https://doi.org/10.1038/s42256-022-00567-4

Low-power object-detection challenge on unmanned aerial vehicles

A design contest for object detection with deep learning on embedded small devices leads to winning hardware–software co-design approaches.

Object detection and tracking is actively explored in various applications such as ambient environment monitoring, precision agriculture and urban planning and, increasingly, on unmanned aerial vehicles (UAVs) to make use of their superior mobility. The task of detecting a single object of interest from captured images can be tackled with deep learning (DL)-based image-processing techniques, but in real-world scenarios involving UAVs, low latency and high throughput are important requirements. Realizing DL-based object detection on UAVs is a representative problem for TinyML, an area that focuses on developing DL algorithms for resource- and power-constrained embedded devices.

In 2018, we founded the System Design Contest (SDC) at the Institute of Electrical and Electronics Engineers (IEEE)/Association for Computing Machinery (ACM) Design Automation Conference (DAC), the flagship conference in the design automation community, which featured a low-power object detection challenge (LPODC) on designing and implementing novel algorithms for object detection on UAVs. Since then, the contest has been held annually, and over 100 teams from more than 10 countries and regions have participated each year. The challenge focuses on UAV applications that have stringent accuracy, real-time detection and power requirements on resource-constrained hardware. In contrast to general computer visual challenges that focus on accuracy, LPODC evaluates overall performance based on a combination of throughput, power and detection accuracy. Three metrics are used[1]: (i) intersection over union (IoU), a metric for object detection accuracy, which is defined as the ratio between the area of the union of the predicted bounding boxes (BB) used to bind the targeting object and the ground-truth BB, and the area of the overlap encompassed by both the predicted BB and the ground-truth BB; (ii) throughput, in frames per second (FPS); and (iii) power consumption, in terms of the average power consumed in the whole evaluation. The final score in LPODC is a combination of the three.

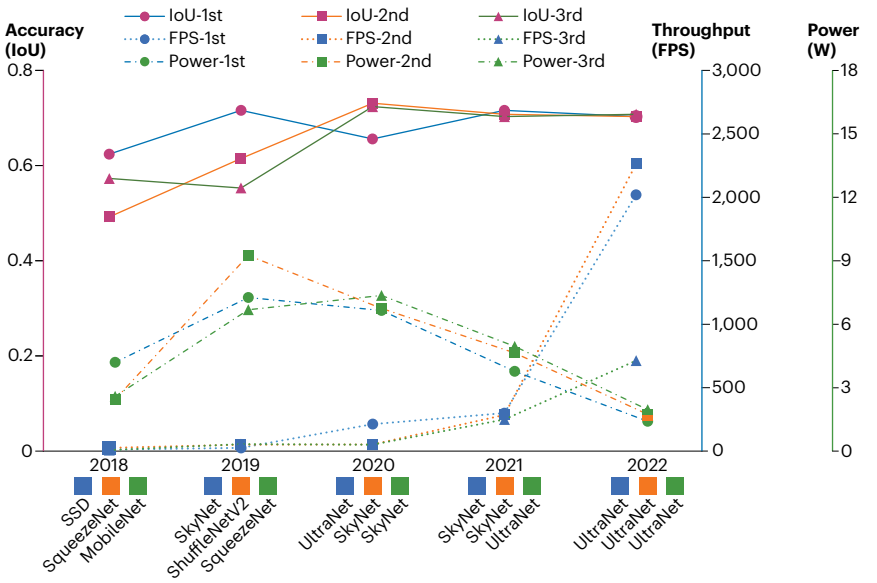

Fig. 1 | IoU, power and throughput performance of each year's top three teams from 2018 to 2022. The neural network backbone used by each team is also included. 1st, 2nd and 3rd denote the first, second and third placed winner team, respectively.

In the first 2 years, LPODC provided two hardware platforms to all participating teams to choose from for their implementations: embedded graphical processing unit (GPU; Jetson TX2 from Nvidia) and system-on-chip field-programmable gate array (SoC FPGA; PYNQ Z-1 board from Xilinx). From 2020 on, only the FPGA SoC platform was used. The images of the dataset, all captured by actual UAVs, reflect real circumstances and challenges encountered in UAV-based applications. The released dataset contains a large quantity of manually annotated training images, while the testing dataset is withheld for evaluation purposes. There are a total of 150,000 images, provided by the industry sponsor, DJI. Participating teams train their models with the training dataset and then send the trained models to the organizers for testing. Evaluations are performed at the end of each month and the details of results and ranking are made public. The final ranking is released at the end of the competition, and the top three entries are invited to present their work at a technical session at DAC.

