

A Trainable Multiplication Layer for Auto-correlation and Co-occurrence Extraction^{*}

Hideaki Hayashi¹[0000-0002-4800-1761] and Seiichi Uchida¹[0000-0001-8592-7566]

Kyushu University, 744 Motooka Nishi-ku Fukuoka, JAPAN
{hayashi, uchida}@ait.kyushu-u.ac.jp

Abstract. In this paper, we propose a trainable multiplication layer (TML) for a neural network that can be used to calculate the multiplication between the input features. Taking an image as an input, the TML raises each pixel value to the power of a weight and then multiplies them, thereby extracting the higher-order local auto-correlation from the input image. The TML can also be used to extract co-occurrence from the feature map of a convolutional network. The training of the TML is formulated based on backpropagation with constraints to the weights, enabling us to learn discriminative multiplication patterns in an end-to-end manner. In the experiments, the characteristics of the TML are investigated by visualizing learned kernels and the corresponding output features. The applicability of the TML for classification and neural network interpretation is also evaluated using public datasets.

Keywords: Neural network · Feature extraction · Auto-correlation · Co-occurrence.

1 Introduction

Typified by the success of convolutional neural networks (CNNs) in recent computer vision research, studies using deep learning-based methods have demonstrated success in several fields [13], [4], [30], [1], [14]. Unlike the traditional machine learning techniques based on handcrafted features, these deep learning-based methods automatically extract features from raw input data via end-to-end network learning.

To exploit of the end-to-end learning capability of the deep neural networks, numerous studies have investigated the development of a network layer that represents a certain function by incorporating a model into a layer structure [27], [20], [5]. For example, Wang *et al.* [27] proposed a trainable structural layer called a global Gaussian distribution embedding network, which involves global Gaussian [17] as an image representation.

Despite such challenging efforts, the majority of layer structures are based on inner products of the input features and weight coefficients. There has been

^{*} Supported by Qdai-jump Research Program and JSPS KAKENHI Grant Number JP17K12752 and JP17K19402.

minimal research attempting to introduce multiplication of the input features. In classical pattern recognition techniques, however, multiplication of the input features is important because it represents auto-correlation or co-occurrence of the input features.

This paper proposes a trainable multiplication layer (TML) for a neural network that can be used to calculate the multiplication between the input features. Taking an image as an input, the TML raises each pixel value to the power of a weight and then multiplies them, thereby extracting the higher-order local auto-correlation from the input image. The TML can also be used to extract co-occurrence from the feature map of a CNN. The training of the TML is formulated based on backpropagation with constraints to the weights, enabling us to learn discriminative multiplication patterns in an end-to-end manner.

The contributions of this work are as follows:

- A trainable multiplication layer is proposed.
- The learning algorithm of the TML based on constrained optimization is formulated.
- Applicability to classification and network interpretation is demonstrated via experiments.

2 Related Work

2.1 Multiplication in a Neural Network Layer

The majority of the existing neural network layers focus on the multiplication of the input features and weight coefficients; there are limited studies that discuss a network layer that calculates the multiplication of the input features. Classically, Sigma-Pi Unit based on the summation of multiplication is applied to a self-organizing map [28], [25]. Incorporation of probabilistic models such as a Gaussian mixture model into the neural network structure consequently leads to the development of a network layer based on multiplication [5], [23], [22]. In recent work, Shih *et al.* [20] proposed a network layer that detects co-occurrence by calculating multiplications between the feature maps of a CNN. Kobayashi [7] achieved a trainable co-occurrence activation by decomposing and approximating a multiplication operation.

Although such multiplication-based layers are not actively studied, they have the potential to create an efficient network structure because they can express auto-correlation and co-occurrence directly. This paper therefore focuses on the development of the multiplication-based neural network layer.

2.2 Higher-Order Local Auto-correlation

Higher-order local auto-correlation (HLAC) is a feature extraction method proposed by Otsu [19]. HLAC is frequently used in the field of image analysis [24],

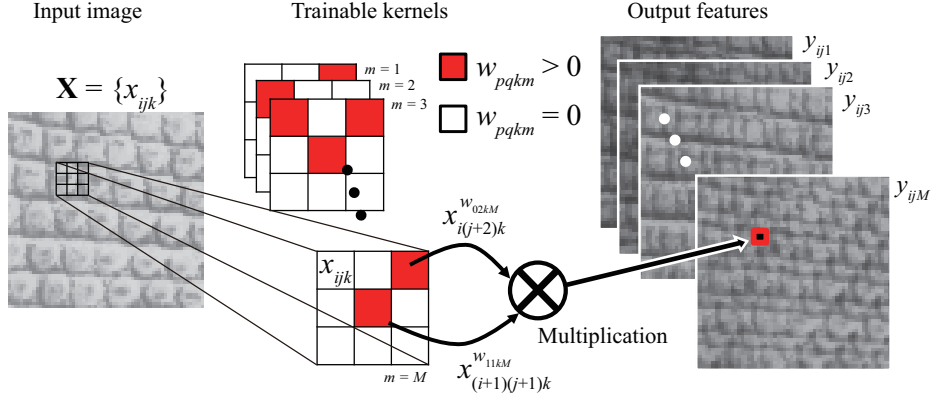


Fig. 1. Overview of forward calculation conducted in the TML. This figure shows a simple case where the number of input channels K is 1 and the kernel size is 3×3 .

