The Retrospect and Prospect of Low-Power Image Recognition Challenge (LPIRC)

Yung-Hsiang Lu, *Professor, Purdue University*

Yiran Chen, *Associate Professor, Duke University*

Alexander C Berg, *Associate Professor, University of North Carolina at Chapel Hill*

Outline

- Motivation
- History
- Metrics
- Development
- Comparison
- Contribution
- Future

Motivation

- **Cameras** have become available in many embedded and mobile systems, including smartphones, wearable devices and aerial robots. It is desirable to have the capability of detecting objects in the images by computers.
- **Energy** is limited in mobile systems so energy conservation is important.



Motivation

- Although many competitions are held each year for various topics, such as image recognition and low power, **Low-Power Image Recognition Challenge (LPIRC)** is the only competition integrating both image recognition and low power.
- Winners will be evaluated based on both high recognition accuracy and low power usage.

Brief Introduction

- LPIRC is Low-Power Image Recognition Challenge
- LPIRC = a competition that combines two important technologies:

Computer Vision + Low-Power Systems

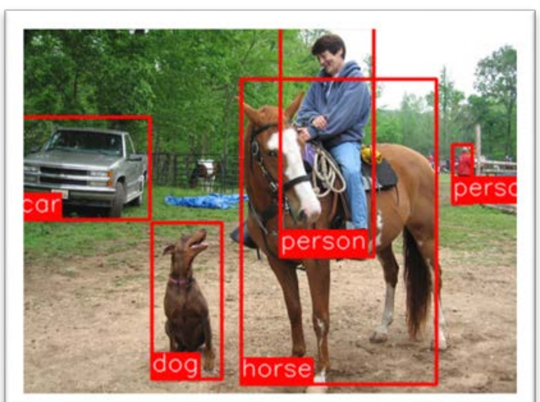

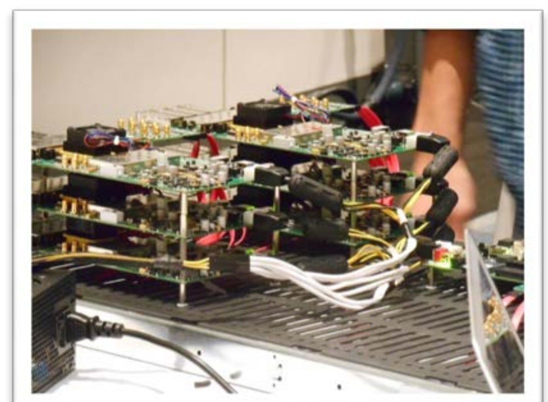

Brief Introduction

- LPIRC has been held annually since 2015 as an on-site competition.
- To encourage innovation, LPIRC intents to have no restrictions on hardware or software platforms.

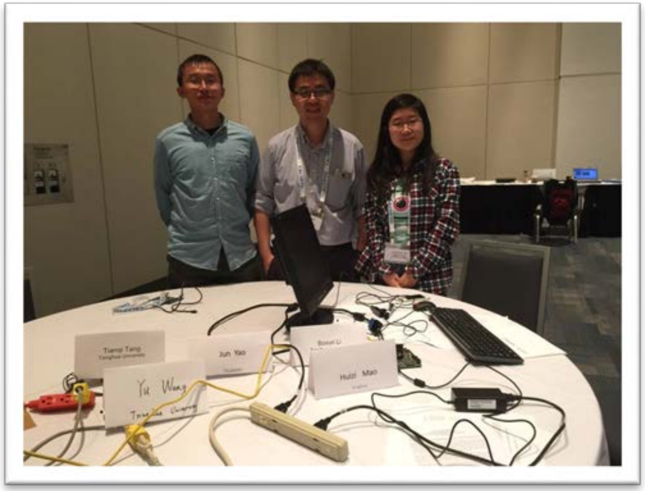

Brief Introduction

- Each team has 10 minutes to recognize the objects in 5000（year2015) or 20,000(year 2016 and 2017) images.

