

Efficient Image Sensor Subsampling for DNN-Based Image Classification

Jia Guo, Hongxiang Gu, and Miodrag Potkonjak
Computer Science Department
University of California, Los Angeles
{jia,hxgu,miodrag}@cs.ucla.edu

ABSTRACT

Today's mobile devices are equipped with cameras capable of taking very high-resolution pictures. For computer vision tasks which require relatively low resolution, such as image classification, subsampling is desired to reduce the unnecessary power consumption of the image sensor. In this paper, we study the relationship between subsampling and the performance degradation of image classifiers that are based on deep neural networks (DNNs). We empirically show that subsampling with the same step size leads to very similar accuracy changes for different classifiers. In particular, we could achieve over 15× energy savings just by subsampling while suffering almost no accuracy lost. For even better energy accuracy trade-offs, we propose AdaSkip, where the row sampling resolution is adaptively changed based on the image gradient. We implement AdaSkip on an FPGA and report its energy consumption.

CCS CONCEPTS

• **Computing methodologies** → **Computer vision; Neural networks**; • **Computer systems organization** → **Embedded systems**;

KEYWORDS

neural network, image sensor, subsampling

ACM Reference Format:

Jia Guo, Hongxiang Gu, and Miodrag Potkonjak. 2018. Efficient Image Sensor Subsampling for DNN-Based Image Classification. In *ISLPED '18: ISLPED '18: International Symposium on Low Power Electronics and Design, July 23–25, 2018, Seattle, WA, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3218603.3218618>

1 INTRODUCTION

High-resolution cameras are prevalent on today's mobile devices. Apple's iPhone X is equipped with a 12-megapixel camera [1] and the latest Sony Xperia XZ has a dazzling 23-megapixel camera [2]. While the increasing consumer demand for high-quality photography justifies the need for a high-resolution camera, image sensors with over 10 million pixels are usually an overkill for many computer vision tasks. Take image classification as an example. For the

ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [20], a popular 1000-class classification competition, competitors usually resize input images to 224×224. For simpler problems, such as the CIFAR-10 10-class classification problem [15], state-of-the-art classifiers achieve over 95% accuracy with input images of size 32×32. As the amount of computation required for a convolutional neural network (CNN) will roughly quadruple if the input image doubles its height and width, it is likely that the input sizes of classifiers will be kept as small as possible. That being said, a 5500×4200 image produced by a 23-megapixel appears disproportionately big compared to what we need for common computer vision tasks.

The question then becomes how do we obtain images with a relatively low resolution from these over-sized image sensors. A straightforward way is to take the full image and interpolate to create a smooth downsized image. However, energy-wise it is not the best solution. There exists an approximately linear relationship between the energy consumption of the image sensor and the number of pixels it samples [8]. Unlike cloud-based DNN applications [10], applications on mobile devices call for low energy consumption. As a result, it is desired to sample as few pixels as possible. But how much should we subsample? If we undersample, the aliasing effect causes distortions that might lead to degraded performances of the computer vision algorithms. If we oversample, we risk wasting energy by sampling unnecessary pixels. It is an important question to hardware and system designers, as it mandates the type of features and APIs that are exposed to application developers. It is also an important question to computer vision app developers, as its answer instructs them on how best to utilize the camera for the overall energy efficiency of mobile computer vision systems.

To bridge the gap between image sensor hardware and DNN-based image classification algorithms, we analyze the cause of errors when we subsample and propose clear and practical strategies to improve it. In particular, we show that subsampling, when done right, does bring about tremendous energy saving with negligible loss of accuracy. We do not intend to provide an exact answer to how much we should subsample. Instead, we show that by studying the performance degradation of just a few classifiers with subsampled input images, we can get a reasonable estimate of the performance degradation of image classification in general. To further reduce the energy consumption of the image sensor, we propose a minimal hardware modification to off-the-shelf image sensors. AdaSkip, the proposed enhancement, adaptively subsamples parts of the image based on the complexity. We show that the proposed method can achieve even better energy accuracy trade-offs than plain subsampling.

