

# PCANet: A Simple Deep Learning Baseline for Image Classification?

Tsung-Han Chan, Kui Jia, Shenghua Gao, Jiwen Lu, Zinan Zeng, and Yi Ma

**Abstract**—In this work, we propose a very simple deep learning network for image classification which comprises only the very basic data processing components: cascaded principal component analysis (PCA), binary hashing, and block-wise histograms. In the proposed architecture, PCA is employed to learn multistage filter banks. It is followed by simple binary hashing and block histograms for indexing and pooling. This architecture is thus named as a PCA network (PCANet) and can be designed and learned extremely easily and efficiently. For comparison and better understanding, we also introduce and study two simple variations to the PCANet, namely the RandNet and LDANet. They share the same topology of PCANet but their cascaded filters are either selected randomly or learned from LDA. We have tested these basic networks extensively on many benchmark visual datasets for different tasks, such as LFW for face verification, MultiPIE, Extended Yale B, AR, FERET datasets for face recognition, as well as MNIST for hand-written digits recognition. Surprisingly, for all tasks, such a seemingly naive PCANet model is on par with the state of the art features, either prefixed, highly hand-crafted or carefully learned (by DNNs). Even more surprisingly, it sets new records for many classification tasks in Extended Yale B, AR, FERET datasets, and MNIST variations. Additional experiments on other public datasets also demonstrate the potential of the PCANet serving as a simple but highly competitive baseline for texture classification and object recognition.

**Index Terms**—Convolution Neural Network, Deep Learning, PCA Network, Random Network, LDA Network, Face Recognition, Handwritten Digit Recognition, Object Classification.



## 1 INTRODUCTION

Image classification based on visual content is a very challenging task, largely because there is usually large amount of intra-class variability, arising from different lightings, misalignment, non-rigid deformations, occlusion and corruptions. Numerous efforts have been made to counter the intra-class variability by manually designing low-level features for classification tasks at hand. Representative examples are Gabor features and local binary patterns (LBP) for texture and face classification, and SIFT and HOG features for object recognition. While the low-level features can be hand-crafted with great success for some specific data and tasks, designing effective features for new data and tasks usually requires new domain knowledge since most hand-crafted features cannot be simply adopted to new conditions [1], [2].

Learning features from the data of interest is considered as a plausible way to remedy the limitation of hand-crafted features. An example of such methods is learning through deep neural networks (DNNs), which draws significant attention recently [1]. The idea of deep learning is to discover multiple levels of representation, with the hope that higher-level features represent more abstract semantics of the data. Such abstract representations learned from a deep network are expected to provide more invariance to intra-class variability. One

key ingredient for success of deep learning in image classification is the use of convolutional architectures [3]–[10]. A convolutional deep neural network (ConvNet) architecture [3]–[5], [8], [9] consists of multiple trainable stages stacked on top of each other, followed by a supervised classifier. Each stage generally comprises of “three layers” – a convolutional filter bank layer, a nonlinear processing layer, and a feature pooling layer. To learn a filter bank in each stage of ConvNet, a variety of techniques has been proposed, such as restricted Boltzmann machines (RBM) [7] and regularized auto-encoders or their variations; see [2] for a review and references therein. In general, such a network is typically learned by stochastic gradient descent (SGD) method. However, learning a network useful for classification critically depends on expertise of parameter tuning and some *ad hoc* tricks.

While many variations of deep convolutional networks have been proposed for different vision tasks and their success is usually justified empirically, arguably the first instance that has led to clear mathematical justification is the wavelet scattering networks (ScatNet) [6], [10]. The only difference there is that the convolutional filters in ScatNet are prefixed – they are simply wavelet operators, hence no learning is needed at all. Somewhat surprisingly, such a pre-fixed filter bank, once utilized in a similar multistage architecture of ConvNet or DNNs, has demonstrated superior performance over ConvNet and DNNs in several challenging vision tasks such as handwritten digit and texture recognition [6], [10]. However, as we will see in this paper, such a prefixed architecture does not generalize so well to tasks

- 
- Tsung-Han Chan, Kui Jia, Shenghua Gao, Jiwen Lu, and Zinan Zeng are with Advanced Digital Sciences Center (ADSC), Singapore.
  - Yi Ma is with the School of Information Science and Technology of ShanghaiTech University, and with the ECE Department of the University of Illinois at Urbana-Champaign

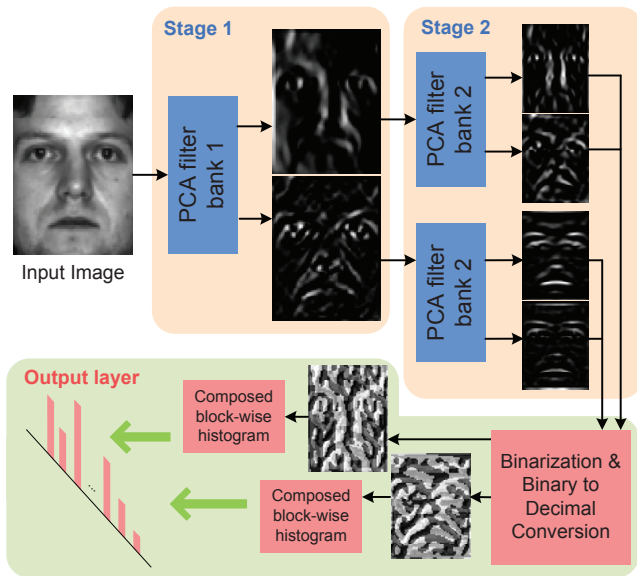


Fig. 1. Illustration of how the proposed PCANet extracts features from an image through three simplest processing components: PCA filters, binary hashing, and histogram.

such as face recognition where the intra-class variability includes significant illumination change and corruption.

### 1.1 Motivations

An initial motivation of our study is trying to resolve some apparent discrepancies between ConvNet and ScatNet. We want to achieve two simple goals: First, we want to design a simple deep learning network which should be very easy, even trivial, to train and to adapt to different data and tasks. Second, such a basic network could serve as a good baseline for people to empirically justify the use of more advanced processing components or more sophisticated architectures for their deep learning networks.

The solution comes as no surprise: We use the most basic and easy operations to emulate the processing layers in a typical (convolutional) neural network mentioned above: The data-adapting convolution filter bank in each stage is chosen to be the most basic PCA filters; the non-linear layer is set to be the simplest binary quantization (hashing); for the feature pooling layer, we simply use the block-wise histograms of the binary codes, which is considered as the final output features of the network. For ease of reference, we name such a data-processing network as a *PCA Network* (PCANet). As example, Figure 1 illustrates how a two-stage PCANet extracts features from an input image.

At least one characteristic of the PCANet model seem to challenge common wisdoms in building a deep learning network such as ConvNet [4], [5], [8] and ScatNet [6], [10]: No nonlinear operations in early stages of the PCANet until the very last output layer where binary hashing and histograms are conducted to compute the output features. Nevertheless, as we will see through

extensive experiments, such drastic simplification does not seem to undermine performance of the network on some of the typical datasets.

A network closely related to PCANet could be two-stage oriented PCA (OPCA), which was first proposed for audio processing [11]. Noticeable differences from PCANet lie in that OPCA does not couple with hashing and local histogram in the output layer. Given covariance of noises, OPCA gains additional robustness to noises and distortions. The baseline PCANet could also incorporate the merit of OPCA, likely offering more invariance to intra-class variability. To this end, we have also explored a supervised extension of PCANet, we replace the PCA filters with filters that are learned from linear discriminant analysis (LDA), called LDANet. As we will see through extensive experiments, the additional discriminative information does not seem to improve performance of the network; see Sections 2.3 and 3. Another, somewhat extreme, variation to PCANet is to replace the PCA filters with totally random filters (say the filter entries are i.i.d. Gaussian variables), called RandNet. In this work, we conducted extensive experiments and fair comparisons of these types of networks with other existing networks such as ConvNet and ScatNet. We hope our experiments and observations will help people gain better understanding of these different networks.

### 1.2 Contributions

Although our initial intention of studying the simple PCANet architecture is to have a simple baseline for comparing and justifying other more advanced deep learning components or architectures, our findings lead to some pleasant but thought-provoking surprises: The very basic PCANet, in *fair* experimental comparison, is already quite on par with, and often better than, the state-of-the-art features (prefixed, hand-crafted, or learned from DNNs) for almost all image classification tasks, including face images, hand-written digits, texture images, and object images. More specifically, for face recognition with one gallery image per person, it achieves 99.58% accuracy in Extended Yale B dataset, and over 95% accuracy for across disguise/illumination subsets in AR dataset. In FERET dataset, it obtains the state-of-the-art average accuracy 97.25% and achieves the best accuracy of 95.84% and 94.02% in Dup-1 and Dup-2 subsets, respectively.<sup>1</sup> In LFW dataset, it achieves competitive 86.28% face verification accuracy under “unsupervised setting”. In MNIST datasets, it achieves the state-of-the-art results for subtasks such as basic, background random, and background image. See Section 3 for more details. Overwhelming empirical evidences demonstrate the effectiveness of the proposed PCANet in learning robust invariant features for various image classification tasks.

