

---

# SKYNET: A HARDWARE-EFFICIENT METHOD FOR OBJECT DETECTION AND TRACKING ON EMBEDDED SYSTEMS

---

Xiaofan Zhang<sup>1</sup> Haoming Lu<sup>1</sup> Cong Hao<sup>1</sup> Jiachen Li<sup>1</sup> Bowen Cheng<sup>1</sup> Yuhong Li<sup>1</sup> Kyle Rupnow<sup>2</sup>  
Jinjun Xiong<sup>3,1</sup> Thomas Huang<sup>1</sup> Honghui Shi<sup>3,1</sup> Wen-mei Hwu<sup>1</sup> Deming Chen<sup>1,2</sup>

## ABSTRACT

Object detection and tracking are challenging tasks for resource-constrained embedded systems. While these tasks are among the most compute-intensive tasks from the artificial intelligence domain, they are only allowed to use limited computation and memory resources on embedded devices. In the meanwhile, such resource-constrained implementations are often required to satisfy additional demanding requirements such as real-time response, high-throughput performance, and reliable inference accuracy. To overcome these challenges, we propose SkyNet, a hardware-efficient neural network to deliver the state-of-the-art detection accuracy and speed for embedded systems. Instead of following the common top-down flow for compact DNN (Deep Neural Network) design, SkyNet provides a bottom-up DNN design approach with comprehensive understanding of the hardware constraints at the very beginning to deliver hardware-efficient DNNs. The effectiveness of SkyNet is demonstrated by winning the competitive System Design Contest for low power object detection in the 56th IEEE/ACM Design Automation Conference (DAC-SDC), where our SkyNet significantly outperforms all other 100+ competitors: it delivers 0.731 Intersection over Union (IoU) and 67.33 frames per second (FPS) on a TX2 embedded GPU; and 0.716 IoU and 25.05 FPS on an Ultra96 embedded FPGA. The evaluation of SkyNet is also extended to GOT-10K, a recent large-scale high-diversity benchmark for generic object tracking in the wild. For state-of-the-art object trackers SiamRPN++ and SiamMask, where ResNet-50 is employed as the backbone, implementations using our SkyNet as the backbone DNN are 1.60X and 1.73X faster with better or similar accuracy when running on a 1080Ti GPU, and 37.20X smaller in terms of parameter size for significantly better memory and storage footprint.

## 1 INTRODUCTION

Edge AI applications require not only high inference accuracy from DNNs, but also aggressive inference speed, throughput, and energy efficiency to meet real-life demands. These applications rely on hardware-efficient DNN designs when they are deployed into embedded systems with extremely limited computation and memory resources. Recently, we have seen intensive studies on DNN accelerators in hardware, which attempt to take advantage of different hardware design styles, such as GPUs, FPGAs and ASICs, to improve the speed and efficiency of DNN inference and training processes (Qiu et al., 2016; Chen et al., 2016; Zhang et al., 2017a; Jouppi et al., 2017; Franklin, 2017; Zhang et al., 2018a; Li et al., 2019b; Chen et al., 2019).

Although hardware accelerators can be helpful, they are still limited by available resources to handle varied real-life applications, especially for embedded systems since most DNNs are not originally designed to be hardware-efficient. As a result, developers started to focus their optimization

efforts on the software side, to compress DNNs for less complexities, lowering computation demands and memory footprints. Recent researches have demonstrated the possibility of using low bit-width data to represent original floating-point parameters, such as using binary and ternary networks (Courbariaux et al., 2016; Rastegari et al., 2016; Li et al., 2016; Tschannen et al., 2018; Wang et al., 2018a; Gope et al., 2019). These solutions are intended to replace the hardware-intensive floating-point multiplications by logical operations, so that DNNs can be more efficient on hardware platforms.

Researchers also investigate the network pruning strategies to reduce the redundancy of DNN structures (Han et al., 2015; 2016; Luo et al., 2017). According to the published pruning strategies, the relatively less important connections between DNN layers are discarded and network retraining is then performed to regain accuracy. Significant reductions can be achieved on the classic DNNs, such as AlexNet (Krizhevsky et al., 2012) and VGG-16 (Simonyan & Zisserman, 2014). Since the major benefit of network compression comes from the fully-connected (FC) layers, to continuously have effective pruning results for latter DNNs (e.g., GoogleNet (Szegedy et al., 2015) and ResNet (He et al., 2016)) with reduced FC layers, more sophisticated algorithms are required to be integrated in network pruning, such as using evolutionary algorithms (Dai et al., 2019),

---

<sup>1</sup>IBM-Illinois Center for Cognitive Computing Systems Research (C3SR), University of Illinois at Urbana-Champaign, USA  
<sup>2</sup>Inspirit IoT, Inc, USA <sup>3</sup>IBM Research, USA. Correspondence to: Xiaofan Zhang <xiaofan3@illinois.edu >.

Open-source code can be found at [this link](#).

*Proceedings of the 3<sup>rd</sup> MLSys Conference*, Austin, TX, USA, 2020. Copyright 2020 by the author(s).

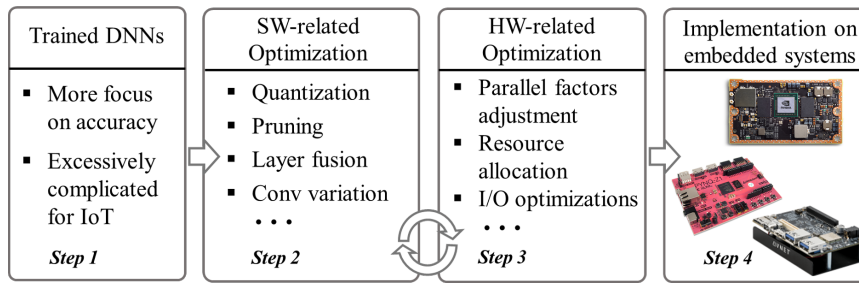


Figure 1. A top-down design flow for hardware-efficient DNN deployment on resource-constrained embedded systems. Challenges appear between *step 2* and *3* where iterative explorations are necessary to balance DNN accuracy and performance on targeted devices.

alternating direction method of multipliers (ADMM) (Ren et al., 2019), and iterative pruning (Ding et al., 2018).

As most of the computations happen inside the convolutional (Conv) layers, previous works also attempt to reduce the computation complexity by using depth-wise separable Conv layers for image classification and ubiquitous keyword-spotting applications (Howard et al., 2017; Zhang et al., 2017b). The depth-wise separable structure can effectively reduce the number of operations and provide more compact DNN designs for resource-constrained hardware. To further improve the DNN deployment on hardware, layer fusion is proposed in (Alwani et al., 2016) to minimize data movements between on-chip and off-chip memory.

In general, a traditional design process for hardware-efficient DNNs can be summarized in Figure 1 with the adoption of above-mentioned technologies. It is a top-down design flow which starts from *step 1*: to select a reference DNN by concentrating on accuracy. Such DNNs are typically too complicated for targeted embedded systems and need to be compressed using software and hardware optimizations in *step 2* and *3*, respectively. Since software compression and hardware implementation are typically carried out in two separate steps, *step 2* and *3* are usually performed in an iterative manner to balance DNN accuracy and hardware performance on targeted devices. Network retraining is also required to regain accuracy after compression before *step 4*. Because of the iterative nature of the process, it is very challenging to cover both inference accuracy in software and deployment efficiency in hardware.

