

Content-Aware Convolutional Neural Networks

Yong Guo^a, Yaofu Chen^a, Mingkui Tan^{a,*}, Kui Jia^a, Jian Chen^{a,*}, Jingdong Wang^a

^aSouth China University of Technology, China

^bMicrosoft Research Asia, China

Abstract

Convolutional Neural Networks (CNNs) have achieved great success due to the powerful feature learning ability of convolution layers. Specifically, the standard convolution traverses the input images/features using a sliding window scheme to extract features. However, not all the windows contribute equally to the prediction results of CNNs. In practice, the convolutional operation on some of the windows (*e.g.*, smooth windows that contain very similar pixels) can be very redundant and may introduce noises into the computation. Such redundancy may not only deteriorate the performance but also incur the unnecessary computational cost. Thus, it is important to reduce the computational redundancy of convolution to improve the performance. To this end, we propose a Content-aware Convolution (CAC) that automatically detects the smooth windows and applies a 1×1 convolutional kernel to replace the original large kernel. In this sense, we are able to effectively avoid the redundant computation on similar pixels. By replacing the standard convolution in CNNs with our CAC, the resultant models yield significantly better performance and lower computational cost than the baseline models with the standard convolution. More critically, we are able to dynamically allocate suitable computation resources according to the data smoothness of different images, making it possible for content-aware computation. Extensive experiments on various computer vision tasks demonstrate the superiority of our method over existing methods.

Keywords: Convolution, Neural Networks, Redundancy Reduction.

*Corresponding author

Email addresses: guo.yong@mail.scut.edu.cn (Yong Guo),
sechenyaofu@mail.scut.edu.cn (Yaofu Chen), mingkuitan@scut.edu.cn (Mingkui Tan),
kuijia@scut.edu.cn (Kui Jia), ellachen@scut.edu.cn (Jian Chen),
jingdw@microsoft.com (Jingdong Wang)

1. Introduction

Recently, convolutional neural networks (CNNs) have achieved remarkable performance in many computer vision tasks, including image classification [1, 2], face recognition [3, 4, 5], semantic segmentation [6, 7, 8], and object detection [9, 10]. Moreover, deep CNNs have also become the workhorse of many other tasks and real-world applications beyond computer vision, such as speech recognition [11, 12] and natural language processing [13, 14].

One of the key factors behind the success of CNNs lies in the powerful feature learning ability of convolution layers. Typically, the standard convolution transforms the input images/features into a set of windows and exploits a sliding window manner to extract features over them [15]. However, not all the windows contribute equally to the prediction results of CNNs. As shown in Figure 1, the input images/features often contain a lot of smooth windows that consist of very similar pixels. These windows may contain very limited information about the data [16, 17] since a similar pattern may also appear in the surrounding areas. As a result, the computation on smooth windows may be very redundant. More critically, performing convolution on these windows may also introduce noises into the computation and thus deteriorate the performance (See results in Tables 1 and 2). Thus, it is important and necessary to reduce the computational redundancy of convolution to improve the performance.

Regarding this issue, existing methods improve the convolutional operation by reducing the redundant communications among different channels of feature maps [18, 19, 20] or reducing the spatial size of some of the redundant channels [21]. Specifically, group convolution [18] and depthwise separable convolution [19, 20] divide the channels of feature maps into multiple groups and perform convolution independently over each group. Recently, Chen *et al.* propose the Octave Convolution (OctConv) [21] method, which downscales the feature maps in some redundant channels into smaller sizes to reduce the computational cost. However, these methods only focus on the redundancy inside the channels of feature maps but ignore the spatial redundancy of the pixels in each window. Moreover, existing methods perform the same computation on the

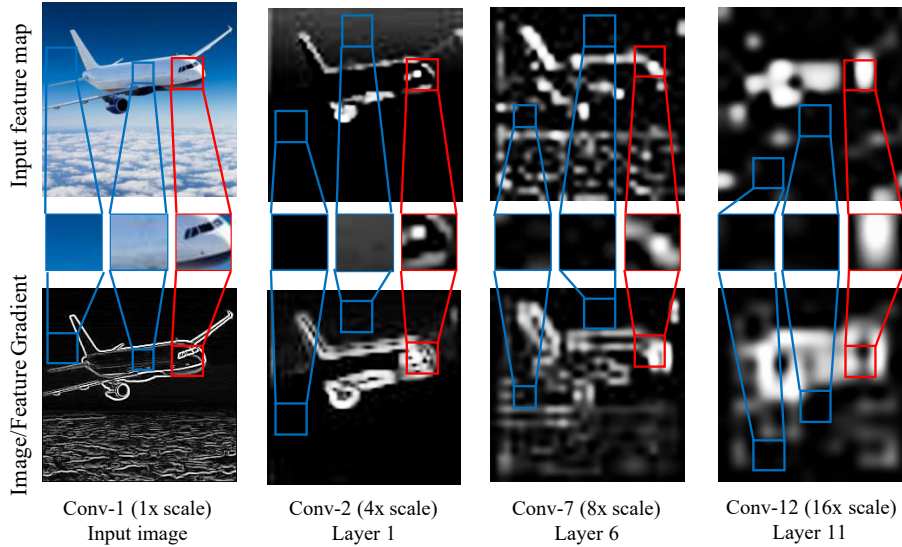


Figure 1: Demonstration of computational redundancy in the input feature maps of different layers of ResNet18 (pretrained on ImageNet). The top row and bottom row show the input feature maps and the corresponding gradient in different layers, respectively. Red boxes denote the sharp windows that contain the main content of the image. Blue boxes denote the smooth windows that contain limited information about the image and require redundant computation.

samples with different spatial redundancy, which makes the prediction not optimal and also very inefficient.

In this paper, we seek to reduce the computational redundancy on smooth windows to improve the performance of convolution. To this end, we propose a Content-aware Convolution (CAC) that uses a 1×1 convolution to replace the computation of original $k \times k$ convolution on smooth windows. In this sense, we are able to effectively avoid the redundant computation to improve the feature learning ability of convolution. To obtain the weights for the 1×1 convolutional kernels, we spatially aggregating the $k \times k$ kernel by summing up all the kernel parameters (See detailed analysis in Section 4.2). In order to automatically detect the smooth windows, we propose an effective training method that seeks for a trade-off between model performance and computational cost. More critically, CAC dynamically allocates computation resources for different samples based on the smoothness of their contents. Therefore, we are able to perform content-aware computation. In practice, our CAC models yield significantly better performance and

lower computational cost than the models with the standard convolution. Extensive experiments on different computer vision tasks demonstrate the superiority of our method over existing methods.

In this paper, we make the following contributions.

- We propose a Content-aware Convolution (CAC) method that replaces the original $k \times k$ kernel with a 1×1 kernel on the smooth windows to improve the performance of convolution. With CAC, we are able to effectively reduce the computational redundancy of convolution and significantly improve the performance.
- We propose an effective training method to automatically detect the smooth windows for each layer. To achieve this goal, we solve a multi-objective optimization problem to find a trade-off between model performance and computational cost.
- Equipped with CAC, the resultant models achieve content-aware computation by dynamically allocating computation resources to different samples according to the data smoothness. Extensive experiments on different computer vision tasks demonstrate the effectiveness of the proposed method.

2. Related Work

Recently, many efforts have been made to reduce the redundancy of deep networks, including channel pruning, network quantization, and energy-efficient model design.

2.1. Channel Pruning

Channel pruning is one of the predominant approaches for deep network compression. Li *et al.* [22] measure the importance of different channels by computing the sum of absolute values of weights to conduct channel selection. Hu *et al.* [23] use the average percentage of zeros (APoZ) to select important channels. Several training based methods [24, 25] have been proposed to automatically identify the redundant channels by introducing a sparsity regularizer in the training objective. The reconstruction methods [26, 27] seek to solve the channel pruning problem by minimizing the reconstruction error between the feature maps of the pretrained model and the compressed

model. Recently, Zhuang *et al.* [28] propose a discrimination-aware channel pruning (DCP) method to choose the channels that contribute to the discriminative power and obtain state-of-the-art results. Based on DCP, Zhuang *et al.* [29] further propose a discrimination-aware kernel pruning (DKP) method by removing the redundant kernels according to the discrimination power. However, these methods only focus on the redundancy in model parameters but ignore the redundancy incurred by the input data.

