

# Linear Context Transform Block

Dongsheng Ruan,<sup>1,2,3</sup> Jun Wen,<sup>1,2</sup> Nenggan Zheng,<sup>1\*</sup> Min Zheng<sup>3</sup>

<sup>1</sup>Qiushi Academy for Advanced Studies, Zhejiang University, Hangzhou, Zhejiang, China

<sup>2</sup>College of Computer Science and Technology, Zhejiang University, Hangzhou, Zhejiang, China

<sup>3</sup>State Key Laboratory for Diagnosis and Treatment of Infectious Diseases, School of Medicine, Zhejiang University, Hangzhou, Zhejiang, China  
{21530003, junwen, zng, minzheng}@zju.edu.cn

## Abstract

Squeeze-and-Excitation (SE) block presents a channel attention mechanism for modeling global context via explicitly capturing dependencies across channels. However, we are still far from understanding how the SE block works. In this work, we first revisit the SE block, and then present a detailed empirical study of the relationship between global context and attention distribution, based on which we propose a simple yet effective module, called Linear Context Transform (LCT) block. We divide all channels into different groups and normalize the globally aggregated context features within each channel group, reducing the disturbance from irrelevant channels. Through linear transform of the normalized context features, we model global context for each channel independently. The LCT block is extremely lightweight and easy to be plugged into different backbone models while with negligible parameters and computational burden increase. Extensive experiments show that the LCT block outperforms the SE block in image classification task on the ImageNet and object detection/segmentation on the COCO dataset with different backbone models. Moreover, LCT yields consistent performance gains over existing state-of-the-art detection architectures, e.g., 1.5~1.7%  $AP^{bbox}$  and 1.0%~1.2%  $AP^{mask}$  improvements on the COCO benchmark, irrespective of different baseline models of varied capacities. We hope our simple yet effective approach will shed some light on future research of attention-based models.

## Introduction

Attention mechanism has achieved remarkable success in a variety of computer visual tasks, e.g., image classification (Wang et al. 2017; Hu, Shen, and Sun 2018), object detection (Wang et al. 2018; Zhang et al. 2018c), and semantic segmentation (Zhang et al. 2018a; Li et al. 2018). The attention module is typically plugged into existing deep networks to improve their representational power (He et al. 2016; Xie et al. 2017; Szegedy et al. 2015; Zagoruyko and Komodakis 2016; Zhang et al. 2018b; Howard et al. 2019). One of the most prominent works is the Squeeze-and-Excitation network (SENet) (Hu, Shen, and Sun 2018), which is channel-

\*Corresponding author

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Visualization of the averaged attention values (first column) and global context features before and after the first SE block (second column) in stage 4 (first row) and stage 5 (second row) on the ImageNet validation set.

attention based and aims to selectively emphasize informative channels and suppress trivial ones through explicitly modeling dependencies across channels. SENet achieves significant performance gains across varied models and has been successfully applied to a variety of computer vision tasks (Sandler et al. 2018; Ma et al. 2018; Yu et al. 2018; Howard et al. 2019).

To dive into this attention mechanism, we are curious of the following two questions: 1) what is the relationship between global context and attention distribution, and 2) which kind of channels are trivial. To answer these questions, we visualize the averaged global context features before and after SE and the corresponding attention activations on the ImageNet validation set. For easier observation, the averaged global context features before SE are sorted in ascending order, as shown in Fig. 1. Interestingly, we observe such a negative correlation in SE that global contexts with larger absolute values tend to be attached with smaller attentions, indicating that these channels are generally trivial. By learning such a correlation, SE effectively suppresses these channels and reduces contextual variations across channels, which enables subsequent filters to extract more primitive semantic features and improve generalization ability. Given this ob-

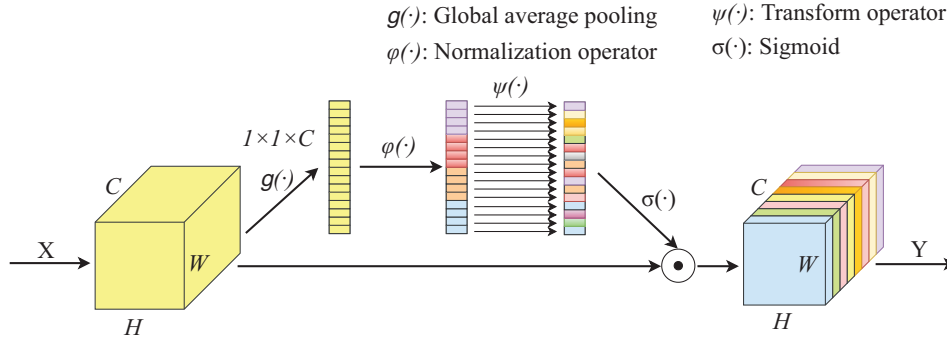


Figure 2: Architecture of the linear context transform block. The input feature maps are defined as  $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$ , where  $C$  is the number of channels and  $H, W$  are the spatial dimensions.  $\mathbf{Y} \in \mathbb{R}^{C \times H \times W}$  denotes the output of the LCT block.  $\odot$  denotes broadcast element-wise multiplication.

servation, a question naturally arises: can we learn such a correlation in a better way?

