Low-Power Image Recognition Challenge

Yung-Hsiang Lu, Alexander C Berg, Yiran Chen

■ *The Low-Power Image Recognition Challenge (LPIRC) has been held annually since 2015. This article summarizes the competition advancements made over the past three years.*

Cameras have become available in many embedded and mobile systems, including vehicles, smartphones, wearable devices, and aerial robots. With the advent of cameras in these systems comes the possibility of detecting objects in the resulting images using the on-board computers. Energy is limited in mobile systems, however, so for this possibility to become a viable opportunity, energy usage must be conservative. The Low-Power Image Recognition Challenge (LPIRC) is the only competition integrating image recognition with low power. LPIRC has been held annually since 2015 as an on-site competition. To encourage innovation, LPIRC has no restriction on hardware or software platforms: the only requirement is that a solution be able to use HTTP to communicate with the referee system to retrieve images and report answers. Each team has 10 minutes to recognize the objects in 5,000 (year 2015) or 20,000 (years 2016 and 2017) images. The score is the ratio of recognition accuracy (measured by mean average precision, MAP) and the total amount of energy consumption (measured by watt-hour). Team are first trained with images from ImageNet. They are then given test images during the competition. Each image may contain one or several objects (such as person, dog, airplane, and table) that belong to 200 predefined categories. To successfully recognize an object in an image, a computer program must correctly determine the category and mark a bounding box around the object. In the first three years of LPIRC, the champions' scores improved significantly, to a ratio of 6.56, as shown in table 1. All teams are from academe; some teams have sponsorships from industry.

Copyright © 2018, Association for the Advancement of Artificial Intelligence. All rights reserved. ISSN 0738-4602

Year	Number of Solutions	mAP	Energy (WH)	Score	Ratio
2015	17	0.02971	1.63405	0.0182	1.00
2016	5	0.03469	0.78910	0.0440	2.42
2017	10	0.24838	2.08188	0.1193	6.56

Table 1. Score Improvement, 2015–2017.

During the three years of LPIRC, participants have brought a wide range of platforms to the competition, including mobile phone, tablet, laptop, desktop, field programmable gate array (FPGA), and experimental boards. The champions of the first three years all used experimental boards with graphics processing units (GPUs). Recognition accuracy is calculated based on the entire set of test data; thus, a winning solution must be able to recognize objects in many images. The 2015 winner used Fast Regions with Convolutional Neural Network (Fast R-CNN) as the foundation running on NVIDIA TK1. The 2016 winner used a convolutional neural network running, together with BING and FAST, on NVIDIA Jetson TX1. The 2017 winner used Tiny YOLO on Jetson TX2. Articles canvassing the LPIRC experience and the solutions of each year's winner are listed in table 2. LPIRC's sponsors include IEEE Rebooting Computing, IEEE GreenICT, the IEEE Council on Superconductivity, the IEEE Council on Electronic Design Automation, NVIDIA, and Xilinx. In 2018, the sponsors included Facebook and Google.

In the first two years, LPIRC had a second track that allowed offloading. A participant's system has two parts: a front end that communicates with the referee system and a back end (another computer or one or more cloud servers) for the computation. Only the power of the front end was measured. In 2015, only one team entered the offloading track, and that team was unable to successfully recognize any image. In 2016, no team entered this track. In 2017, this track was no longer offered. In 2016, a third track was added: the images were acquired by a camera, rather than as files through the network. The purpose was to use the camera to simulate the human eye. Only one team entered this track and the score was substantially lower than the same team's solution using the network. In 2017, this track, too, was abandoned. In 2017, only one track was offered: images were acquired through the network and offloading was disallowed. The scores reported in table 1 are from this main track since it was constant over the three years.

In 2018, LPIRC will include two different tracks. The first track keeps the same rules as the first three LPIRCs: the competition is on-site at CVPR and each team can bring any hardware or software platform. The second track will provide a software development kit for a chosen platform. In this second track, participants submit their solutions before CVPR and their solutions are evaluated in this predetermined platform.[1]

Note

1. More details about LPIRC can be found at rebootingcomputing.ieee.org/lpirc

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 an ACM distinguished scientist, an ACM distinguished speaker, and a member in the organizing committee of the IEEE Rebooting Computing Initiative. Lu is the lead organizer of the Low-Power Image Recognition Challenge, and the chair of the Multimedia Communication Systems Interest Group in the IEEE Multimedia Communications Technical Committee.

Alexander C. Berg is an associate professor in the Department of Computer Science at the University of North Carolina at Chapel Hill. His research concerns computational visual recognition. He has worked on general object recognition in images, action recognition in video, human pose identification in images, image parsing, face recognition, image search, and machine learning for computer vision. He co-organizes the ImageNet Large-Scale Visual Recognition Challenge.

Yiran Chen is an associate professor at Duke University and serves as the codirector of the Duke Center of Evolutionary Intelligence. His research focuses on emerging memory and storage systems, deep learning acceleration and neuromorphic computing, and mobile computing. He is a recipient of an NSF CAREER award and the Outstanding New Faculty Award from the ACM Special Interest Group on Design Automation.

Year	Article
2017	CNN-Based Object Detection Solutions for Embedded Heterogeneous Multicore SoCs. Asia and South Pacific Design Automation Conference.
2017	Low-Power Image Recognition Challenge. Asia and South Pacific Design Automation Conference.
2016	Towards Real-Time Object Detection on Embedded Systems. *IEEE Transactions on Emerging Topics in Computing*, August.
2015	Rebooting Computing and Low-Power Image Recognition Challenge. IEEE/ACM International Conference on Computer-Aided Design.

Table 2. Articles Detailing the LPIRC Experience and Winners' Solutions.