2.2. Network Quantization

Network quantization aims to convert the pretrained full-precision convolution networks into the low-precision versions to reduce the computational cost. Recently, Song *et al.* [30] propose a three-stage deep compression pipeline, including pruning, trained quantization, and Huffman coding. DoReFa-Net [31] seeks to quantize the full precision weights, activations, and gradients to the low bit ones for deep networks. In ternary weight networks (TWNs) [32, 33], the parameters are constrained to +1, 0, and -1, and the model achieves higher accuracy than binary neural networks. Similar to channel pruning, network quantization also reduces the redundancy in model parameters and may yield limited performance.

2.3. Energy-efficient Model Design

Many energy-efficient modules have been proposed to reduce the computational cost of deep networks. Specifically, sparse convolution [34, 35] zeros out a large number of parameters to reduce the model size. Group convolution [18] and depthwise separable convolution [19] divide the input channels into groups to reduce the redundant communications among different groups. NAT [36, 37] replaces redundant operations with identity mapping or directly remove them to obtain efficient models.

Related to our method, Li *et al.* [38] propose a Region Convolution (RC) that reduces the computational redundancy for semantic segmentation models. Specifically, given an input feature map, RC performs convolutions on the regions of hard pixels and discards the easy pixels according to the confidence of the predicted mask. However, it has some underlying limitations. First, RC relies on the predicted confidence of all pixels to construct the mask and cannot be applied to tasks without dense prediction,

e.g., image classification. Second, RC completely discards the easy regions and may influence the features learned in the deeper layers. Unlike RC, our CAC preserves the information in all regions/pixels and can be applied to most computer vision tasks, *e.g.*, image classification, semantic segmentation, and object detection.

Very recently, Chen *et al.* propose the octave convolution (OctConv) [21] method, which reduces the spatial resolution of some low-frequency feature maps to reduce the computational complexity. However, it has two major limitations. First, the low-frequency feature maps are predefined before training rather than detected according to the input data. Second, OctConv only considers the redundancy in the channels of the feature maps but ignores the redundancy in the data content, *e.g.*, pixels. Compared to OctConv, our CAC automatically detects the sharp/smooth windows from the input data to achieve content-aware computation. Moreover, CAC considers the pixel level redundancy caused by the input data.

3. Notations and Problem Definition

In this paper, we assume that the convolution has the stride of 1 and is performed with padding to guarantee that the output feature maps have the same size as the input feature maps. In this paper, we consider the squared input images and feature maps. For simplicity, we consider one single-channel input feature map $\mathbf{X} \in \mathbb{R}^{n \times n}$ and one convolutional kernel $\mathbf{W} \in \mathbb{R}^{k \times k}$, where n and k denote the sizes of the feature maps and kernels, respectively. Thus, the standard convolution can be written as

$$\mathbf{Y} = \mathbf{X} \otimes \mathbf{W}, \quad (1)$$

where \otimes denotes the convolutional operator.

In practice, the computation of convolution is often converted to the matrix-matrix or matrix-vector multiplication [39]. Given a convolution with a $k \times k$ kernel and an input feature map $\mathbf{X} \in \mathbb{R}^{n \times n}$, there are n^2 windows that are convolved by the kernel [40]. In this sense, we can represent \mathbf{X} by a set of windows

$$\Psi := \{\mathbf{Q}_i \in \mathbb{R}^{k \times k} \mid i = 1, \dots, n^2\}, \quad (2)$$

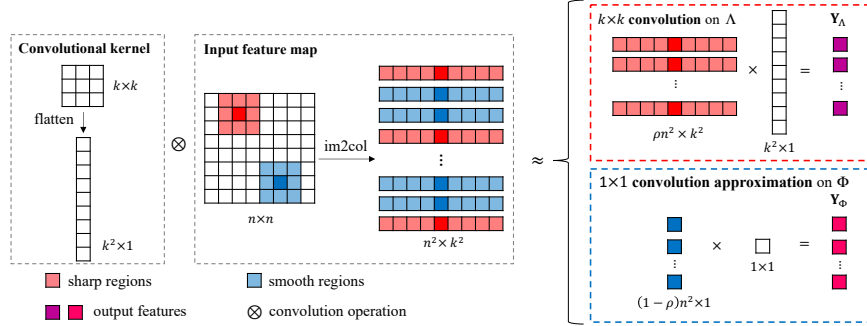


Figure 2: The computation method of Content-aware Convolution. We first divide the input feature maps into two parts, namely the sharp windows (red boxes) and smooth windows (blue boxes). Then, we perform $k \times k$ convolution on sharp windows and 1×1 convolution on smooth windows.

where \mathbf{Q}_i denotes the i -th window in Ψ . For any window \mathbf{Q}_i , we can reshape it into a vector $\mathbf{p}_i = \text{vec}(\mathbf{Q}_i) \in \mathbb{R}^{k^2}$. For convenience, we define $\mathbf{P} := [\mathbf{p}_1, \dots, \mathbf{p}_{n^2}] \in \mathbb{R}^{k^2 \times n^2}$ and $\mathbf{w} := \text{vec}(\mathbf{W}) \in \mathbb{R}^{k^2}$. Note that the transformation from \mathbf{X} to \mathbf{P} is often called `im2col` (See Fig. 2). Thus, Eqn. (1) can be written as a matrix-vector multiplication:

$$\mathbf{Y} = \text{vec2mat}(\mathbf{P}^T \mathbf{w}), \quad (3)$$

where the function `vec2mat`(\cdot) denotes the operation to reshape a vector to a matrix.

Clearly, the complexity of the matrix-vector multiplication is $O(n^2 k^2)$, which can be very expensive when n and/or k are very large. However, some of the windows are very smooth and only contain limited information. More critically, performing convolution on smooth windows may also introduce noises into the computation and thus hamper the performance. Thus, it is necessary and important to reduce the computational redundancy on smooth windows to improve the performance of convolution.

4. Proposed Method

In this paper, we propose a Content-aware Convolution (CAC) that replaces the original kernel with a 1×1 kernel to perform convolution on smooth windows. Moreover, we present an effective method to automatically detect smooth windows. We show the overall scheme in Fig. 2 and the detailed computation method of CAC in Algorithm 1.

4.1. Motivation

The standard convolution convolves the input images/features by transforming them into a set of windows and adopts a sliding window scheme to extract features [40]. However, not all the windows contribute equally to the prediction results of deep networks. As shown in Fig. 1, there are a large number of smooth windows in the feature maps of each layer. These windows often contain very similar pixels and come with very limited information about the data. As a result, performing convolution on smooth windows can be redundant. More critically, the computational redundancy on these smooth windows may also introduce noises into the computation and thus deteriorate the performance.

Instead of convolving all the windows using the same kernel, we seek to perform different convolutional computations on the windows according to their smoothness. Specifically, we first recognize the sharp and smooth windows. Then, we perform the standard convolution on the sharp windows and perform the convolution with a smaller kernel of 1×1 on a single pixel of the smooth windows to reduce the computational redundancy. Since the computation depends on the content of the input data in each layer, we call our method **Content-aware Convolution (CAC)**.

4.2. Content-aware Convolution

Given an input feature map or image, we divide the whole window set Ψ into two disjoint subsets, namely the **sharp window set** Λ and the **smooth window set** Φ . We will illustrate how to detect smooth/sharp windows from Ψ in Section 4.3. To reduce the computational redundancy on smooth windows, we seek to use a 1×1 kernel to replace the original large kernel. Given an input window $\mathbf{Q}_i \in \mathbb{R}^{k \times k}$ and a single-channel kernel $\mathbf{W} \in \mathbb{R}^{k \times k}$, the output of a convolution layer y_i can be computed by

$$y_i = \mathbf{Q}_i \otimes \mathbf{W} = \mathbf{p}_i^\top \mathbf{w} = \sum_{j=1}^{k^2} p_j w_j, \quad (4)$$

where \mathbf{p}_i denotes vector presentation of \mathbf{Q}_i and p_j denotes the j -th element of \mathbf{p}_i . If \mathbf{Q}_i is a smooth window, it implies that all the elements of the window should have very

Algorithm 1: Content-Aware Convolution (CAC).