[3], [18], [6]. For L displacements $\mathbf{a}_1, \dots, \mathbf{a}_L$, the L th-order autocorrelation function is defined as

$$R_L(\mathbf{a}_1, \dots, \mathbf{a}_L) = \sum_{D'} f(\mathbf{r})f(\mathbf{r} + \mathbf{a}_1) \cdots f(\mathbf{r} + \mathbf{a}_L), \quad (1)$$

$$D' = \{\mathbf{r} | \mathbf{r} + \mathbf{a}_l \in D, \forall l \in \{1, 2, \dots, L\}\}, \quad (2)$$

where $f(\mathbf{r})$ is the pixel value of the input image at a coordinate \mathbf{r} , D is a set of coordinates of the input image.

In HLAC, the patterns of the displacement must be prepared manually. The number of displacement patterns increases explosively based on the mask size and order; hence, they are limited practically.

3 Trainable Multiplication Layer

3.1 Layer Structure

Fig. 1 presents an overview of the forward calculation in the TML. The main idea behind this layer is to achieve the multiplication of the input values chosen by the kernels, which is different from the well-known convolutional layer that conducts multiplication of the input values and kernels.

Given an input image or feature map of a convolutional layer $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2 \times K}$, where N_1 , N_2 , and K are the number of rows, columns, and channels, respectively, the forward pass of the TML is defined as follows:

$$y_{ijm} = \prod_{k=0}^{K-1} \prod_{p=0}^{H-1} \prod_{q=0}^{W-1} x_{(i+p)(j+q)k}^{w_{pqkm}}, \quad (3)$$

where x_{ijk} ($i = 0, \dots, N_1 - 1$, $j = 0, \dots, N_2 - 1$, $k = 0, \dots, K - 1$) is the (i, j, k) -th element of \mathbf{X} and $w_{pqkm} \geq 0$ is the (p, q) -th weight of the m -th kernel for channel

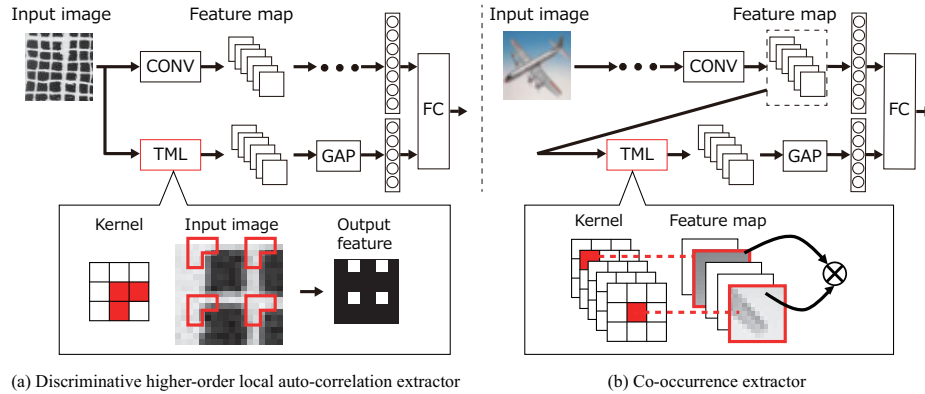


Fig. 2. Two usages of the TML in a CNN: (a) used as discriminative higher-order local auto-correlation extractor and (b) used as co-occurrence extractor. The abbreviations GAP, CONV, and FC denote global average pooling, convolutional layer, and fully connected layer, respectively.

k ($p = 0, \dots, H - 1$, $q = 0, \dots, W - 1$, $m = 0, \dots, M - 1$; M is the number of kernels). Since any number to the zero power is one (we define the value of zero to the power of zero to also be one), the forward pass of the proposed layer is regarded as the multiplication of the input values, where the value of the kernel at the corresponding coordinate is greater than zero.

In practice, (3) is calculated in the logarithmic form to prevent under/overflow. Assuming that $x_{ijk} > 0$, since \mathbf{X} is an image or a feature map passed through a ReLU function, (3) is rewritten as follows:

$$z_{ijmpqk} = \log (x_{(i+p)(j+q)k} + \epsilon), \quad (4)$$

$$y_{ijm} = \exp \left(\sum_{k=0}^{K-1} \sum_{p=0}^{H-1} \sum_{q=0}^{W-1} w_{pqkm} z_{ijmpqk} \right), \quad (5)$$

where ϵ is a small positive quantity to avoid $\log 0$.