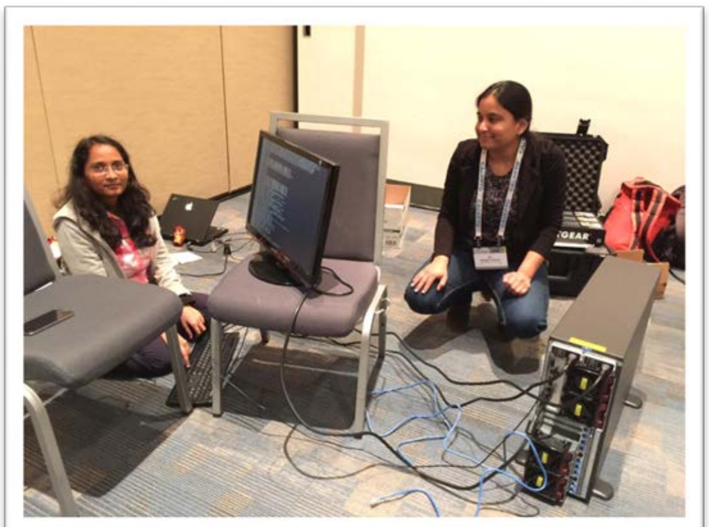

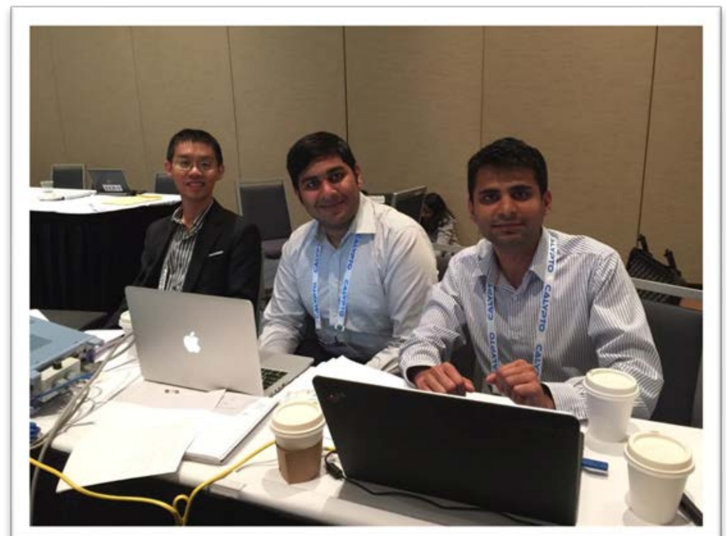

Brief Introduction

- Team are first trained with ImageNet's labels.
- They are then given test images during the competition from our own dataset.
- Each image may contain one or several objects (such as person, dog, airplane, and table) that belong to 200 predefined categories. (more details later)
- To successfully recognize an object in an image, a computer program must correctly determine the category and mark a bounding box around the object.

Brief Introduction

Sample Output

(identified objects are enclosed by bounding boxes)