1. The results were obtained by following FERET standard training CD, and could be marginally better when the PCANet is trained on MultiPIE database.

The method hardly contains any deep or new techniques and our study so far is entirely empirical.<sup>2</sup> Nevertheless, a thorough report on such a baseline system has tremendous value to the deep learning and visual recognition community, sending both *sobering and encouraging* messages: On one hand, for future study, PCANet can serve as a simple but surprisingly competitive baseline to empirically justify any advanced designs of multistage features or networks. On the other hand, the empirical success of PCANet (even that of RandNet) confirms again certain remarkable benefits from cascaded feature learning or extraction architectures. Even more importantly, since PCANet consists of only a (cascaded) linear map, followed by binary hashing and block histograms, it is now amenable to mathematical analysis and justification of its effectiveness. That could lead to fundamental theoretical insights about general deep networks, which seems in urgent need for deep learning nowadays.

## 2 CASCADED LINEAR NETWORKS

### 2.1 Structures of the PCA Network (PCANet)

Suppose that we are given  $N$  input training images  $\{\mathcal{I}_i\}_{i=1}^N$  of size  $m \times n$ , and assume that the patch size (or 2D filter size) is  $k_1 \times k_2$  at all stages. The proposed PCANet model is illustrated in Figure 2, and only the PCA filters need to be learned from the input images  $\{\mathcal{I}_i\}_{i=1}^N$ . In what follows, we describe each component of the block diagram more precisely.

#### 2.1.1 The first stage: PCA

Around each pixel, we take a  $k_1 \times k_2$  patch, and we collect all (overlapping) patches of the  $i$ th image; i.e.,  $\mathbf{x}_{i,1}, \mathbf{x}_{i,2}, \dots, \mathbf{x}_{i,mn} \in \mathbb{R}^{k_1 k_2}$  where each  $\mathbf{x}_{i,j}$  denotes the  $j$ th vectorized patch in  $\mathcal{I}_i$ . We then subtract patch mean from each patch and obtain  $\bar{\mathbf{X}}_i = [\bar{\mathbf{x}}_{i,1}, \bar{\mathbf{x}}_{i,2}, \dots, \bar{\mathbf{x}}_{i,mn}]$ , where  $\bar{\mathbf{x}}_{i,j}$  is a mean-removed patch. By constructing the same matrix for all input images and putting them together, we get

$$\mathbf{X} = [\bar{\mathbf{X}}_1, \bar{\mathbf{X}}_2, \dots, \bar{\mathbf{X}}_N] \in \mathbb{R}^{k_1 k_2 \times Nmn}. \quad (1)$$

Assuming that the number of filters in layer  $i$  is  $L_i$ , PCA minimizes the reconstruction error within a family of orthonormal filters, i.e.,

$$\min_{\mathbf{V} \in \mathbb{R}^{k_1 k_2 \times L_1}} \|\mathbf{X} - \mathbf{V}\mathbf{V}^T \mathbf{X}\|_F^2, \text{ s.t. } \mathbf{V}^T \mathbf{V} = \mathbf{I}_{L_1}, \quad (2)$$

where  $\mathbf{I}_{L_1}$  is identity matrix of size  $L_1 \times L_1$ . The solution is known as the  $L_1$  principal eigenvectors of  $\mathbf{X}\mathbf{X}^T$ . The PCA filters are therefore expressed as

$$\mathbf{W}_l^1 \doteq \text{mat}_{k_1, k_2}(\mathbf{q}_l(\mathbf{X}\mathbf{X}^T)) \in \mathbb{R}^{k_1 \times k_2}, \quad l = 1, 2, \dots, L_1, \quad (3)$$

where  $\text{mat}_{k_1, k_2}(\mathbf{v})$  is a function that maps  $\mathbf{v} \in \mathbb{R}^{k_1 k_2}$  to a matrix  $\mathbf{W} \in \mathbb{R}^{k_1 \times k_2}$ , and  $\mathbf{q}_l(\mathbf{X}\mathbf{X}^T)$  denotes the  $l$ th principal eigenvector of  $\mathbf{X}\mathbf{X}^T$ . The leading principal

2. We would be surprised if something similar to PCANet or variations to OPCA [11] have not been suggested or experimented with before in the vast learning literature.

eigenvectors capture the main variation of all the mean-removed training patches. Of course, similar to DNN or ScatNet, we can stack multiple stages of PCA filters to extract higher level features.

#### 2.1.2 The second stage: PCA

Almost repeating the same process as the first stage. Let the  $l$ th filter output of the first stage be

$$\mathcal{I}_i^l \doteq \mathcal{I}_i * \mathbf{W}_l^1, \quad i = 1, 2, \dots, N, \quad (4)$$

where  $*$  denotes 2D convolution, and the boundary of  $\mathcal{I}_i$  is zero-padded before convolving with  $\mathbf{W}_l^1$  so as to make  $\mathcal{I}_i^l$  having the same size of  $\mathcal{I}_i$ . Like the first stage, we can collect all the overlapping patches of  $\mathcal{I}_i^l$ , subtract patch mean from each patch, and form  $\bar{\mathbf{Y}}_i^l = [\bar{\mathbf{y}}_{i,l,1}, \bar{\mathbf{y}}_{i,l,2}, \dots, \bar{\mathbf{y}}_{i,l,mn}] \in \mathbb{R}^{k_1 k_2 \times mn}$ , where  $\bar{\mathbf{y}}_{i,l,j}$  is the  $j$ th mean-removed patch in  $\mathcal{I}_i^l$ . We further define  $\mathbf{Y}^l = [\bar{\mathbf{Y}}_1^l, \bar{\mathbf{Y}}_2^l, \dots, \bar{\mathbf{Y}}_N^l] \in \mathbb{R}^{k_1 k_2 \times Nmn}$  for the matrix collecting all mean-removed patches of the  $l$ th filter output, and concatenate  $\mathbf{Y}^l$  for all the filter outputs as

$$\mathbf{Y} = [\mathbf{Y}^1, \mathbf{Y}^2, \dots, \mathbf{Y}^{L_1}] \in \mathbb{R}^{k_1 k_2 \times L_1 Nmn}. \quad (5)$$

The PCA filters of the second stage are then obtained as

$$\mathbf{W}_\ell^2 \doteq \text{mat}_{k_1, k_2}(\mathbf{q}_\ell(\mathbf{Y}\mathbf{Y}^T)) \in \mathbb{R}^{k_1 \times k_2}, \quad \ell = 1, 2, \dots, L_2. \quad (6)$$

For each input  $\mathcal{I}_i^l$  of the second stage, we will have  $L_2$  outputs, each convolves  $\mathcal{I}_i^l$  with  $\mathbf{W}_\ell^2$  for  $\ell = 1, 2, \dots, L_2$ :

$$\mathcal{O}_i^\ell \doteq \{\mathcal{I}_i^l * \mathbf{W}_\ell^2\}_{\ell=1}^{L_2}. \quad (7)$$

The number of outputs of the second stage is  $L_1 L_2$ . One can simply repeat the above process to build more (PCA) stages if a deeper architecture is found to be beneficial.

#### 2.1.3 Output stage: hashing and histogram

For each of the  $L_1$  input images  $\mathcal{I}_i^l$  for the second stage, it has  $L_2$  real-valued outputs  $\{\mathcal{I}_i^l * \mathbf{W}_\ell^2\}_{\ell=1}^{L_2}$  from the second stage. We binarize these outputs and get  $\{H(\mathcal{I}_i^l * \mathbf{W}_\ell^2)\}_{\ell=1}^{L_2}$ , where  $H(\cdot)$  is a Heaviside step (like) function whose value is one for positive entries and zero otherwise.

Around each pixel, we view the vector of  $L_2$  binary bits as a decimal number. This converts the  $L_2$  outputs in  $\mathcal{O}_i^\ell$  back into a single integer-valued “image”:

$$\mathcal{T}_i^l \doteq \sum_{\ell=1}^{L_2} 2^{\ell-1} H(\mathcal{I}_i^l * \mathbf{W}_\ell^2), \quad (8)$$

whose every pixel is an integer in the range  $[0, 2^{L_2} - 1]$ . The order and weights of for the  $L_2$  outputs is irrelevant as we here treat each integer as a distinct “word.”