In this paper, we address the hardware-efficient DNN design problem by proposing SkyNet, which is designed following a bottom-up DNN design approach with comprehensive awareness of hardware constraints. The main contributions of this paper are summarized as follows:

- We survey the latest low power object detectors for embedded systems and identify the potential obstacles of using top-down DNN design flows, which may prevent improved DNN accuracy and hardware efficiency.
- We propose a bottom-up design strategy of hardware-efficient DNNs for both embedded GPU and embedded FPGA; using such a design method, we propose

SkyNet, which has comprehensive awareness of hardware limitations to overcome the challenges of top-down design flow.

- We demonstrate SkyNet in DAC-SDC’19 using both TX2 GPU and Ultra96 FPGA with the state-of-the-art accuracy. SkyNet achieved the highest overall score regarding accuracy, throughput, and energy-efficiency, and won the first place winner award for both GPU and FPGA tracks.
- We extend SkyNet for object tracking. By using SkyNet as the backbone DNN, SiamRPN++ and SiamMask obtain 1.60X and 1.73X speedup with better or similar accuracy, and 37.20X smaller parameter size compared to using the original ResNet-50 backbone when running on a 1080Ti GPU.

## 2 RELATED WORK

Recent state-of-the-art object detectors feature DNN backbones to extract input features. Researchers initially propose a two-stage approach which first outputs multiple region proposals for object candidates and then generates more accurate regions with corresponding class labels (Dai et al., 2016; Lin et al., 2017a; He et al., 2017; Cheng et al., 2018a;b; Li et al., 2019c). To improve the detection speed, some one-stage approaches are proposed to simultaneously regress object locations and classes (Sermanet et al., 2014; Redmon et al., 2016; Liu et al., 2016; Lin et al., 2017b; Shen et al., 2019; Tian et al., 2019). Object tracking also relies on the features extracted from DNN backbones, and we have seen recent Siamese network based trackers formulate trackers as feature between the exemplar image and search region (Tao et al., 2016; Valmadre et al., 2017; Wang et al., 2018b; Li et al., 2019a; Wang et al., 2019). These emerging methods make real-time object detection and tracking possible using desktop GPUs but still need aggressive compression before deploying onto embedded systems.

### 2.1 Low-Power Object Detectors

Nowadays, much attention has been paid to delivering hardware-efficient designs for object detection instead of simply pursuing higher inference quality. To address the design difficulties of real-life applications, a low power object detection challenge in DAC-SDC is proposed to target

Table 1. DAC-SDC winning entries from both GPU and FPGA tracks. They follow a top-down approach, from choosing reference DNNs to applying optimization strategies on software and hardware sides, so that they compress DNNs with improved hardware efficiency. Optimizations include: ① input resizing, ② network pruning, ③ data quantization, and ④ TensorRT (Vanholder, 2016) on software, and ⑤ CPU-FPGA task partition, ⑥ double-pumped DSP, ⑦ fine-grained pipeline, ⑧ clock gating, and ⑨ multithreading on hardware.

RANK	GPU-TRACK	REFERENCE DNN	OPTIMIZATIONS	
'19 2ND	THINKER (XIONG ET AL., 2019)	SHUFFLENET + RETINANET	① ② ③	⑨
'19 3RD	DEEPZS (DENG ET AL., 2019)	TINY YOLO	NOT CLEAR	⑨
'18 1ST	ICT-CAS (LU ET AL., 2018)	TINY YOLO	① ② ③ ④	NOT CLEAR
'18 2ND	DEEPZ (DENG & ZHUO, 2018)	TINY YOLO	NOT CLEAR	⑨
'18 3RD	SDU-LEGEND (ZANG ET AL., 2018)	YOLOV2	① ② ③	⑨
RANK	FPGA-TRACK	REFERENCE DNN	OPTIMIZATIONS	
'19 2ND	XJTU TRIPLER (ZHAO ET AL., 2019)	SHUFFLENETV2 + YOLO	② ③	⑤ ⑥ ⑧
'19 3RD	SYSTEMSETHZ (KARA & ALONSO, 2019)	SQUEEZE NET + YOLO	① ② ③	⑦
'18 1ST	TGIIF (ZENG ET AL., 2018)	SSD	① ② ③	⑤ ⑥
'18 2ND	SYSTEMSETHZ (KARA ET AL., 2018)	SQUEEZE NET + YOLO	① ② ③	⑦
'18 3RD	ISMARK2 (HAO ET AL., 2018)	MOBILENET + YOLO	① ② ③	⑤ ⑦

unmanned aerial vehicle (UAV) applications using embedded platforms, such as NVIDIA TX2 GPU, Ultra96 FPGA, and Xilinx Pynq-Z1 FPGA (Xu et al., 2019). By examining the winning entries, we notice that all of them adopt one-stage detectors and share similar top-down DNN design approaches in Figure 1. As shown in Table 1, most of them start from well-established hardware-efficient DNNs, such as ShuffleNet (Zhang et al., 2018b), SqueezeNet (Iandola et al., 2016), and MobileNet (Howard et al., 2017), and replace the image classifier with YOLO (Redmon et al., 2016; Redmon & Farhadi, 2017) or RetinaNet (Lin et al., 2017b) back-end for object detection. Other solutions directly adopt the object detection algorithms, such as SSD (Liu et al., 2016) and YOLO. To deliver hardware-efficient DNNs, they employ input resizing and network pruning to lower the network complexity. Some of the GPU entries use half-precision data format (16-bit) and TensorRT for improved throughput. More aggressive compression is found in FPGA designs because of even tighter resource budgets. DNN parameters are quantized to around 8 bits and some even down to 1 bit. They also cover task partitioning (between host CPU and FPGA), double-pumped DSP (with doubled working frequency in DSP units), tailored pipeline, multithreading, and clock gating to boost hardware performance and energy-efficiency.

## 2.2 Hardware-Aware Neural Network Search

To deliver DNNs for edge devices, there have been growing interests in using neural architecture search (NAS) (Tan et al., 2019; Wu et al., 2019; Howard et al., 2019; Cai et al., 2018) to automatically find resource constrained DNN targeting edge-platforms. To find the efficient networks for a specific platform, (Tan et al., 2019) uses real-time latency by running models on the targeted device instead of latency proxy. Limited by the number of available physical devices, (Wu et al., 2019; Cai et al., 2018) use look-up-table (LUT) to approximate the run-time of models on a specific device.

To incorporate human knowledge, (Howard et al., 2019) uses platform-aware NAS to search DNNs for a platform and manually adjust the structure to make it more efficient. Compared to previous hardware-aware NAS methods that target a specific platform, SkyNet can target both embedded GPU and embedded FPGA platforms and capture hardware limitations by using the realistic hardware performance feedback instead of using LUT approximation.

## 3 MOTIVATIONS

To deliver an even better solution, we investigate the potential obstacles in the top-down design flow (Figure 1) which may hinder further improvements on DNN accuracy and efficiency. We summarize two challenges as follows:

- 1) It is difficult to balance the sensitivities of DNN configurations on software and hardware during model compression following the top-down approach.
- 2) It is difficult to select appropriate reference DNNs at the very beginning of the top-down flow because of the uncertain accuracy variations for a given real-life task.