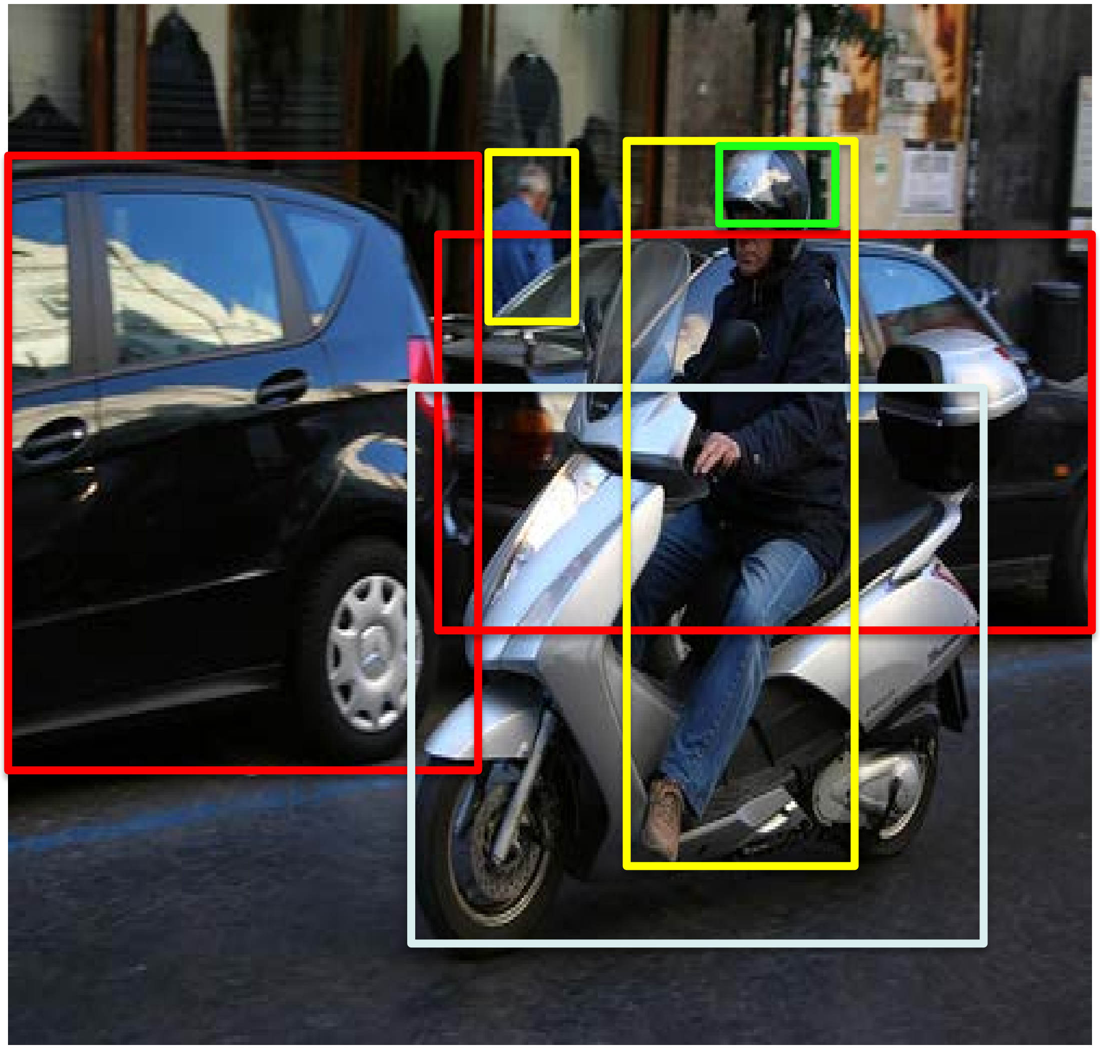

Person
Car
Motorcycle
Helmet

Organization Team

- **Yung-Hsiang Lu** is a professor in the School of Electrical and Computer Engineering and (by courtesy) the Department of Computer Science of Purdue University.
- He is the leader organizer of LPIRC 2015-2017. He is an ACM distinguished scientist and ACM distinguished speaker. He is a member in the organizing committee of the IEEE Rebooting Computing Initiative.
- He is the lead organizer of Low-Power Image Recognition Challenge, the chair of the Multimedia Communication Systems Interest Group in IEEE Multimedia Communications Technical Committee.

PURDUE UNIVERSITY

Organization Team

- **Yiran Chen** is an associate professor at Duke University. He is the associate editor of IEEE TCAD, IEEE D&T, IEEE ESL, ACM JETC, and ACM SIGDA E-newsletter, and served on the technical and organization committees of around 40 international conferences.

- **Alexander C Berg** is an associate professor at UNC. His research concerns computational visual recognition. He has worked on general object recognition in images, human pose identification in images, image parsing, face recognition, and machine learning for computer vision.

- **Bo Chen** is a software engineer at Google Inc. He received his PhD in Computation and Neural Systems from the California Institute of Technology.

LPIRC rule and metric

- Each team has 10 minutes.
- Classify the class and recognize objects.
- Draw the bounding box and report to referee system.