The SENet has shown the effectiveness of explicit dependency modelling across channels. However, a potential problem for SENet is that when the number of feature channel becomes higher, it will be much more difficult to capture the dependency across all channels to learn such a correlation stably because lots of irrelevant information from other channels can be introduced. An alternative approach is to boost the capacity of context feature transform module as in the GENet (Hu et al. 2018), but which brings significant increase of model complexity.

In this paper, we propose a simpler and robust approach to learning the above negative correlation with a novel module, called Linear Context Transform (LCT) block, which is extremely lightweight and brings negligible parameters and computational burden increase. Specifically, LCT achieves context feature transform with the following two cheap operators: *normalization* and *transform*. To enable stable context modelling, we divide channels into different groups and normalize the global context within each group using the *normalization* operation. With the *transform* operation, we then linearly transform the normalized global contexts for each channel independently. With varied architectures, we investigate the difference between the SE block and our LCT block in terms of attention distribution and global context feature, and find that the combined *normalization* and *transform* operators play a similar role as the fully connected (FC) layers of SE in learning the negative correlation while with smaller fluctuations (Fig. 3). In summary, our main contributions can be summarized as follows:

- We present an empirical study of the relationship between global context and attention distribution of the SENet, and find a negative correlation between them two, which help researchers better understand the mechanism of channel-wise attention and shed light on future research of attention-based models.
- We propose a novel light-weight attention block (LCT) for global context modeling by combining simple group normalization and linear transform. To our best knowl-

edge, this is the first work to model global context for each channel independently.

- Comprehensive experiments with three visual tasks (image classification on the ImageNet and object detection/segmentation on the COCO) consistently demonstrate the superiority and generalization ability of our attention model.

## Related Work

**Normalization** Batch normalization (BN) (Ioffe and Szegedy 2015) is a milestone technique that normalizes the statistics for each training mini-batch to stabilize the distributions of layer inputs, which enables deep networks to train faster and more stably. However, the property that depends on the mini-batch size leads to a rapid decline in network performance when the batch size becomes smaller. A series of normalization methods (Ba, Kiros, and Hinton 2016; Ulyanov, Vedaldi, and Lempitsky 2016; Wu and He 2018; Salimans and Kingma 2016) have been proposed to address this issue caused by inaccurate batch statistics estimation. Layer normalization (LN) (Ba, Kiros, and Hinton 2016) computes the statistics along the channel dimension and is well suited for recurrent neural network. Instance normalization (Ulyanov, Vedaldi, and Lempitsky 2016) proposes to perform the normalization across spatial locations. Group normalization (GN) divides features into different groups and normalizes them within each group (Wu and He 2018; Wen et al. 2019). Since GN does not exploit the batch dimension, it is still able to achieve high accuracy even in small batch size.

The design of LCT is inspired by GN. Instead of stabilizing the distribution of layer inputs, LCT is essentially a channel-wise attention mechanism that aims to model global context dependency with group normalization.

**Attention modules** Recently, several attention modules (Chen et al. 2017; Wang et al. 2018; Chen et al. 2018; Fu et al. 2019; Huang et al. 2018) have been proposed to exploit global contextual information to enhance the representational power of the networks. In particular, SENet

(Hu, Shen, and Sun 2018) develops a lightweight attention block to recalibrate feature channels by exciting the aggregated contexts from original features. Further, GENet (Hu et al. 2018) proposes a gather-excite framework for better context exploitation and yields further performance gains at the expense of increasing parameters. GCNet (Cao et al. 2019) combines simplified non-local block (Wang et al. 2018) and SE block (Hu, Shen, and Sun 2018) to effectively model the global context via addition fusion. In addition to channel attention, CBAM (Woo et al. 2018) and BAM (Park et al. 2018) exploit both spatial and channel-wise information to yield further performance gains. SKNet (Li et al. 2019) proposes a dynamic selection mechanism that enables the network to adaptively adjust receptive field. More recently, Li *et al.* (Li, Hu, and Yang 2019) introduce a spatial group-wise enhance module to spatially enhance the semantic expression in each group, showing excellent performance in image classification and object detection.

Our work builds on the idea developed in the SE block. However, different from SE, LCT implicitly captures channel-wise dependencies and linearly models the global context of each channel, which is more lightweight and robust.

## Method

In this section, we first review the SE block, and then present the proposed linear context transform (LCT) block.