3.2 Usage of the Layer in a Convolutional Neural Network

The TML is used as a part of a CNN-based deep neural network. Fig. 2 shows two different usages of the TML in a CNN. Specifically, Fig. 2(a) shows the case that the TML is the use as a discriminative higher-order local auto-correlation (DHLAC) extractor; Fig. 2(b) shows the use as a co-occurrence extractor.

The former usage is proposed to achieve the discriminative learning of the displacement patterns in HLAC. In this usage, the TML is inserted immediately behind the input image. The output features calculated by the TML are then passed through a global average pooling (GAP) [12]. GAP calculates the average of each feature map and the resulting vector is input to the next layer. This GAP

is placed to simulate the integral computation of HLAC. Because the displacement patterns are represented by the power functions to be differentiable, they are trainable as kernel patterns.

The latter usage is designed to calculate the co-occurrence between feature maps. In the deeper layer of a CNN, feature maps calculated by a convolutional layer involve abstracted information of the input. In this usage, the TML determines the co-occurrence from this abstracted information by calculating the multiplication between the feature maps. This calculation allows the extraction of non-Markovian co-occurrence that is useful for classification. Furthermore, we can use the TML to extract not only co-occurrence between features of a single network, but also co-occurrence between two or more networks.

3.3 Learning Algorithm

Given a set of images $\mathbf{X}^{(n)}$ ($n = 1, \dots, N$) for training with the teacher vector $\mathbf{t}^{(n)}$, the training process of the network into which the TML is incorporated involves minimizing the energy function E defined as

$$\text{minimize } E = \frac{1}{N} \sum_{n=1}^N \text{Loss}(\mathbf{O}^{(n)}, \mathbf{t}^{(n)}) + \lambda \|\mathbf{w}\|_1 \quad (6)$$

$$\text{subject to } \sum_{k=0}^{K-1} \sum_{p=0}^{H-1} \sum_{q=0}^{W-1} w_{pqkm} = C_1, \quad (7)$$

$$0 \leq w_{pqkm} \leq C_2, \quad (8)$$

where $\text{Loss}(\mathbf{O}^{(n)}, \mathbf{t}^{(n)})$ is the loss function defined by the final network output vector $\mathbf{O}^{(n)}$ corresponding to $\mathbf{X}^{(n)}$ and the teacher vector $\mathbf{t}^{(n)}$. The operator $\|\cdot\|_1$ indicates the L_1 norm, \mathbf{w} is all of the kernels presented in a vectorized form, and λ is a constant. The constants C_1 and C_2 are hyperparameters that influence the training results, and the details are described in the next section. The L_1 regularization is recruited expecting to obtain sparse kernels based on the concept of the TML. The first constraint shown in (7) sets the total value of each kernel and prevents overflow. The second constraint shown in (8) prevents each kernel from being overly sparse because each kernel must have two or more nonzero elements to be a meaningful kernel. It could appear that the L_1 regularization in (6) is redundant as the L_1 norm is fixed to a constant value by both constraints. However, the L_1 regularization influences the gradient vectors to obtain a sparse solution during the network learning.

Algorithm 1 shows the procedure to solve the above optimization, where \mathbf{w}_m is the vectorized m -th kernel. In this algorithm, gradient-based weight updating to decrease the energy function E and weight modification to maintain the constraints are calculated alternately. Although this is an approximated approach different from the well-known Lagrange multiplier, we employ this approach to strictly satisfy the constraints during the network learning.

Algorithm 1 Algorithm of weight updating

Require: Parameters C_1 , C_2 , and λ , training image set $\mathbf{X}^{(n)}$, teacher vector $\mathbf{t}^{(n)}$.**Ensure:** Trained network.

- 1: Initialize the kernels \mathbf{w} and other network weights $\boldsymbol{\theta}$.
 - 2: **while** \mathbf{w} and $\boldsymbol{\theta}$ have not converged **do**
 - 3: Calculate E .
 - 4: Calculate gradients of E with respect to \mathbf{w} and $\boldsymbol{\theta}$.
 - 5: Update \mathbf{w} and $\boldsymbol{\theta}$ using gradient-based updating.
 - 6: $\mathbf{w} \leftarrow \text{clip}(\mathbf{w}, 0, C_2)$
 - 7: $\mathbf{w}_m \leftarrow C_1 \mathbf{w}_m / \text{sum}(\mathbf{w}_m)$ for $m = 0, \dots, M - 1$
 - 8: **end while**
-

To calculate backpropagation, the partial derivative of E with respect to each kernel w_{pqkm} is required. Since the TML consists of multiplication and power functions, the partial derivative is calculated simply as follows:

$$\begin{aligned}
 \frac{\partial E}{\partial w_{pqkm}} &= \frac{\partial E}{\partial y_{ijm}} \frac{\partial y_{ijm}}{\partial w_{pqkm}} \\
 &= \frac{\partial E}{\partial y_{ijm}} \prod_{k=0}^{K-1} \prod_{p=0}^{H-1} \prod_{q=0}^{W-1} x_{(i+p)(j+q)k}^{w_{pqkm}} \log x_{(i+p)(j+q)k} \\
 &= \frac{\partial E}{\partial y_{ijm}} y_{ijm} \log x_{(i+p)(j+q)k},
 \end{aligned} \tag{9}$$

where the form of $\frac{\partial E}{\partial y_{ijm}}$ depends on the layer connected between the TML and the output.