The first challenge causes tedious iterative explorations between software and hardware optimizations. With the similar hardware performance (e.g., throughput and latency), DNNs may have different accuracy results as the compression technologies are applied to different network components. We take data quantization as an example. In Figure 2 (a), the accuracy results vary significantly between parameter and intermediate feature map (FM) quantization. In this figure, the coordinates of the bubble center represent accuracy and model compression ratio, while the area of a bubble shows data size in megabyte (MB). We scale-up the FM bubble for better graphic effect. By compressing the model from Float32 to fixed point, we reduce 22X parameter size (237.9MB→10.8MB) and 16X FM size (15.7MB→0.98MB), respectively. In this case, inference accuracy is more sensitive to the FM precision.




Figure 2. (a) Accuracy results for the same AlexNet under different compression schemes: blue for parameter compression and green for FM compression. The legend shows the quantization details: each quantized model (#2~#5) is denoted as bit precision  $p_1$  for FMs across all layers,  $p_2$  for 1<sup>st</sup> Conv parameters,  $p_3$  for the parameters in 2<sup>nd</sup> ~ 5<sup>th</sup> Convs,  $p_4$  for the parameters in 1<sup>st</sup> ~ 2<sup>nd</sup> FCs, and  $p_5$  for parameters in the 3<sup>rd</sup> FC in  $p_1$ - $p_2$  $p_3$  $p_4$  $p_5$  format. (b) BRAM usages of accelerators with the same architecture but 12~16-bit quantization for feature maps (FM12~FM16) and different image resize factors. (c) DSP utilization of accelerator using different quantizations between weights (W) and feature maps (FMs) with the numbers indicating bits allocated.

Table 2. Accuracy on DAC-SDC dataset using ResNet, VGG, and SkyNet backbones and the same back-end for object detection.

BACKBONE DNN	# OF PARAMETER	IOU
RESNET-18	11.18M	0.61
RESNET-34	21.28M	0.26
RESNET-50	23.51M	0.32
VGG-16	14.71M	0.25
SKYNET (OURS)	<b>0.44M</b>	<b>0.73</b>

On the other hand, DNNs with similar accuracy may cause differences in hardware. To provide a quantitative analysis, Figure 2 (b) shows the BRAM (on-chip memory in FPGA) usages with different input sizes and FM quantizations. By reducing the resize factor from 1.00 to 0.78, we can maintain nearly the same DNN accuracy (<1.0% drop), but save half memory when the factor is smaller than 0.9. Similarly, Figure 2 (c) indicates small changes may lead to diverse DSP utilization. By taking the 16-bit FM (FM16) as an example, the required DSPs reduce from 128 to 64 when weights are changed from 15-bit (W15) to 14-bit (W14).

For the second challenge, it is difficult to select a reference DNN with relatively high accuracy upper bound on a given task. The DNNs with impressive accuracy on published datasets (e.g., CIFAR-10/100 and ImageNet) may not be always suitable. We evaluate the accuracy of popular DNNs on DAC-SDC object detection dataset and list the results in Table 2. With the same box regression part, these DNNs show no clear connection between their parameter size and inference accuracy after adequate training. Thus, it is not easy to select a promising reference model for a given task.

## 4 A BOTTOM-UP DESIGN APPROACH

Motivated by the discussed challenges in Section 3, we propose a bottom-up approach to leverage the hardware-efficient DNN design for embedded systems. It is a three-stage approach as shown in Figure 3.

### 4.1 Stage 1: Bundle Selection and Evaluation

This flow starts with building the hardware-aware basic blocks, called Bundles. From a software perspective, a Bundle is a set of sequential DNN layers, which can be repeatedly stacked for constructing DNNs. From the hardware perspective, a Bundle is a set of IPs to be implemented on hardware. To capture the hardware constraints, Bundles need to be evaluated on targeted embedded systems for collecting realistic latency (for both FPGA and GPU) and resource utilization (for FPGA) results.

In the first stage, we enumerate DNN components (such as Conv, pooling, activation layers, etc.) and assemble them into Bundle 1 ~  $n$ . Each Bundle is then implemented and evaluated in targeted hardware devices for hardware performance metrics. To get their potential accuracy contribution, we build DNN sketches with fixed front- and back-end structures based on given tasks, and respectively insert one type of Bundle (with replications) in the middle. We limit one type of Bundle for one DNN sketch to guarantee its hardware efficiency. Then, DNN sketches are fast trained using targeted datasets to find out the ones with relatively high accuracy. By targeting the object detection task, for example, we can concatenate a input resizing unit (front-end) and a bounding box regression (back-end) with the selected Bundle to build a DNN sketch. The number of training epochs may vary from different datasets as a 20-epoch-training can distinguish sketches using the DAC-SDC dataset (with 100K images), while 5 epochs are enough if using Cifar-10 dataset. We have also seen similar strategies in (Jiang et al., 2019) distinguish candidates by a 25-epoch-training on a subset of ImageNet. At last, the most promising Bundles located in the Pareto curve are selected for the next stage.

### 4.2 Stage 2: Hardware-Aware DNN Search

During DNN search, the inputs include the software and hardware metrics (e.g., DNN accuracy and throughput per-

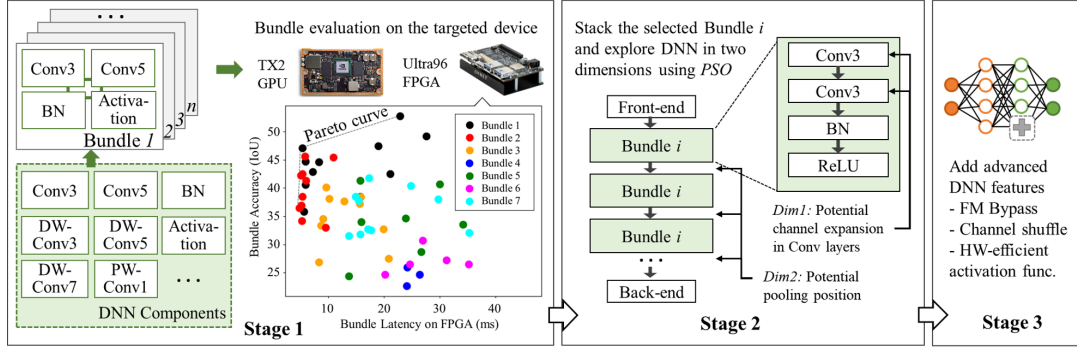


Figure 3. The proposed bottom-up DNN design flow to deliver hardware-efficient DNNs for embedded systems in three stages.

formance) and the targeted hardware platforms while the outputs are DNN candidates which meet the software and hardware requirements. To solve such a multi-objective optimization problem, we propose a group-based particle swarm optimization (PSO) evolutionary algorithm to discover proper DNN candidates since literature has demonstrated the validity of using evolutionary methods to discover DNNs with state-of-the-art accuracy (Real et al., 2019; Elsken et al., 2019). From the design methodology perspective, SkyNet can be extended to support other optimization algorithms and meet the needs of different scenarios.