### Revisiting the SE block

The SE aims to emphasize informative features and suppress trivial ones by modeling the channel-wise relationship. To obtain contextual information, SE proposes to squeeze global spatial information. Specifically, it aggregates global context information across spatial dimension through global average pooling operation. Further, to fully capture channel-wise dependencies, the SE block excites the aggregated contexts using two fully-connected layers. Here we define  $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$  as the input feature maps of SE, where  $C$  is the number of channels and  $H, W$  are the spatial dimensions. The SE block can be formulated as:

$$\mathbf{Y} = \mathbf{X} \cdot \sigma(f(g(\mathbf{X}))) = \mathbf{X} \cdot \sigma(\mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 g(\mathbf{X}))), \quad (1)$$

where  $\cdot$  denotes channel-wise multiplication and  $g(\cdot)$  global average pooling to generate channel-wise statistics.  $\mathbf{W}_1$  and  $\mathbf{W}_2$  denote the weights of FC layers and  $\sigma(\cdot)$  the *sigmoid* function.

As shown in Fig. 1, the SE performs a non-linear transform to learn a negative correlation between global contexts and attention values by explicitly capturing the dependencies across channels. However, this negative correlation is learned from all channels, which may bring in each channel irrelevant information from other channels and make the global context modeling unstable, resulting in incorrect mapping. To tackle this problem, we propose the novel LCT block.

### Linear context transform block

In this section, we introduce the proposed LCT block in detail, which is illustrated in Fig. 2. As summarized in GCNet (Cao et al. 2019), global context modeling framework can be abstracted as the following three modules: (a) context aggregation; (b) context feature transform; (c) feature fusion, which framework is also followed by the LCT.

**Context aggregation** Context aggregation aims to help the network capture long-range dependencies by exploiting information beyond the local receptive fields of each filter. A number of aggregation strategies can be chosen to aggregate contextual information, such as second-order attention pooling (Chen et al. 2018), global attention pooling (Hu et al. 2018; Cao et al. 2019), and global average pooling (Hu, Shen, and Sun 2018). Complex aggregation operators can be used to improve performance of the LCT block, but which are not the focus of our work. Hence we simply employ global average pooling to aggregate the global context features of each sample across spatial dimensions generating a channel descriptor as  $\mathbf{z} = \{z_k = \frac{1}{H \times W} \sum_{i=1}^W \sum_{j=1}^H \mathbf{X}_k(i, j) : k \in \{1, \dots, C\}\}$ .

**Context feature transform** To effectively and efficiently model the context feature, the LCT introduces a pair of lightweight operators: a normalization operator, which normalizes the global context features within each group, and a transform operator, which takes in the normalized global contexts to produce the importance scores. Specifically, we first divide the descriptor  $\mathbf{z}$  into groups and then normalize it within each group along channel dimension. More formally, we define  $\mathbf{v}^i = \{z_{mi+1}, \dots, z_{m(i+1)}\}$  as the  $i$ -th local context group, where  $i \in \{0, \dots, G-1\}$  and  $G$  are the index and the number of groups, respectively.  $m = C/G$  is the number of channels per group. The normalization operator  $\varphi$  can be formulated as:

$$\hat{\mathbf{v}}^i = \varphi(\mathbf{v}^i) = \frac{1}{\sigma^i}(\mathbf{v}^i - \mu^i), \quad (2)$$

where  $\mu^i$  and  $\sigma^i$  are the mean and standard deviation of the  $i$ -th group, respectively, and can be computed as:

$$\mu^i = \frac{1}{m} \sum_{n \in \mathcal{S}_i} z_n, \sigma^i = \sqrt{\frac{1}{m} \sum_{n \in \mathcal{S}_i} (z_n - \mu^i)^2 + \epsilon}. \quad (3)$$

Here  $\epsilon$  is a small constant.  $\mathcal{S}_i$  is the set of the  $i$ -th group of channel index.

The normalization operator plays two crucial roles in context feature transform. First, it enables each channel to adjust its own context feature by perceiving context information within each group, implicitly capturing dependencies across channels. Second, it can effectively eliminate the inconsistency of the context feature distribution caused by different samples, which stabilizes the distribution of global context features.

Next, we define the transform operator to be a function  $\psi: \mathbb{R}^C \rightarrow \mathbb{R}^C$  that maps the gathered context features  $\hat{\mathbf{z}}$  to the importance scores  $\mathbf{a}$ , formulated as:

$$\mathbf{a} = \psi(\hat{\mathbf{z}}) = \mathbf{w} \cdot \hat{\mathbf{z}} + \mathbf{b}, \quad (4)$$

where  $\hat{\mathbf{z}} = [\hat{\mathbf{v}}^0, \hat{\mathbf{v}}^1, \dots, \hat{\mathbf{v}}^{G-1}]$ .  $\mathbf{w}$  and  $\mathbf{b}$  are trainable gain and bias parameters of the same dimension as  $\hat{\mathbf{z}}$ . Note that the transform operator  $\psi$  is a channel-wise linear transform, which means that information from other channels is not taken into account in the context transform process. In addition, it only introduces the parameters of  $\mathbf{w}$  and  $\mathbf{b}$ , which are almost negligible compared to the entire network.

Interestingly, the composition of two operators can be regarded as a special case of GN where the spatial height  $H$  and width  $W$  are 1. In the case of  $G = 1$ , it is equivalent to LN. But it is worth noting that the transform operator in the LCT block is designed to transform the global context features, not to compensate for the potential loss of representational ability caused by normalization, which is essentially different from other normalization methods.