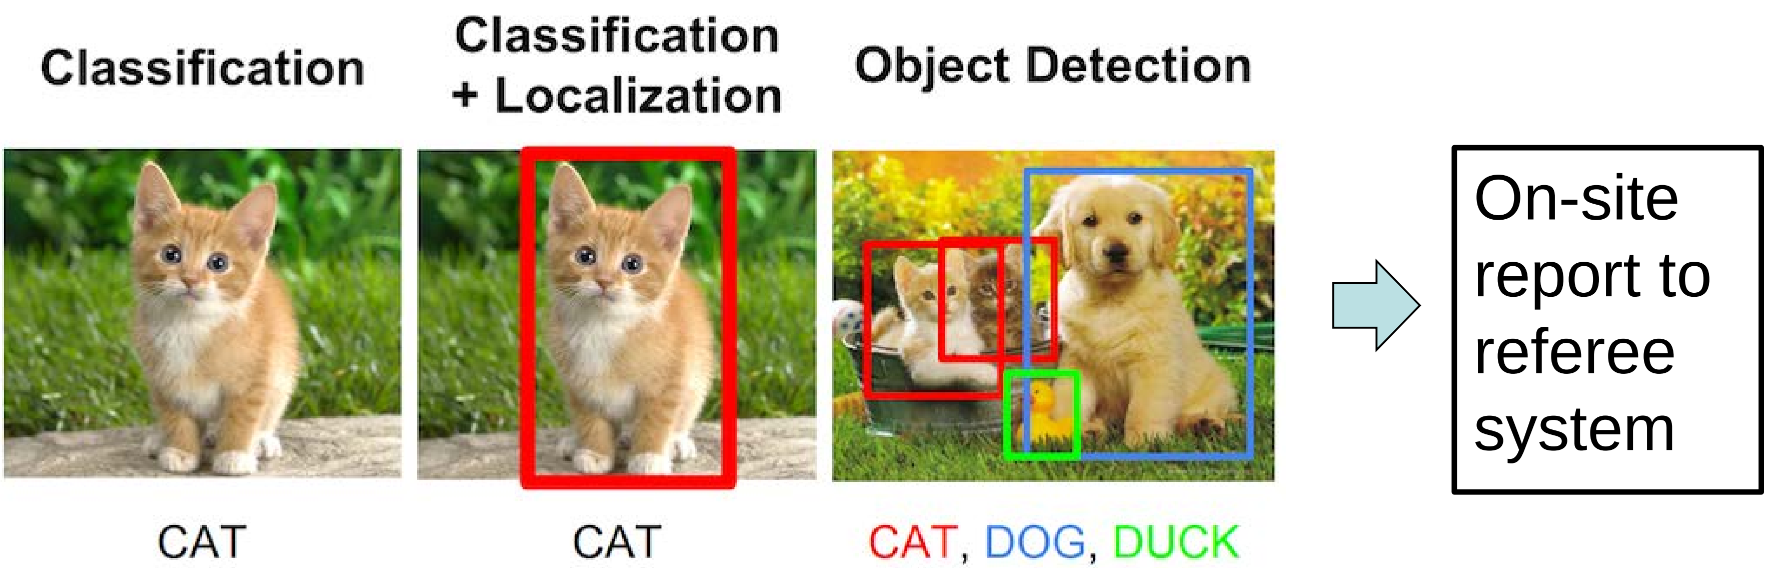

LPIRC rule and metric

- There are total 200 object classes in the competition.
- These object classes are the same as ImageNet's.
- But, contestant's training set is not restricted to ImageNet's.
- **The participants can use any datasets to Train!**

1. airplane
2. apple
3. banana
4. basketball
5. bee
6. bicycle
7. bird
8. bus
9. car
10. chair
11. dog
12. drum
13. hammer
14. laptop
15. orange
16. rabbit
17. snake
18. sofa
19. tennis ball
20. …

200 categories from ILSVRC, http://image-net.org/challenges/LSVRC/2015/browse-det-synsets

LPIRC rule and metric

- The score is the ratio between recognition accuracy and the total amount of energy consumption.
- We adopt mean average precision (mAP) as recognition accuracy.
- Energy consumption part measures the participant's hardware power usage without its own battery. We provide AC/DC power supply.

$$\text{Score} = \frac{\text{mean average precision}}{\text{Energy } consumption}$$