In the proposed group-based PSO algorithm, each individual DNN is regarded as a particle, and all active DNNs during the search contribute to the swarm. Since we only use one type of Bundle in each DNN, DNNs composed by the same type of Bundle are considered as a particle group. In order to maintain evolution stability, a DNN only evolves within its own group. We label the group optimal position as  $P_{group}^i$  within the  $i$ -th group, meaning such DNN has the best fitness value evaluated under given conditions. We denote a DNN particle  $j$  within group  $i$  as  $n_j^i$  and each  $n_j^i$  has a pair of feature vectors ( $fv1$ ,  $fv2$ ) to illustrate two hyper-parameters regarding DNN structure.  $fv1$  represents the number of channels of each Bundle replication; and  $fv2$  describes the pooling position between Bundles. Both feature vectors with dimension equal to the number of stacked Bundles in  $n_j^i$ , and both of them affect accuracy and hardware performance. To locate the best DNN candidates, we propose Algorithm 1 with the following major components:

**Population generation.** An initial network population  $P$  (a set of DNN candidates) is generated with  $\mathcal{M}$  groups and  $\mathcal{N}$  networks for each group. The search contains  $\mathcal{I}$  iterations and in the  $itr$ -th iteration, all networks are fast trained for  $e_{itr}$  epochs, where  $e_{itr}$  increases with  $itr$ .

**Latency estimation.** We perform a platform-specific latency estimation. For GPUs, we directly measure the inference latency on the training GPU, and scale latency to the targeted GPU for deployment if the target GPU is different from the training one. For FPGAs, we follow a predefined IP-based DNN accelerator template (Hao et al., 2019) for

hardware performance evaluation. Layer-specific IPs are implemented in hardware and shared by corresponding DNN layers. To maximize the performance, IPs are configured to fully consume the available resources. We then collect the end-to-end performance and resource overhead of each DNN from an FPGA high level synthesis tool.

**Fitness value.** After network training and latency estimation, we calculate the fitness value for each network  $n_j^i$  as:

$$Fit_j^i = Acc_j^i + \alpha \cdot (Est(n_j^i) - Tar) \quad (1)$$

where  $Acc_j^i$  is the validation accuracy of  $n_j^i$  and  $Est(n_j^i)$  represents the latency on hardware;  $Tar$  is the targeted latency. Parameters  $\alpha$  ( $\alpha < 0$ ) is used to balance between network accuracy and hardware performance.

**Velocity calculation and particle update.** In standard PSO, the updated velocity of a particle is calculated every iteration based on the current velocity, the velocities toward the local and the global best positions. Particles can move to a better position with assigned probabilities following the updated velocity. Similarly, in our case, DNNs in the same group update their positions (meaning network structures represented by feature vectors) based on the current design, the local best design (the best one across all passing iterations), and the group best design. To determine the velocity toward the local best  $V_{local}$  and the group best  $V_{group}$ , we compute the differences between positions of current and the local/group best designs. Since each position is represented by ( $fv1$ ,  $fv2$ ), position differences can be captured by the mismatch of layer expansion factors  $fv1$  and pooling spots  $fv2$ , respectively. Then, with the velocities known, we start evolving the current network by updating its position toward the local and the group best by a random percentage.

### 4.3 Stage 3: Feature Addition

More advanced DNN design features are added if hardware metrics allow. For example, we can include a bypass from low-level features to high-level features along with feature map reordering (Redmon & Farhadi, 2017) to improve small object detection. We can also replace ReLU with ReLU6 (Sandler et al., 2018) to enhance hardware efficiency. More discussions are provided in the next section.

**Algorithm 1** The bottom-up DNN design with PSO

```

P ← InitialPopulation(M, N)
while itr < I do
  FastTraining(P, eitr)
  Fitji ← GetFitnessVal(P) //evaluate all candidates
  for each group i do
    GroupRank(i) //rank candidates in group i
    Ngroupi ← GroupBest(i) //select the best one in group i
    //get the group best position
    Pgroupi(fv1, fv2) ← GetPosition(Ngroupi)
    for each candidate nji(itr) in group i do
      //rank nji across all passing iterations
      LocalRank(i, j)
      Nlocalij ← LocalBest(i, j)
      //get the local best position
      Plocalij(fv1, fv2) ← GetPosition(Nlocalij)
      //get the current position
      Pji(fv1, fv2) ← GetPosition(nji(itr))
      //get the velocity toward the local and the group best
      Vlocal ← GetV(Pji, Plocalij)
      Vgroup ← GetV(Pji, Pgroupi)
      nji(itr + 1) ← Evolve(nji(itr), Vlocal, Vgroup)
    end
  end
end
end

```

## 5 SKYNET

### 5.1 SkyNet Architecture for object detection

Following the proposed flow, the best Bundle is selected as a combination of  $3 \times 3$  depth-wise Conv layer (DW-Conv3 (Howard et al., 2017)),  $1 \times 1$  point-wise Conv layer (PW-Conv1), batch normalization layer (BN (Ioffe & Szegedy, 2015)), and ReLU6. By repeatedly stacking this Bundle, we generate three backbones shown in Table 3 for object detection in DAC-SDC. These networks share the same chain structure but with different configurations of feature map bypass. For model A, no bypass is included; while for the model B and C, output feature maps of Bundle #3 are fed into the Bundle #6. SkyNet also adapts the YOLO detector head by removing the classification output and use two anchors for bounding box regression.

### 5.2 Feature Map Bypass, Reordering, and ReLU6

By examining the DAC-SDC training data, we keep a record of the size ratio between the output bounding box and the input image and present a distribution diagram in Figure 6. It clearly shows that 91% of the objects to be detected are less than 9% of the original input image size, and 31% of them are even smaller than 1% of the input image size. It means the majority of objects inside this dataset are small objects and we need to provide additional DNN features accordingly. So, we add feature map bypass and reordering to enhance the ability of detecting small object (model B and C). The bypass helps to keep small object features in the later part (closer to the output layer) of the DNN by adding low-level high-resolution feature maps. Also, it is beneficial

Table 3. The SkyNet architecture with number of channels shown in the bracket. Each convolutional layer except the last one is followed by a BN and a ReLU (omitted for conciseness).

CONFIGURATIONS OF SKYNET			
A	B	C	BUNDLE
INPUT ( $3 \times 160 \times 360$ COLOR IMAGE)			
	DW-CONV3 (3)		#1
	PW-CONV1 (48)		
2x2 MAX-POOLING			
	DW-CONV3 (48)		#2
	PW-CONV1 (96)		
2x2 MAX-POOLING			
	DW-CONV3 (96)		#3
	PW-CONV1 (192)		
	[Bypass Start] FM Reordering (768)		
2x2 MAX-POOLING			
	DW-CONV3 (192)		#4
	PW-CONV1 (384)		
	DW-CONV3 (384)		#5
	PW-CONV1 (512)		
	[Bypass End] FM Concatenated	[Bypass End] FM Concatenated	#6
	DW-CONV3 (512+768)	DW-CONV3 (512+768)	
	PW-CONV1 (48)	PW-CONV1 (96)	
PW-CONV1 (10)			
BACK-END FOR BOUNDING BOX REGRESSION			

to have multiple feature maps (from different layers) before generating the bounding boxes. Since the bypass crosses a pooling layer (highlighted in Figure 4), we use reordering (shown in Figure 5) to align the size of original feature map (generated by the Bundle #5) and the low-level feature without losing information. The other feature to improve hardware efficiency is the ReLU6, which clips output range to  $[0, 6]$ . Since ReLU6 generates much smaller data range compared to the original ReLU ( $[0, +\infty)$ ), less bits are required to represent intermediate FMs. It also helps to better implement lower-precision floating point in embedded GPUs and fixed-point data format in embedded FPGAs.