Figure 1 demonstrates the trend of the performance metrics for each year's top three teams and the corresponding network backbone choice in the past 5 years. Between 2018 and 2020, the detection accuracy and throughput improved dramatically, at the cost of slightly increased power needs. In 2018, the champion, from Tsinghua University in Beijing[2], focused on machine learning model optimization, and downsized the single-shot detection (SSD) network topology by removing the last two convolutional layers. Pruning and quantization of network parameters were also applied. The top team in 2019, from the

University of Illinois at Urbana–Champaign in the United States, started to explore the hardware design space through a bottom-up deep neural network (DNN) model-design strategy together with a top-down flow for accelerator design[3]. Using a lightweight SkyNet as the backbone[4], this enables a joint optimization of both DNN models and their deployment configurations and achieves much higher IoU and more than twice the throughput (>200% of the original value) of previous year's first place team's approaches. In 2020, the top winner, from Beijing University of Technology, used a learnable parameter, soft clipping full-integer quantization (LSFQ), as well as full parallelization of multiplications based on another lightweight architecture, UltraNet[5].

Starting in 2021, little improvement in detection accuracy has been achieved, and the winning strategy has shifted towards hardware efficiency, with a great throughput increase and power reduction seen each year. The 2021 champion, from Shanghai Tech, adopted tuneable activation imbalance transfer (TAIT) for quantization based on SkyNet, and further exploited the power of the hardware through fine-grained multithreading and parallelization[6]. Finally, this year, all the top three teams adopted UltraNet. The top team, from Southeast University in Nanjing, China, optimized the data flow in programmable logic to increase the data computation reuse rate, which eventually approached the theoretical performance boundary of the hardware in terms of throughput[7].

LPODC at SDC-DAC has been a success these last 5 years, and we expect that it will continue to be a premier contest in low-power object detection. LPODC is only the starting point for TinyML, which deploys artificial intelligence on small hardware platforms with constrained resources. As evidenced by the LPODC winners, for TinyML it is crucial to deploy a hardware–software co-design approach[8] rather than optimizing software and hardware separately. Beyond accelerator and neural architecture design, there are also opportunities in communication, compiler or even device optimization[9] in future contests to push forward the field of TinyML and make it usable, reliable and prevalent in daily applications.

The source code from the top three teams in each of the past 5 years can be found on Github[10].

Zhenge Jia[1], Xiaowei Xu[2], Jingtong Hu[3] & Yiyu Shi[1]

[1]Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN, USA. [2]Guangdong Provincial People's Hospital, Guangdong Academy of Medical Sciences, Guangzhou, China. [3]Department of Electrical and Computer Engineering, University of Pittsburgh, Pittsburgh, PA, USA.

e-mail: yshi4@nd.edu

Published online: 25 November 2022

References

1. Xu, X. et al. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 392–403 (2019).
2. Wang, Y. et al. in *Low-Power Computer Vision: Improve the Efficiency of Artificial Intelligence* (Thiruvathukal, G. K., Lu, Y.-H., Kim, J., Chen, Y. & Chen, B.) pp. 35–64 (CRC Press, 2022).
3. Zhang, X. et al. in *36th International Conference on Machine Learning (ICML) Workshop on ODML-CDNNR* (2019).
4. Zhang, X. et al. in *Conference on Machine Learning and Systems (MLSys)* (2020).
5. Bao, Z., Zhan, K., Zhang, W. & Guo, J. in *2021 IEEE Symposium in Low-Power and High-Speed Chips (COOL CHIPS)* https://doi.org/10.1109/COOLCHIPS52128.2021.9410327 (IEEE, 2021).
6. Jiang, W., Yu, H., Liu, X., Sun, H., Li, R. & Ha, Y. in *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pp. 1027–1032 (IEEE, 2021).
7. Guoqing, L. Zhang, M., Li, J., Lv, F. &Tong, G. *Pattern Recognit.* **109**, 107610 (2021).
8. Jiang, W. et al. *IEEE Trans. Comp. Aided Des. Integr. Circuits Syst.* **39**, 4805–4815 (2020).
9. Jiang, W. et al. *IEEE Trans. Comput.* **70**, 595–605 (2021).
10. Goeders, J., Hao, C. & Zhuo, C. Contest Winning Designs; https://byuccl.github.io/dac_sdc_2022/info/ (2022).

Competing interests

The authors declare no competing interests.