Overview of on-site testing

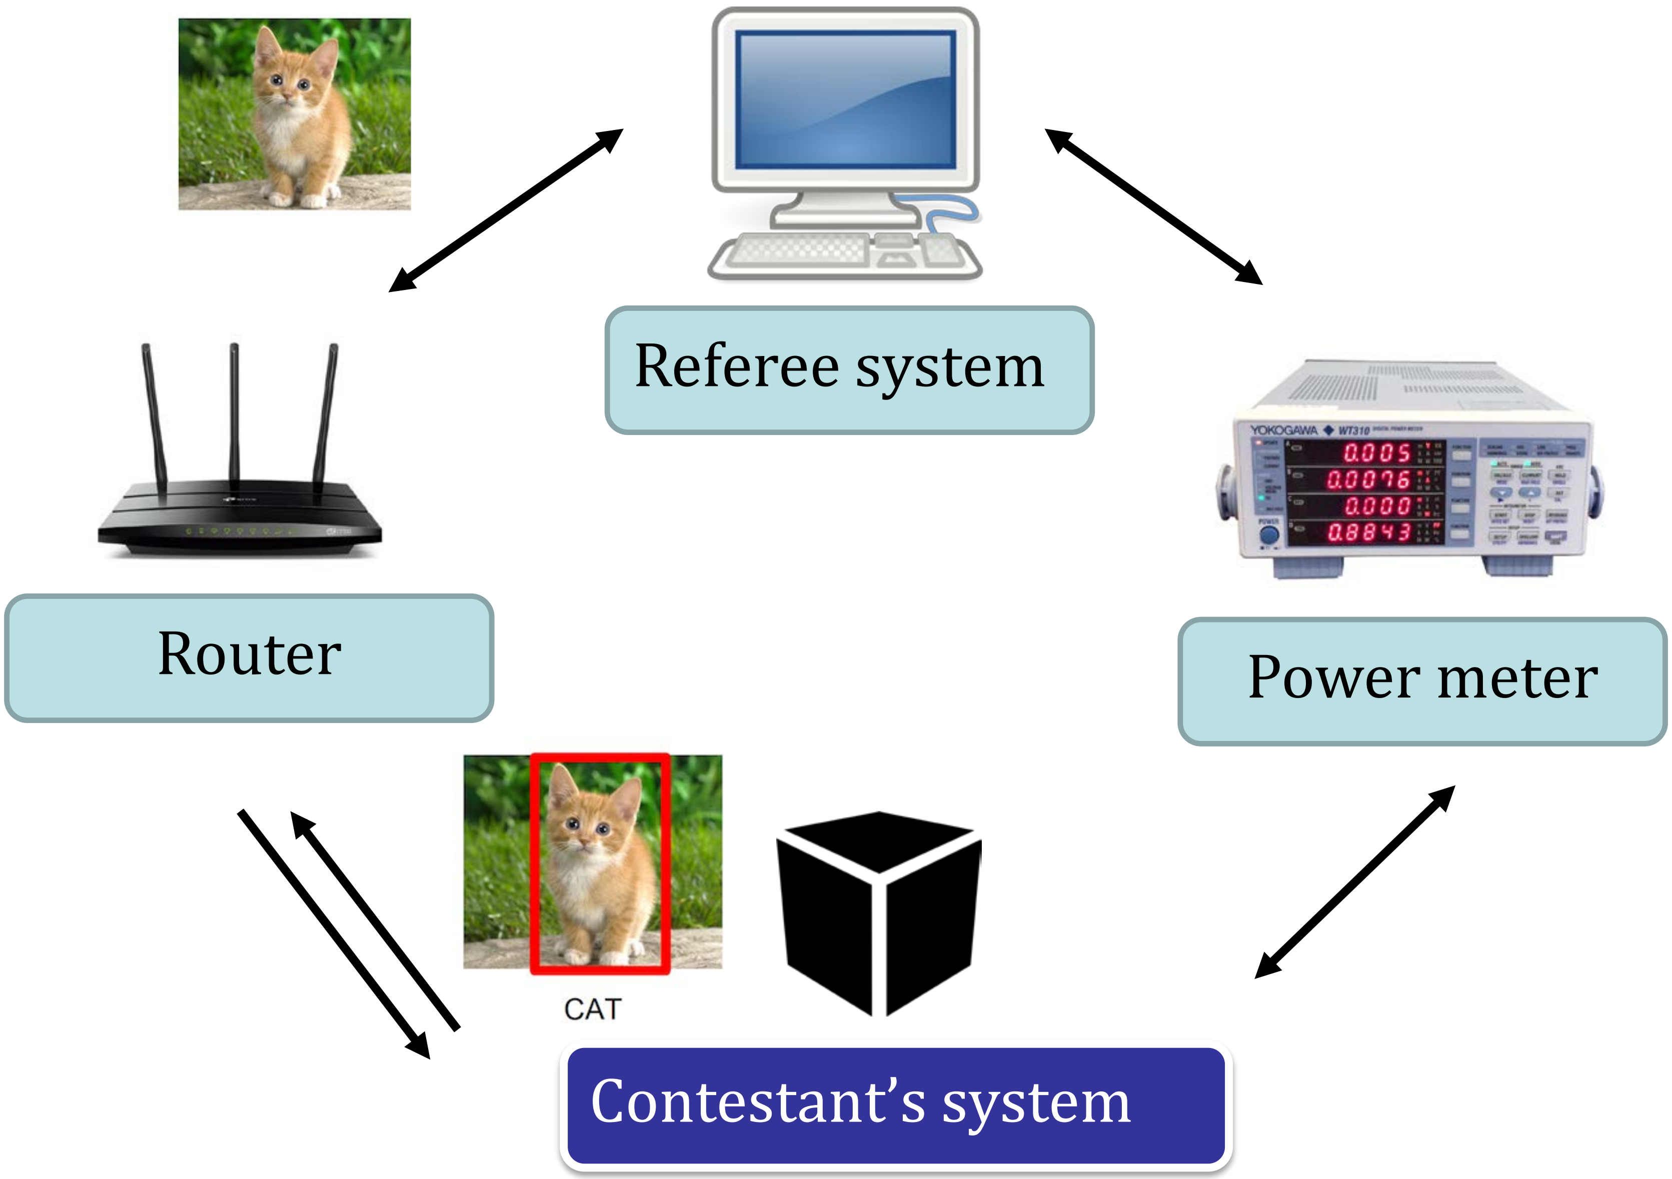

LPIRC rule and metric

Reported Bounding Box

Correct Answer

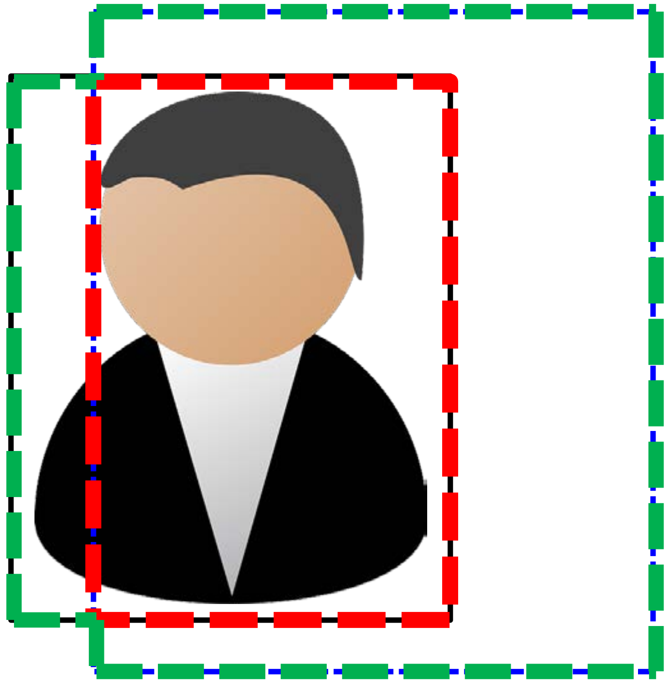

Correct Detection if $\frac{\text{Correct} \cap \text{Reported}}{\text{Correct} \cup \text{Reported}} > 0.5$

Development (LPIRC 2015)

- Low-Power Image Recognition Challenge (LPIRC) June 7 2015 in Austin Texas (with Design Automation Conference)
- 34 participants of 10 teams came from 4 countries.

Development (LPIRC 2015)

- **Competition Tracks:**
- Track 1: Computation offloading to another device or server is not allowed.
- Track 2: Computation offloading is allowed.
- **2015 Winner:** Tsinghua University won the first LPIRC in 2015. The team used Fast R-CNN (Regions with Convolutional Neural Network) as a baseline solution. Additional modifications on the Fast R-CNN method are made to fit the specific platform and achieve trade-off between speed and accuracy on embedded systems. The system used the Jetson TK1 platform and achieved the best score in 2015.
- Only one team participated the 2nd track but was not able to successfully recognize any images.

Development (LPIRC 2016)