- Require:** Input feature map $\mathbf{X} \in \mathbb{R}^{n \times n}$;
Convolutional kernel $\mathbf{W} \in \mathbb{R}^{k \times k}$;
Learnable parameters kernel γ and β ;
Set of input windows $\Psi := \{\mathbf{Q}_i \in \mathbb{R}^{k \times k} \mid i = 1, \dots, n^2\}$.
- 1: Compute the 1×1 kernel w_Φ using Eqn. (6);
 - 2: Compute the average feature map using Eqn. (9);
 - 3: Compute the gradient of feature maps \mathbf{G} using Eqn. (10);
 - 4: Compute the score map based on \mathbf{G} :
 $\mathbf{M} = \text{Sigmoid}(\gamma \mathbf{G} + \beta)$;
 - 5: Obtain the set of sharp windows:
 $\Lambda = \{\mathbf{Q}_i \mid M_i > 0.5\}$;
 - 6: Obtain the set of smooth windows:
 $\Phi = \Psi \setminus \Lambda$;
 - 7: Perform convolution on Λ :
 $\mathbf{Y}_\Lambda = \text{Conv}(\Lambda; \mathbf{W})$;
 - 8: Perform convolution on Φ :
 $\mathbf{Y}_\Phi = \text{Conv}(\Phi; w_\Phi)$;
 - 9: Combine \mathbf{Y}_Λ and \mathbf{Y}_Φ to obtain the final output:
 $\mathbf{Y} = \text{Combine}(\mathbf{Y}_\Lambda, \mathbf{Y}_\Phi)$.
-

similar values, *i.e.*, for $\forall m, n \in \{1, \dots, k^2\}, p_m \approx p_n$. Therefore, it follows that

$$\mathbf{p}_i^\top \mathbf{w} = \sum_{j=1}^{k^2} p_j w_j \approx \bar{p} \cdot \sum_{j=1}^{k^2} w_j, \quad (5)$$

where \bar{p} can be the average value of all p_j ($\bar{p} = \frac{1}{k^2} \sum_{j=1}^{k^2} p_j$) or any element of this window. In this paper, we choose the center element to compute \bar{p} . Relying on Eqn. (5), we approximate the original $k \times k$ convolution kernel using a 1×1 kernel:

$$w_\Phi = \sum_{j=1}^{k^2} w_j. \quad (6)$$

Note that the computation on similar pixels may introduce noises into the computation. Our CAC performs convolution on a single pixel in a window and thus effectively reduces the impact of the noisy information. In this sense, the CAC based models are often more robust than the models built with the standard convolution (See results in Table 5). Moreover, the computational cost of CAC on smooth windows can be reduced to $1/k^2$ of the cost with the $k \times k$ kernels.

Given the sharp windows Λ and smooth windows Φ , we obtain the output of CAC by performing a $k \times k$ convolution and a 1×1 convolution on Λ and Φ , respectively. Let $\mathbf{W} \in \mathbb{R}^{k \times k}$ be the parameters of a single-channel convolutional kernel, w_Φ be the 1×1 kernel obtained by Eqn. (6). The computation on Λ and Φ can be formulated by

$$\mathbf{Y}_\Lambda = \text{Conv}(\Lambda; \mathbf{W}), \quad \mathbf{Y}_\Phi = \text{Conv}(\Phi; w_\Phi). \quad (7)$$

Then, we combine \mathbf{Y}_Λ and \mathbf{Y}_Φ to obtain the final output according to the relative positions of the windows.

$$\mathbf{Y} = \text{Combine}(\mathbf{Y}_\Lambda, \mathbf{Y}_\Phi). \quad (8)$$

4.3. Sharp and Smooth Window Recognition

Based on the smoothness of windows, we propose an effective method to automatically detect the sharp/smooth windows. In this paper, we measure the data smoothness using the gradients of the input images or features.

For any layer of a deep network, not all the channels are useful and some of them are noisy or irrelevant to the final prediction results [28]. As a result, the features in these channels can be very noisy and thus may hamper the final prediction results [41]. Regarding this issue, we seek to compute the average feature map over different channels to alleviate the influence of noisy features:

$$\bar{\mathbf{X}} = \frac{1}{m} \sum_{i=1}^m \mathbf{X}_i, \quad (9)$$

where m denotes the number of channels. Then, we compute the gradient of the averaged feature map to compare the sharpness of different windows. In this paper, we use the Sobel operator [42] to compute the gradient by performing two 1-d convolutions along

the x- and y-axis, respectively:

$$\mathbf{G}_x = \bar{\mathbf{X}} \otimes \begin{bmatrix} -1 & 0 & +1 \end{bmatrix} \otimes \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}, \quad \mathbf{G}_y = \bar{\mathbf{X}} \otimes \begin{bmatrix} 1 & 2 & 1 \end{bmatrix} \otimes \begin{bmatrix} -1 \\ 0 \\ +1 \end{bmatrix}.$$

Thus, the total gradient becomes

$$\mathbf{G} = \sqrt{\mathbf{G}_x^2 + \mathbf{G}_y^2}. \quad (10)$$

Given the feature map with m channels, computing gradients on the averaged feature map only yields $1/m$ cost of the computation on all the channels. Compare to the cost of convolution, the cost of computing gradients can be negligible in practice.

Based on the computed gradients, we may divide the windows into sharp and smooth windows according to some threshold. However, such a threshold has to be carefully selected for each layer, making it very time consuming and labor-intensive. To address this issue, we propose a learnable module to automatically discriminate the sharp windows from the smooth windows. Specifically, we exploit an affine function to transform the gradients and then apply the Sigmoid function to compute the probability of a window being sharp. The probability map \mathbf{M} can be computed by

$$\mathbf{M} = \text{Sigmoid}(\gamma \mathbf{G} + \beta), \quad (11)$$

where γ and β are trainable parameters.

Here, we consider the windows with the probability larger than 0.5 as sharp windows and the other windows as smooth windows. Formally, the sets of sharp and smooth windows can be represented by:

$$\Lambda = \{\mathbf{Q}_i \mid M_i > 0.5\}, \quad \Phi = \Psi \setminus \Lambda, \quad (12)$$

where M_i is the score of the window \mathbf{Q}_i and Ψ denote the set of all the windows. By changing a hard threshold manner to a learnable scheme, the model is able to adjust γ

and β to find the optimal number of smooth windows to perform 1×1 convolution. We will show the training method for γ and β in Section 4.4.

4.4. Training Method

Note that CAC seeks to find a number of smooth windows to perform 1×1 convolution to improve the performance. Although we can reduce the computational cost, performing 1×1 convolution on too many windows may also hamper the performance (See results in Table 6). To find a good trade-off between model performance and computational cost, we propose to solve a multi-objective optimization problem.

Let M be the CAC-based model to be trained, M^b be the baseline model with the standard convolution, and $c(*)$ be the function to measure the computational cost of deep models, *e.g.*, the number of multiply-adds (MAdds). To train CAC-based models, we use the weighted product method¹ to build the objective:

$$L = \ell(M) \left(\frac{c(M)}{c(M^b)} \right)^\lambda, \quad (13)$$

where $\ell(M)$ denotes the standard loss function w.r.t. M (*e.g.*, the cross-entropy loss for classification models) and $\lambda \geq 0$ is a constant weight factor. When $\lambda = 0$, the objective is reduced to the standard loss for a specific task. When $\lambda > 0$, we seek to find a promising trade-off between model performance and computational cost. In this paper, to obtain a good balance, we use λ to control the importance of computational cost (See discussions on λ in Section 7.2).