## 6 EXPERIMENTS ON DAC-SDC

DAC-SDC features a single object detection challenge for embedded systems, which include embedded GPUs (NVIDIA TX2) and FPGAs (Pynq-Z1 and Ultra96) with very low energy consumption. The goal is to consider the most appropriate needs of UAV applications, such as capability of real-time processing, energy efficiency, and detection accuracy. To better reflect real-life challenges, images of the dataset are captured by UAVs in the real environment. The whole dataset is divided by two parts: the training dataset with 100,000 images with objects of interest across 12 main categories and 95 sub-categories, and the hidden test set for official evaluation with 50,000 images that only the contest organizers could access (DJI, 2018). Results generated by SkyNet are shown in Figure 7. In DAC-SDC'19, 52 GPU teams and 58 FPGA teams participated worldwide creating

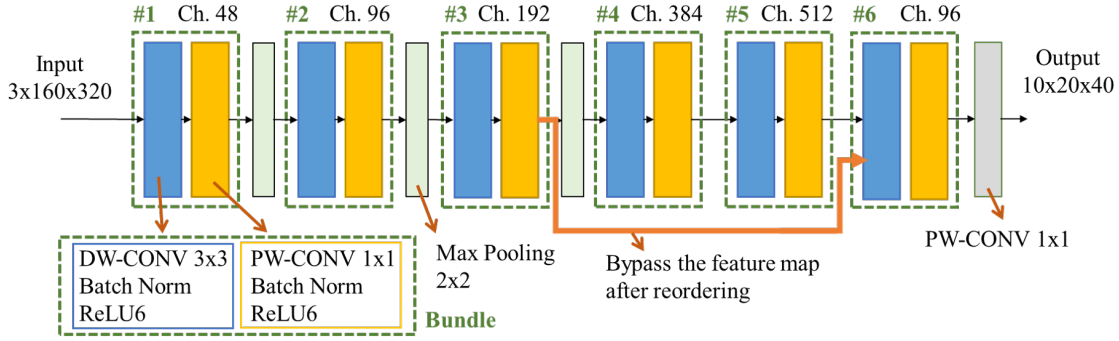


Figure 4. SkyNet backbone (model C in Table 3) generated by stacking six of the selected Bundle (circled by green dashed line) with DNN components as: DW-Conv3, PW-Conv1, BN, and ReLU6. The number of output channels is listed on top of each Bundle denoted as Ch. Three  $2 \times 2$  pooling layers are inserted. The bypass is highlighted in orange, which passes feature maps generated by the Bundle #3 directly to the last Bundle. The feature map reordering is also performed along with the bypass.




Figure 5. Feature map reordering from  $1 \times 4 \times 4$  to  $4 \times 2 \times 2$  with shrunken width and height but expanded number of channels. There is no information loss compared to pooling operation. In addition, this reorder pattern also ensures larger receptive field.




Figure 6. The distribution of bounding box relative size in DAC-SDC training dataset. We capture the bounding box relative size by computing the ratio of output bounding box size divided by the input image size. The green bars show the ratio distribution, and the blue curve shows the corresponding cumulative distribution.

a very intense competition. Our SkyNet design has successfully delivered the best inference accuracy and total score for both GPU and FPGA tracks.

### 6.1 Ablation Study

We perform an ablation study on DAC-SDC dataset to analyze these three configurations of SkyNet (Model A, B, and C listed in Table 3). By combining two activation functions (ReLU and ReLU6), six configurations of SkyNet are evaluated. We train these models in an end-to-end fashion using multi-scale training with the learning rate starting from  $1e-4$  to  $1e-7$ . We apply stochastic gradient descent (SGD) to update parameters. To further enrich the training data, we use data augmentations to distort, jitter, crop, and resize inputs

Table 4. Validation accuracy of SkyNet.

DNN MODEL	PARAMETER SIZE	IoU
SKYNET A - ReLU	1.27 MB	0.653
SKYNET A - ReLU6	1.27 MB	0.673
SKYNET B - ReLU	1.57 MB	0.685
SKYNET B - ReLU6	1.57 MB	0.703
SKYNET C - ReLU	1.82 MB	0.713
SKYNET C - ReLU6	1.82 MB	<b>0.741</b>

with size  $160 \times 320$ . The accuracy results are presented in Table 4, where SkyNet C - ReLU6 reaches the highest IoU (0.741) on the validation set. Therefore, we use this model as the proposed design for the following experiments.

### 6.2 Evaluation Criteria

Comprehensive evaluations are introduced in DAC-SDC, covering detection accuracy (IoU), throughput (FPS), and energy consumption. To identify the best design, a total score is calculated following Equation 2 to 5. Assuming there are  $I$  registered teams and  $K$  images in the test set, the IoU score for team  $i$ , denoted as  $R_{IoU_i}$ , is calculated as:

$$R_{IoU_i} = \frac{\sum_{k=1}^K IoU_{i,k}}{K} \quad (2)$$

For energy,  $\bar{E}_I$  is denoted as the average energy consumption of all  $I$  entries when performing DNN inference on the test dataset (Equation 3). The energy score of team  $i$  ( $ES_i$ ) is then computed using Equation 4 relating to the ratio between average energy and the energy consumed by this team.  $x$  is set to 2 and 10 for FPGA track and GPU track, respectively. Eventually, the total score, denoted as  $TS_i$ , is calculated in Equation 5 including both inference accuracy ( $R_{IoU_i}$ ) and energy consumption ( $ES_i$ ).

$$\bar{E}_I = \frac{\sum_{i=1}^I E_i}{I} \quad (3)$$

$$ES_i = \max\{0, 1 + 0.2 \times \log_x \frac{\bar{E}_I}{E_i}\} \quad (4)$$



Figure 7. Object detection results generated by SkyNet on DAC-SDC dataset. Challenges include to detect small objects and distinguish multiple similar objects (e.g., images in the first row).

$$TS_i = R_{IoU_i} \times (1 + ES_i) \quad (5)$$

### 6.3 GPU Implementation

For the TX2 GPU implementation, we keep all network parameters using Float32 to maintain the best inference accuracy. Since most of the compute-intensive parts of DNN inference are handled by NVIDIA cuDNN, which leaves little space for customized improvement, we start optimizing our design on a system-level.