- Low-Power Image Recognition Challenge (LPIRC) June 5 2016 in Austin Texas (with Design Automation Conference).

Development (LPIRC 2016)

- **Competition Tracks:**
- Track 1: Computation offloading to another device or server is not allowed.
- Track 2: Computation offloading is allowed.
- Track 3: The team can choose offloading or not. The images are acquired by a camera, rather than as files through the network. only one team participates each time.
- **2016 Winner**: The team from Chinese Academy of Science won 2016 LPIRC. The team evaluated two different object detection architectures (BING+FAST-RCNN and Faster-RCNN) on Jetson TX1. The feature-extracting CNN dominated the speed, accuracy and power consumption of the systems; thus, the team explored the design space to find the most energy effective network implementation.
- The score of this year's winner is more than double the score of last year.

Development (LPIRC 2017)

- LPIRC 2017 Competition in Honolulu, Hawaii, USA on July 21, 2017.
- **Competition Track:**
- <u>Only one track</u>: images were acquired through the network and offloading was disallowed.

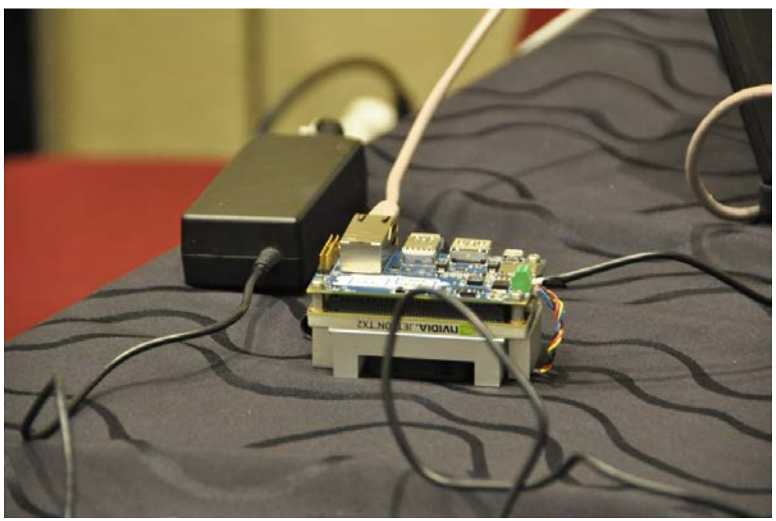

Summary of Development

- The score of 2016 winner is more than double the score of 2015 winner.