5. More Discussions

In this section, we first analyze the computational complexity of the proposed CAC convolution in Section 5.1. Then, we discuss the differences between the proposed methods and existing methods in Section 5.2.

¹We use the weighted product method because it is easy to customize for different models. The weighted sum method is also appropriate.

5.1. Computational Complexity Analysis

To analyze the computational complexity of the proposed CAC, we consider the more general case in which a convolution layer contains multiple channels. Let $\mathcal{X} \in \mathbb{R}^{n \times n \times c_{in}}$ be the input feature maps of a convolution layer and the convolutional kernel be $\mathcal{W} \in \mathbb{R}^{k \times k \times c_{in} \times c_{out}}$, where c_{in} and c_{out} are the numbers of input and output channels, respectively. Then, the convolution in a standard convolution layer can be computed by

$$\mathcal{O} = \mathcal{X} \otimes \mathcal{W}, \quad \mathcal{O} \in \mathbb{R}^{n \times n \times c_{out}}, \quad (14)$$

where \otimes denotes the convolution operation. The number of multiply-adds (MAdds) required by the standard convolution is given by:

$$\Omega_{\text{Conv}} = c_{in} \cdot c_{out} \cdot k \cdot k \cdot n \cdot n. \quad (15)$$

Given a specific proportion of sharp windows (denoted by ρ), the computational complexity consists of three parts. **First**, there are $\rho \cdot n^2$ sharp windows in Λ where we perform the standard convolution using the $k \times k$ kernel. Thus, the complexity of the first part becomes $\rho \Omega_{\text{conv}}$. **Second**, we perform 1×1 convolution on $(1-\rho) \cdot n^2$ smooth windows in Φ , each of which only requires $1/k^2$ complexity of the standard $k \times k$ convolution. Therefore, the computational complexity of the second part is $(1-\rho)\Omega_{\text{conv}}/k^2$. **Third**, we perform two 1-d convolutions with a single output channel to compute the gradients along the x- and y-axis, respectively. Thus, there are a total of four 1×3 or 3×1 convolutions to compute the gradient.

Compared to the standard convolution with c_{in} input channels and c_{out} output channels, the complexity of computing gradients in the third part is

$$4 \cdot \underbrace{\frac{3}{k^2} \cdot \frac{1}{c_{in}c_{out}} \Omega_{\text{Conv}}}_{\text{computing gradient}} + \underbrace{\frac{1}{k^2 c_{in} c_{out}} \Omega_{\text{conv}}}_{\text{linear transformation}} = \frac{13}{k^2 c_{in} c_{out}} \Omega_{\text{conv}}. \quad (16)$$

As a result, the total computational complexity of the CAC convolution becomes:

$$\Omega_{\text{CAC}} = \left(\underbrace{\rho}_{k \times k \text{ conv}} + \underbrace{\frac{1-\rho}{k^2}}_{1 \times 1 \text{ conv}} + \underbrace{\frac{13}{k^2 c_{in} c_{out}}}_{\text{computing score map}} \right) \Omega_{\text{Conv}}. \quad (17)$$

To accelerate the computation of the convolutions, according to Eqn. (17), we have to satisfy the condition such that $\Omega_{\text{CAC}}/\Omega_{\text{Conv}} \leq 1$. In this sense, we can obtain the upper bound of the ratio ρ :

$$\rho \leq \bar{\rho} = 1 - \frac{13}{(k^2 - 1) \cdot c_{in} \cdot c_{out}}. \quad (18)$$

Specifically, for a 3×3 convolution (*i.e.*, $k = 3$), the upper bound is $\bar{\rho} = 1 - \frac{13}{8 \cdot c_{in} \cdot c_{out}}$. Taking ResNet [1] as an example, the number of channels c_{out} ranges from 16 to 512. In this sense, the ratio only needs to be $\rho < 99.3\%$ when we substitute the smallest value $c_{in} = c_{out} = 16$ into Eqn. (18). Thus, there is considerable potential to accelerate the computation of the standard convolution.

5.2. Differences from Existing Methods

The proposed CAC method has several essential differences from existing methods. **First**, the standard convolution performs convolution using a general kernel on all the windows and ignores the inherent redundancy, which may hamper the performance (See results in Tables 1 and 2). In contrast, the proposed CAC performs different convolutions on these windows to reduce the computational redundancy and improve the performance. **Second**, existing methods perform the same computation on the samples with different spatial redundancy. Unlike these methods, our CAC adopts a content-aware computation scheme that dynamically allocates suitable computational resources for different samples according to their data smoothness. It is worth noting that our CAC is more robust to the samples with adversarial perturbations than existing convolution methods. The main reason is that CAC replaces the large kernel convolution with a 1×1 convolution on smooth windows, which effectively reduces the influence incurred by the noises/attacks in these windows.

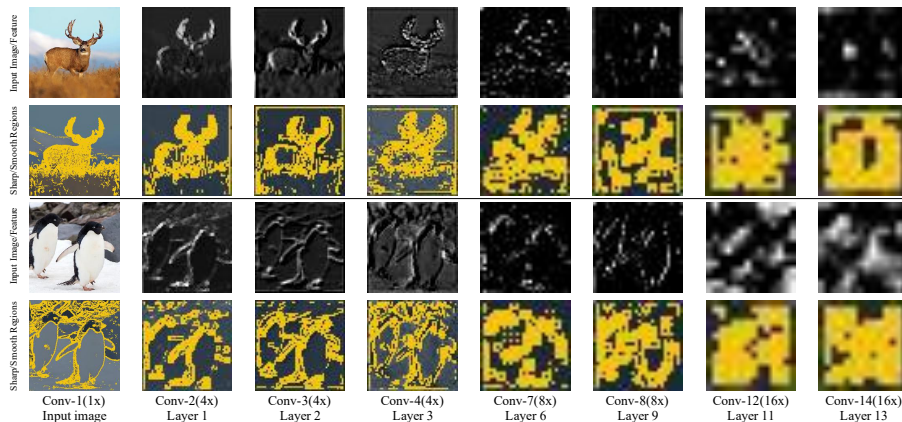


Figure 3: Visualization of the feature maps of different layers in CAC-ResNet18 on ImageNet. For each image, the top row shows the feature map of different layers and the bottom row shows the corresponding map of sharp windows detected by CAC. In the bottom row, yellow regions denote the sharp windows to perform 3×3 convolution and the dark regions denote the smooth ones to perform 1×1 convolution. We scale the feature maps of different layers to the same spatial size for better visualization.

6. Experiments

In this section, we use CAC to accelerate two popular convolution methods, namely the standard convolution and the octave convolution (OctConv) [21]. We apply CAC to various architectures and demonstrate the performance on three computer vision tasks, including image classification, semantic segmentation, and object detection. All implementations are based on PyTorch².

We organize the experiments as follows. First, we show the visual interpretation of each CAC layer in Section 6.1. Second, we evaluate our CAC on image classification tasks in Section 6.2. Third, we apply our CAC method to semantic segmentation models and evaluate the proposed method in Section 6.3. Fourth, we conduct experiments to show the effectiveness of our CAC method on object detection tasks in Section 6.4. Finally, we investigate the effect of the hyperparameter λ in Section 7.2.

²The implementation of the proposed CAC method is available at <https://github.com/guoyongcs/CAC>.

6.1. Visual Interpretation of CAC Convolution

To better understand the proposed CAC method, we visualize the feature maps and the corresponding masks of the sharp windows of different CAC layers inside deep networks. In this experiment, we take the CAC-ResNet18 model as an example and show the results in Fig. 3.

From Fig. 3, the sharp windows (marked in yellow) are often located at the edges and contain the main information about the object. However, the smooth windows (marked in black) are often very smooth areas that only contain little information. As discussed in Section 4, when the input windows are very smooth, it is not necessary to use a large kernel to perform convolution. Thus, we can reduce the computational cost of convolution on the smooth windows using a 1×1 kernel to approximate the original output. In this way, our CAC method can greatly reduce the computational complexity without a loss of information. More critically, since the input images or features may have different numbers of smooth windows, the resultant CAC models can dynamically allocate suitable computation power to different input images. Thus, our CAC models can perform content-aware computation to improve the performance of convolution.