For each of the  $L_1$  images  $\mathcal{T}_i^l, l = 1, \dots, L_1$ , we partition it into  $B$  blocks. We compute the histogram (with  $2^{L_2}$  bins) of the decimal values in each block, and concatenate all the  $B$  histograms into one vector and denote as  $\text{Bhist}(\mathcal{T}_i^l)$ . After this encoding process, the






Fig. 2. The detailed block diagram of the proposed (two-stage) PCANet.

“feature” of the input image  $\mathcal{I}_i$  is then defined to be the set of block-wise histograms; i.e.,

$$f_i \doteq [\text{Bhist}(\mathcal{T}_i^1), \dots, \text{Bhist}(\mathcal{T}_i^{L_1})]^T \in \mathbb{R}^{(2^{L_2})L_1 B}. \quad (9)$$

The local blocks can be either overlapping or non-overlapping, depending on applications. Our empirical experience suggests that non-overlapping blocks are suitable for face images, whereas the overlapping blocks are appropriate for hand-written digits, textures, and object images. Furthermore, the histogram offers some degree of translation invariance in the extracted features, as in hand-crafted features (e.g., scale-invariant feature transform (SIFT) [12] or histogram of oriented gradients (HOG) [13]), learned features (e.g., bag-of-words (BoW) model [14]), and average or maximum pooling process in ConvNet [3]–[5], [8], [9].

The model parameters of PCANet include the filter size  $k_1, k_2$ , the number of filters in each stage  $L_1, L_2$ , the number of stages, and the block size for local histograms in the output layer. PCA filter banks require that  $k_1 k_2 \geq L_1, L_2$ . In our experiments in Section 3, we always set  $L_1 = L_2 = 8$  inspired from the common setting of Gabor filters [15] with 8 orientations, although some fine-tuned  $L_1, L_2$  could lead to marginal performance improvement. Moreover, we have noticed empirically that two-stage PCANet is in general sufficient to achieve good performance and a deeper architecture does not necessarily lead to further improvement. Also, larger block size for local histograms provides more translation invariance in the extracted feature  $f_i$ .

#### 2.1.4 Comparison with ConvNet and ScatNet

Clearly, PCANet shares some similarities with ConvNet [5]. The patch-mean removal in PCANet is reminiscent of local contrast normalization in ConvNet.<sup>3</sup> This operation moves all the patches to be centered around the origin of the vector space, so that the learned PCA filters can better catch major variations in the data. In addition, PCA

3. We have tested the PCANet without patch-mean removal and the performance degrades significantly.

can be viewed as the simplest class of auto-encoders, which minimizes reconstruction error.

The PCANet contains no non-linearity process between/in stages, running contrary to the common wisdom of building deep learning networks; e.g., the absolute rectification layer in ConvNet [5] and the modulus layer in ScatNet [6], [10]. We have tested the PCANet with an absolute rectification layer added right after the first stage, but we did not observe any improvement on the final classification results. The reason could be that the use of quantization plus local histogram (in the output layer) already introduces sufficient invariance and robustness in the final feature.

The overall process prior to the output layer in PCANet is completely linear. One may wonder what if we merge the two stages into just one that has an equivalently same number of (PCA-like) filters and size of receptive field. To be specific, one may be interested in how the single-staged PCANet with  $L_1 L_2$  filters of size  $(k_1 + k_2 - 1) \times (k_1 + k_2 - 1)$  could perform, in comparison to the two-staged PCANet we described in Section 2.1. We have experimented with such settings on various image classification tasks including face and handwritten digit and observed that the two-staged PCANet consistently outperforms this single-stage alternative. Note that within the capability of learning, it is more desirable to use a larger receptive field, since it contains more holistic observations of the objects in images, and learning invariance from it can essentially capture more semantic information. Our comparative experiments suggest that hierarchical architectures with multiple stacked stages are more efficient in learning semantically related representations from relatively large receptive fields of input images, which coincides with what have been observed in [7].

## 2.2 Computational Complexity

The components for constructing the PCANet are extremely basic and computationally efficient. To see how light the computational complexity of PCANet would be, let us take the two-staged PCANet as an example. In

each stage of PCANet, forming the patch-mean-removed matrix  $\mathbf{X}$  costs  $k_1k_2 + k_1k_2mn$  flops; the inner product  $\mathbf{X}\mathbf{X}^T$  has complexity of  $2(k_1k_2)^2mn$  flops; and the complexity of eigen-decomposition is  $\mathcal{O}((k_1k_2)^3)$ . The PCA filter convolution takes  $L_ik_1k_2mn$  flops for stage  $i$ . In the output layer, the conversion of  $L_2$  binary bits to a decimal number costs  $2L_2mn$ , and the naive histogram operation is of complexity  $\mathcal{O}(mnBL_2\log 2)$ . Assuming  $mn \gg \max(k_1, k_2, L_1, L_2, B)$ , the overall complexity of PCANet is easy to be verified as

$$\mathcal{O}(mnk_1k_2(L_1 + L_2) + mn(k_1k_2)^2).$$

The above computational complexity applies to training phase and testing phase of PCANet, as the extra computational burden in training phase from testing phase is the eigen-decomposition, whose complexity is ignorable when  $mn \gg \max(k_1, k_2, L_1, L_2, B)$ .

In comparison to ConvNet, the SGD for filter learning is also a simple gradient-based optimization solver, but the overall training time is still much longer than PCANet. For example, training PCANet on around 100,000 images of  $80 \times 60$  pixel dimension took only half a hour, but CNN-2 took 6 hours, excluding the fine-tuning process; see Section 3.1.1.D for details.

### 2.3 Two Variations: RandNet and LDANet

The PCANet is an extremely simple network, requiring only minimum learning of the filters from the training data. One could immediately think of two possible variations of the PCANet towards two opposite directions:

- 1) We could further eliminate the necessity of training data and replace the PCA filters at each layer with random filters of the same size. Be more specific, for random filters, i.e., the elements of  $\mathbf{W}_l^1$  and  $\mathbf{W}_l^2$ , are generated following standard Gaussian distribution. We call such a network *Random Network*, or RandNet as a shorthand. It is natural to wonder how much degradation such a randomly chosen network would perform in comparison with PCANet.
- 2) If the task of the learned network is for classification, we could further enhance the supervision of the learned filters by incorporating the information of class labels in the training data and learn the filters based on the idea of multi-class linear discriminant analysis (LDA). We called so composed network *LDA Network*, or LDANet for ease of reference. Again we are interested in how much the enhanced supervision would help improve the performance of the network.

To be more clear, we here describe with more details how to construct the LDANet. Suppose that the  $N$  training images are classified into  $C$  classes  $\{\mathcal{I}_i\}_{i \in S_c}$ ,  $c = 1, 2, \dots, C$  where  $S_c$  is the set of indices of images in class  $c$ , and the mean-removed patches associated with each image of distinct classes  $\bar{\mathbf{X}}_i \in \mathbb{R}^{k_1k_2 \times mn}$ ,  $i \in S_c$  (in the same spirit of  $\bar{\mathbf{X}}_i$  in (1)) are given. We can first

compute the class mean  $\Gamma_c$  and the intra-class variability  $\Sigma_c$  for all the patches as follows,

$$\Gamma_c = \sum_{i \in S_c} \bar{\mathbf{X}}_i / |S_c|, \quad (10)$$

$$\Sigma_c = \sum_{i \in S_c} (\bar{\mathbf{X}}_i - \Gamma_c)(\bar{\mathbf{X}}_i - \Gamma_c)^T / |S_c|. \quad (11)$$

Each column of  $\Gamma_c$  denotes the mean of patches around each pixel in the class  $c$ , and  $\Sigma_c$  is the sum of all the patch-wise sample covariances in the class  $c$ . Likewise, the inter-class variability of the patches is defined as

$$\Phi = \sum_{c=1}^C (\Gamma_c - \Gamma)(\Gamma_c - \Gamma)^T / C, \quad (12)$$

where  $\Gamma$  is the mean of class means. The idea of LDA is to maximize the ratio of the inter-class variability to sum of the intra-class variability within a family of orthonormal filters; i.e.,

$$\max_{\mathbf{V} \in \mathbb{R}^{k_1k_2 \times L_1}} \frac{\text{Tr}(\mathbf{V}^T \Phi \mathbf{V})}{\text{Tr}(\mathbf{V}^T (\sum_{c=1}^C \Sigma_c) \mathbf{V})}, \text{ s.t. } \mathbf{V}^T \mathbf{V} = \mathbf{I}_{L_1}, \quad (13)$$

where  $\text{Tr}(\cdot)$  is the trace operator. The solution is known as the  $L_1$  principal eigenvectors of  $\tilde{\Phi} = (\sum_{c=1}^C \Sigma_c)^\dagger \Phi$ , where the superscript  $\dagger$  denotes the pseudo-inverse. The pseudo inverse is to deal with the case when  $\sum_{c=1}^C \Sigma_c$  is not of full rank, though there might be another way of handling this with better numeric stability [16]. The LDA filters are thus expressed as  $\mathbf{W}_l^1 = \text{mat}_{k_1, k_2}(\mathbf{q}_l(\tilde{\Phi})) \in \mathbb{R}^{k_1 \times k_2}$ ,  $l = 1, 2, \dots, L_1$ . A deeper network can be built by repeating the same process as above. .