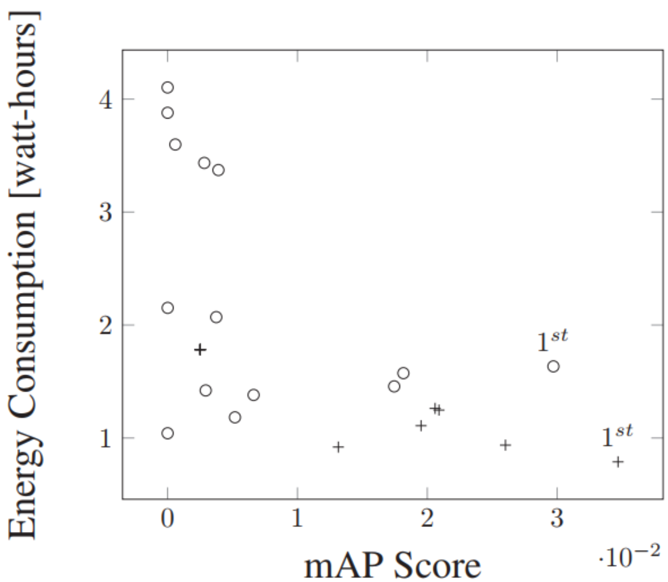

LPIRC 2016 (“+”) and 2015 (“o”) mAP versus power consumption results.

2015 & 2016 Summary		
Year	All Score Mean	Top-5 Score Mean
2016	0.02280	.0245
2015	0.00393	.0102
Improvement	5.80215	2.4020

Summary of results from LPIRC 2016 and 2015

Summary of Development

- The accuracy (mAP) improves 8.53 times from 2015 to 2017.
- The final score improves 6.56 times.

Year	mAP	Energy (WH)	Score	Ratio
2015	0.02971	1.634	0.0182	1.00
2016	0.03469	0.789	0.0440	2.42
2017	0.24838	2.082	0.1193	6.56

- 32 different solutions have been presented since 2015.

Year	Date	Conference	# Solutions
2015	June 7	Design Automation Conference	17
2016	June 5	Design Automation Conference	5
2017	July 21	Computer Vision and Pattern Recognition	10

Summary of Development

- Winners = low energy + high speed + high accuracy
- Winners of the previous three years propose their expectations, which are:
- A system consumes at most 0.1 Watt.
- The system can process more than 100 images (or video frames) at high resolutions (12 MP or higher) per second.
- More than 100 objects appear in each image (or video frame) and these objects belong to 1,000 different categories.
- The recognition accuracy is more than 99.99%.

Successes and Confusions

- **Successes:**
  - The most flexibility: participants tried out a wide range of hardware, including mobile phones, tablets, laptops, FPGA, etc.
- **Confusions:**
  - Whether it is potentially beneficial to process only some images and stop early in order to “game” the evaluation criteria? No

Successes and Confusions

- **Answer of confusions:**
  - The ground truth used for computing mAP includes the entire test set, regardless of how many images a solution processes.
  - No matter how many images the participants' systems processed, mAP is based on all the images in the test set.

$$\text{mAP} = \frac{\text{mean average precision} * throughpht}{\text{number of images in the test set}}$$

$$\text{Score} = \frac{\text{mAP}}{\text{Energy } consumption}$$

Comparison to ILSVRC

- ImageNet Large Scale Visual Recognition Challenge (ILSVRC) has been run annually from 2010 to present, attracting participations from more than fifty institutions.
- Different in three ways:
  - LPIRC is an on-site competition. Contestants must bring their systems to compete. ILSVRC does not require contestants compete on-site.
  - In LPIRC, each solution has 10 minutes. ILSVRC has no time limit.
  - Most important, LPIRC considers the energy consumption.

Comparison to System Design Contest (DAC)

- The Design Automation Conference (DAC) is recognized as the premier conference for design and automation of electronic systems.
- San Francisco, CA. June 24-28, 2018
- Target Platforms
  - Nvidia Jetson TX2
  - PynQ Z-1 board
- Features embedded system implementation of neural network based object detection for drones.
- Two categories
  - GPU
  - FPGA

DAC's Evaluation

The evaluation for the design is based on the accuracy, throughput, and energy consumption.

- Intersection over Union (IoU):

  IoU is an evaluation metric used to measure the accuracy of an object detector on a particular dataset (without threshold).

$$IoU = \frac{Area\ of\ Overlap}{Area\ of\ Union} = \frac{DetectionResult \cap GroundTruth}{DetectionResult \cup GroundTruth}$$

- Throughput: The minimum speed requirement is 20FPS in this competition.
- Energy: Energy consumption for a detector to process all the images.

DAC's Evaluation

- Let R_{IoU_i} be the IoU score for team i:

$$R_{IoU_i} = \frac{\sum_{k=1}^{K} IoU_{i_k}}{K}$$

  - K: total number of images in the dataset
  - I: total number of registered teams

- Let ES_i be the energy consumption score for team i.

$$ES_i = \max\left\{0, 1 + 0.2 * \log_2 \frac{\bar{E}_i}{E_i}\right\}$$

  - E_i: energy consumption for team i.
  - $\bar{E}_i$: average energy consumption of I teams.