6.2. Experiments on Image Classification

In this experiment, we consider two popular convolution methods as the baseline methods, namely, the standard convolution and the OctConv [21]. We apply the proposed CAC method to various image classification models, including ResNet [1], DenseNet [43], and ShuffleNetV2 [44].

6.2.1. Datasets and Implementation Details

We conduct experiments on two benchmark image classification datasets, including CIFAR-10 [45], and ImageNet [46]. CIFAR-10 consists of 50k training samples and 10k testing images with 10 classes. ImageNet contains 123k training samples and 50k testing images for 1,000 classes.

We follow the settings in [1] and use SGD with nesterov [47] for the optimization. The momentum and weight decay are set to 0.9 and 0.0001, respectively. On CIFAR-10, we train the models for 400 epochs using a mini-batch size of 128. The learning rate is

Table 1: Comparisons of different convolutions in terms of both computational complexity and testing error based on various architectures on CIFAR-10. “\” denotes the missing results of the models to which OctConv cannot be applied.

Conv Type	ResNet20		ResNet32		ResNet56		DenseNet121		ShuffleNetV2	
	#MAdds (M)	Error (%)	#MAdds (M)	Error (%)	#MAdds (M)	Error (%)	#MAdds (M)	Error (%)	#MAdds (M)	Error (%)
Standard Conv	40.93	8.75	69.12	7.51	126.08	6.97	888.51	4.78	45.04	7.25
CAC	25.79	8.37	46.58	7.25	87.27	6.51	733.82	4.60	41.33	7.13
OctConv [21]	26.33	8.77	44.02	8.07	61.72	7.58	403.80	5.52	\	
CAC-OctConv	19.49	8.60	34.39	7.63	52.12	6.99	356.63	4.68		

initialized to 0.1 and divided by 10 at epochs 160 and 240, respectively. On ImageNet, we train the models for 90 epochs with a mini-batch size of 256. The learning rate is started at 0.1 and divided by 10 at epochs 30 and 60, respectively. We train the CAC based models with $\lambda = 0.3$. We use the number of multiply-adds (MAdds) to measure the computational complexity of deep models. Based on sharp window set Λ and the set of all the windows Ψ , we compute the ratio of sharp windows inside a convolution layer by $\rho = |\Lambda|/|\Psi|$, where $|\cdot|$ denotes the cardinality of a set. In general, a lower ratio implies that the more windows would be convolved with a 1×1 kernel.

6.2.2. Comparisons on CIFAR-10

In this experiment, we evaluate our CAC method on a small dataset CIFAR-10. From Table 1, the proposed CAC method greatly accelerates ResNet and DenseNet models equipped with different convolution types. Specifically, for both the models with the standard convolution and OctConv, our CAC consistently yields significantly better performance and lower computational cost. Moreover, we also show the ratios of sharp windows ρ of each layer based on several models in Fig. 4. From this figure, deep layers tend to have larger ratios than shallow layers due to their better representation ability and smaller feature map, yielding a smaller risk of containing smooth windows. These results show that the proposed CAC removes the redundancy caused by the smooth windows in each layer.

We also consider very compact models, *e.g.*, ShuffleNetV2, which mainly consists of 1×1 group convolution. However, since OctConv requires information exchange among groups, it would destroy the computation of group convolution and thus cannot be directly applied to ShuffleNetV2. Thus, we only compare the performance of the

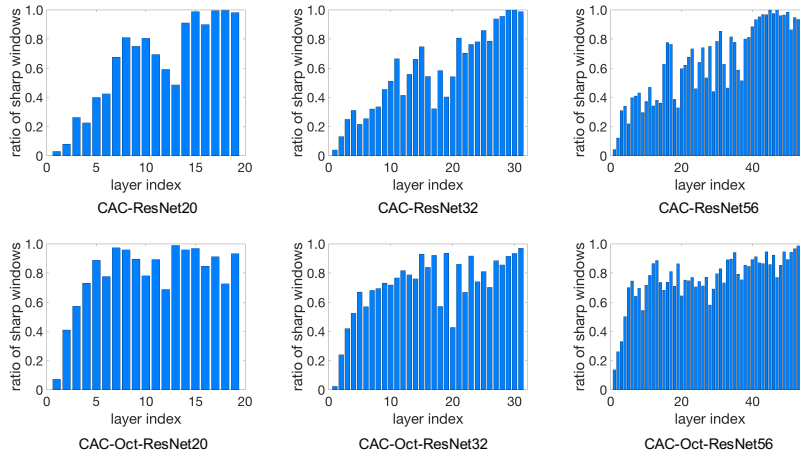


Figure 4: Visualization of the ratios of sharp windows for different layers in ResNet20, ResNet32, and ResNet56 on CIFAR-10. We compute the ratios by averaging the ratios over 10,000 testing samples.

models equipped with the standard convolution and the proposed CAC. Even with such a compact model, our CAC further improves the validation accuracy and reduces the redundancy in the model. These results demonstrate the effectiveness of our method.

6.2.3. Comparisons on ImageNet

We also evaluate our method on a large-scale dataset ImageNet. Similar to the experiments on CIFAR-10, we apply our CAC to improve both the standard convolution and OctConv. In this experiment, we consider ResNet, DenseNet and ShuffleNetV2 as the baseline model. The results are shown in Table 2.

From Table 2, our CAC based models significantly outperform the baseline models with different architectures in terms of Top-1 and Top-5 error. More critically, the resultant models often have lower computational cost and thus become more compact. These results demonstrate the superiority of the proposed CAC method over the existing methods. We also show the ratios ρ of each layer for ImageNet models in Fig. 5. From this figure, due to the training difficulty on a large-scale dataset, ImageNet models are often hard to compress without performance degradation [25, 27, 28], yielding larger ratios of the intermediate layers than the CIFAR-10 models.

Table 2: Comparisons of different convolutions in terms of both computational complexity and validation error on ImageNet. “\” denotes the missing results of the models to which OctConv cannot be applied.

Conv Type	ResNet18			ResNet50			DenseNet121			ShuffleNetV2		
	#MAAdds (G)	Error (%)		#MAAdds (G)	Error (%)		#MAAdds (G)	Error (%)		#MAAdds (G)	Error (%)	
		Top-1	Top-5		Top-1	Top-5		Top-1	Top-5		Top-1	Top-5
Conv	1.81	30.36	11.02	4.09	24.01	7.07	2.83	25.35	7.83	0.15	30.64	11.68
CAC	1.41	30.19	10.87	3.75	23.79	6.81	2.52	24.49	7.37	0.13	30.13	11.27
OctConv [21]	1.14	29.64	10.48	2.37	23.27	6.55	1.37	25.68	7.90	\	\	\
CAC-OctConv	0.96	29.43	10.27	2.19	23.05	6.37	1.28	24.82	7.61	\	\	\

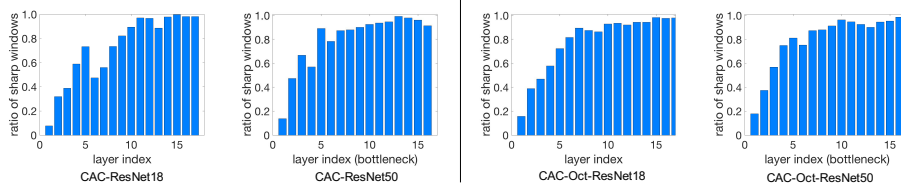


Figure 5: Visualization of the ratios of sharp windows for different layers in ResNet18 and ResNet50 on ImageNet. The layer index for ResNet50 denotes the index of the bottleneck block. We compute the ratios for different layers by averaging the samples in the validation set.