The contribution of this paper is neither proposing new image sensor architecture nor building more efficient computer vision

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISLPED '18, July 23–25, 2018, Seattle, WA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5704-3/18/07...\$15.00

<https://doi.org/10.1145/3218603.3218618>

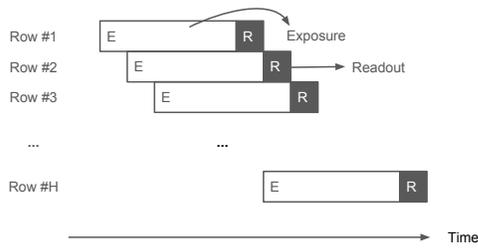


Figure 1: CMOS image sensor operation with rolling shutter.

algorithms. Rather, we take existing and widely-used designs from both sides and optimize from a system’s perspective. To the best of our knowledge, we are the first to a) systematically study the relationship between subsampling and performance degradation of DNN-based image classification; b) propose a low-footprint hardware extension to off-the-shelf CMOS image sensors for more efficient subsampling in the context of efficient image classification.

We organize the rest of the paper as follows. Section 2 introduces prior art from the images sensor and deep DNN optimization communities. Section 3 briefly describes the operation of CMOS image sensors and our energy model. Section 4 analyzes the negative effect subsampling has on DNNs and motivates the problem. Section 5 details both of our proposed methods. The experimental setup is described in Section 6 and the results in Section 7.

2 RELATED WORK

CMOS Image Sensors. It is well known that the analog readout chain consumes a significant amount of power. Thus there has been a continuous effort in creating *specialized* image sensors that take reduced amount of sampling. Notable approaches include the use of DCT proposed by Kawahito *et al.* [7], predictive coding proposed by Leon-Salas *et al.* [17], Quadtree Decomposition algorithm proposed by Artyomov *et al.* [3], and compressive sensing proposed by Oike *et al.*[19]. Kemeny *et al.* also proposed an image sensor that allows programmable readout of pixels at arbitrary locations and various resolutions [14]. LiKamwa *et al.* studied the energy proportionality of off-the-shelf image sensors and proposed optimization strategies for simple computer vision tasks [8].

In our work, we look at the effect of subsampling on the particular computer vision task of image classification rather than image quality. To reduce sensor energy consumption in vision tasks, one way is to derive a separate image sensor pipeline along side the normal one[4]. Instead, we base our techniques solely on subsampling by skipping pixels, a feature that is widely available on today’s mobile image sensors.

Neural Network Design. DNNs are extremely power hungry. Pruning, quantization and factorization have already become widely adopted techniques to improve efficiency [11][6]. Meanwhile, the quadratic relationship between computation and input size can also plays a critical role, as evident in the multi-scale Densenet architecture [13]. Other researchers have also proposed to reduce the computation of neural networks by using different image sensor designs. Chen *et al.* have proposed to approximate the first layer of CNNs using image sensors with angle sensitive pixels [5]. LiKamwa

et al. proposed to move early layers in CNNs to the analog domain to save energy [18].

In comparison to the designs mentioned above, our work targets a broader audience. We base our experiments on generic neural network models. And the main target of our study is how the most widely used image sensor architecture could best be utilized in order to serve need of neural networks.

3 CMOS IMAGE SENSORS

In this section, we describe the architecture and operation of the most commonly used mobile CMOS image sensors. We then describe the model based on previous energy characterizations of off-the-shelf units [8].

3.1 Architecture and Operation

Today’s mobile CMOS image sensors usually consist of pixel arrays, an analog signal chain and various digital processing and I/O logic. The analog chain contains ADCs, amplifiers, and bias adjusting circuits and it is known to be the major source of power consumption. A common design choice for image sensors is to use *column-parallel readout*, where one pixel column share one signal chain. In this architecture, the pixels are usually read out row by row.

Rolling electronic shutters are widely used to control the exposure. During the operation of a rolling shutter, only one row or column of pixels read out at a time. The operation fits nicely with the column-parallel readout architecture. When one row finished exposure, the readout logic starts the processing of that row. At the same time, the subsequent row is in the process of exposure, waiting for its turn for pixel readout. The whole process is depicted in Figure 1.