The whole procedure of running SkyNet contains four steps as: 1) input fetching from the flash storage in a unit of batch; 2) image pre-process which includes input resizing and normalization; 3) DNN inference; and 4) post-process to generate bounding boxes and buffer results in DDR memory. The most straightforward way is to execute these steps in serial but with the cost of low resource utilization and poor throughput performance. In our design, we first merge step 1 and 2 in pre-process and enable multithreading technology to execute these steps in a pipelined fashion as shown in Figure 10. We use NVIDIA System Profiler (L4T) to capture the latency results. In average, the proposed system-level optimizations enable a 3.35X speedup compared to the original serial design and help our design reach the highest throughput performance, peaking at 67.33 FPS.

### 6.4 FPGA Implementation

To implement DNNs on FPGA, we suffer even scarcer resource budgets, as the theoretical peak performance provided by Ultra96 FPGA (144 GOPS @200MHz) is much lower than the TX2 GPU (665 GFLOPS @1300MHz). By using the proposed bottom-up design flow, hardware limitations have already captured by the Bundle design and the Bundle is instantiated on FPGA as a single customized hardware IP. Since the proposed network is structured by the same type of Bundle, this IP can be shared across different layers to cope with the resource constraints. Still, we need more optimizations to further enhance the performance.

#### 6.4.1 Quantization, Batch Process, and Tiling

Since fixed-point representation is more favorable in FPGA design, we quantize the FMs and weights from Float32 to

fixed point and explore different quantization schemes in Table 7. After quantization, the SkyNet backbone suffers different levels of accuracy drop from 1.4% to 6.1% in scheme 1 to 4. We finally pick scheme 1 as accuracy has higher weight in the total score calculation (Equation 5).

Since network parameters can not be accommodated by the FPGA on-chip memory (which is BRAM with only 0.95 MB available), we have to store them in the external memory (DRAM), which easily makes the memory access bandwidth a bottleneck. To mitigate the bandwidth demand, input batch process is applied to exploit data reuse opportunities, where a certain number (which equals to the batch size) of input images are assembled before sending to FPGA for DNN inference, so that task size (the number of images being processed at one time) increases while consuming the same amount of network parameters from DRAM.

With larger batch size, the process of network inference asks for larger amount of FPGA on-chip memory to buffer intermediate FMs. Since our implementation is based on an IP-shared structure, buffers instantiated on FPGA are shared by different layers, which means the buffer may not be large enough for the FMs generated by the first few layers while too large for the last few layers as FMs get smaller after pooling. To solve this problem, we propose an input tiling and batch scheme as shown in Figure 9. Four inputs are stitched to form a larger input which can be processed as an entirety. With the tiling and batch process, it is possible to use one shared buffer across different layers without changing its size. The proposed solution inherits the benefit from batch process to allow better reuse of DNN weights and it eliminates the possible waste of unused buffer space.

#### 6.4.2 Layer Fusion, Memory Hierarchy, and Task Partitioning

To avoid dealing with the floating-point operations (e.g., inverse-square root) in BN layer, we use layer fusion to merge both parameters from Conv and its successive BN offline. So, there are no separated BN layers nor expensive floating-point operations required during DNN inference.

With hardware resources shared by DNN layers, the inter-



Table 5. GPU final results from DAC-SDC'19 and '18 using the hidden test set with 50K images, evaluated by a TX2 GPU.

TEAM NAME	IoU	FPS	POWER(W)	TOTAL SCORE
RESULTS FROM 2019				
<b>SKYNET (OURS)</b>	<b>0.731</b>	<b>67.33</b>	13.50	<b>1.504</b>
THINKER (XIONG ET AL., 2019)	0.713	28.79	<b>8.55</b>	1.442
DEEPZS (DENG ET AL., 2019)	0.723	26.37	15.12	1.422
RESULTS FROM 2018				
ICT-CAS (LU ET AL., 2018)	0.698	24.55	12.58	1.373
DEEPZ (DENG & ZHUO, 2018)	0.691	25.30	13.27	1.359
SDU-LEGEND (ZANG ET AL., 2018)	0.685	23.64	10.31	1.358

Table 6. FPGA final results in DAC-SDC'19 and '18 using the hidden test set with 50K images. Designs in 2019 are evaluated on a Ultra96 FPGA while designs in 2018 use a Pynq-Z1 FPGA.

TEAM NAME	IoU	FPS	POWER (W)	TOTAL SCORE
RESULTS IN 2019				
<b>SKYNET (OURS)</b>	<b>0.716</b>	25.05	7.26	<b>1.526</b>
XJTU_TRIPLER (ZHAO ET AL., 2019)	0.615	50.91	9.25	1.394
SYSTEMSETHZ (KARA & ALONSO, 2019)	0.553	<b>55.13</b>	6.69	1.318
RESULTS IN 2018				
TGIIF (ZENG ET AL., 2018)	0.624	11.96	4.20	1.267
SYSTEMSETHZ (KARA ET AL., 2018)	0.492	25.97	<b>2.45</b>	1.179
iSMART2 (HAO ET AL., 2018)	0.573	7.35	2.59	1.164



Figure 8. Object tracking results generated by SkyNet on GOT-10K dataset.

Table 7. Validation accuracy results regarding different quantization schemes during FPGA implementation

Scheme	Feature Map	Weight	Accuracy (IoU)
0	Float32	Float32	0.741
1	9 bits	11 bits	0.727
2	9 bits	10 bits	0.714
3	8 bits	11 bits	0.690
4	8 bits	10 bits	0.680

mediate results need to be swapped in/out between on-chip and external memory. To boost the performance, we instantiate the selected Bundle on hardware and implement a five-stage pipeline with Load, EXE\_CONV3, EXE\_CONV1, EXE.Pooling, and WriteBack stages. By using Ping-pong buffers between memory and computation units, data transfer (in Load and WriteBack stages) can be fully overlapped by computation latency. Regarding the data transfer be-




Figure 9. The proposed batch and tiling design to increase the data reuse opportunity and avoid on-chip memory waste.

tween adjacent execution stages (with “EXE” prefix), we keep data on-chip without going through external memory.

To fully utilize the available computational resource, we also implement task partitioning on the Ultra96. The whole design is shown in Figure 10, which is highly similar to our GPU design. Workloads are distributed to both CPU and FPGA and creating a system-level pipeline. With all three




Figure 10. Task partitioning in SkyNet implementation on TX2 GPU and Ultra96 FPGA.

tasks (pre-process, SkyNet inference, and post-process) overlapped, our FPGA design can reach 25.05 FPS.

## 6.5 Result Comparison

After implementing the proposed DNN on GPU and FPGA following the strategies mentioned in Section 6.3 and 6.4, our designs are evaluated by the DAC-SDC organizers using the hidden test set. As shown in Table 5 and 6, we present the comparison results with the top-3 teams in DAC-SDC’19 and ’18. In our GPU design, SkyNet outperforms all other competitors by delivering the best accuracy (0.731), throughput performance (67.33), and total score (1.504). In terms of the FPGA design, SkyNet also reaches the best accuracy and gets the highest total score.