**Feature fusion** Finally, the feature fusion module modulates the input features by conditioning on the transformed contexts. Specifically, the output  $\mathbf{Y} \in \mathbb{R}^{C \times H \times W}$  of the LCT block is obtained by rescaling the original response  $\mathbf{X}$  according to the attention activations  $\sigma(\mathbf{a})$  and can be expressed as:

$$\mathbf{Y} = \mathbf{X} \cdot \sigma(\mathbf{a}). \quad (5)$$

**Relationship to SE block** LCT shares the same context aggregation module and feature fusion module with SE. The main difference between them is the context transform module, which reflects different perspectives of two blocks for global context modeling. First, SE makes use of global information from other channels to help model the global contexts, which actually increases the complexity of context transform. In comparison, our LCT block is more lightweight and simplifies global context modeling by independently transforming the global contexts of each channel. The number of parameters in the SE block is  $2C^2/r$ , while the number of parameters in the LCT block is  $C$ , where  $r$  is the reduction ratio. It is apparent that LCT has significantly decreased parameters. Second, SE explicitly captures channel-wise dependencies using two FC layers, while our approach implicitly captures dependencies within each group through group normalization operator. The results in Table. 2 show that the normalization operator can effectively capture channel dependencies within each group.

## Experiments

In this section, we first evaluate the proposed LCT block on the task of image classification on ImageNet-1K (Rusakovsky et al. 2015), and then conduct extensive ablation studies. Finally, we experiment on the COCO 2017 dataset (Lin et al. 2014) to demonstrate the general applicability of the LCT block.

### Image classification on ImageNet

The ImageNet 2012 dataset contains 1.28 million training images and 50K validation images with 1000 classes.

**Implementation details** We train all models from scratch on 4 GPUs for 100 epochs, using synchronous SGD optimizer with a weight decay of 0.0001 and momentum 0.9. The initial learning rate is set to 0.1, and decreases by a

Backbone	Params	FLOPs	Top-1 (%)	Top-5 (%)
ResNet50	25.56M	4.122G	76.15	92.87
+SE	28.09M	4.130G	77.31	93.68
+LCT	25.59M	4.127G	<b>77.45</b>	<b>93.71</b>
ResNet101	44.55M	7.849G	77.37	93.56
+SE	49.33M	7.863G	78.49	94.19
+LCT	44.61M	7.858G	<b>78.55</b>	<b>94.26</b>

Table 1: Classification accuracies on the ImageNet validation set. Params denotes the number of parameters. FLOPs denotes the number of multiply-adds.

G	1	4	8	16	32	64	128
Top-1	77.37	77.36	77.44	77.34	77.32	<b>77.45</b>	-
Top-5	93.66	93.57	93.56	93.54	93.52	<b>93.71</b>	-

Table 2: Classification accuracies (%) of LCT-ResNet50 with different group numbers  $G$  on the ImageNet validation set. - denotes that the network can not converge.

factor of 0.1 every 30 epochs. The weight initialization is adopted in (He et al. 2015). For ResNet50 backbone, the total batch size is set as 256. For ResNet101 backbone, we reduce the batch size to 220 due to the limited GPU memory. The standard data augmentation is performed for training: a  $224 \times 224$  crop is randomly sampled from a  $256 \times 256$  image or its horizontal flip using the scale and aspect ratio augmentation. Input images are normalized using the channel means and standard deviations.

As is widely practiced in (Hu, Shen, and Sun 2018; Woo et al. 2018), our LCT blocks are inserted into each residual block of ResNet. We use 0 and 1 to initialize all  $\mathbf{w}$  and  $\mathbf{b}$  parameters respectively.  $G$  is set as 64 by default. To make a fair comparison, the baseline models are reproduced in the same training settings. We report the top-1 and top-5 classification accuracy on the single  $224 \times 224$  center crop in the validation set.

**Classification results** Table 1 presents the main results of our experiments. We observe that LCT performs better than SE with fewer parameters and less computational burden, irrespective of the depth of the backbone. Compared to ResNet, our LCT block adds few parameters and computations, but achieves significant performance gains ( $> 1.0\%$   $\uparrow$  on Top-1 accuracy) even in deeper ResNet101. Remarkably, LCT-ResNet50 is able to outperform ResNet101, which indicates that the improvements brought by LCT exceed the benefits of increased network depth (51 layers). These results demonstrate effective of LCT on image classification.