## 3 EXPERIMENTS

We now evaluate the performance of the proposed PCANet and the two simple variations (RandNet and LDANet) in various tasks, including face recognition, face verification, hand-written digits recognition, texture discrimination, and object recognition in this section.

### 3.1 Face Recognition on Many Datasets

We first focus on the problem of *face recognition with one gallery image per person*. We use part of MultiPIE dataset to learn PCA filters in PCANet, and then apply such trained PCANet to extract features of new subjects in MultiPIE dataset, Extended Yale B, AR, and FERET datasets for face recognition.

#### 3.1.1 Training and Testing on MultiPIE Dataset.

MultiPIE dataset [17] contains 337 subjects across simultaneous variation in pose, expression, and illumination. Of these 337 subjects, we select the images of 129 subjects that enrolled all the four sessions. The images of a subject under all illuminations and all expressions at pose  $-30^\circ$  to  $+30^\circ$  with step size  $15^\circ$ , a total of 5 poses, were collected. We manually select eye corners as the ground truth for registration, and down-sample




Fig. 3. Recognition accuracy of PCANet on MultiPIE cross-illumination test set for varying number of filters in the first stage. (a) PCANet-1; (b) PCANet-2 with  $L_2 = 8$ .

the images to  $80 \times 60$  pixels. The distance between the two outer eye corners is normalized to be 50 pixels. This generic training set comprises around 100,000 images, and all images are converted to gray scale.

We use these assembled face images to train the PCANet and together with data labels to learn LDANet, and then apply the trained networks to extract features of the new subjects in Multi-PIE dataset. As mentioned above, 129 subjects enrolling all four sessions are used for PCANet training. The remaining 120 new subjects in Session 1 are used for gallery training and testing. Frontal view of each subject with neutral expression and frontal illumination is used in gallery, and the rest is for testing. We classify all the possible variations into 7 test sets, namely cross illumination, cross expression, cross pose, cross expression-plus-pose, cross illumination-plus-expression, cross illumination-plus-pose, and cross illumination-plus-expression-and-pose. The cross-pose test set is specifically collected over the poses  $-30^\circ$ ,  $-15^\circ$ ,  $+15^\circ$ ,  $+30^\circ$ .

*A. Impact of the number of filters.* Before comparing RandNet, PCANet, and LDANet with existing methods on all the 7 test sets, we first investigate the impact of the number of filters of these networks on the cross-illumination test set only. The filter size of the networks is  $k_1 = k_2 = 5$  and their non-overlapping blocks is of size  $8 \times 6$ . We vary the number of filters in the first stage  $L_1$  from 2 to 12 for one-staged networks. When considering two-staged networks, we set  $L_2 = 8$  and vary  $L_1$  from 4 to 24. The results are shown in Figure 3. One can see that PCANet-1 achieves the best results for  $L_1 \geq 4$  and PCANet-2 is the best for all  $L_1$  under test. Moreover, the accuracy of PCANet and LDANet (for both one-staged and two-staged networks) increases for larger  $L_1$ , and the RandNet also has similar performance trend. However, some performance fluctuation is observed for RandNet due to the filters' randomness.

*B. Impact of the the block size.* We next examine the impact of the block size (for histogram computation) on robustness of PCANet against image deformations. We use the cross-illumination test set, and introduce artificial deformation to the testing image with a translation, in-plane rotation or scaling; see Figure 4. The parameters




Fig. 4. Original image and its artificially deformed images.

TABLE 1  
Face recognition rates (%) of PCANet on MultiPIE cross-illumination test set, with respect to different amount of generic training images ( $S$ ).

$S$	100	500	1,000	5,000	10,000	50,000
PCANet-1	98.01	98.44	98.61	98.65	98.70	98.70
PCANet-2	100.00	100.00	100.00	100.00	100.00	100.00

of PCANet are set to  $k_1 = k_2 = 5$  and  $L_1 = L_2 = 8$ . Two block sizes  $8 \times 6$  and  $12 \times 9$  are considered. Figure 5 shows the recognition accuracy for each artificial deformation. It is observed that PCANet-2 achieves more than 90 percent accuracy with translation up to 4 pixels in all directions, up to  $8^\circ$  in-plane rotation, or with scale varying from 0.9 to 1.075. Moreover, the results suggest that PCANet-2 with larger block size provides more robustness against various deformations, but a larger block side may sacrifice some performance for PCANet-1.

*C. Impact of the number of generic training samples.* We also report the recognition accuracy of the PCANet for different number of the generic training images. Again, we use cross-illumination test set. We randomly select  $S$  images from the generic training set to train the PCANet, and varies  $S$  from 10 to 50,000. The parameters of PCANet are set to  $k_1 = k_2 = 5$ ,  $L_1 = L_2 = 8$ , and block size  $8 \times 6$ . The results are tabulated in Table 1. Surprisingly, the accuracy of PCANet is less-sensitive to the number of generic training images. The performance of PCANet-1 gradually improves as the number of generic training samples increases, and PCANet-2 keeps perfect recognition even when there are only 100 generic training samples.

*D. Comparison with state of the arts.* We compare the RandNet, PCANet, and LDANet with Gabor<sup>4</sup> [15], LBP<sup>5</sup> [18], and two-staged ScatNet (ScatNet-2) [6]. We set the parameters of PCANet to the filter size  $k_1 = k_2 = 5$ , the number of filters  $L_1 = L_2 = 8$ , and  $8 \times 6$  block size, and the learned PCANet filters are shown in Figure 6. The number of scales and the number of orientations

4. Each face is convolved with a family of Gabor kernels with 5 scales and 8 orientations. Each filter response is down-sampled by a  $3 \times 3$  uniform lattice, and normalized to zero mean and unit variance.

5. Each face is divided into several blocks, each of size the same as PCANet. The histogram of 59 uniform binary patterns is then computed, where the patterns are generated by thresholding 8 neighboring pixels in a circle of radius 2 using the central pixel value.



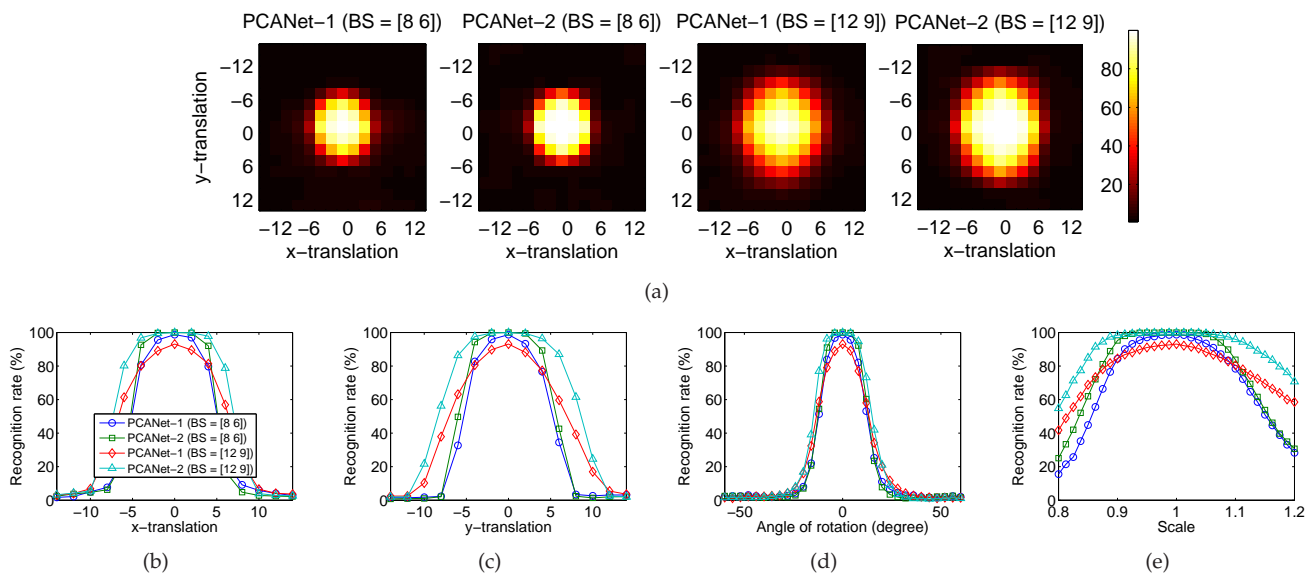


Fig. 5. Recognition rate of PCANet on MultiPIE cross-illumination test set, for different PCANet block size and deformation to the test image. Two block sizes [8 6] and [12 9] for histogram aggregation are tested. (a) Simultaneous translation in  $x$  and  $y$  directions. (b) Translation in  $x$  direction. (c) Translation in  $y$  direction. (d) In-plane rotation. (e) Scale variation.