As outlined above, the TML incorporated in the deep network calculates DHLAC of the input image or co-occurrence between the feature maps. The kernels are trained in an end-to-end manner via backpropagation with constraints.

4 Investigation of the Layer Characteristics

Before beginning the classification experiments, we investigated the characteristics of the TML by training a relatively shallow network with the TML inserted, and visualizing the learned kernels and corresponding feature maps.

4.1 Relationships between Hyperparameters and Learned Kernels

The TML has important hyperparameters including C_1 , C_2 for the learning constraints, and the kernel size. We investigated the changes in the learned kernels and the corresponding features according to the parameter variation.

In this trial, LeNet-5 [11] was used as the basic structure. The TML was inserted immediately behind the input. The ReLU activation function [16] was used for the convolution layers and the sigmoid function was used for the fully connected layer.




Fig. 3. Changes in learned kernels according to C_2 . The kernel size and other parameters are fixed as 3×3 , $C_1 = 1.0$, and $\lambda = 0.01$.

For the training data, we used the MNIST dataset [11]. This dataset includes ten classes of handwritten binary digit images with a size of 28×28 ; it contains 60,000 training images and 10,000 testing images. The network was trained with all the training data for each hyperparameter.

Following the network training, we visualized the learned kernels and the responses for the testing images. First, we observed the changes in the learned kernels according to the parameters for the constraints. Because the ratio of C_1 to C_2 influences the training results, we varied the value of C_2 as $C_2 = 1.0, 0.5, 0.33, 0.25$ for fixed a C_1 of 1.0. The kernel size and λ were set as 3×3 and $\lambda = 0.01$, respectively.

Fig. 3 shows the changes in the learned kernels according to C_2 . In this figure, the values of the learned kernels normalized by C_2 are displayed in a grayscale heat map, and therefore black and white pixels show “0” and C_2 , respectively. The number of nonzero elements in each kernel increases according to the decrease in the value of C_2 . Since the total value of elements in each kernel is fixed to C_1 and the upper limit of each element is suppressed by C_2 , the number of nonzero elements approximates to C_1/C_2 if the L_1 regularization functions appropriately. These results demonstrate that the number of nonzero elements, which is equivalent to the order in HLAC, can be controlled by changing the ratio of C_1 to C_2 .

Secondly, the kernel size was varied as 3×3 , 5×5 , 7×7 , and 9×9 for the fixed value of $C_1 = 1.0$, $C_2 = 0.5$, and $\lambda = 0.01$. Fig. 4 shows the learned kernels and the corresponding output features according to the kernel size. The kernel pattern and output feature pattern placed at the same location correspond to each other. The learned kernel values displayed in the left panels are also normalized in the same manner as Fig. 3. The values of the output features shown in the right panels are rescaled from $[0, 1]$ to $[0, 255]$ before visualization.

In terms of kernel size variation, the characteristics of the learned kernels and output features changed depending on the kernel size. In Fig. 4(a), the number of nonzero elements is two in each kernel and these elements adjoin each other in the majority of the kernels. Kernels with neighboring nonzero elements extract rather local correlation, and the output features are virtually the same as the input images. However, kernels with nonzero elements apart from each other extract different output features according to the kernel pattern. In Fig. 4(b), the number of kernels with nonzero elements apart from each other increased




Fig. 4. Learned kernels (left panels) and corresponding output features (right panels) for each kernel size. Other parameters are fixed as $C_1 = 1.0$, $C_2 = 0.5$, and $\lambda = 0.01$.

and hence a richer variation of output features was obtained compared to Fig. 4(a). Fig. 4(c) and (d) include kernels that have gray pixels. This means that three or more values are multiplied in a calculation and therefore high-order auto-correlation is extracted. These results indicate that the number of nonzero elements in each kernel can be controlled to a certain extent by the ratio of C_1 to C_2 , and the variety of the output features changes according to the kernel size.

4.2 Characteristics as a DHLAC Extractor

We verified if the TML can make discriminative responses by observing the learned kernel patterns and the corresponding output features for the synthetic texture patterns.

