

Low-power image recognition

To accelerate the development of energy-efficient and intelligent machines, Yung-Hsiang Lu and organizers launched a challenge for low-power approaches to image recognition.

In 2012, IEEE Future Directions started a new initiative called Rebooting Computing to think about computer designs for the year 2040 and beyond. In the inaugural summit, co-chair Elie Track asked each attendee to volunteer for one task after the meeting. David Kirk from Nvidia and I suggested the idea of organizing a competition that would require technologies far beyond the capabilities of existing computers. This competition would demonstrate an intelligent machine using only ambient energy such as sunlight, vibrations or wind. The first challenge is to define ‘intelligence’. Around that time a new era of computer vision began; deep neural networks achieved impressive progress in the ImageNet competition. I recruited Alexander Berg, one of the organizers of ImageNet, to design a new competition. We decided to build on ImageNet’s success using image recognition to judge intelligence, adding low energy consumption as a second factor. The score is simply the recognition accuracy divided by energy consumption.

After two years of preparation, the first IEEE Low-Power Image Recognition Challenge (LPIRC; <https://rebootingcomputing.ieee.org/lpirc>) was held in June 2015, co-located with the Design Automation Conference. A contestant’s system is connected to the referee system through an intranet. The contestant’s system issues HTTP GET commands to retrieve images and issues HTTP POST commands to send the answers. Meanwhile, a power meter measures the energy consumption of the system. Each image may contain one or several distinct objects. Each team has ten minutes to recognize objects and mark their locations in the images. The objects belong to 200 predefined categories, such as humans, cars, tables and dogs. For each recognized object, its category and location are reported. The location is marked by a rectangle, called the bounding box, enclosing this object. The bounding box must have at least 50% overlap with the correct answer (also called the ground truth, marked by the organizers).

LPIRC has no restrictions on hardware or software, so it is an onsite competition. Contestants have brought a wide range of systems: laptops, phones, desktop with GPUs, embedded computers, reconfigurable systems, tablets and so on. To encourage



The LPIRC 2017 competitors and organizers. Credit: IEEE

more participation, I recruited Bo Chen and Yiran Chen to help create two new tracks in 2018 with preselected hardware (TensorFlow model for the Pixel 2 XL phone and Nvidia TX2). Contestants can submit their solutions online without the need to travel, and more than 200 solutions were submitted. The onsite competition is still held so that contestants can bring their (possibly proprietary) hardware. The accuracy is measured in the same way as ImageNet (mean average precision, or MAP). In 2015, 5,000 images were used and no team finished all 5,000 images. The champion obtained 0.02971 MAP with an energy consumption of 1.634 Wh; their score was 0.0182. Since 2016, 20,000 images have been used. The 2018 champion finished 20,000 images within ten minutes and obtained 0.1832 MAP with 0.412 Wh energy consumption. Their score was 0.4446, an over 24-fold improvement since 2015. As a reference, the winner of the 2017 ImageNet competition (without time or energy restrictions) had an accuracy of 0.731 MAP.

LPIRC is only the beginning of a long journey toward highly intelligent machines with extremely low power consumption. Detecting objects in images is the first step in understanding the meaning of visual data. LPIRC will retire in 2019 and be replaced with a new, much more challenging

competition for low-power computer vision. Intelligent machines need to understand action, intention, emotion and implication in images or videos. Processors optimized for machine learning have been introduced by several vendors recently and these new systems may appear in future competitions. Significant efforts would be needed to create training and testing data. Meanwhile, the energy consumption must be reduced by 99.9999% before these intelligent machines can run perpetually using ambient energy. IEEE sets the goal of creating new benchmarks for computers in 2040. We have 20 more years to work towards this. □

Yung-Hsiang Lu

*Purdue University, West Lafayette, IN, USA.
e-mail: yunglu@purdue.edu*

Published online: 9 April 2019
<https://doi.org/10.1038/s42256-019-0041-4>

Acknowledgements

LPIRC sponsors include IEEE Rebooting Computing, IEEE Council on Electronic Design Automation, IEEE Council on Superconductivity, IEEE Circuits and Systems Society, IEEE GreenICT, Google, Facebook, Nvidia, Xilinx and Mediatek. Many students from the three universities have contributed to the creation and management of LPIRC. More information is available at <https://rebootingcomputing.ieee.org/lpirc>. Alexander Berg, University of North Carolina at Chapel Hill; Bo Chen, Google; Yiran Chen, Duke University.