Fig. 6. The PCANet filters learned on MultiPIE dataset. Top row: the first stage. Bottom row: the second stage.

in ScatNet-2 are set to 3 and 8, respectively. We use the nearest neighbor (NN) classifier with the chi-square distance for RandNet, PCANet, LDANet and LBP, or with the cosine distance for Gabor and ScatNet. The NN classifier with different distance measure is to secure the best performances of respective features.

We also compare with CNN. Since we could not find any work that successfully applies CNN to the same face recognition tasks, we use Caffe framework [19] to pre-train a two-staged CNN (CNN-2) on the generic training set. The CNN-2 is fully-supervised network with filter size  $5 \times 5$ ; 20 channels for the first stage and 50 channels for the second stage. Each convolution output is followed by a rectified linear function  $relu(x) = \max(x, 0)$  and  $2 \times 2$  max-pooling. The output layer is a softmax classifier. After pre-training the CNN-2 on the generic training set, the CNN-2 is also fine-tuned on the 120 gallery images for 500 epochs.

The performance of all methods are given in Table 2. Except for cross-pose test set, the PCANet yields the best precision. For all test sets, the performance of RandNet and LDANet is inferior to that of PCANet, and LDANet does not seem to take advantage of discriminative information. One can also see that whenever there is illumination variation, the performance of LBP drops

significantly. The PCANet overcomes this drawback and offers comparable performance to LBP for cross-pose and cross-expression variations. As a final note, ScatNet and CNN seem not performing well. This is the case for all face-related experiments below, and therefore ScatNet and CNN are not included for comparison in these experiments. We also do not include RandNet and LDANet in the following face-related experiments, as they did not show performance superior over PCANet.

Another issue worth mentioning is the efficiency of the PCANet. Training PCANet-2 on the generic training set (i.e., around 100,000 face images of  $80 \times 60$  pixel dimension) took only half a hour, but CNN-2 took 6 hours, excluding the fine-tuning process.

### 3.1.2 Testing on Extended Yale B Dataset.

We then apply the PCANet with the PCA filters learned from MultiPIE to Extended Yale B dataset [20]. Extended Yale B dataset consists of 2414 frontal-face images of 38 individuals. The cropped and normalized  $192 \times 168$  face images were captured under various laboratory-controlled lighting conditions. For each subject, we select frontal illumination as the gallery images, and the rest for testing. To challenge ourselves, in the test images, we also simulate various levels of contiguous occlusion, from 0 percent to 80 percent, by replacing a randomly located square block of each test image with an unrelated image; see Figure 7 for example. The size of non-overlapping blocks in the PCANet is set to  $8 \times 8$ . We compare with LBP [18] and LBP of the test images being processed by illumination normalization, P-LBP [21]. We use the NN classifier with the chi-square distance measure.

TABLE 2  
Comparison of face recognition rates (%) of various methods on MultiPIE test sets.

Test Sets	Illum.	Exps.	Pose	Exps.+Pose	Illum.+Exps.	Illum.+Pose	Illum.+Exps.+Pose
Gabor [15]	68.75	94.17	84.17	64.70	38.09	39.76	25.92
LBP [18]	79.77	98.33	<b>95.63</b>	86.88	53.77	50.72	40.55
ScatNet-2 [6]	20.88	66.67	71.46	54.37	14.51	15.00	14.47
CNN-2 [8]	46.71	75.00	73.54	57.50	23.38	25.05	18.74
RandNet-1	80.88	98.33	87.50	75.62	46.57	42.80	31.85
RandNet-2	97.64	97.50	83.13	75.21	63.87	53.50	42.47
PCANet-1	98.70	<b>99.17</b>	94.17	<b>87.71</b>	72.40	65.76	53.80
PCANet-2	<b>100.00</b>	<b>99.17</b>	93.33	87.29	<b>87.89</b>	<b>75.29</b>	<b>66.49</b>
LDANet-1	99.95	98.33	92.08	82.71	77.89	68.55	57.97
LDANet-2	96.02	99.17	93.33	83.96	65.78	60.14	46.72




Fig. 7. Illustration of varying level of an occluded test face image.

TABLE 3  
Recognition rates (%) on Extended Yale B dataset.

Percent occluded	0%	20%	40%	60%	80%
LBP [18]	75.76	65.66	54.92	43.22	18.06
P-LBP [21]	96.13	91.84	84.13	70.96	41.29
PCANet-1	97.77	96.34	93.81	84.60	<b>54.38</b>
PCANet-2	<b>99.58</b>	<b>99.16</b>	<b>96.30</b>	<b>86.49</b>	51.73

The experimental results are given in Table 3. One can see that the PCANet outperforms the P-LBP for different levels of occlusion. It is also observed that the PCANet is not only illumination-insensitive, but also robust against block occlusion. Under such a single sample per person setting and various difficult lighting conditions, the PCANet surprisingly achieves almost perfect recognition 99.58%, and still sustains 86.49% accuracy when 60% pixels of every test image are occluded! The reason could be that the learned PCA filters in PCANet provide sufficient robustness to occlusion. Specifically, each PCA filter can be seen as a detector with the maximum response for patches from a face. In other words, the contribution from occluded patches would somehow be ignored after PCA filtering and are not passed onto the output layer of the PCANet, thereby yielding striking robustness to occlusion.

### 3.1.3 Testing on AR Dataset.

We further evaluate the ability of the MultiPIE-learned PCANet to cope with real possibly malicious occlusions using AR dataset [22]. AR dataset consists of over 4,000 frontal images for 126 subjects. For each individual, 26 pictures were taken in two separate sessions. These

TABLE 4  
Recognition rates (%) on AR dataset.

Test sets	Illum.	Exps.	Disguise	Disguise + Illum.
LBP [18]	93.83	81.33	91.25	79.63
P-LBP [21]	97.50	80.33	93.00	88.58
PCANet-1	98.00	<b>85.67</b>	95.75	92.75
PCANet-2	<b>99.50</b>	85.00	<b>97.00</b>	<b>95.00</b>

images include different facial expressions, illumination conditions and disguises. In the experiment, we chose a subset of the data consisting of 50 male subjects and 50 female subjects. The images are cropped with dimension  $165 \times 120$  and converted to gray scale. For each subject, we select the face with frontal illumination and neutral expression in the gallery training, and the rest are all for testing. The size of non-overlapping blocks in the PCANet is set to  $8 \times 6$ . We also compare with LBP [18] and P-LBP [21]. We use the NN classifier with the chi-square distance measure.

The results are given in Table 4. For test set of illumination variations, the recognition by PCANet is again almost perfect, and for cross-disguise related test sets, the accuracy is more than 95%. The results are consistent with that on MultiPIE and Extended Yale B datasets: PCANet is insensitive to illumination and robust to occlusions. To the best of our knowledge, no single feature with a simple classifier can achieve such performances, even if in extended representation-based classification (ESRC) [23]!

### 3.1.4 Testing on FERET Dataset.

We finally apply the MultiPIE-learned PCANet to the popular FERET dataset [24], which is a standard dataset used for facial recognition system evaluation. FERET contains images of 1,196 different individuals with up to 5 images of each individual captured under different lighting conditions, with non-neural expressions and over the period of three years. The complete dataset is partitioned into disjoint sets: gallery and probe. The probe set is further subdivided into four categories: *Fb* with different expression changes; *Fc* with different lighting conditions; *Dup-I* taken within the period of three to four months; *Dup-II* taken at least one and a half year



TABLE 5  
Recognition rates (%) on FERET dataset.

Probe sets	<i>Fb</i>	<i>Fc</i>	<i>Dup-I</i>	<i>Dup-II</i>	Avg.
LBP [18]	93.00	51.00	61.00	50.00	63.75
DMMA [25]	98.10	98.50	81.60	83.20	89.60
P-LBP [21]	98.00	98.00	90.00	85.00	92.75
POEM [26]	99.60	99.50	88.80	85.00	93.20
G-LQP [27]	<b>99.90</b>	<b>100</b>	93.20	91.00	96.03
LGBP-LGXP [28]	99.00	99.00	94.00	93.00	96.25
sPOEM+POD [29]	99.70	<b>100</b>	94.90	94.00	97.15
GOM [30]	99.90	<b>100</b>	95.70	93.10	97.18
PCANet-1 (Trn. CD)	99.33	99.48	88.92	84.19	92.98
PCANet-2 (Trn. CD)	99.67	99.48	<b>95.84</b>	<b>94.02</b>	97.25
PCANet-1	99.50	98.97	89.89	86.75	93.78
PCANet-2	99.58	<b>100</b>	95.43	<b>94.02</b>	<b>97.26</b>

apart. We use the gray-scale images, cropped to image size of  $150 \times 90$  pixels. The size of non-overlapping blocks in the PCANet is set to  $15 \times 15$ . To compare fairly with prior methods, the dimension of the PCANet features are reduced to 1000 by a whitening PCA (WPCA),<sup>6</sup> where the projection matrix is learned from the features of gallery samples. The NN classifier with cosine distance is used. Moreover, in addition to PCANet trained from MultiPIE database, we also train PCANet on the FERET generic training set, consisting of 1,002 images of 429 people listed in the FERET standard training CD.