The texture dataset used in this experiment contained six classes of 32×32 artificially generated images from the following procedure (examples are displayed in the leftmost panels in Fig. 5): First, six stripe patterns, different for each class, were drawn on a 1024×1024 black image. Then, uniform random noise in the range of $[0, 1]$ was added to the images. Finally, a set of randomly cropped 32×32 -sized images were used as the dataset. We generated 100 training images for each class (600 training images in total). After training, we observed the learned kernels and the corresponding output features to the testing samples that were generated independently from the training images.

The network structure used in this experiment was LeNet-5, as in Section 4.1. The parameters were set as $C_1 = 1.0$, $C_2 = 0.5$, $\lambda = 0.01$, and the kernel size was 9×9 .

Fig. 5 shows the output features for each combination of the input image and learned kernel. The row and column correspond to the class number and kernel number, respectively. As with Fig. 4, the features are rescaled $[0, 1]$ to $[0, 255]$.

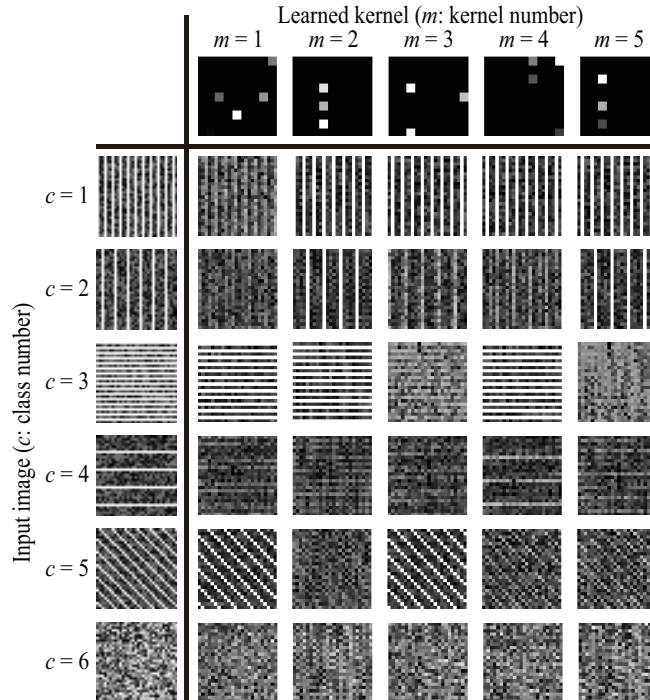


Fig. 5. Output features for each combination of the input image and learned kernel.

Fig. 5 indicates that different output features are obtained according to the combination of the class and the kernel. For example, for the input image of $c = 5$, the slanting white stripe remains only in the output features from $m = 1$ and $m = 3$. This is because the kernels $m = 1$ and $m = 3$ have nonzero elements apart diagonally from each other. For the other classes, the pattern remains in the output features if the direction or interval of the pattern corresponds to those of the kernel patterns. This means that the TML learned kernels that can extract discriminative features from the input image.

4.3 Characteristics as a Co-occurrence Extractor

To investigate the capability of the TML for extracting co-occurrence features, we visualized the co-occurring regions detected by the TML.

For the use as a co-occurrence extractor, the TML was connected to the fully connected layer as shown in Fig. 2(b). In this experiment, the TML was inserted between the second convolutional layer and the fully connected layer of the LeNet-5. Because this fully connected layer maps the input vector into the likelihood for each class, the weights of this layer represent the relevance between each dimension of the input vector and each class. Based on this fact, we extracted the most relevant kernel to the target category. We then defined



Fig. 6. Visualization of the co-occurrence features. The top panels show examples of the original MNIST images. The bottom panels are the corresponding co-occurrence features highlighted over the original images.

the co-occurrence features as the input features to the TML that were activated by the kernel relevant to the target category.

Fig. 6 shows the visualization results of the co-occurrence features. In this figure, it appears that the distinctive features for each category are highlighted. For example, the bilateral lines of “0” are highlighted. In “7” the upper left edge and vertical line are highlighted. These results demonstrate the capability of the TML for extracting co-occurrence features that are useful for classification from CNN feature maps.

5 Classification Experiments

To evaluate the applicability of the TML, we conducted a classification experiment using public datasets.

5.1 Dataset

We used the following datasets in this experiment.

MNIST: As described in the previous section, this dataset includes ten classes of handwritten binary digit images with a size of 28×28 . We used 60,000 images as training data and 10,000 images as testing data.

Fashion-MNIST: Fashion-MNIST [29] includes ten classes of binary fashion images with a size of 28×28 . It includes 60,000 images for training data and 10,000 images for testing data.

CIFAR10: CIFAR10 [9] is labeled subsets of the 80 million tiny images dataset. This dataset consists of 60,000 32×32 color images in 10 classes, with 6,000 images per class. There are 50,000 training images and 10,000 test images.