6.3. Experiments on Semantic Segmentation

We further apply the proposed CAC to semantic segmentation models, *e.g.*, fully convolutional network (FCN) [48]. We compare the performance of the models with and without CAC based on a benchmark dataset PSACAL VOC 2012 [49].

6.3.1. Compared Methods

We adopt FCN as the baseline model and apply our CAC to show the effectiveness of CAC. In this experiment, we compare our CAC with a strong baseline region convolution (RC). Moreover, we also consider several semantic segmentation methods as the baselines, including BONN-SVR [50], O₂P [51], SDS [52], and MSRA-CFM [53].

6.3.2. Datasets and Implementation Details

We conduct experiments on the benchmark semantic segmentation dataset PASCAL VOC 2012, which consists of 1,464 training images and 1,449 validation images. We measure the performance using the commonly used metric, *i.e.*, the mean intersection over union (mIoU), which computes the percent between the intersection and union of the ground truth segmentation mask and the prediction mask.

In this experiment, we use the ImageNet pretrained model as the backbone model, *e.g.*, ResNet18. Following the setting in [54], we make some modifications to adapt the

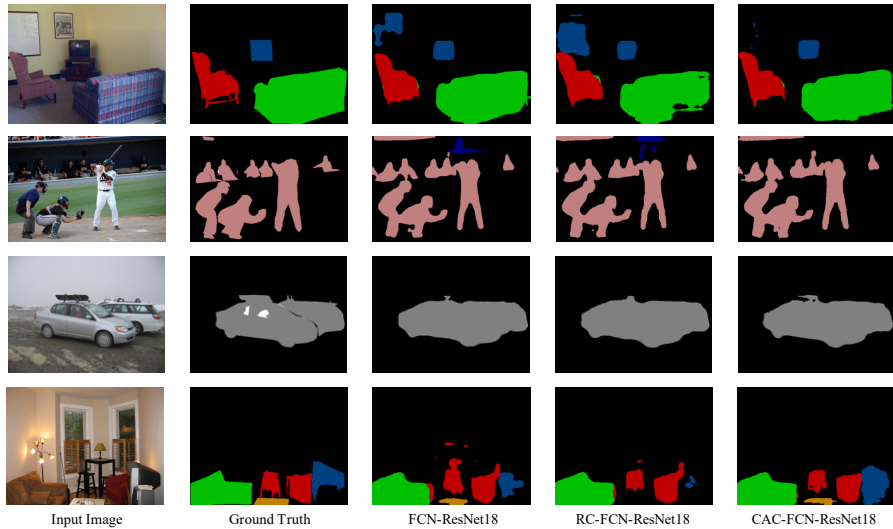


Figure 6: Visual comparison of the segmentation masks produced by different methods. The first and the second columns show the input images and the corresponding ground truth segmentation masks, respectively. The last three columns show the segmentation masks predicted by different models.

Table 3: Comparisons of different models on each class for semantic segmentation. We adopt the FCN-ResNet18 model as the baseline model. “-” denotes the results that are not reported.

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mIoU	#MAdds (G)
BONN-SVR [50]	54.3	23.9	39.5	35.3	42.6	65.4	53.5	46.1	15.0	47.4	30.1	33.9	48.8	54.4	46.4	28.8	51.3	26.2	44.9	37.2	43.3	-
O ₂ P [51]	64.0	27.3	54.1	39.2	48.7	56.6	57.7	52.5	14.2	54.8	29.6	42.2	58.0	54.8	50.2	36.6	58.6	31.6	48.4	38.6	47.8	-
SDS [52]	63.3	25.7	63.0	39.8	59.2	70.9	61.4	54.9	16.8	45.0	48.2	50.5	51.0	57.7	63.3	31.8	58.7	31.2	55.7	48.5	51.6	-
MSRA-CFM [53]	75.7	26.7	69.5	48.8	65.6	81.0	69.2	73.3	30.0	68.7	51.5	69.1	68.1	71.7	67.5	50.4	66.5	44.4	58.9	53.5	61.8	-
FCN [48]	83.5	29.7	68.7	59.8	50.2	80.3	71.9	71.8	28.6	59.6	44.4	61.2	59.0	66.7	80.3	40.8	63.4	42.5	74.4	66.9	61.7	45.0
RC-FCN [38]	82.8	28.9	64.4	58.6	50.3	80.5	70.5	71.6	26.7	58.3	45.5	58.4	58.1	67.8	79.6	40.9	58.6	42.7	73.9	63.6	60.6	38.2
CAC-FCN	84.1	29.9	69.1	60.5	51.9	82.2	73.1	71.5	28.9	59.7	48.0	59.5	60.2	69.2	80.2	41.8	63.1	44.1	76.9	67.0	62.5	37.2

model to the semantic segmentation task. First, we remove the last two subsampling layers in ResNet18 to upscale the size of the output feature map by $4\times$. Second, we replace the last four convolution layers with dilated convolutions. Third, we replace the linear layer with an interpolated upsampling layer. In the training, we first train the model on MS-COCO dataset and then finetune it on PASCAL VOC 2012. We make some modifications to adapt it to the semantic segmentation task. We apply data augmentation by randomly scaling the input images (from 0.5 to 2.0) and randomly left-right flipping in training. The input images are resized to 480×480 in testing. We finetune 30 epochs using mini-batch SGD with a weight decay of 0.0001 and a momentum of 0.9. The learning rate is started at 0.01. In all experiments, we train the CAC based models with $\lambda = 0.3$ and $C = 0.5$.

6.3.3. Performance Comparison

In this section, we compare the proposed CAC method with the standard convolution and region convolution (RC) on semantic segmentation tasks. For convenience, we use FCN, RC-FCN, and CAC-FCN to represent the models with the standard convolution, RC, and CAC. We show the results in Table 3 and Fig. 6.

From Table 3, our CAC-FCN outperforms the baseline models with the standard convolution and RC on most of the categories. For the average performance in terms of mIoU, the CAC-FCN model yields significant performance improvement. We also provide a visual comparison of different models in Fig. 6. From this figure, our CAC-FCN produces more accurate segmentation masks than the FCN and RC-FCN baselines.

6.4. Experiments on Object Detection

In this section, we apply the proposed CAC method to object detection models. We evaluate the CAC-based models on the benchmark dataset MS COCO [55].

6.4.1. Compared Methods

In this experiment, we adopt the widely used model Faster-RCNN network as the baseline model. We compare our CAC based model with several state-of-the-art object detection models, including Fast-RCNN [56], ION [57], YOLOv2 [58], SSD300 [59] and SSD512 [59].

6.4.2. Datasets and Implementation Details

We conduct experiments on the MS COCO dataset which contains 117k training images and 5k validation images with 80 classes. We use the ImageNet pretrained ResNet18-based FPN network for comparing different methods on MS COCO dataset. We follow the setting in [60]. The networks are optimized for 13 epochs using SGD with a weight decay of 0.0001 and a momentum of 0.9. The learning rate is initialized with 0.02 and divided by 10 at 8 and 11 epochs.

We evaluate different object detection models using the COCO’s standard metric, namely mAP@0.5 and mAP@0.75. These two metrics represent the mean average precision (mAP) computed at the IoU thresholds of 0.5 and 0.75, respectively. We also

Table 4: Comparisons of different object detection models on MS COCO dataset. We use the ResNet18 model as the backbone. “-” denotes results that are not reported.

Model	mmAP	mAP@0.5	mAP@0.75	#MAdds (G)
Fast-RCNN [56]	18.9	38.6	-	-
ION [57]	23.6	43.2	23.6	-
YOLOv2 [58]	21.6	44.0	27.8	-
SSD300 [59]	23.2	41.2	23.4	-
SSD512 [59]	26.8	46.5	27.8	-
Faster-RCNN [60]	28.6	49.2	29.6	23.82
CAC-Faster-RCNN	29.9	49.9	31.1	20.49



Figure 7: Visual comparison of different object detection methods. The first column shows the input images with the corresponding ground truth bounding boxes. The second and third columns show the bounding boxes predicted by different models.

compute mmAP for different models by averaging multiple mAP values with the IoU thresholds ranging from 0.5 to 0.95 with the step of 0.05. In all experiments, we train the CAC based models with $\lambda = 0.3$ and $C = 0.5$.