There are usually two supported modes of subsampling: *skipping* and *binning*. In skipping mode, the sensor skips the entire pixel readout chain of pixels at certain locations. In binning mode, pixel values are averaged before output. We mainly consider the skipping mode, since it is the mode that enables the most energy savings.

3.2 Energy Model

We mainly use LiKamWa *et al.*’s energy characterization on commercial image sensors[8]. They proposed to estimate the energy consumption of image sensors using the following equation:

$$E = P_{idle}t_{idle} + P_{active}t_{active}$$

P_{idle} and t_{idle} represents the power and time when no pixel readout is occurring, i.e. when the first row is going through exposure. Once the pixel readout starts, the system switches to the active mode with higher power consumption. The length of the readout t_{active} is linear to the number of pixels sampled.

4 PROBLEMS WITH SUBSAMPLING

Given a neural network that requires inputs of size $H \times W$, it is most desirable if we directly subsample the pixel arrays to create an $H \times W$ image. However, images naturally contain a variety of high-frequency components in its original form. Subsampling images can easily give rise to strong *aliasing* effect. Neural networks have very unpredictable behaviors when they encounter input data derived from distributions that are different from what they are trained on, as in the case of adversarial examples [9]. From the neural

Table 1: Effects of direct subsampling on 10-class classifiers with different input sizes

Subsampling method	Accuracy loss (%)
32 × 32, Direct subsampling	-48.6 ± 6.4
32 × 32, Smoothed (radius=0.3)	-36.6 ± 5.9
32 × 32, Smoothed (radius=0.65)	-9.3 ± 1.3
32 × 32, Smoothed (radius=0.9)	-50.2 ± 5.8
64 × 64, Direct subsampling	-31.7 ± 8.3
64 × 64, Smoothed (radius=0.55)	-6.1 ± 1.3

networks' perspective, the process of subsampling creates input data of a slightly different distribution. It is thus likely that neural networks do not behave correctly in these situations.

To verify our conjecture, we tested direct subsampling on 10-class image classifiers with two different input sizes. Section 6 explains the experimental setup in more details. To create an image of the target input size, we directly subsample (by skipping pixels) from a 512 × 512 image. Then we feed it to 10 different classifiers, with classification accuracies averaging 86.9% and 90.6% for 32 × 32 and 64 × 64 classifiers respectively. Fearing that the extreme variation in neighboring pixels might be adversely affecting the performance of the classifiers, we further added a low pass Gaussian filter to smooth out the images. We used binary search to find a radius that performs relatively well. Table 1 shows the mean accuracies and corresponding 95% confidence intervals. We observe substantial accuracy drop when direct subsampling is applied. Smoothing helps, but in the best case, we still see an average of 9.3% accuracy drop for 32 × 32 classifiers and 6.1% accuracy drop for 64 × 64 classifiers.

To further study reasons behind the misclassifications, we study the effect each pixel has on the final loss function. To quantify the effect, we define the following measurement given loss function L , bilinearly resampled image X and smoothed subsampled image X' :

$$I = \frac{\partial L}{\partial X'} \odot (X' - X)$$

$\frac{\partial L}{\partial X'}$ represents how sensitive the loss function to the change of an individual pixel. It could be positive or negative, where increasing pixel value causes loss increase or decrease, respectively. \odot is the Hadamard (element-wise) product. $(X' - X)$ represents the change from the standard image to the subsampled. Intuitively, large changes in the sensitive pixels will inflate the loss and cause the classifier to misclassify.

We studied many misclassified samples and selected an example image of a bus as shown in Figure 2. Although area inside the windshield in the subsampled image appears somewhat jagged, it is not contributing much in terms of error. It is the edge of the windshield, where the color abruptly changes, that is contributing the most to the classification error. Apparently failing to sample the right color during subsampling distorted the appearance of the window frame and caused the classifier to misclassify. Above is a typical example where smoothing cannot recover the information lost during subsampling. The *only* way to mitigate the problem is to sample at a higher resolution to include more information.