- The total score TS_i for team i is:

$$TS_i = R_{IoU_i} * (1 + ES_i)$$

Comparison to 2017 EDLDC (ESWeek)

- Embedded deep learning design contest (EDLDC) is a competition for the energy-efficient design of a deep learning application under a given set of constraints.

- The contest targets object detection which is one of the most popular deep learning applications on embedded systems, e.g., self-driving car, AI secretary, robot, surveillance, etc.

EDLDC Introduction

- Problem description
  - Given a set of images, the participant's design needs to find as many correct bounding boxes of objects as possible.
- Evaluation metric
  - Mean average precision (mAP)
  - Energy consumption
  - Number of processed images.

Tracks and Winners

- Tracks
  - Track A: Any hardware platform having power connector (110V or 220V), except NVIDIA TX2, can be used. The evaluation metric is mAP*(#images in 5min)/energy.
    - 1st Seoul National University, 2nd Samsung Electronics.
  - Track B: Only NVIDIA TX2 is used. The evaluation metric is mAP*(#images in 5min).
    - 1st Seoul National University, 2nd Seoul National University.

LPIRC's Contributions

Focus on both accuracy and energy consumption:

- There exist many possible uses of low-power devices which can detect objects in images. i.e. smartphones, electronic glass, autonomous robots. These systems use batteries and energy conservation is essential.
- LPIRC is the only known competition to consider both accuracy of object detection and power consumption before 2017.



LPIRC's Contributions

Power and time efficienct design

- Energy consumption is a major constraint
  - Limited battery power
  - Reduced RAM
  - Less capable CPU/GPU
- Time is limited
  - Recognition completed in a short time

This year: LPIRC 2018

- LPIRC 2018 Salt Lake City, Utah, June 18, 2018
- https://rebootingcomputing.ieee.org/lpirc

This year : LPIRC 2018

- **Competition Tracks:**
- **Track 1:** Participants are required to use the specific software platform for neural networks. They will submit their **models** in TfLite format before the competition. This track focuses on **accuracy** and **execution time** on a fixed compute platform.
- Participants focus on improving the accuracy of vision solutions within a time budget. The participant with the highest accuracy wins the competition.
- Since this track focuses solely on neural network architectures, it makes the competition accessible to people that do not have expertise on hardware or access to top-tier systems.

This year : LPIRC 2018

- **Track 2:** Participants use the same hardware platform, Nvidia TX2. They can submit their programs before the competition. The organizers will execute the programs, measure the **accuracy and energy consumption** and give feedback to participants for reference.
- Both accuracy and energy are considered: the winner must have <u>high accuracy and low energy consumption</u>. This track encourages system-level improvements (e.g., better cache performance or voltage scaling) that are not possible in the first track.

This year : LPIRC 2018

- **Track 3:** The third track gives participants complete freedom and they can use any hardware platforms.
- This is an on-site competition: participants must bring their systems to CVPR.
- This is **the same competition** as the first three years of L PIRC.

Comparison

	System Design Contest (DAC)		EDLDC (ESWeek)		LPIRC		
	Track 1	Track 2	Track 1	Track 2	Track 1	Track 2	Track 3
metric	$R_{IoU_i} * (1 + ES_i)$		mAP*(#images in 5min)/energy	mAP*(#images in 5min)	(mean average precision)/(energy consumption)		
Object detection	drones		embedded systems		embedded and other systems		
Target platforms	NVIDIA Jetson TX2	PynQ Z-1 board	Any hardware platform except for NVIDIA TX2	NVIDIA TX2	Fixed compute platform	NVIDIA TX2	Any hardware platforms

Future plan

- Expand LPRIC to a broader and bigger prize competition such as XPRIZE
- Stronger connection between researcher and industry. For example,
  - Now we have a track for Tensorflow, we expect to create more tracks for specific frameworks and SDKs to encourage usage.
  - Add more tracks on solving different problems to encourage participants from different areas. For example, add tracks for NLP, trade prediction.
- Showcase of the relevant technologies.

Acknowledgements

- Thanks to the following sponsor.

IEEE CSC Council on Superconductivity

IEEE rebooting COMPUTING

Google

CEDA® IEEE Council on Electronic Design Automation

IEEE GREENICT

See you at CVPR 2018!

Contact: Yung-Hsiang Lu
yunglu@purdue.edu
Yiran Chen
yiran.chen@duke.edu

CVPR 2018
Salt Lake City

Thanks to Dave Cheng, Xin Liu and Jingchi Zhang's efforts for the slides.

Q & A

CVPR 2018

Salt Lake City