**Analysis and discussion** To gain some insights into the channel attention mechanism, we investigate the relationship between global context features and attention distribution. Specifically, we first compute the averaged global context features before and after attention blocks and the corresponding attention activations across 1000 classes on ImageNet validation set. Then we sort the averaged global context features in ascending order for better observation. Fig. 3 shows the results of the first attention blocks at different stages. In order to observe the difference more intuitively,



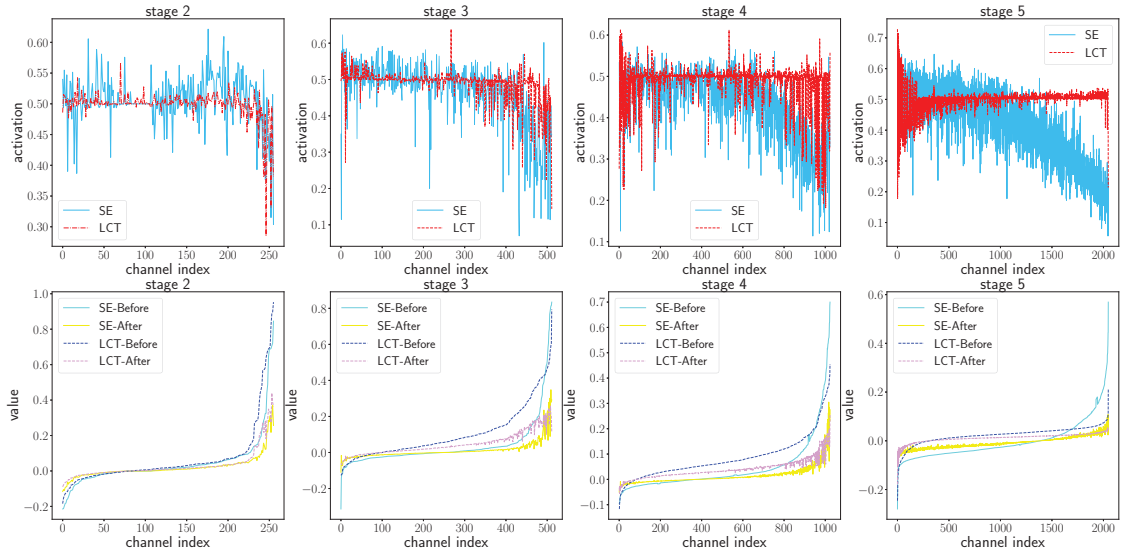


Figure 3: Visualizations of the averaged attention values and averaged global context features before and after the first attention blocks at different stages on the ImageNet validation set. The backbone network is ResNet50. Top row: averaged attention valued. Bottom row: averaged global context features.




Figure 4: Visualizations of the absolute context variations before and after attention blocks at different stages.

we also visualize the  $\Delta$ value that represents the absolute context variation before and after attention block, shown in Fig. 4.

We observe that both SE and LCT learn a negative correlation that global context features with larger absolute values tend to be assigned smaller activations, which suggests that channels with these context features are generally less useful. This is reasonable to some extent, since a large amount of noise is more likely to exist in these channels. When the magnitude of the features of some channels is dramatically larger than that of other channels, subsequent filters will pay more attention on these trivial channels, leading to unstable semantic representation learning. By performing fea-

Normalization	w/	w/o
Top-1 (%)	77.45	76.89
Top-5 (%)	93.71	93.33
Transform	w/	w/o
Top-1 (%)	77.45	76.82
Top-5 (%)	93.71	93.32

Table 3: Classification accuracies of LCT-ResNet50 with and without normalization/transform operator on the ImageNet validation set.

	LCT	SE	SE+
Top-1 (%)	<b>77.45</b>	77.31	77.37
Top-5 (%)	93.71	93.68	<b>93.73</b>

Table 4: Effects of inserting a normalization operator before the two FC layers of the SE block. The backbone is ResNet50.

ture recalibration, both blocks effectively suppress the influence of these channels and reduce the contextual differences across channels, which enables subsequent filters to capture robust semantics of each channel. In a sense, global contexts act like an indicator of which channels need to be suppressed.

While SE and LCT learn similar attention distributions, there are still several differences. First, the attention distribution learned by LCT is more stable because no other channel information is introduced in the transform operator. Second, LCT does not over-suppress the original feature responses, thus retaining important semantic information. These findings provide explanations for the effectiveness of the LCT block.

		Initialization	
w	b	Top-1 (%)	Top-5 (%)
0	0	77.36	93.60
0	1	<b>77.45</b>	<b>93.71</b>
1	0	77.24	93.54

Table 5: Results of different initializations with LCT-ResNet50 on the ImageNet validation set.

## Ablation study

**Number of groups** In this experiment, we assess the effect of group number on the performance of the LCT block. As shown in Table. 2, LCT is not sensitive to the variation of group number, which is reasonable because the mean and variance do not change significantly with the number of channels per group. We observe that when  $G = 128$ , the network has failed to converge since too many groups may lead to incorrect statistical estimation. When  $G = 64$ , the performance is slightly higher than that of other settings, indicating that introducing too much information from other irrelevant channels may not be helpful. By default, we set  $G = 64$  for LCT. Moreover, LCT consistently outperforms SE for all  $G$  values, which indicates that the normalization operator can well capture the dependency across channels, even in the extreme case of  $G = 1$ .