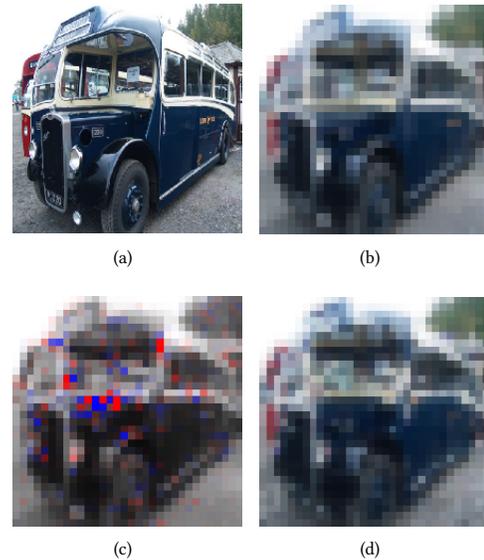


Figure 2: Errors caused by subsampling. (a) The original image of a bus; (b) 32 × 32 resampled image using bilinear kernel; (c) visualized source of error I in channel “R”. Red dots are pixels that increase loss, and blue dots are those that decrease loss. The less transparent the color, the larger impact a pixel has; (d) 32 × 32 smoothed subsampled image.

REMARK 1. *Subsampling may distort key features that image classifiers use, causing unrecoverable classification performance degradation.*

5 PROPOSED METHOD

To avoid confusion, we are going to use the term "two-step subsampling" to describe the process of subsampling to create input images of specific sizes for a classifier. We formally define this process in the next subsection. Following the definition, we describe our proposed adaptive sampling strategy, AdaSkip. AdaSkip requires minimal hardware support. We discuss an alternative hardware design in the last subsection.

5.1 Two-Step Subsampling

Let $H_o \times W_o$ be the resolution of the image sensor. The image sensor also supports a list of K subsampling steps, $\mathbb{S} = \{s_1, s_2, \dots, s_K\}$, where s_i represents the number of pixels among which one is sampled. We use $\text{subsample}(X, s)$ to represent the process of sampling with step s , both horizontally and vertically. X is of size $H_o \times W_o$, so the resulting image is of size $\frac{H_o}{s} \times \frac{W_o}{s}$. Instead of subsampling the whole image X , we use the same expression to represent the subsampling of one row x , $\text{subsample}(x, s)$.

To subsample to create images of specific size $H_t \times W_t$, we first command the image sensor to subsample with step s , where $\frac{H_o}{s} > H_t, \frac{W_o}{s} > W_t$. Then, taking the output from the sensor, we apply resampling (in software) using kernel function f $\text{resample}(x, f)$

to create image of the target size. We use bilinear kernel in our experiments.

5.2 AdaSkip

Since undersampling distorts essential image features, we should sample at a rate that is high enough to keep the features in tact. However, not all parts of the image need a high sampling rate. It's more desirable to adaptively change the rate of sampling. The process of progressive exposure and readout in rolling shutters creates slacks after each row. We could use the slacks for making decisions on the step size used for the subsequent rows, based on the information that we obtain from the current row. Using this information, we could speculate on whether to sample them using high resolution or low resolution. With the varying granularity within one frame, the resulting image will look very undesirable aesthetically. However, deep learning algorithm might be able to look at the resulting image in a different light. Since the process of deciding what resolution to sample is adaptive, we coined the name AdaSkip, which stands for adaptively skipping rows.

In computer vision, the *gradient* of an image is a very fundamental element in feature engineering. It is usually defined in two terms, the gradient along the x direction G_x and the gradient along the y direction G_y .

$$G_x(\mathbf{X}) = \left| \frac{\partial \mathbf{X}}{\partial x} \right|$$

$$G_y(\mathbf{X}) = \left| \frac{\partial \mathbf{X}}{\partial y} \right|$$

Image gradients have an extensive range of usages. For one, it is commonly used as building blocks for edge features. As a direct feature, it is also used to determine the complexity of different parts of the images in seam carving. As a simple indicator of the complexity of the images, it is adopted in our algorithm to decide the sampling resolution. However, since we only have the information about the current row sampled, we will only use the gradient along the x-direction G_x . Algorithm 1 details our approach. We use $G_x(\mathbf{x}_i)$ to represent the gradient in the row \mathbf{x}_i .

