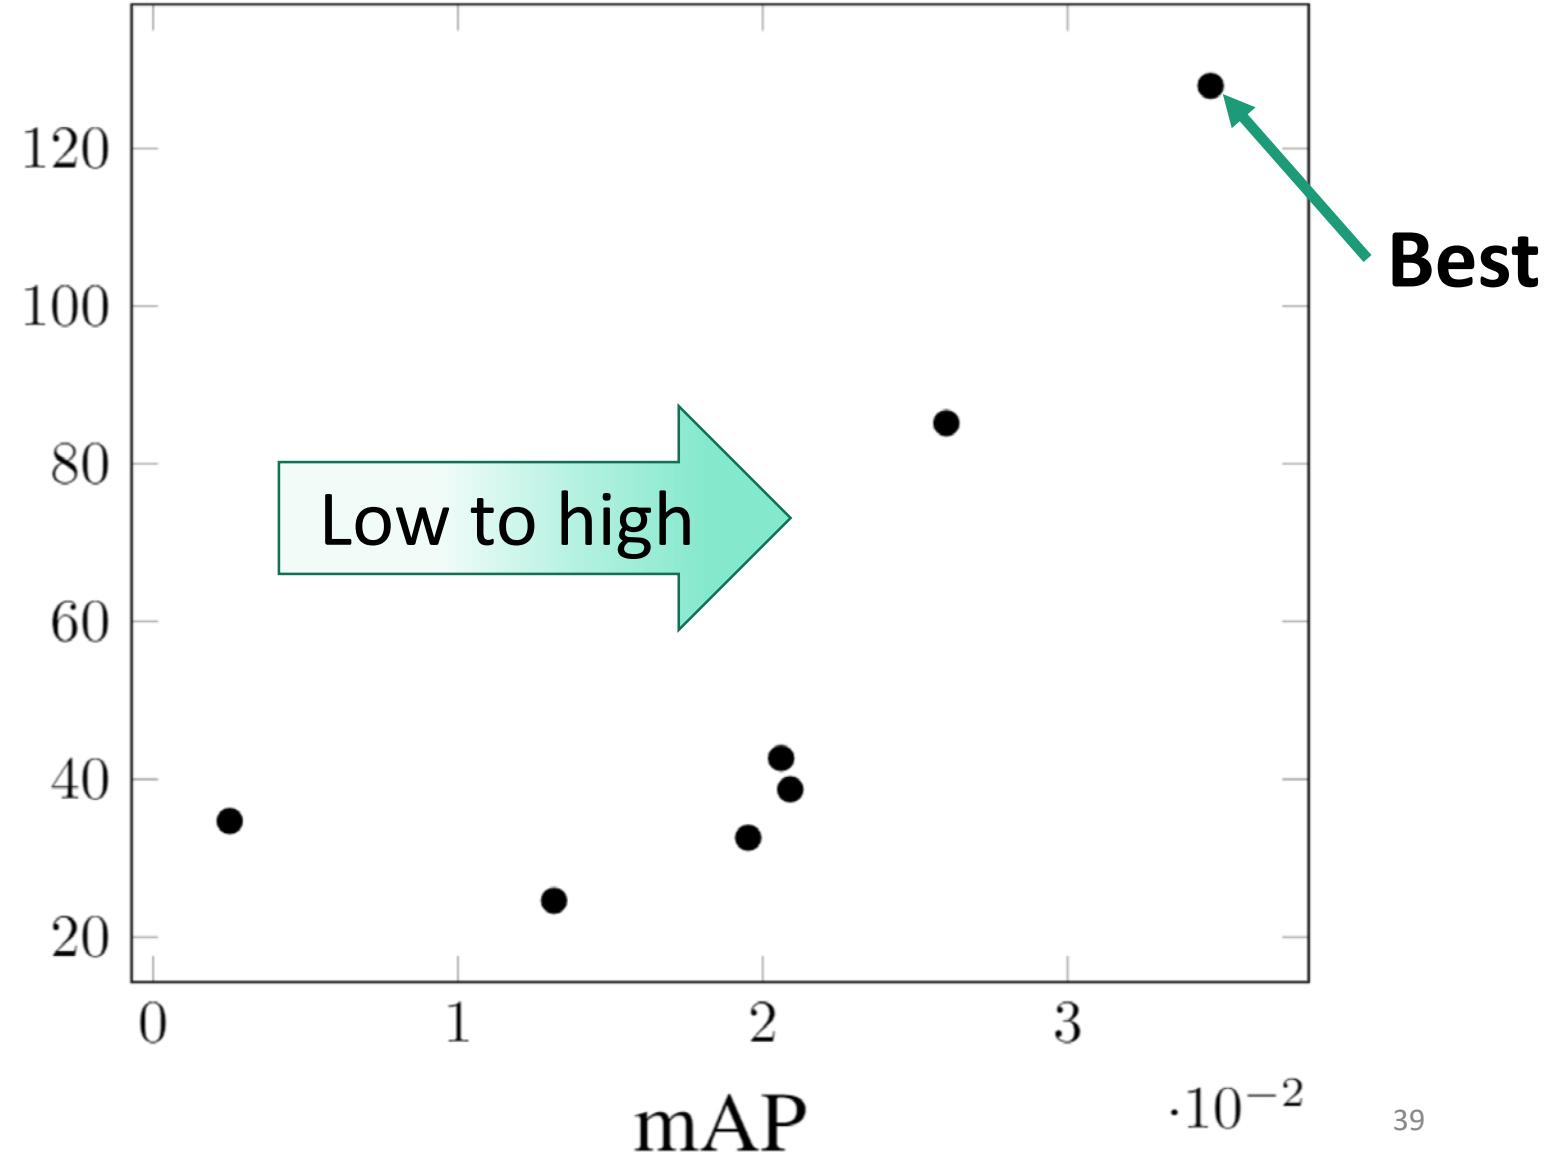
Energy Consumption [watt-hours]