Kylberg: The Kylberg texture dataset [10] contains unique texture patches for each class. We used its small subset provided by the author which includes six classes with 40 samples each. We divided the original 576×576 patches into nine 64×64 nonoverlapping patches and considered each patch as one sample; thus 2,160 samples were available. We conducted 10-fold cross-validation and calculated the classification accuracy.

Brodatz: The Brodatz texture dataset [2] contains 112 texture classes with a size of 640×640 . We divided each image into 100 64×64 nonoverlapping

Table 1. Comparison of recognition rates (%).

	MNIST	Fashion	CIFAR10	Kylberg	Brodatz
Baseline CNN	99.27	92.44	81.43	99.12	91.69
Baseline CNN + HLAC	99.31	92.09	82.23	99.31	91.33
Baseline CNN + TML (DHLAC)	99.39	92.45	81.50	99.02	92.51
Baseline CNN + TML (Co-occurrence)	99.27	92.54	81.49	99.35	91.56

patches, and considered each patch as one sample; thus 11,200 samples were available. We conducted 10-fold cross-validation to calculate the accuracy.

5.2 Experimental Setup

As the baseline, we used a simple CNN (called “baseline CNN” hereinafter) to clarify the effect of the TML. The baseline CNN consisted of five convolutional layers with four maxpooling layers between them; it also included two fully connected layers (the structure is illustrated in the supplemental file). In the network, dropout was conducted after the first and second maxpooling and the first fully connected layer.

The effectiveness of the TML was then examined by connecting the TML to the baseline CNN, both as a DHLAC extractor and a co-occurrence extractor. The kernel size of the TML was set as 5×5 . We also compared the results to the results using HLAC instead of the TML.

5.3 Results

Table 1 shows the recognition rates for each dataset. In all datasets, applying the TML to the CNN improved the classification performance. In particular, the use as DHLAC demonstrated remarkable performance for the Brodatz dataset. This is because texture is based on a repeated structure, and hence auto-correlation information is effective for the classification. Moreover, the discriminative training of the HLAC features conducted by the TML functioned efficiently. These results confirm the applicability of the TML for classification.

6 Applications

6.1 Interpretation of the Network

In this section, we demonstrate that the TML can be used to interpret on which part of an input image the CNN focuses. In the usage as a co-occurrence extractor, the TML is connected between the feature maps and the fully connected layer (Fig. 2(b)). By tracing this structure in the reverse direction in the trained network, it is possible to extract features having a strong influence for classification. It should be emphasized that the purpose of this experiment is not to improve classification accuracy, rather, it is to improve interpretability.




Fig. 7. Visualization of network interpretation results using the TML. The areas having a strong influence for classification extracted by the TML are highlighted.

Specifically, we interpret a trained network in the following procedure: 1. Perform a forward calculation for a certain input image; 2. Calculate the largest weight among the weights of the fully connected layer between the unit that outputs the posterior probability of the target class and the features calculated by the TML; 3. Extract the kernel of the TML connecting the largest weight calculated in the previous step; 4. Extract CNN feature maps connecting to nonzero elements of the kernel; and 5. Visualize the extracted feature maps by upsampling them to the size of the input image and overlapping them on the input image. The property of the TML that the learned kernel acquires sparsity allows this calculation.

We used the Caltech-UCSD Birds-200-2011 dataset (CUB-200-2011) [26]. Each species of birds has unique characteristics such as feather pattern and body shape. This fact makes the interpretation of visualization results easier. This dataset contains 200 species of bird images and consists of 5,994 images for training data and 5,794 images for test data. Although this dataset is frequently used for fine-grained classification and segmentation, we used this dataset for only visualization in this experiment.

The basic network structure used in this experiment was the LeNet-5. The TML was inserted between the second convolutional layer and the fully connected layer. The parameters were set as $C_1 = 1.0$, $C_2 = 0.5$, $\lambda = 0.01$, and 1×1 for the kernel size.

Fig. 7 shows examples of the visualization results. In Fig. 7, images of the same species are arranged in each column. The feather patterns or body shapes unique for each species are highlighted. For example, in the far left panels, yellowish green patterns on the head and the wing are highlighted. These results indicate the applicability of the proposed TML for interpreting a neural network.