The high-level idea of the algorithm is extremely straightforward: use a lower resolution if the current row is not complex, else use a higher resolution. There are a few details that are not reflected in the algorithm. First, in order to make sure that the *thd* is the same for two resolutions (two rows of different sizes), the denser row should be subsampled in the calculation of $G_x(\mathbf{x}_i)$ to match the sparser one. Second, since the subsampled rows are of different lengths, we duplicate pixels in rows with the lower resolution to create rows of equal lengths. Those rows are later duplicated to fill in the skipped row in order to retain the aspect ratio.

5.3 AdaSkip Hardware Extension

Since the AdaSkip algorithm varies sampling rate per row, it can only be implemented as part of the images sensor hardware. The core logic is relatively simple, and it mainly involves line 7 in Algorithm 1. The calculation of gradient involves adders and converting two's complement representations. Comparators are next for finding the maximum value. Both components are already present in current image sensor design. Adders exist in digital gain adjustment circuits, and comparators in SAR based ADCs. To reduce the energy

Algorithm 1 The AdaSkip algorithm

Input: \mathbf{X} , the original image, of size $H_o \times W_o$; s_h, s_l , the number of pixels to skip, $s_h < s_l$; *thd*, the threshold to determine which subsampling mode to use;

Output: \mathbf{X}' , the image of desired size $H_t \times W_t$

```

1: step  $\leftarrow s_h$ 
2: Initialize counter
3: for each row  $\mathbf{x}'_i$  do            $\triangleright$  The main algorithm in hardware
4:   if counter = step then      $\triangleright$  Sample and decide resolution
5:      $\mathbf{x}'_i \leftarrow \text{subsample}(\mathbf{x}_i, \textit{step})$ 
6:     counter  $\leftarrow 1$ 
7:     if  $\max(G_x(\mathbf{x}'_i)) > \textit{thd}$  then
8:       step  $\leftarrow s_h$ 
9:     else
10:      step  $\leftarrow s_l$ 
11:    end if
12:  else                            $\triangleright$  Skip the row
13:    counter  $\leftarrow \textit{counter} + 1$ 
14:  end if
15: end for
16:  $\mathbf{X}' \leftarrow$  collections of rows  $\mathbf{x}_i$ 
17: return  $\mathbf{X}'$ 

```

of the design, we could simply lower the comparator precision by taking the most significant bits (in our case 4 bits) of the pixel value.

6 EXPERIMENTAL SETUP

The ImageNet dataset used for Large Scale Visual Recognition Challenge (ILSVRC) contains 1,461,406 images from 1000 classes [20]. It contains object classes of very different nature and granularity, including classes ranging from "warplane" to a specific dog breed called "Irish water spaniel". The median number of pixels in the images is around 200,000. Due to resolution limitations of the dataset, it is not possible to conduct experiments to reflect the performance on today's high resolution cameras. We decide to simulate the behavior of an image sensor with resolution 512×512 . Images are bilinear resampled and reshaped to the resolution, and we assume that each pixel in the reshaped image corresponds to a sensor pixel. For subsampling, we assume that when the supported skip steps are $s = 1, 2, 4$ and 8 . For simplicity, we simulate outdoor scenarios with relatively less exposure time and thus less *t_{idle}* [8]. In this case, the majority power consumption occurs during the active mode. The image sensor energy is estimated based on the model described in Section 3.2.

To reenact a reasonable ratio between the image sensor resolution and classifier input size, we experiment on classifiers with 32×32 , 48×48 , and 64×64 input images. The resized images are created using bilinear resampling¹. We randomly select 10, 20 or 30 classes from the 1000 classes, and train 5 or 10 classifiers to form the population of classifiers with the corresponding configuration. The randomly sampled classes have varying levels semantic similarities. Using these classifiers of very different nature as well as

¹To our knowledge, all of today's classifiers are use smooth natural and smooth images that are interpolated at least linearly when downsizing. In our experience, classifiers trained on images with strong aliasing suffer more from overfitting and exhibit inferior performance.