Better is high

Objects-per-second



Person 0.54

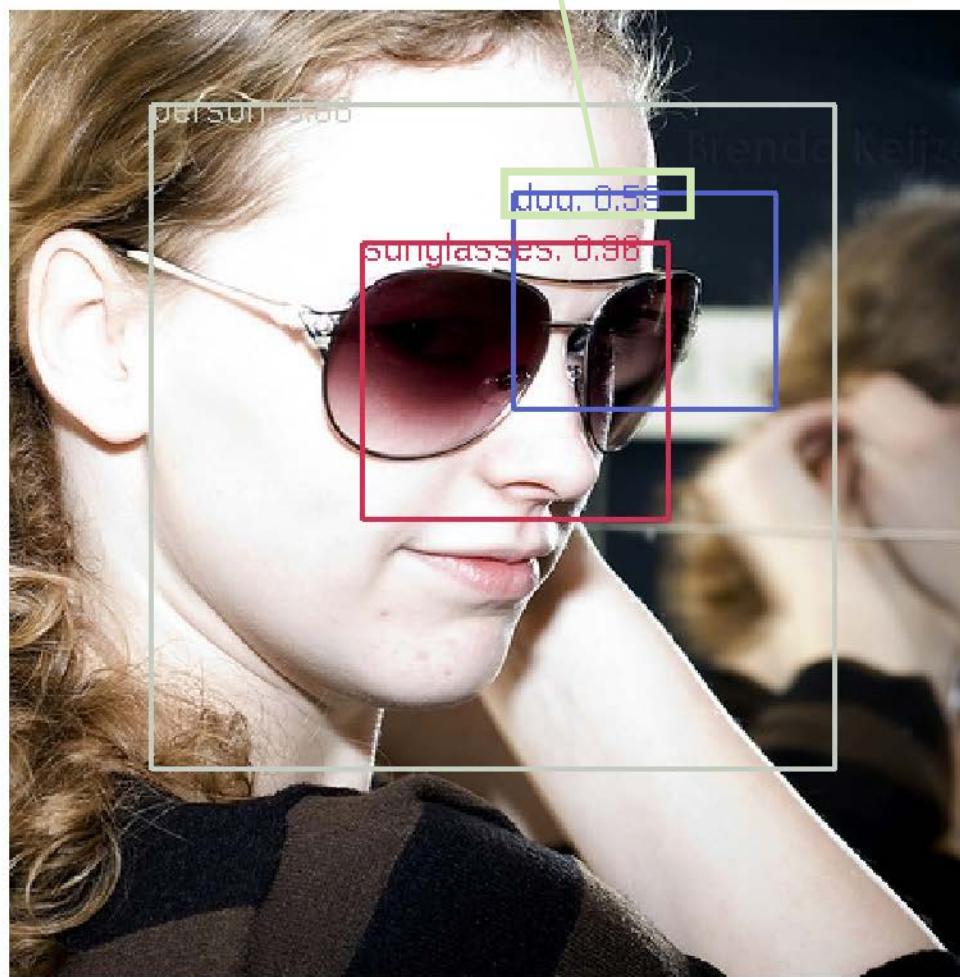
Person 0.92



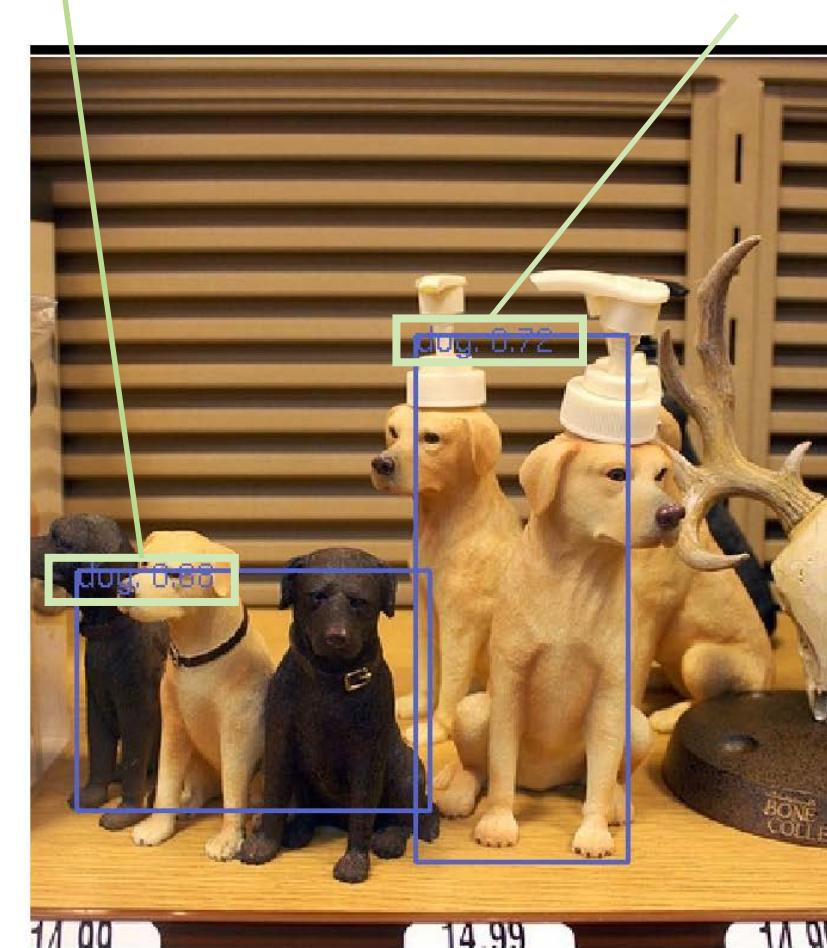
Dog 0.99



Dog 0.55



Dog 0.68

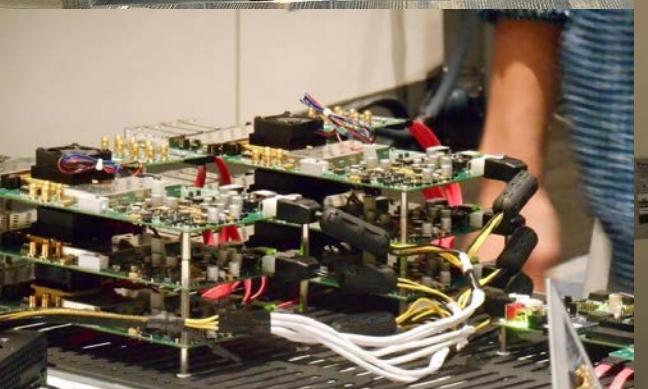
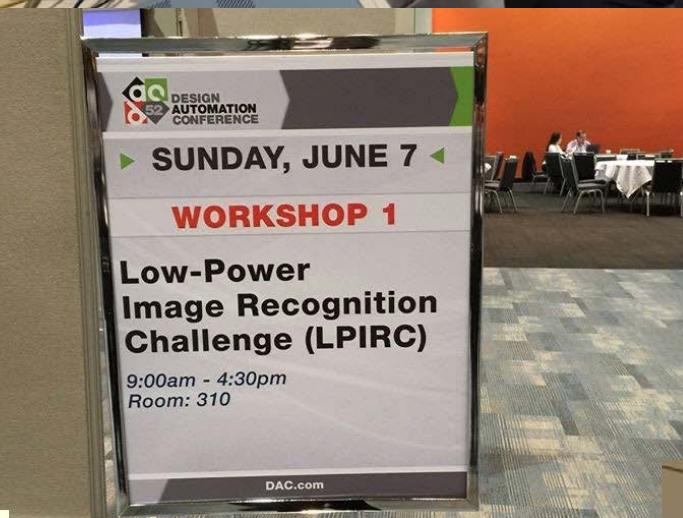
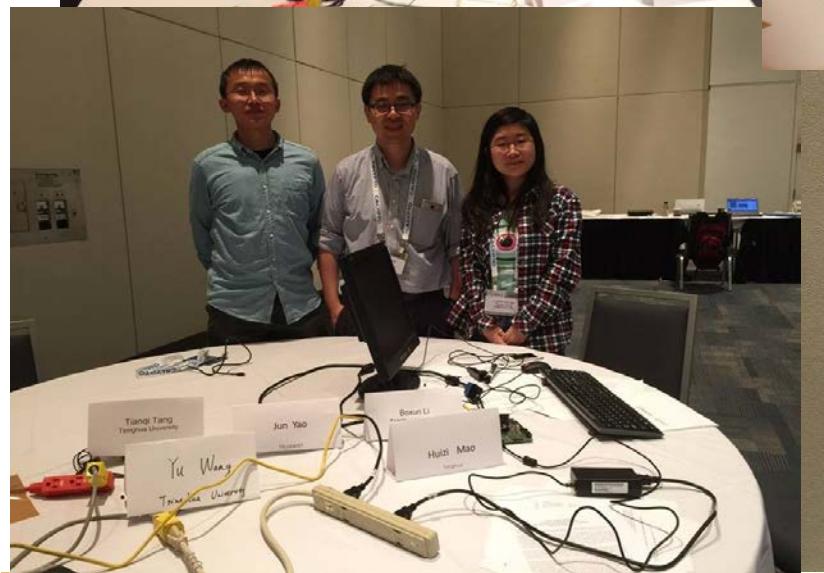


Dog 0.72

- Yung-Hsiang Lu et al, Rebooting Computing and Low-Power Image Recognition Challenge, *IEEE/ACM International Conference on Computer-Aided Design (ICCAD) : digest of technical papers, 2015.*
- Jingyu Liu et al, Fast Object Detection at Constrained Energy IEEE Transactions on Emerging Topics in Computing (Volume: PP, Issue: 99) Date of Publication: 07 June 2016.
- Huizi Mao et al, Towards Real-Time Object Detection on Embedded Systems, *IEEE Transactions on Emerging Topics in Computing* : Volume: PP Issue: 99, 2015.







- July 2017 in Hawaii before CVPR (Computer Vision and Pattern Recognition) |
- More teams: please encourage your colleagues and students to participate
- Student travel grants will be available
- 2015/2016 Prizes were \$1000, \$500, \$200
- Higher prizes in 2017
- Registration is coming soon (email yunlu@purdue.edu or gauenk@purdue.edu)

- Vital to development of practical image processing applications
- Users perform “object detection” on ImageNet data
- Final score is the image processing score divided by the energy consumption
- Over x2 improvement from 2015 to 2016
- Registration for the July LPIRC 2017 is coming soon (email yunlu@purdue.edu or gauenk@purdue.edu)