The results of the PCANet and other state-of-the-art methods are listed in Table 5. Surprisingly, both simple MultiPIE-learned PCANet-2 and FERET-learned PCANet-2 (with Trn. CD in a parentheses) achieve the state-of-the-art accuracies 97.25% and 97.26% on average, respectively. As the variations in MultiPIE database are much richer than the standard FERET training set, it is nature to see that the MultiPIE-learned PCANet slightly outperforms FERET-learned PCANet. More importantly, PCANet-2 breaks the records in *Dup-I* and *Dup-II*.

A prominent message drawn from the above experiments in sections 3.1.1, 3.1.2, 3.1.3, and 3.1.4 is that training PCANet from a face dataset can be very effective to capture the abstract representation of new subjects or new datasets. After the PCANet is trained, extracting PCANet-2 feature for one test face only takes 0.3 second in Matlab. We can anticipate that the performance of PCANet could be further improved and moved toward practical use if the PCANet is trained upon a wide and deep dataset that collect sufficiently many inter-class and intra-class variations.

### 3.2 Face Verification on LFW Dataset

Besides tests with laboratory face datasets, we also evaluate the PCANet on the LFW dataset [31] for unconstrained face verification. LFW contains 13,233 face images of 5,749 different individuals, collected from the web with large variations in pose, expression, illumination, clothing, hairstyles, etc. We consider “unsupervised setting”, which is the best choice for evaluating

6. The PCA projection directions are weighted by the inverse of their corresponding square-root energies, respectively.

the learned features, for it does not depend on metric learning and discriminative model learning. The aligned version of the faces, namely LFW-a, provided by Wolf *et al.* [32] is used, and the face images were cropped into  $150 \times 80$  pixel dimensions. We follow the standard evaluation protocol, which splits the View 2 dataset into 10 subsets with each subset containing 300 intra-class pairs and 300 inter-class pairs. We perform 10-fold cross validation using the 10 subsets of pairs in View 2. In PCANet, the filter size, the number of filters, and the (non-overlapping) block size are set to  $k_1 = k_2 = 7$ ,  $L_1 = L_2 = 8$ , and  $15 \times 13$ , respectively. The performances are measured by averaging the 10-fold cross validation. We project the PCANet features onto 400 and 3,200 dimensions using WPCA for PCANet-1 and PCANet-2, respectively, and use NN classifier with cosine distance.

Table 6 tabulates the results.<sup>7</sup> Note that PCANet followed by sqrt in a parentheses represents the PCANet feature taking square-root operation. One can see that the square-root PCANet outperforms PCANet, and this performance boost from square-root operation has also been observed in other features for this dataset [33]. Moreover, the square-root PCANet-2 that achieves 86.28% accuracy is quite competitive to the current state-of-the-art methods. This shows that the proposed PCANet is also effective in learning invariant features for face images captured in less controlled conditions.

In preparation of this paper, we are aware of two concurrent works [34], [35] that employ ConvNet for LFW face verification. While both works achieve very impressive results on LFW, their experimental setting differs from ours largely. These two works require some outside database to train the ConvNet and the face images have to be more precisely aligned; e.g., [34] uses 3-dimensional model for face alignment and [35] extracts multi-scale features based on detected landmark positions. On the contrary, we only trained PCANet based on LFW-a [32], an aligned version of LFW images using the commercial alignment system of face.com.

TABLE 6  
Comparison of verification rates (%) on LFW under unsupervised setting.

Methods	Accuracy
POEM [26]	82.70±0.59
High-dim. LBP [36]	84.08
High-dim. LE [36]	84.58
SFRD [37]	84.81
I-LQP [27]	86.20±0.46
OCLBP [33]	<b>86.66±0.30</b>
PCANet-1	81.18 ± 1.99
PCANet-1 (sqrt)	82.55 ± 1.48
PCANet-2	85.20 ± 1.46
PCANet-2 (sqrt)	86.28 ± 1.14

7. For fair comparison, we only report results of single descriptor. The best known LFW result under unsupervised setting is 88.57% [33], which is inferred from four different descriptors.

TABLE 7

Error rates (%) of PCANet-2 on *basic* dataset for varying block overlap ratios (BORs).

BOR	0.1	0.2	0.3	0.4	0.5	0.6	0.7
RandNet-2	1.31	1.35	1.23	1.34	1.18	1.14	1.24
PCANet-2	<b>1.12</b>	<b>1.12</b>	<b>1.07</b>	1.06	<b>1.02</b>	<b>1.02</b>	<b>1.05</b>
LDANet-2	1.14	1.14	1.11	<b>1.05</b>	<b>1.05</b>	1.05	1.06

### 3.3 Digit Recognition on MNIST Datasets

We now move forward to test the proposed PCANet, along with RandNet and LDANet, on MNIST [4] and MNIST variations [38], a widely-used benchmark for testing hierarchical representations. There are 9 classification tasks in total, as listed in Table 8. All the images are of size  $28 \times 28$ . In the following, we use *MNIST basic* as the dataset to investigate the influence of the number of filters or different block overlap ratios for RandNet, PCANet and LDANet, and then compare with other state-of-the-art methods on all the MNIST datasets.

#### 3.3.1 Impact of the number of filters

We vary the number of filters in the first stage  $L_1$  from 2 to 12 for one-staged networks. Regarding two-staged networks, we set  $L_2 = 8$  and change  $L_1$  from 4 to 24. The filter size of the networks is  $k_1 = k_2 = 7$ , block size is  $7 \times 7$ , and the overlapping region between blocks is half of the block size. The results are shown in Figure 8. The results are consistent with that for MultiPIE face database in Figure 3; PCANet outperforms RandNet and LDANet for almost all the cases.

#### 3.3.2 Impact of the block overlap ratio

The number of filters is fixed to  $L_1 = L_2 = 8$ , and the filter size is again  $k_1 = k_2 = 7$  and block size is  $7 \times 7$ . We only vary the block overlap ratio (BOR) from 0.1 to 0.7. Table 7 tabulates the results of RandNet-2, PCANet-2, and LDANet-2. Clearly, PCANet-2 and LDANet-2 achieve their minimum error rates for BOR equal to 0.5 and 0.6, respectively, and PCANet-2 performs the best for all conditions.

#### 3.3.3 Comparison with state of the arts

We compare RandNet, PCANet, and LDANet with ConvNet [5], 2-staged ScatNet (ScatNet-2) [6], and other existing methods. In ScatNet, the number of scales and the number of orientations are set to 3 and 8, respectively. Regarding the parameters of PCANet, we set the filter size  $k_1 = k_2 = 7$ , the number of PCA filters  $L_1 = L_2 = 8$ ; the block size is tuned by a cross-validation for *MNIST*, and the validation sets for MNIST variations<sup>8</sup>. The overlapping region between blocks is half of the block size. Unless otherwise specified, we use linear SVM classifier

8. Using either cross-validation or validation set, the optimal block size is obtained as  $7 \times 7$  for *MNIST*, *basic*, *rec-img*,  $4 \times 4$  for *rot*, *bg-img*, *bg-rnd*, *bg-img-rot*,  $14 \times 14$  for *rec*, and  $28 \times 28$  for *convex*.

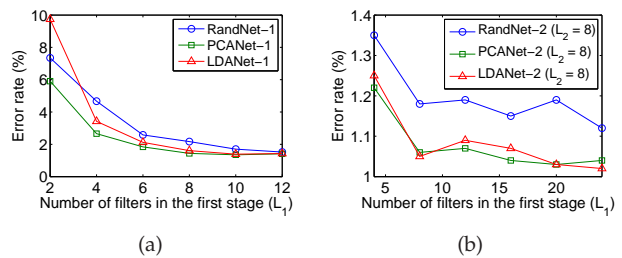


Fig. 8. Error rate of PCANet on MNIST basic test set for varying number of filters in the first stage. (a) PCANet-1; (b) PCANet-2 with  $L_2 = 8$ .

for ScatNet and RandNet, PCANet and LDANet for the 9 classification tasks.