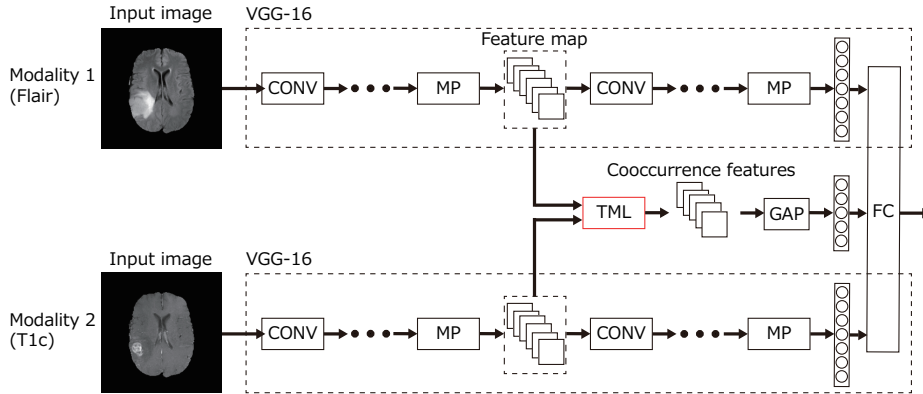


Fig. 8. Network architecture for MRI classification. The abbreviations MP, GAP, CONV, and FC denote max pooling, global average pooling, convolutional layer, and fully connected layer, respectively. Each VGG-16 model takes different modality images as input. The co-occurrence features between two networks are extracted by the TML.

6.2 Co-occurrence Extraction between Two Networks for Multimodal Data

In this experiment, we applied the TML to multimodal data classification. As mentioned in Section 3.2, the TML can also be used to extract co-occurrence between two networks. By extracting the co-occurrence between the features of two CNNs that take input data from different modalities, it is expected that the two networks complement each other.

The problem we addressed in this experiment is tumor detection in magnetic resonance (MR) images. In MR imaging, there are different modalities depending on a setting of pulse sequences. Because each modality has a different response, extracting co-occurrence between multiple modalities could possibly improve tumor detection accuracy.

We prepared an MR image dataset containing two modalities (Flair and T1c) from the multimodal brain tumor image segmentation benchmark (BraTS) [15]. This dataset consisted of three-dimensional brain MR images with tumor annotation. We created two-dimensional axial (transverse) slice images and separated them into tumor and non-tumor classes based on the annotation. The dataset contained 220 subjects; we randomly divided these into 154, 22, and 44 for training, validation, and testing, respectively. Because approximately 60 images were extracted from each subject, we obtained 8,980, 1,448, and 2,458 images for training, validation, and testing, respectively. We resized the images from 240×240 pixels to 224×224 to fit the network input size.

Fig. 8 illustrates the network architecture used in this experiment. The network was constructed based on VGG-16 [21]. VGG-16 has three fully connected layers following five blocks consisting of convolutional layers and a max pooling layer. We applied the TML to the CNN features after the b -th max pooling layer

Table 2. Comparison of classification accuracies for the MRI dataset.

Method	Modality	Accuracy (%)
Single VGG	Flair	94.22
Single VGG	T1c	89.86
Concatenated two VGGs	Flair and T1c	94.47
Ours ($b = 1$)	Flair and T1c	95.16
Ours ($b = 2$)	Flair and T1c	95.57
Ours ($b = 3$)	Flair and T1c	96.14
Ours ($b = 4$)	Flair and T1c	95.85
Ours ($b = 5$)	Flair and T1c	95.77

from the top. For comparison, we calculated the classification accuracy for a single VGG-16 with each modality and two VGG-16’s concatenated at the first fully connected layer.

Table 2 shows the results of the MR image classification. This confirms that the TML is effective for improving classification performance. In particular, the improvement is greatest for $b = 3$. One possible explanation for this is that the information necessary for classification is extracted to the extent that position information is not lost in the middle of the network, and extracting co-occurrence from such information is effective for multimodal data classification. This result demonstrates the effectiveness of co-occurrence extraction using the TML.

7 Conclusion

In this paper, we proposed a trainable multiplication layer (TML) for a neural network that can be used to calculate the multiplication between the input features. Taking an image as an input, the TML raises each pixel value to the power of a weight and then multiplies them, thereby extracting the higher-order local auto-correlation from the input image. The TML can also be used to extract co-occurrence from the feature map of a convolutional network. The training of the TML is formulated based on backpropagation with constraints to the weights, enabling us to learn discriminative multiplication patterns in an end-to-end manner. In the experiments, the characteristics of the TML were investigated by visualizing learned kernels and the corresponding output features. The applicability of the TML for classification was also evaluated using public datasets. Applications such as network interpretation and co-occurrence extraction between two neural networks were also demonstrated.

In future work, we plan to investigate the practical applications of the proposed TML. We will also expand the layer to the 3D structure in the same manner as cubic HLAC [8]. Application to time-series data analysis is also expected by constructing a one-dimensional structure.