## 7 SKYNET EXTENSION ON GOT-10K

Since SkyNet can deliver real-time object detection on embedded systems, we setup experiments on the GOT-10k benchmark (Huang et al., 2019) to demonstrate its potential on object tracking. GOT-10k is a large high-diversity database for generic object tracking with rich motion trajectory and wide coverage of object classes. Models are evaluated with two metrics in GOT-10k as average overlap (AO) and success rate (SR). AO is defined as the mean of IoU between prediction and ground truth bounding boxes, while SR is defined as the proportion of predictions where the IoU is beyond some threshold. During evaluation, Got-10K only provides the ground truth bounding box in the first frame and expect trackers to keep tracking on the same object for subsequent frames by predicting bounding boxes. The predictions will then be evaluated by the Got-10K server. In this section, we integrate the SkyNet backbone with two of the state-of-the-art trackers (SiamRPN++ and SiamMask) and evaluate its capability of real-time tracking.

### 7.1 Evaluation Using SiamRPN++

Siamese network is one of the most popular network structures for building object trackers. The Siamese trackers locate the object by the correlation between features extracted from the exemplar image and search image, where DNN-based feature extraction plays an important role. SiamRPN++ (Li et al., 2019a) is the first Siamese tracker that has been proven to profit from using DNN backbones with different capacities as long as they are properly trained.

Table 8. Performance of SiamRPN++ trackers on GOT-10k with different backbones evaluated on single NVIDIA 1080Ti.

Backbone	AO	SR <sub>0.50</sub>	SR <sub>0.75</sub>	FPS
AlexNet	0.354	0.385	0.101	52.36
ResNet-50	0.365	0.411	0.115	25.90
SkyNet	0.364	0.391	0.116	41.22

Table 9. Performance of SiamMask trackers on GOT-10k with different backbones evaluated on single NVIDIA 1080Ti.

Backbone	AO	SR <sub>0.50</sub>	SR <sub>0.75</sub>	FPS
ResNet-50	0.380	0.439	0.153	17.44
SkyNet	0.390	0.442	0.158	30.15

To evaluate the performance of different backbones, we train three SiamRPN++ trackers with AlexNet, ResNet-50, and SkyNet backbones on GOT-10k. We maintain the size of exemplar and search images as  $127 \times 127$  and  $255 \times 255$  ( $128 \times 128$  and  $256 \times 256$  for SkyNet for better implementation efficiency), respectively, and we set the learning rates start from  $1e-3$  to  $1e-5$ . Results are shown in Table 8 where SkyNet achieves nearly the same quality (AO and SR) as the ResNet-50 backbone but much better speed (1.59X faster).

### 7.2 Evaluation Using SiamMask

SiamMask (Wang et al., 2019) is another Siamese tracker which outperforms SiamRPN++ by incorporating image segmentation for object tracking tasks. Since information of the segmentation is not provided, it cannot be directly trained with GOT-10k dataset. Instead, we perform training using Youtube-VOS dataset (Xu et al., 2018) and apply object tracking on Got-10K to compare the performance of different backbones using the SiamMask structure. We maintain the same input size setup as Section 7.1 and apply the learning rates from  $1e-3$  to  $1e-4$ . As shown in Table 9, the proposed SkyNet backbone outperforms ResNet-50 in all metrics when using SiamMask tracker with better tracking quality and 1.73X speedup.

## 8 CONCLUSIONS

In this paper, we proposed SkyNet, as well as a hardware-efficient method to generate compact DNNs for object detection running on embedded GPUs and embedded FPGAs. SkyNet design methodology is a novel bottom-up DNN design flow which can capture hardware limitations using realistic hardware feedbacks and deliver DNNs with great balance between software and hardware metrics, such as DNN inference accuracy and throughput performance. SkyNet was demonstrated on the 56th IEEE/ACM DAC-SDC low power object detection challenge and won the first place winner award for both GPU and FPGA tracks. We also extended SkyNet to handle object tracking task and it delivered 1.60X and 1.73X higher FPS, and 37.20X smaller parameter size with comparable accuracy when compared to the state-of-the-art Siamese trackers with ResNet-50 backbone.

---

## 9 ACKNOWLEDGMENTS

This work was partly supported by the IBM-Illinois Center for Cognitive Computing System Research (C<sup>3</sup>SR) – a research collaboration as part of IBM AI Horizons Network.

## REFERENCES

- Alwani, M., Chen, H., Ferdman, M., and Milder, P. Fused-layer CNN accelerators. In *Proceedings of the International Symposium on Microarchitecture*, 2016.
- Cai, H., Zhu, L., and Han, S. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332*, 2018.
- Chen, Y., He, J., Zhang, X., Hao, C., and Chen, D. Cloud-DNN: An open framework for mapping DNN models to cloud FPGAs. In *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2019.
- Chen, Y.-H., Krishna, T., Emer, J., and Sze, V. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. In *IEEE International Solid-State Circuits Conference (ISSCC)*, 2016.
- Cheng, B., Wei, Y., Shi, H., Feris, R., Xiong, J., and Huang, T. Decoupled classification refinement: Hard false positive suppression for object detection. *arXiv preprint arXiv:1810.04002*, 2018a.
- Cheng, B., Wei, Y., Shi, H., Feris, R., Xiong, J., and Huang, T. Revisiting RCNN: On awakening the classification power of faster RCNN. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018b.
- Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., and Bengio, Y. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv preprint arXiv:1602.02830*, 2016.
- Dai, J., Li, Y., He, K., and Sun, J. R-FCN: Object detection via region-based fully convolutional networks. In *Advances in neural information processing systems*, 2016.
- Dai, X., Yin, H., and Jha, N. K. Nest: A neural network synthesis tool based on a grow-and-prune paradigm. *IEEE Transactions on Computers*, 68(10):1487–1497, 2019.
- Deng, J. and Zhuo, C. a. DAC-SDC’18 2nd place winner in GPU track. <https://github.com/jndeng/DACSDC-DeepZ>, 2018. Accessed: 2020-02-28.
- Deng, J., Shen, T., Yan, X., Chen, Y., Zhang, H., Wang, R., Zhou, P., and Zhuo, C. DAC-SDC’19 3rd place winner in GPU track, 2019.
- Ding, X., Ding, G., Han, J., and Tang, S. Auto-balanced filter pruning for efficient convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- DJI. The DAC-SDC dataset for low power object detection. <http://www.cse.cuhk.edu.hk/byu/2019-DAC-SDC/index.html>, 2018. Accessed: 2020-02-28.
- Elsken, T., Metzen, J. H., and Hutter, F. Efficient multi-objective neural architecture search via lamarckian evolution. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Franklin, D. NVIDIA Jetson TX2 delivers twice the intelligence to the edge. *NVIDIA Accelerated Computing—Parallel For all*, 2017.
- Gope, D., Dasika, G., and Mattina, M. Ternary hybrid neural-tree networks for highly constrained iot applications. 2019.
- Han, S., Pool, J., Tran, J., and Dally, W. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, 2015.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2016.
- Hao, C., Li, Y., Huang, S. H., Zhang, X., Gao, T., Xiong, J., Rupnow, K., Yu, H., Hwu, W.-M., and Chen, D. DAC-SDC’18 3rd place winner in FPGA track. <https://github.com/onioncc/iSmartDNN>, 2018. Accessed: 2020-02-28.
- Hao, C., Zhang, X., Li, Y., Huang, S., Xiong, J., Rupnow, K., Hwu, W.-m., and Chen, D. FPGA/DNN co-design: An efficient design methodology for IoT intelligence on the edge. In *Proceedings of the 56th Annual Design Automation Conference 2019 (DAC)*. ACM, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016.
- He, K., Gkioxari, G., Dollár, P., and Girshick, R. Mask R-CNN. In *Proceedings of the IEEE international conference on computer vision (CVPR)*, 2017.
- Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al. Searching for mobilenetv3. In *Proceedings of the International Conference on Computer Vision*, 2019.