6.4.3. Performance Comparison

In this experiment, we use CAC to replace the standard convolution layers in the Faster-RCNN model. We show the quantitative and visual results in Table 4 and Fig. 7.

From Table 4, CAC can obtain very compact models with better performance and lower computational cost than the baseline models equipped with the standard convolution. Specifically, the resultant CAC-Faster-RCNN has a lower computational cost and consistently outperforms the baseline model on all the considered metrics, including mmAP, mAP@0.5, and mAP@0.75. We also show the visual comparison of different object detection models in Fig. 7. From this figure, our CAC model generates more accurate bounding boxes than the Faster-RCNN baseline. The results demonstrate the effectiveness of the proposed CAC method on object detection tasks.

Table 5: Comparisons of robustness of deep models with different convolutions on CIFAR-10.

Model	Method	Error on Adversarial Examples (%)			
		FGSM [61]	MI-FGSM [62]	PGD-10 [63]	PGD-100 [63]
ResNet20	Standard Conv	29.82	54.89	57.99	65.41
	CAC	28.97	53.67	57.13	64.74
ResNet56	Standard Conv	26.91	51.91	55.33	64.81
	CAC	25.58	51.03	54.72	63.75

7. Further Experiments

In this section, we first compare the accuracy of our CAC method with the standard convolution method under adversarial perturbations to investigate the robustness of our method. Then, we conduct more experiments to investigate the effect of the hyperparameter λ .

7.1. Comparisons of Robustness

We investigate the robustness of the proposed CAC by comparing the accuracy on adversarial samples generated by four different attack methods, including FGSM [61], MI-FGSM [62], PGD10 [63] and PGD100 [63]. We train the standard convolution models (namely ResNet20 and ResNet56) and the CAC based models with adversarial samples on CIFAR-10. We report the adversarial accuracy that is evaluated on adversarial samples. The higher adversarial accuracy the model has, the more adversarially robust the model will be.

Following the setting in [63], we train all the models for 200 epochs and use an SGD optimizer with a momentum of 0.9 and a weight decay of 0.0005. The initialized learning rate is set to 0.1 and divided by 10 at epochs 90, 140 and 160, respectively. All attacks are with a total perturbation scale of $8/255$ (0.03) and a step size of $2/255$ (0.01). We set the number of attack iterations to 10, 10 and 100 for MIFGSM [62], PGD10 [63] and PGD100 [63], respectively. From Table 5, our CAC based models achieve higher adversarial accuracy than the standard convolution ones under four different attack perturbations. These results demonstrate that the proposed CAC convolution is more robust than the standard convolution. The main reason is that our CAC replaces the original large kernel with a 1×1 kernel to perform convolution on smooth windows.

Table 6: Comparisons of the CAC-ResNet20 models with different values of λ on CIFAR-10.

Method	Standard Conv	CAC			
λ	-	0.3	0.5	0.7	1.0
Error (%)	8.75	8.37	9.92	10.07	11.81
#MAdds \downarrow (%)	0	36.98	40.38	51.68	69.39

In this sense, we are able to effectively reduce the influence incurred by the noises and attacks in these windows.

7.2. Effect of λ on CAC

We investigate the effect of the hyperparameter λ in Eqn. (13) on the performance of CAC models. In this experiment, we use ResNet20 as the baseline model and train the models with different values of $\lambda \in \{0.3, 0.5, 0.7, 1.0\}$ on CIFAR-10. We show the results in Table 6.

From Table 6, the reduction of MAdds would increase when we gradually increase λ . However, due to the redundancy incurred by the smooth windows, it is possible to simultaneously reduce the computational cost and improve the performance, *e.g.*, when setting $\lambda = 0.3$. If we further increase the value of λ , the objective in Eqn. (13) encourages the model to focus more on the computational cost but compromise the performance. For example, when λ is set to 1.0, CAC reduces the computational cost by 69.39% but incur significant performance degradation. Thus, we set $\lambda = 0.3$ to train CAC models in the experiments.

8. Conclusion

In this paper, we have proposed a Content-aware Convolution (CAC) to reduce the computational redundancy incurred by the smooth windows when performing convolution. To reduce the computational redundancy and improve the performance, CAC replaces the original $k \times k$ kernel with a 1×1 kernel to perform convolutions on the smooth windows. Moreover, we propose an efficient algorithm to automatically recognize the sharp and smooth windows. Given different samples, the resultant CAC models could allocate different computation resources according to their data smoothness, which makes it possible for content-aware computation. Extensive results on

image classification, semantic segmentation, and object detection tasks show that our CAC based models yield significantly better performance and lower computational cost than the baseline models with the standard convolution.

9. Acknowledgements

This work was partially supported by the Key-Area Research and Development Program of Guangdong Province (Grant No. 2018B010107001), National Natural Science Foundation of China (NSFC) 61836003 (key project), Guangdong Project (Grant No. 2017ZT07X183), Fundamental Research Funds for the Central Universities (Grant No. D2191240), the National Natural Science Foundation of China (Grant No. 62072186), the Guangdong Basic and Applied Basic Research Foundation (Grant No. 2019B1515130001), the Guangdong Special Branch Plans Young Talent with Scientific and Technological Innovation (Grant No. 2016TQ03X445), the Guangzhou Science and Technology Planning Project (Grant No. 201904010197), and the Opening Project of Guangdong Key Laboratory of Big Data Analysis and Processing.

References

- [1] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [2] Y. Guo, J. Chen, Q. Du, A. Van Den Hengel, Q. Shi, M. Tan, Multi-way backpropagation for training compact deep neural networks, *Neural networks*.
- [3] F. Schroff, D. Kalenichenko, J. Philbin, Facenet: A Unified Embedding for Face Recognition and Clustering, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 815–823.
- [4] Y. Sun, X. Wang, X. Tang, Deeply Learned Face Representations are Sparse, Selective, and Robust, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 2892–2900.

- [5] S. Ozawa, S. L. Toh, S. Abe, S. Pang, N. Kasabov, Incremental learning of feature space and classifier for face recognition, *Neural Networks* 18 (5-6) (2005) 575–584.
- [6] E. Shelhamer, J. Long, T. Darrell, Fully convolutional networks for semantic segmentation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39 (4) (2017) 640–651.
- [7] L. Liu, J. Cao, M. Liu, Y. Guo, Q. Chen, M. Tan, Dynamic extension nets for few-shot semantic segmentation, in: *Proceedings of the 28th ACM international conference on Multimedia*, 2020.
- [8] N. Ibtehaz, M. S. Rahman, Multiresunet: Rethinking the u-net architecture for multimodal biomedical image segmentation, *Neural Networks* 121 (2020) 74–87.
- [9] S. Ren, K. He, R. B. Girshick, J. Sun, Faster R-CNN: towards real-time object detection with region proposal networks, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39 (6) (2017) 1137–1149.
- [10] H. Wang, L. Dai, Y. Cai, X. Sun, L. Chen, Salient object detection based on multi-scale contrast, *Neural Networks* 101 (2018) 47–56.
- [11] M. D. Skowronski, J. G. Harris, Automatic speech recognition using a predictive echo state network classifier, *Neural networks* 20 (3) (2007) 414–423.
- [12] B. Schrauwen, M. D’Haene, D. Verstraeten, J. Van Campenhout, Compact hardware liquid state machines on fpga for real-time speech recognition, *Neural networks* 21 (2-3) (2008) 511–523.
- [13] A. Gross, D. Murthy, Modeling virtual organizations with latent dirichlet allocation: A case for natural language processing, *Neural networks* 58 (2014) 38–49.
- [14] W. Duch, P. Matykiewicz, J. Pestian, Neurolinguistic approach to natural language processing with applications to medical text analysis, *Neural Networks* 21 (10) (2008) 1500–1510.
- [15] C. S. Burrus, T. Parks, and *Convolution Algorithms*, Citeseer, 1985.