Table 2: Classifiers used in our experiments

Input size	# Classes	Quantity	Accuracy (%)
32 × 32	10	10	86.9 ± 3.3
32 × 32	20	10	81.1 ± 2.0
32 × 32	30	5	78.9 ± 2.3
48 × 48	10	10	87.3 ± 4.2
48 × 48	20	5	86.0 ± 3.3
48 × 48	30	5	84.9 ± 2.7
64 × 64	10	10	90.6 ± 2.1
64 × 64	20	5	88.0 ± 3.5

Table 3: Accuracy loss caused by two-step subsampling with different step sizes on different classifiers

Energy savings (step) \ Classifier	56.0× (s = 8)	15.5× (s = 4)	3.97× (s = 2)
32 × 32, 10-class	-1.6 ± 0.6	-0.1 ± 0.2	0
32 × 32, 20-class	-1.3 ± 0.4	-0.1 ± 0.1	0
32 × 32, 30-class	-2.7 ± 0.4	-0.0 ± 0.2	0
48 × 48, 10-class	-3.0 ± 0.9	-0.4 ± 0.3	0.0 ± 0.1
48 × 48, 20-class	-4.4 ± 1.8	-0.2 ± 0.1	0
48 × 48, 30-class	-4.1 ± 0.7	-0.4 ± 0.3	0
64 × 64, 10-class	-30.5 ± 9.0	-0.2 ± 0.2	0
64 × 64, 20-class	-36.5 ± 9.5	-0.4 ± 0.5	0

using different input sizes are all parts of our efforts to make our conclusions more generalizable.

All of the classifiers use the ResNet34 architecture [12], trained with over 150 epochs on an Nvidia GTX 1080Ti. We use SGD with a momentum of 0.9. The learning rate starts at 0.1 and is reduced to 1/10 the value every 50 epochs. Table 2 lists the classifiers. In addition, we filter out images with original confidence score below 0.8 when we compare different subsampling methods ².

7 EVALUATION

We start the evaluation by presenting the performance characteristics of all our classifiers when two-step subsampling is used. Next, we demonstrate the AdaSkip algorithm by first visualizing it, and then show its performance with respect to energy and accuracy.

7.1 Two-Step Subsampling

Table 3 shows the accuracy loss using different steps s for subsampling. Different columns represent different step sizes and on the top row we also showed the relative energy savings. The gray areas represent cases where after subsampling, the image is less than twice (in both width and height) the size of the required input. Comparing to Table 1, even just doubling the sampling rate substantially reduces the accuracy loss. Further, we observe that the accuracy loss is correlated much more with the step size than with what classifier it is. It appears that as long it is at least 2× larger than

²In these cases, the class scores before the softmax layer are almost uniformly distributed. A small change in the input may trigger an arbitrary change in the prediction. These predictions are usually discarded in a practical setting.



Figure 3: Visualizing AdaSkip-based subsampling. (a) The original 512 × 512 image; (b) the output of the AdaSkip algorithm, where the greyscale parts are rows sampled with $s_l = 8$ and the colored part are rows sampled with $s_h = 4$; (c) 32 × 32 output images from two-step subsampling with $s = 4$ and bilinear resampling; (d) 32 × 32 images obtained by bilinear resampling the output in (b).

the target resolution of the classifiers, the same resolution after subsampling causes similar accuracy loss to all classifiers. These findings echo with the analysis from Section 4, where we argue that information loss at the subsampling step plays a crucial role in determining the performance degradation of image classifiers.

This characteristic has an interesting implication for system designers. Given a subsampling step size s , one could estimate the accuracy loss of potentially all classifiers using just a few of them. Based on this characteristics, hardware providers, for example, could present guidelines to app developers on the best subsampling strategy to use for different target image sizes needed by classifiers.

REMARK 2. Regardless of what classifier it is, the performance degradation of a classifier given subsampled input images is largely decided by the subsampling step size.