- 
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Huang, L., Zhao, X., and Huang, K. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- Iandola, F. N., Han, S., Moskewicz, M. W., Ashraf, K., Dally, W. J., and Keutzer, K. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 MB model size. *arXiv preprint arXiv:1602.07360*, 2016.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Jiang, W., Zhang, X., Sha, E. H.-M., Yang, L., Zhuge, Q., Shi, Y., and Hu, J. Accuracy vs. efficiency: Achieving both through FPGA-implementation aware neural architecture search. In *Proceedings of the 56th Annual Design Automation Conference 2019 (DAC)*, 2019.
- Jouppi, N. P., Young, C., Patil, N., Patterson, D., Agrawal, G., Bajwa, R., Bates, S., Bhatia, S., Boden, N., Borchers, A., et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of International Symposium on Computer Architecture (ISCA)*, 2017.
- Kara, K. and Alonso, G. DAC-SDC'19 3rd place winner in FPGA track, 2019.
- Kara, K., Zhang, C., and Alonso, G. DAC-SDC'18 2nd place winner in FPGA track. <https://github.com/fpgasystems/spooNN>, 2018. Accessed: 2020-02-28.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 2012.
- Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., and Yan, J. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019a.
- Li, F., Zhang, B., and Liu, B. Ternary weight networks. *arXiv preprint arXiv:1605.04711*, 2016.
- Li, Q., Zhang, X., Xiong, J., Hwu, W.-m., and Chen, D. Implementing neural machine translation with bi-directional GRU and attention mechanism on FPGAs using HLS. In *Proceedings of the 24th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2019b.
- Li, Y., Chen, Y., Wang, N., and Zhang, Z. Scale-aware trident networks for object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019c.
- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., and Belongie, S. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2017a.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, 2017b.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., and Berg, A. C. SSD: Single shot multibox detector. In *Proceedings of the European conference on computer vision (ECCV)*, 2016.
- Lu, H., Cai, X., Zhao, X., and Wang, Y. DAC-SDC'18 1st place winner in GPU track. <https://github.com/lvhao7896/DAC2018>, 2018. Accessed: 2020-02-28.
- Luo, J.-H., Wu, J., and Lin, W. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, 2017.
- Qiu, J., Wang, J., Yao, S., Guo, K., Li, B., Zhou, E., Yu, J., Tang, T., Xu, N., Song, S., et al. Going deeper with embedded FPGA platform for convolutional neural network. In *Proceedings of International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2016.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Proceedings of European Conference on Computer Vision*, 2016.
- Real, E., Aggarwal, A., Huang, Y., and Le, Q. V. Regularized evolution for image classifier architecture search. In *Proceedings of the AAAI conference on artificial intelligence*, 2019.
- Redmon, J. and Farhadi, A. YOLO9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2017.
- Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016.
- Ren, A., Zhang, T., Ye, S., Li, J., Xu, W., Qian, X., Lin, X., and Wang, Y. Admm-nn: An algorithm-hardware co-design framework of dnns using alternating direction

- 
- methods of multipliers. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2019.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., and LeCun, Y. Overfeat: Integrated recognition, localization and detection using convolutional networks. 2014.
- Shen, Z., Shi, H., Yu, J., Phan, H., Feris, R., Cao, L., Liu, D., Wang, X., Huang, T., and Savvides, M. Improving object detection from scratch via gated feature reuse. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2019.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2015.
- Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., and Le, Q. V. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Tao, R., Gavves, E., and Smeulders, A. W. Siamese instance search for tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016.
- Tian, Z., Shen, C., Chen, H., and He, T. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- Tschannen, M., Khanna, A., and Anandkumar, A. Strassenets: Deep learning with a multiplication budget. 2018.
- Valmadre, J., Bertinetto, L., Henriques, J., Vedaldi, A., and Torr, P. H. End-to-end representation learning for correlation filter based tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Vanholder, H. Efficient inference with tensorrt, 2016.
- Wang, J., Lou, Q., Zhang, X., Zhu, C., Lin, Y., and Chen, D. Design flow of accelerating hybrid extremely low bit-width neural network in embedded FPGA. In *Proceedings of the International Conference on Field Programmable Logic and Applications (FPL)*, 2018a.
- Wang, Q., Teng, Z., Xing, J., Gao, J., Hu, W., and Maybank, S. Learning attentions: residual attentional siamese network for high performance online visual tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018b.
- Wang, Q., Zhang, L., Bertinetto, L., Hu, W., and Torr, P. H. Fast online object tracking and segmentation: A unifying approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., and Keutzer, K. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Xiong, F., Yin, S., Fan, Y., and Ouyang, P. DAC-SDC'19 2nd place winner in GPU track, 2019.
- Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., and Huang, T. Youtube-vos: A large-scale video object segmentation benchmark. In *European Conference on Computer Vision (ECCV)*, 2018.
- Xu, X., Zhang, X., Yu, B., Hu, X. S., Rowen, C., Hu, J., and Shi, Y. Dac-sdc low power object detection challenge for uav applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- Zang, C., Liu, J., Hao, Y., Li, S., Yu, M., Zhao, Y., Li, M., Xue, P., Qin, X., Ju, L., Li, X., Zhao, M., and Dai, H. DAC-SDC'18 3rd place winner in GPU track. <https://github.com/xiaoyuuyuuu/dac-hdc-2018-object-detection-in-Jetson-TX2>, 2018. Accessed: 2020-02-28.
- Zeng, S., Chen, W., Huang, T., Lin, Y., Meng, W., Zhu, Z., and Wang, Y. DAC-SDC'18 1st place winner in FPGA track. <https://github.com/hirayaku/DAC2018-TGIIF>, 2018. Accessed: 2020-02-28.
- Zhang, X., Liu, X., Ramachandran, A., Zhuge, C., Tang, S., Ouyang, P., Cheng, Z., Rupnow, K., and Chen, D. High-performance video content recognition with long-term recurrent convolutional network for FPGA. In *Proceedings of the International Conference on Field Programmable Logic and Applications (FPL)*, 2017a.

---

Zhang, X., Wang, J., Zhu, C., Lin, Y., Xiong, J., Hwu, W.-m., and Chen, D. DNNBuilder: an automated tool for building high-performance DNN hardware accelerators for FPGAs. In *Proceedings of International Conference on Computer-Aided Design (ICCAD)*, 2018a.

Zhang, X., Zhou, X., Lin, M., and Sun, J. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018b.

Zhang, Y., Suda, N., Lai, L., and Chandra, V. Hello edge: Keyword spotting on microcontrollers. *arXiv preprint arXiv:1711.07128*, 2017b.

Zhao, B., Zhao, W., Xia, T., Chen, F., Fan, L., Zong, P., Wei, Y., Tu, Z., Zhao, Z., Dong, Z., and Ren, P. DAC-SDC'19 2nd place winner in FPGA track, 2019.