The testing error rates of the various methods on *MNIST* are shown in Table 9. For fair comparison, we do not include the results of methods using augmented training samples with distortions or other information, for that the best known result is 0.23% [39]. We see that RandNet-2, PCANet-2, and LDANet-2 are comparable with the state-of-the-art methods on this standard *MNIST* task. However, as *MNIST* has many training data, all methods perform very well and very close – the difference is not so statistically meaningful.

Accordingly, we also report results of different methods on MNIST variations in Table 10. To the best of our knowledge, the PCANet-2 achieves the state-of-the-art results for four out of the eight remaining tasks: *basic*, *bg-img*, *bg-img-rot*, and *convex*. Especially for *bg-img*, the error rate reduces from 12.25% [40] to 6.27%.

Furthermore, we also draw the learned PCANet filters in Figure 9 and Figure 10. An intriguing pattern is observed in the filters of *rect* and *rect-img* datasets. For *rect*, we can see both horizontal and vertical stripes, for these patterns attempt to capture the edges of the rectangles. When there is some image background in *rect-img*, several filters become low-pass, in order to secure the responses from background images.

TABLE 9

Comparison of error rates (%) of the methods on *MNIST*, excluding methods that augment the training data.

Methods	<i>MNIST</i>
HSC [41]	0.77
K-NN-SCM [42]	0.63
K-NN-IDM [43]	0.54
CDBN [7]	0.82
ConvNet [5]	0.53
Stochastic pooling ConvNet [44]	0.47
Conv. Maxout + Dropout [3]	0.45
ScatNet-2 (SVM <sub>rbf</sub> ) [6]	<b>0.43</b>
RandNet-1	1.32
RandNet-2	0.63
PCANet-1	0.94
PCANet-2	0.66
LDANet-1	0.98
LDANet-2	0.62

TABLE 8  
Details of the 9 classification tasks on MNIST and MNIST variations.

Data Sets	Description	Num. of classes	Train-Valid-Test
<i>MNIST</i>	Standard MNIST	10	60000-0-10000
<i>basic</i>	Smaller subset of MNIST	10	10000-2000-50000
<i>rot</i>	MNIST with rotation	10	10000-2000-50000
<i>bg-rand</i>	MNIST with noise background	10	10000-2000-50000
<i>bg-img</i>	MNIST with image background	10	10000-2000-50000
<i>bg-img-rot</i>	MNIST with rotation and image background	10	10000-2000-50000
<i>rect</i>	Discriminate between tall and wide rectangles	2	1000-200-50000
<i>rect-img</i>	Dataset <i>rect</i> with image background	2	10000-2000-50000
<i>convex</i>	Discriminate between convex and concave shape	2	6000-2000-50000

TABLE 10  
Comparison of testing error rates (%) of the various methods on MNIST variations.

Methods	<i>basic</i>	<i>rot</i>	<i>bg-rand</i>	<i>bg-img</i>	<i>bg-img-rot</i>	<i>rect</i>	<i>rect-img</i>	<i>convex</i>
CAE-2 [45]	2.48	9.66	10.90	15.50	45.23	1.21	21.54	-
TIRBM [46]	-	<b>4.20</b>	-	-	35.50	-	-	-
PGBM + DN-1 [40]	-	-	<b>6.08</b>	12.25	36.76	-	-	-
ScatNet-2 [6]	1.27	7.48	18.40	12.30	50.48	<b>0.01</b>	<b>8.02</b>	6.50
RandNet-1	1.86	14.25	15.97	18.81	51.82	0.21	15.94	6.78
RandNet-2	1.25	8.47	11.65	13.47	43.69	0.09	17.00	5.45
PCANet-1	1.44	10.55	11.11	6.77	42.03	0.15	25.55	5.93
PCANet-2	1.06	7.37	10.95	<b>6.19</b>	<b>35.48</b>	0.24	14.08	<b>4.36</b>
LDANet-1	1.61	11.40	13.03	7.16	43.86	0.15	23.63	6.89
LDANet-2	<b>1.05</b>	7.52	12.42	6.81	38.54	0.14	16.20	7.22



Fig. 9. The PCANet filters learned on MNIST dataset. Top row: the first stage. Bottom row: the second stage.



Fig. 11. The PCANet filters learned on CURET database. Top row: the first stage. Bottom row: the second stage.

### 3.4 Texture Classification on CURET Dataset

The CURET texture dataset contains 61 classes of image textures. Each texture class has images of the same material with different pose and illumination conditions. Other than the above variations, specularities, shadowing and surface normal variations also make this classification challenging. In this experiment, a subset of the dataset with azimuthal viewing angle less than 60 degrees is selected, thereby yielding 92 images in each class. A central  $200 \times 200$  region is cropped from each of the selected images. The dataset is randomly split into a training and a testing set, with 46 training images for each class, as in [47]. The PCANet is trained with filter size  $k_1 = k_2 = 5$ , the number of filters  $L_1 = L_2 = 8$ , and block size  $50 \times 50$ . We use linear SVM classifier. The testing error rates averaged over 10 different random splits are shown in Table 11. We see that the PCANet-1 outperforms ScatNet-1, but the improvement from PCANet-1 to PCANet-2 is not as large as that of ScatNet. Note that ScatNet-2 followed by a PCA-based classifier gives the best result [6].

TABLE 11  
Comparison of error rates (%) on CURET.

Methods	Error rates
Textons [48]	1.50
BIF [49]	1.40
Histogram [50]	1.00
ScatNet-1 (PCA) [6]	0.50
ScatNet-2 (PCA) [6]	<b>0.20</b>
RandNet-1	0.61
RandNet-2	0.46
PCANet-1	0.45
PCANet-2	0.39
LDANet-1	0.69
LDANet-2	0.54

### 3.5 Object Recognition on Caltech101 Dataset

We finally evaluate the performance of PCANet on Caltech101 dataset for object recognition. Caltech101 dataset contains 9,144 images in 101 classes, including animals, vehicles, flowers, etc, with significant variance in shape, and a background class. The number of images per category varies from 31 to 800. All the images were converted to gray scale and resized to be no larger than  $300 \times 300$  pixels with preserved aspect ratio. We randomly partitioned the whole dataset into 15 and 30



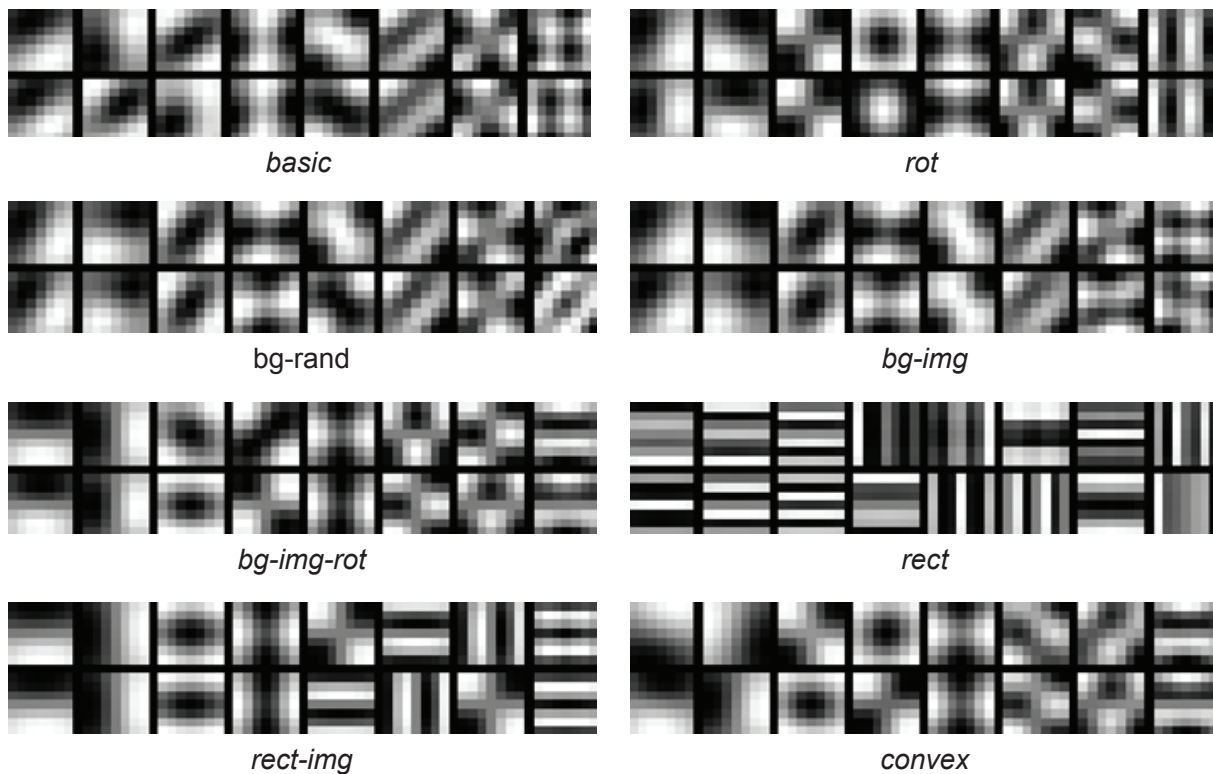


Fig. 10. The PCANet filters learned on various MNIST datasets. For each dataset, the top row shows the filters of the first stage; the bottom row shows the filters of the second stage.

training images per class, and the rest are for testing. The PCANet are trained with filter size  $k_1 = k_2 = 7$ , the number of filters  $L_1 = L_2 = 8$ , and block size equal to a quarter of the image size. Also, we set the overlapping region between blocks to half of the block size. The histogram dimension of each block is reduced from 256 ( $2^{L_2}$ ) to 64 by WPCA. The performance is averaged over 5 drawings of the training set, and for each draw, the accuracy is averaged over the 102 classes (including background class). We do not compare with LDANet because the image sizes in Caltech101 are different, incapable of measuring inter-class and intra-class variabilities in LDA calculation.