7.2 AdaSkip Subsampling

7.2.1 Visualization. Figure 3 visualizes the simulated subsampling procedure of the AdaSkip algorithm. The image belongs to the class `black_stork`. For AdaSkip, we set the threshold for switching between the higher resolution $s_h = 4$ (equivalent to resolution 128 × 128), the lower resolution $s_l = 8$ (equivalent to resolution 64 × 64), and threshold to 96 ³. To visualize the effect of using different subsampling resolutions, we resort to using grayscale and colored parts. As a result, as shown in Figure 3(b), a few rows on the top depicting grass and roughly the entire lower half are sampled

³In our algorithm the threshold does not vary with respect to images.

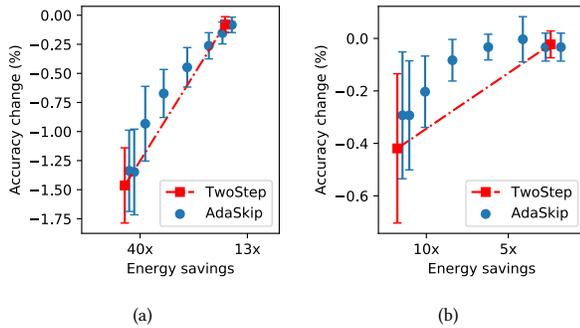


Figure 4: Energy accuracy trade-offs using AdaSkip. (a) 32×32 classifiers with $s_l = 8, s_h = 4$; (b) 48×48 classifiers with $s_l = 4, s_h = 2$

using the lower resolution. Figure 3(c) depicts the output image from a two-step subsampling procedure. The first step is the hardware subsampling with $s = 4$ (resulting in a 128×128 image), and the second step is software bilinear resampling. Figure 3(d) shows the results from bilinear resampling the image in 3(b). Comparing 3(d) with 3(c), we barely see any difference. The resulting image is indeed correctly classified by the classifier.

7.2.2 Hardware Energy Consumption. We assume our hardware support is an auxiliary to an existing image sensor. Our implementation on a Spartan-6 LX45 field-programmable-gate-array (FPGA) shows the power overhead of the 128-pixel ($s_l = 4$) AdaSkip hardware support is approximately 13.5 mW. We claim that this overhead can be significantly reduced if AdaSkip hardware support is implemented on ASIC or co-optimized with the image sensor. Kuon *et al.* estimate the power gap between ASIC and FPGA implementations of logic only designs could be as large as $5.7 \times -52 \times$ [16]. We estimate the power consumption of AdaSkip hardware support to be 2.4 mW or lower on ASIC, indicating less than 1% additional power in our energy model comparing to an image sensor without AdaSkip. In an actual image sensor, we argue that there is no need for more than 600-pixel AdaSkip support, as rarely will a classification problem need that large of a resolution. Without considering the possibility of using memory and serializing, that amounts to less than 5% of additional power. We use that as the overhead in the following discussion of energy.

7.2.3 Energy Accuracy Trade-Off. Simply by two-step subsampling, we already saved significant amount of energy. To take a step further, we apply AdaSkip in this setting. Figure 4 shows the accuracy changes using different thresholds in AdaSkip. In the figure, the y-axis is the accuracy change, and the x-axis is the relative energy consumption. Blue dots represent the results from AdaSkip, where each point represents a threshold (multiples of 32). Red square represents the baselines using two-step subsampling for s_l (left) and s_h (right) respectively.

When the threshold is 0, AdaSkip degrades to two-step but with additional hardware overhead, thus falls below the baseline. However, for almost all the other points, AdaSkip constantly achieves

better energy accuracy trade-offs. In the case of Figure 4(b), sampling at $s = 2$ sustained an average of 0 accuracy lost. That is equivalent to $4 \times$ energy savings. Using AdaSkip, we can boost the energy savings to around $7 \times$ without losing any accuracy.

8 CONCLUSION

We present a study on how subsampling affects the performance of DNN-based image classifiers. We first demonstrate that one could achieve over $15 \times$ energy savings just by subsampling while suffering almost no accuracy lost. Then we empirically show that the accuracy lost can be predicted as an approximate function of subsampling step size, regardless of what classifier it is. To achieve better accuracy when subsampling aggressively, we propose the AdaSkip algorithm. We implement AdaSkip on an FPGA and estimate the overhead in a real image sensor to be less than 5%.