**Normalization operator** To investigate the influence of normalization in the LCT block, we conduct experiments by removing the normalization operator from LCT. Table. 3 shows the results. It is clear that the LCT block without normalization operator suffers considerable performance degradation. This comparison shows that global context can not be effectively transformed using transform operator alone. It also demonstrates that normalization operator can effectively eliminate the inconsistency of context feature distribution and captures dependencies between channels well.

We have seen that normalization operator can improve the performance of the LCT block and would like to explore whether normalization operator can also help SE block yield further performance gains. For this purpose, we insert a normalization operator before the FC layers of the SE block. We refer to this block as SE+.  $G$  is set to 64. The results are shown in Table. 4. We find that normalization operator does not bring significant gain to the SE block. The top-1 accuracy of the SE+ block is slightly inferior to ours. Based on these results, we can draw the following conclusions: 1) The two FC layers in SE not only can transform the global context features, but also effectively prevent the inconsistency of feature distribution caused by different samples, which is surprisingly similar to two operators in LCT. The difference is that LCT decomposes the roles of two FC layers into two independent operators, each of which performs its own function. 2) After normalization, a per-channel linear transform is sufficient to transform the global contexts. Introducing information from other channels complicates context feature transform. These findings provide an explanation for the effectiveness of the LCT block.

**Transform operator** We study the effect of transform operator. To this end, we retain the normalization operator and remove the transform operator from LCT. The results are shown in Table. 3. We observe that performance is noticeably reduced and is slightly worse than that without normalization operator, suggesting that transform operator is vitally important for global context transform. The reason is that normalization operator can not learn the negative correlation between global context features and attention distribution. We also find that the LCT block with two operators achieves the best performance, which indicates that two operators are complementary and indispensable for global context modeling.

**Initialization** Table 5 shows the ablation results of initialization. Different from the initialization in GN, IN and LN, we find that it is more appropriate to initialize  $w$  and  $b$  to 0 and 1 respectively, which is consistent with the finding in SGE (Li, Hu, and Yang 2019). Initializing  $w$  and  $b$  to 0 gets suboptimal results. As shown in Fig. 3, we observe that most of the attention values fluctuate around 0.5 for both SE and LCT. Hence a possible explanation is that initializing  $w$  to 0 makes  $\sigma(0 \sim 1)$  around 0.5, which is conducive to the learning of attention distribution. When  $w = 1$  and  $b = 0$ , LCT achieves the worst results, because the transform operator is designed to transform the context features rather than compensate for the lost of representational ability caused by normalization.

## Object detection and segmentation on COCO

In this section, we evaluate our block with object detection and instance segmentation tasks on the COCO-2017 dataset (Lin et al. 2014). We train using 118k train images and evaluate on 5k val images. The COCO-style average precisions at different boxes and the mask IoUs are reported.

**Implementation details** All experiments are implemented with *mmdetection* framework (Chen et al. 2019). The input images are resized such that the long edge and short edge are 1333 and 800 pixels respectively. We train on 4 GPUs with 1 images per each for 12 epochs. All models are trained using synchronized SGD with a weight decay of  $1e-4$  and momentum of 0.9. According to the linear scaling rule (Goyal et al. 2017), the initial learning rate is set to 0.005, which is decreased by 10 at the 9th and 12th epochs. The backbones of all models are pretrained on ImageNet. We finetune all layers except for c1 and c2 with FPN (Lin et al. 2017), detection and segmentation heads. During finetuning the BatchNorm layers are frozen. Other hyper-parameters follow the default settings of the *mmdetection* framework. The backbone is ResNet101 in all experiments.

**Object detection** We evaluate the LCT block on the object detection task. To this end, we insert LCT into four state-of-the-art detection frameworks, including Faster RCNN (Ren et al. 2015), Mask RCNN (He et al. 2017), Cascade RCNN (Cai and Vasconcelos 2018) and Cascade Mask RCNN (Chen et al. 2019). The results on val set are given in Table 6. We observe that our approach is better than SE with fewer parameters and less computations, irrespective of dif-