References

1. Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Cheng, Q., Chen, G., et al.: Deep speech 2: End-to-end speech recognition in English and Mandarin. In: International Conference on Machine Learning. pp. 173–182. IMLS, New York (2016)
2. Brodatz, P.: Textures: A photographic album for artists and designers. Dover Publications (1966)
3. Fujino, K., Mitani, Y., Fujita, Y., Hamamoto, Y., Sakaida, I.: Liver cirrhosis classification on M-mode ultrasound images by higher-order local auto-correlation features. *Journal of Medical and Bioengineering* **3**(1), 29–32 (2014)
4. Greenspan, H., van Ginneken, B., Summers, R.M.: Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. *IEEE Transactions on Medical Imaging* **35**(5), 1153–1159 (2016)
5. Hayashi, H., Shibasaki, T., Shima, K., Kurita, Y., Tsuji, T.: A recurrent probabilistic neural network with dimensionality reduction based on time-series discriminant component analysis. *IEEE Transactions on Neural Networks and Learning Systems* **26**(12), 3021–3033 (2015)
6. Hu, E., Nosato, H., Sakanashi, H., Murakawa, M.: A modified anomaly detection method for capsule endoscopy images using non-linear color conversion and higher-order local auto-correlation (HLAC). In: Proceedings of the 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 5477–5480. IEEE, Osaka, Japan (2013)
7. Kobayashi, T.: Trainable co-occurrence activation unit for improving ConvNet. In: Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1273–1277. IEEE, Calgary, Canada (2018)
8. Kobayashi, T., Otsu, N.: Action and simultaneous multiple-person identification using cubic higher-order local auto-correlation. In: Proceedings of the 17th International Conference on Pattern Recognition (ICPR). vol. 4, pp. 741–744. IEEE, Cambridge, UK (2004)
9. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto (2009)
10. Kylberg, G.: The kylberg texture dataset v. 1.0. External report (Blue series) 35, Centre for Image Analysis, Swedish University of Agricultural Sciences and Uppsala University, Uppsala, Sweden (2011)
11. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
12. Lin, M., Chen, Q., Yan, S.: Network in network. In: Proceedings of the International Conference on Learning Representations (ICLR). Banff, Canada (2014)
13. Ma, L., Lu, Z., Li, H.: Learning to answer questions from image using convolutional neural network. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. pp. 3567–3573. Phoenix, Arizona (2016)
14. Majumder, N., Poria, S., Gelbukh, A., Cambria, E.: Deep learning-based document modeling for personality detection from text. *IEEE Intelligent Systems* **32**(2), 74–79 (2017)
15. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Transactions on Medical Imaging* **34**(10), 1993–2024 (2015)

16. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: Proceedings of the 27th International Conference on Machine Learning (ICML). pp. 807–814. Haifa, Israel (2010)
17. Nakayama, H., Harada, T., Kuniyoshi, Y.: Global Gaussian approach for scene categorization using information geometry. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2336–2343. IEEE, San Francisco, CA (2010)
18. Nosato, H., Sakanashi, H., Takahashi, E., Murakawa, M.: An objective evaluation method of ulcerative colitis with optical colonoscopy images based on higher order local auto-correlation features. In: Proceedings of the 11th International Symposium on Biomedical Imaging (ISBI). pp. 89–92. IEEE, Beijing, China (2014)
19. Otsu, N., Kurita, T.: A new scheme for practical flexible and intelligent vision systems. In: Proceedings of the IAPR Workshop on Computer Vision. pp. 431–435. IAPR, Tokyo, Japan (1988)
20. Shih, Y.F., Yeh, Y.M., Lin, Y.Y., Weng, M.F., Lu, Y.C., Chuang, Y.Y.: Deep co-occurrence feature learning for visual object recognition. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Honolulu, HI (2017)
21. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Proceedings of the International Conference on Learning Representations (ICLR). San Diego, CA (2015)
22. Tsuji, T., Bu, N., Fukuda, O., Kaneko, M.: A recurrent log-linearized Gaussian mixture network. *IEEE Transactions on Neural Networks* **14**(2), 304–316 (2003)
23. Tsuji, T., Fukuda, O., Ichinobe, H., Kaneko, M.: A log-linearized Gaussian mixture network and its application to EEG pattern classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* **29**(1), 60–72 (1999)
24. Uehara, K., Sakanashi, H., Nosato, H., Murakawa, M., Miyamoto, H., Nakamura, R.: Object detection of satellite images using multi-channel higher-order local autocorrelation pp. 1339–1344 (2017)
25. Valle-Lisboa, J.C., Reali, F., Anastasia, H., Mizraji, E.: Elman topology with sigma-pi units: An application to the modeling of verbal hallucinations in schizophrenia. *Neural networks* **18**(7), 863–877 (2005)
26. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Tech. rep. (2011)
27. Wang, Q., Li, P., Zhang, L.: G2DeNet: Global Gaussian distribution embedding network and its application to visual recognition. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2730–2739. IEEE, Honolulu, HI (2017)
28. Weber, C., Wermter, S.: A self-organizing map of sigma-pi units. *Neurocomputing* **70**(13), 2552–2560 (2007)
29. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747 (2017)
30. Zhang, T., Kahn, G., Levine, S., Abbeel, P.: Learning deep control policies for autonomous aerial vehicles with MPC-guided policy search. In: Proceedings of the International Conference on Robotics and Automation (ICRA). pp. 528–535. IEEE, Stockholm, Sweden (2016)