9 ACKNOWLEDGMENTS

This work was supported in part by the NSF award CNS-1513306.

REFERENCES

- [1] 2018. iPhone X - Technical Specifications. <https://www.apple.com/iphone-x/specs/>. (2018). Accessed: 2018-03-01.
- [2] 2018. Xperia™XZ Specifications - Sony Mobile. <https://www.sonymobile.com/us/products/phones/xperia-xz/specifications/>. (2018). Accessed: 2018-03-01.
- [3] Evgeny Artyomov and Orly Yadid-Pecht. 2006. Adaptive Multiple-Resolution CMOS Active Pixel Sensor. *IEEE Trans. Circuits Syst.* 53-I, 10 (2006), 2178–2186.
- [4] Mark Buckler, Suren Jayasuriya, and Adrian Sampson. 2017. Reconfiguring the Imaging Pipeline for Computer Vision. In *ICCV*. 975–984.
- [5] Huaijin G. Chen, Suren Jayasuriya, Jiyue Yang, Judy Stephen, Sriram Sivaramakrishnan, Ashok Veeraraghavan, and Alyosha C. Molnar. 2016. ASP Vision: Optically Computing the First Layer of Convolutional Neural Networks Using Angle Sensitive Pixels. In *CVPR*. 903–912.
- [6] Jeng-Hau Lin et al. 2017. Binarized Convolutional Neural Networks with Separable Filters for Efficient Hardware Acceleration. In *CVPR Workshops*. 344–352.
- [7] Kawahito Shoji et al. 1997. A CMOS image sensor with analog two-dimensional DCT-based compression circuits for one-chip cameras. *J. Solid-State Circuits* 32, 12 (1997), 2030–2041.
- [8] Robert LiKamWa et al. 2013. Energy characterization and optimization of image sensing toward continuous mobile vision. In *MobiSys*. 69–82.
- [9] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *ICLR*.
- [10] Hongxiang Gu and Viswanathan Swaminathan. 2018. From Thumbnails to Summaries - A single Deep Neural Network to Rule Them All. In *ICME*. To appear.
- [11] Jia Guo and Miodrag Potkonjak. 2017. Pruning ConvNets Online for Efficient Specialist Models. In *CVPR Workshops*. 430–437.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. 770–778.
- [13] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q. Weinberger. 2018. Multi-Scale Dense Convolutional Networks for Efficient Prediction. In *ICLR*.
- [14] Sabrina E. Kemeny, Roger Panicacci, Bedabrata Pain, Larry H. Matthies, and Eric R. Fossum. 1997. Multiresolution image sensor. *IEEE Trans. Circuits Syst. Video Techn.* 7, 4 (1997), 575–583.
- [15] Alex Krizhevsky and Geoffrey Hinton. 2009. *Learning multiple layers of features from tiny images*. Master’s thesis. Department of Computer Science, University of Toronto.
- [16] Ian Kuon and Jonathan Rose. 2007. Measuring the Gap Between FPGAs and ASICs. *IEEE Trans. on CAD of Integrated Circuits and Systems* 26, 2 (2007), 203–215.
- [17] Walter D. Leon-Salas, Sina Balkir, Khalid Sayood, Michael W. Hoffman, and Nathan Schemm. 2006. A CMOS imager with focal plane compression. In *ISCAS*.
- [18] Robert LiKamWa, Yunhui Hou, Yuan Gao, Mia Polansky, and Lin Zhong. 2016. RedEye: Analog ConvNet Image Sensor Architecture for Continuous Mobile Vision. In *ISCA*. 255–266.
- [19] Yusuke Oike and Abbas El Gamal. 2013. CMOS Image Sensor With Per-Column $\Sigma\Delta$ ADC and Programmable Compressed Sensing. *J. Solid-State Circuits* 48, 1 (2013), 318–328.
- [20] Olga Russakovsky et al. 2015. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Computer Vision* 115, 3 (2015), 211–252.