Detector	Backbone	$\Delta$ Params	$\Delta$ FLOPs	$AP_{0.5:0.95}^{bbox}$	$AP_{0.5}^{bbox}$	$AP_{0.75}^{bbox}$	$AP_{small}^{bbox}$	$AP_{media}^{bbox}$	$AP_{large}^{bbox}$
Faster R-CNN	baseline	-	-	38.5	60.5	41.8	22.3	43.2	49.8
	+SE	+4.78M	+0.191G	39.8(+1.3)	61.9	43.1	23.9	43.8	<b>51.5</b>
	+LCT	+0.06M	+0.187G	<b>40.0(+1.5)</b>	<b>62.8</b>	<b>43.4</b>	<b>24.8</b>	<b>44.4</b>	50.9
Mask R-CNN	baseline	-	-	39.4	61.0	43.3	23.1	43.7	51.3
	+SE	+4.78M	+0.191G	40.7(+1.3)	62.7	44.3	24.5	44.8	52.7
	+LCT	+0.06M	+0.187G	<b>40.9(+1.5)</b>	<b>63.1</b>	<b>44.6</b>	<b>25.0</b>	<b>45.1</b>	<b>52.9</b>
Cascade R-CNN	baseline	-	-	42.0	60.3	45.9	23.2	46.0	56.3
	+SE	+4.78M	+0.191G	43.4(+1.4)	62.2	47.4	24.7	47.4	57.0
	+LCT	+0.06M	+0.187G	<b>43.6(+1.6)</b>	<b>62.4</b>	<b>47.6</b>	<b>25.4</b>	<b>47.6</b>	57.3
Cascade Mask R-CNN	baseline	-	-	42.6	60.7	46.7	23.8	46.4	56.9
	+SE	+4.78M	+0.191G	43.7(+1.1)	61.8	47.5	24.3	47.5	58.6
	+LCT	+0.06M	+0.187G	<b>44.1(+1.5)</b>	<b>62.4</b>	<b>48.3</b>	<b>25.0</b>	<b>47.7</b>	<b>59.3</b>

Table 6: Comparisons based on ResNet101 backbone on the task of **object detection**.  $\Delta$ Params denotes the change amount of parameters.  $\Delta$ FLOPs denotes the change amount of computations. The numbers in brackets denote the improvements over the baseline backbone.

Detector	Backbone	$AP_{0.5:0.95}^{mask}$	$AP_{0.5}^{mask}$	$AP_{0.75}^{mask}$	$AP_{small}^{mask}$	$AP_{media}^{mask}$	$AP_{large}^{mask}$
Mask R-CNN	baseline	35.9	57.7	38.4	19.2	39.7	49.7
	+SE	36.9(+1.0)	59.4	39.2	20.0	<b>40.8</b>	50.3
	+LCT	<b>37.0(+1.1)</b>	<b>59.6</b>	<b>39.3</b>	<b>20.5</b>	<b>40.8</b>	<b>50.5</b>
Cascade Mask R-CNN	baseline	37.0	58.0	39.9	19.1	40.5	51.4
	+SE	37.7(+0.7)	59.0	40.5	19.4	41.1	52.4
	+LCT	<b>38.1(+1.1)</b>	<b>59.5</b>	<b>41.3</b>	<b>19.9</b>	<b>41.3</b>	<b>53.2</b>

Table 7: Comparisons based on ResNet101 backbone on the task of **instance segmentation**. The results show that LCT outperforms SE.

ferent detectors, which indicates that modeling global context for each channel independently is also effective on the task of object detection. In addition, compared to the baselines, LCT consistently yields 1.5 ~ 1.6%  $AP_{0.5:0.95}^{bbox}$  points with neglectable extra parameters and computations, suggesting that our approach is widely applicable across various detector architectures. We also find that LCT greatly improves the detection performance of Faster RCNN, Mask RCNN and Cascade RCNN for small objects with the gain exceeding 1.9%  $AP_{small}^{mask}$ . For Cascade Mask RCNN, the detection performance of large objects is significantly boosted (2.4%  $\uparrow AP_{large}^{mask}$ ).

**Instance segmentation** Finally, we explore the applicability to the instance segmentation task. We select two popular frameworks, Mask RCNN and Cascade Mask RCNN. As can be seen in Table 7, LCT also outperforms SE, which is consistent with the results in image classification and object detection. When adopting stronger detector Cascade Mask RCNN, the improvements achieved by LCT are still significant, suggesting that our approach is complementary to the capacity of current model. Compared to the baselines, the LCT block can boost performance by 1.1 %  $AP_{0.5:0.95}^{mask}$  regardless of the strength of the detectors. These results suggest the generalization and effectiveness of our approach.

## Conclusion

In this paper, we presented an empirical study of the relationship between global context and attention distribution of

SENet. Then we considered the question of how to effectively learn the correlation between them. To this end, we introduced a simple yet effective channel attention architecture, the LCT block, to explore this question and provided experimental evidence that demonstrates the effectiveness and generalization of our approach across multiple visual tasks. In further work, we plan to develop more efficient algorithms to exploit feature context, which may provide new insights into channel attention mechanism.

## Acknowledgement

This work is supported by the Zhejiang Provincial Natural Science Foundation (LR19F020005), National Natural Science Foundation of China (61572433, 61972347), 13-5 State S&T Projects of China (2018ZX1030206) and thanks for a gift grant from Baidu inc.