- [16] A. Bar-Hillel, D. Weinshall, Efficient learning of relational object class models, *International Journal of Computer Vision* 77 (1-3) (2008) 175–198.
- [17] R. Fergus, P. Perona, A. Zisserman, Object class recognition by unsupervised scale-invariant learning, in: *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings., Vol. 2, IEEE, 2003*, pp. II–II.
- [18] A. Krizhevsky, I. Sutskever, G. Hinton, Imagenet classification with deep convolutional neural networks, in: *Advances in Neural Information Processing Systems, 2012*, pp. 1106–1114.
- [19] L. Sifre, S. Mallat, Rigid-motion scattering for image classification, Ph.D. thesis, Citeseer (2014).
- [20] F. Chollet, Xception: Deep learning with depthwise separable convolutions, in: *The IEEE Conference on Computer Vision and Pattern Recognition, 2017*, pp. 1800–1807.
- [21] Y. Chen, H. Fan, B. Xu, Z. Yan, Y. Kalantidis, M. Rohrbach, S. Yan, J. Feng, Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution, in: *The IEEE International Conference on Computer Vision, 2019*, pp. 3434–3443.
- [22] H. Li, A. Kadav, I. Durdanovic, H. Samet, H. P. Graf, Pruning filters for efficient convnets, in: *International Conference on Learning Representations, 2017*.
- [23] H. Hu, R. Peng, Y.-W. Tai, C.-K. Tang, Network trimming: A data-driven neuron pruning approach towards efficient deep architectures, *ArXiv abs/1607.03250*.
- [24] J. M. Alvarez, M. Salzmann, Learning the number of neurons in deep networks, in: *Advances in Neural Information Processing Systems, 2016*, pp. 2270–2278.
- [25] Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, C. Zhang, Learning efficient convolutional networks through network slimming, in: *The IEEE International Conference on Computer Vision, 2017*, pp. 2755–2763.

- [26] Y. He, X. Zhang, J. Sun, Channel pruning for accelerating very deep neural networks, in: The IEEE International Conference on Computer Vision, 2017, pp. 1398–1406.
- [27] J.-H. Luo, J. Wu, W. Lin, Thinet: A filter level pruning method for deep neural network compression, in: The IEEE International Conference on Computer Vision, 2017, pp. 5068–5076.
- [28] Z. Zhuang, M. Tan, B. Zhuang, J. Liu, Y. Guo, Q. Wu, J. Huang, J. Zhu, Discrimination-aware channel pruning for deep neural networks, in: Advances in Neural Information Processing Systems, 2018, pp. 883–894.
- [29] J. Liu, B. Zhuang, Z. Zhuang, Y. Guo, J. Huang, J. Zhu, M. Tan, Discrimination-aware network pruning for deep model compression, arXiv preprint arXiv:2001.01050.
- [30] S. Han, H. Mao, W. J. Dally, Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding, in: International Conference on Learning Representations, 2016.
- [31] S. Zhou, Y. Wu, Z. Ni, X. Zhou, H. Wen, Y. Zou, Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients, Arxiv abs/1606.06160.
- [32] F. Li, B. Zhang, B. Liu, Ternary weight networks, ArXiv abs/1605.04711.
- [33] C. Zhu, S. Han, H. Mao, W. J. Dally, Trained ternary quantization, in: International Conference on Learning Representations, 2017.
- [34] B. Liu, M. Wang, H. Foroosh, M. Tappen, M. Pensky, Sparse convolutional neural networks, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 806–814.
- [35] B. Graham, M. Engelcke, L. van der Maaten, 3d semantic segmentation with sub-manifold sparse convolutional networks, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 9224–9232.

- [36] Y. Guo, Y. Zheng, M. Tan, Q. Chen, J. Chen, P. Zhao, J. Huang, Nat: Neural architecture transformer for accurate and compact architectures, in: *Advances in Neural Information Processing Systems*, 2019, pp. 737–748.
- [37] Y. Guo, Y. Zheng, M. Tan, Q. Chen, Z. Li, J. Chen, P. Zhao, J. Huang, Towards accurate and compact architectures via neural architecture transformer, *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- [38] X. Li, Z. Liu, P. Luo, C. Change Loy, X. Tang, Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade, in: *The IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6459–6468.
- [39] A. Vasudevan, A. Anderson, D. Gregg, Parallel multi channel convolution using general matrix multiplication, in: *2017 IEEE 28th International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, IEEE, 2017, pp. 19–24.
- [40] J. Ludwig, *Image convolution*, Portland State University.
- [41] L. Wang, W. Ouyang, X. Wang, H. Lu, Visual tracking with fully convolutional networks, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 3119–3127.
- [42] N. Kanopoulos, N. Vasanthavada, R. L. Baker, Design of an image edge detection filter using the sobel operator, *IEEE Journal of solid-state circuits* 23 (2) (1988) 358–367.
- [43] G. Huang, Z. Liu, L. van der Maaten, K. Q. Weinberger, Densely connected convolutional networks, in: *The IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2261–2269.
- [44] N. Ma, X. Zhang, H.-T. Zheng, J. Sun, Shufflenet v2: Practical guidelines for efficient cnn architecture design, in: *The European Conference on Computer Vision*, 2018, pp. 122–138.

- [45] A. Krizhevsky, G. Hinton, Learning multiple layers of features from tiny images, Tech Report.
- [46] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2009.
- [47] Y. Nesterov, A method of solving a convex programming problem with convergence rate $o(1/k^2)$, in: Soviet Math. Dokl, Vol. 27, 1983, pp. 372–376.
- [48] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 3431–3440.
- [49] M. Everingham, S. M. A. Eslami, L. V. Gool, C. K. I. Williams, J. M. Winn, A. Zisserman, The pascal visual object classes challenge: A retrospective, International Journal of Computer Vision 111 (1) (2015) 98–136.
- [50] J. Carreira, F. Li, C. Sminchisescu, Object recognition by sequential figure-ground ranking, International Journal of Computer Vision 98 (3) (2012) 243–262.
- [51] J. Carreira, R. Caseiro, J. Batista, C. Sminchisescu, Semantic segmentation with second-order pooling, in: The European Conference on Computer Vision, 2012, pp. 430–443.
- [52] B. Hariharan, P. Arbeláez, R. Girshick, J. Malik, Simultaneous detection and segmentation, in: The European Conference on Computer Vision, 2014, pp. 297–312.
- [53] J. Dai, K. He, J. Sun, Convolutional feature masking for joint object and stuff segmentation, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 3992–4000.
- [54] L. Chen, G. Papandreou, F. Schroff, H. Adam, Rethinking atrous convolution for semantic image segmentation, Arxiv abs/1706.05587.

- [55] T. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft COCO: common objects in context, in: The European Conference on Computer Vision, 2014, pp. 740–755.
- [56] R. Girshick, Fast r-cnn, in: The IEEE International Conference on Computer Vision, 2015, pp. 1440–1448.
- [57] S. Bell, C. Lawrence Zitnick, K. Bala, R. Girshick, Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 2874–2883.
- [58] J. Redmon, A. Farhadi, Yolo9000: better, faster, stronger, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 7263–7271.
- [59] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, A. C. Berg, Ssd: Single shot multibox detector, in: The European Conference on Computer Vision, Springer, 2016, pp. 21–37.
- [60] S. Ren, K. He, R. B. Girshick, J. Sun, Faster R-CNN: towards real-time object detection with region proposal networks, in: Advances in Neural Information Processing Systems, 2015, pp. 91–99.
- [61] I. J. Goodfellow, J. Shlens, C. Szegedy, Explaining and harnessing adversarial examples, in: Y. Bengio, Y. LeCun (Eds.), International Conference on Learning Representations, 2015.
- [62] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, J. Li, Boosting adversarial attacks with momentum, in: The IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 9185–9193.
- [63] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, A. Vladu, Towards deep learning models resistant to adversarial attacks, in: International Conference on Learning Representations, 2018.