## References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Cai, Z., and Vasconcelos, N. 2018. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6154–6162.
- Cao, Y.; Xu, J.; Lin, S.; Wei, F.; and Hu, H. 2019. Gcnet: Non-local networks meet squeeze-excitation networks and beyond. *arXiv preprint arXiv:1904.11492*.
- Chen, L.; Zhang, H.; Xiao, J.; Nie, L.; Shao, J.; Liu, W.; and Chua, T.-S. 2017. Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning. In *Proceedings of*

- the *IEEE conference on computer vision and pattern recognition*, 5659–5667.
- Chen, Y.; Kalantidis, Y.; Li, J.; Yan, S.; and Feng, J. 2018. A<sup>2</sup>-nets: Double attention networks. *Advances in Neural Information Processing Systems* 352–361.
- Chen, K.; Wang, J.; Pang, J.; Cao, Y.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Xu, J.; Zhang, Z.; Cheng, D.; Zhu, C.; Cheng, T.; Zhao, Q.; Li, B.; Lu, X.; Zhu, R.; Wu, Y.; Dai, J.; Wang, J.; Shi, J.; Ouyang, W.; Loy, C. C.; and Lin, D. 2019. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*.
- Fu, J.; Liu, J.; Tian, H.; Li, Y.; Bao, Y.; Fang, Z.; and Lu, H. 2019. Dual attention network for scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3146–3154.
- Goyal, P.; Dollár, P.; Girshick, R.; Noordhuis, P.; Wesolowski, L.; Kyrola, A.; Tulloch, A.; Jia, Y.; and He, K. 2017. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *ICCV* 1026–1034.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. *computer vision and pattern recognition* 770–778.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- Howard, A.; Sandler, M.; Chu, G.; Chen, L.-C.; Chen, B.; Tan, M.; Wang, W.; Zhu, Y.; Pang, R.; Vasudevan, V.; et al. 2019. Searching for mobilenetv3. *arXiv preprint arXiv:1905.02244*.
- Hu, J.; Shen, L.; Albanie, S.; Sun, G.; and Vedaldi, A. 2018. Gather-excite: Exploiting feature context in convolutional neural networks. *Advances in Neural Information Processing Systems* 9401–9411.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. *computer vision and pattern recognition* 7132–7141.
- Huang, Z.; Wang, X.; Huang, L.; Huang, C.; Wei, Y.; and Liu, W. 2018. Ccnet: Criss-cross attention for semantic segmentation. *arXiv preprint arXiv:1811.11721*.
- Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *international conference on machine learning* 448–456.
- Li, H.; Xiong, P.; An, J.; and Wang, L. 2018. Pyramid attention network for semantic segmentation. *arXiv preprint arXiv:1805.10180*.
- Li, X.; Wang, W.; Hu, X.; and Yang, J. 2019. Selective kernel networks. *Proceedings of the IEEE conference on computer vision and pattern recognition* 510–519.
- Li, X.; Hu, X.; and Yang, J. 2019. Spatial group-wise enhance: Enhancing semantic feature learning in convolutional networks. In *arXiv preprint arXiv:1905.09646*.
- Lin, T.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. *europaean conference on computer vision* 740–755.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Ma, N.; Zhang, X.; Zheng, H.; and Sun, J. 2018. Shufflenet v2: Practical guidelines for efficient cnn architecture design. *europaean conference on computer vision* 122–138.
- Park, J.; Woo, S.; Lee, J.; and Kweon, I. S. 2018. Bam: Bottleneck attention module. *british machine vision conference* 147.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3):211–252.
- Salimans, T., and Kingma, D. P. 2016. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in Neural Information Processing Systems* 901–909.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. *computer vision and pattern recognition* 4510–4520.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S. E.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. *computer vision and pattern recognition* 1–9.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2016. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*.
- Wang, F.; Jiang, M.; Qian, C.; Yang, S.; Li, C.; Zhang, H.; Wang, X.; and Tang, X. 2017. Residual attention network for image classification. *computer vision and pattern recognition* 6450–6458.
- Wang, X.; Girshick, R. B.; Gupta, A.; and He, K. 2018. Non-local neural networks. *computer vision and pattern recognition* 7794–7803.
- Wen, J.; Liu, R.; Zheng, N.; Zheng, Q.; Gong, Z.; and Yuan, J. 2019. Exploiting local feature patterns for unsupervised domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 5401–5408.
- Woo, S.; Park, J.; Lee, J.; and Kweon, I. S. 2018. Cbam: Convolutional block attention module. *europaean conference on computer vision* 3–19.
- Wu, Y., and He, K. 2018. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 3–19.
- Xie, S.; Girshick, R. B.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. *computer vision and pattern recognition* 5987–5995.
- Yu, C.; Wang, J.; Peng, C.; Gao, C.; Yu, G.; and Sang, N. 2018. Learning a discriminative feature network for semantic segmentation. *computer vision and pattern recognition* 1857–1866.
- Zagoruyko, S., and Komodakis, N. 2016. Wide residual networks. *british machine vision conference*.
- Zhang, H.; Dana, K. J.; Shi, J.; Zhang, Z.; Wang, X.; Tyagi, A.; and Agrawal, A. K. 2018a. Context encoding for semantic segmentation. *computer vision and pattern recognition* 7151–7160.
- Zhang, X.; Zhou, X.; Lin, M.; and Sun, J. 2018b. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6848–6856.
- Zhang, X.; Wang, T.; Qi, J.; Lu, H.; and Wang, G. 2018c. Progressive attention guided recurrent network for salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 714–722.