The results are shown in Table 12.<sup>9</sup> For the setting with 15 training samples per class, PCANet-2 achieves the best performance. For the setting with 30 training samples, PCANet-2 outperforms most of deep learning techniques, only inferior to hierarchical sparse coding (HSC) [41]. We noticed that the representation learning from raw pixels are in general inferior to those based on SIFT features. Applying PCANet to SIFT features will be considered as our future study.

#### 4 CONCLUSION

In this paper, we have proposed arguably the simplest convolutional deep learning network— PCANet.

<sup>9</sup>. Only the methods which learn representations from gray-scale image pixels are listed.



Fig. 12. The PCANet filters learned on Caltech101 database. Top row: the first stage. Bottom row: the second stage.

TABLE 12  
Comparison of accuracy (%) of the methods based on raw pixels of Caltech101 dataset.

Training Size	15	30
CDBN [7]	57.70±1.50	65.40±0.50
ConvNet [9]	57.60±0.40	66.30±1.50
DeconvNet [51]	58.60±0.70	66.90±1.10
Chen <i>et al.</i> [52]	58.20±1.20	65.80±0.60
Zou <i>et al.</i> [53]	-	66.50
HSC [41]	-	<b>74.00</b>
RandNet-1	46.81 ± 1.69	53.27 ± 1.13
RandNet-2	57.16 ± 1.35	63.68 ± 1.63
PCANet-1	51.10±0.53	56.40±0.98
PCANet-2	<b>61.46±0.76</b>	68.56±1.01

The network processes input images by cascaded PCA, binary hashing, and block histograms. Like the most ConvNet models, the network parameters such as the number of layers, the filter size, and the number of filters have to be given to PCANet. Once the param-

eters are fixed, training PCANet is extremely simple and efficient, for the filter learning in PCANet does not involve regularized parameters and does not require numerical optimization solver. Moreover, building the PCANet comprises only a cascaded linear map, followed by a nonlinear output stage. Such a simplicity offers an alternative and yet refreshing perspective to convolutional deep learning networks, and could further facilitate mathematical analysis and justification of its effectiveness.

A couple of simple extensions of PCANet; that is, RandNet and LDANet, have been introduced and tested together with PCANet on many image classification tasks, including face, hand-written digit, texture, and object. Extensive experimental results have consistently shown that the PCANet outperforms RandNet and LDANet, and is generally on par with and often superior to ScatNet and variations of ConvNet. Furthermore, the performance of PCANet is closely comparable and often better than highly engineered hand-crafted features (such as LBP and LQP). In tasks such as face recognition, PCANet also demonstrates remarkable robustness to corruption and ability to transfer to new datasets. Regardless, extensive experiments given in this paper sufficiently establish one fact: The PCANet is a simple and valuable baseline for studying advanced deep learning architectures for image classification tasks.

Our experiments also convey that as long as the images in databases are somehow well prepared; i.e., images are roughly aligned and do not exhibit diverse scales or poses, PCANet is able to eliminate the image variability and gives reasonably competitive accuracy. In complex image databases such as ImageNet, PCANet might not be sufficient to handle the variability and some preprocessing of pose alignment and scale normalization might be needed for good performance guarantee. To accommodate these issues, a more complicated (say more sophisticated filters) or deeper (more number of stages) PCANet may be inevitable. The current bottleneck that keeps PCANet from growing deeper (e.g., more than two stages) is that the dimension of the resulted feature would increase exponentially with the number of stages, if the number of filters are fixed in all stages. An improved version of PCANet that replaces the 2-dimensional convolution filters with tensor-like filters would be more scalable and will be investigated for future study. We will also leave as future work to augment PCANet with a simple, scalable baseline classifier, readily applicable to much larger scale datasets or problems.

## REFERENCES

- [1] G. Hinton, S. Osindero, and Y.-W. Teh, "A fast learning algorithm for deep belief nets," *Neural Computation*, vol. 18, no. 7, pp. 1527–1554, 2006.
- [2] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: a review and new perspectives," *IEEE TPAMI*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [3] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio, "Maxout networks," in *ICML*, 2013.
- [4] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [5] K. Jarrett, K. Kavukcuoglu, M. Ranzato, and Y. LeCun, "What is the best multi-stage architecture for object recognition," in *ICCV*, 2009.
- [6] J. Bruna and S. Mallat, "Invariant scattering convolution networks," *IEEE TPAMI*, vol. 35, no. 8, pp. 1872–1886, 2013.
- [7] H. Lee, R. Grosse, R. Rananth, and A. Ng, "Convolutional deep belief networks for scalable unsupervised learning of hierarchical representation," in *ICML*, 2009.
- [8] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convolutional neural network," in *NIPS*, 2012.
- [9] K. Kavukcuoglu, P. Sermanet, Y. Boureau, K. Gregor, M. Mathieu, and Y. LeCun, "Learning convolutional feature hierarchies for visual recognition," in *NIPS*, 2010.
- [10] L. Sifre and S. Mallat, "Rotation, scaling and deformation invariant scattering for texture discrimination," in *CVPR*, 2013.
- [11] C. J. C. Burges, J. C. Platt, and S. Jana, "Distortion discriminant analysis for audio fingerprinting," *IEEE TSAP*, vol. 11, no. 3, pp. 165–174, 2003.
- [12] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110.
- [13] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *IEEE CVPR*, 2001.
- [14] L. Fei-Fei and P. Perona, "A Bayesian hierarchical model for learning natural scene categories," in *IEEE CVPR*, 2005.
- [15] C. Liu and H. Wechsler, "Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition," *IEEE TIP*, vol. 11, no. 4, pp. 467–476, 2002.
- [16] H. Yu and J. Yang, "A direct LDA algorithm for high-dimensional data— with application to face recognition," *Pattern Recognition*, vol. 34, no. 10, pp. 2067–2069, 2001.
- [17] R. Gross, I. Matthews, and S. Baker, "Multi-pie," in *IEEE Conference on Automatic Face and Gesture Recognition*, 2008.
- [18] T. Ahonen, A. Hadid, and M. Pietikainen, "Face description with local binary patterns: application to face recognition," *IEEE TPAMI*, vol. 28, no. 12, pp. 2037–2041, 2006.
- [19] Y. Jia, "Caffe: An open source convolutional architecture for fast feature embedding," <http://caffe.berkeleyvision.org/>, 2013.
- [20] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE TPAMI*, vol. 23, no. 6, pp. 643–660, June 2001.
- [21] X. Tan and B. Triggs, "Enhanced local texture feature sets for face recognition under difficult lighting conditions," *IEEE TIP*, vol. 19, no. 6, pp. 1635–1650, 2010.
- [22] A. Martinez and R. Benavente, *CVC Technical Report 24*, 1998.
- [23] W. Deng, J. Hu, and J. Guo, "Extended SRC: Undersampled face recognition via intraclass variant dictionary," *IEEE TPAMI*, vol. 34, no. 9, pp. 1864–1870, Sept. 2012.
- [24] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss, "The FERET database and evaluation procedure for face-recognition algorithms," *Image Vision Comput.*, vol. 16, no. 5, pp. 295–306, 1998.
- [25] J. Lu, Y.-P. Tan, and G. Wang, "Discriminative multi-manifold analysis for face recognition from a single training sample per person," *IEEE TPAMI*, vol. 35, no. 1, pp. 39–51, 2013.
- [26] N.-S. Vu and A. Caplier, "Enhanced patterns of oriented edge magnitudes for face recognition and image matching," *IEEE TIP*, vol. 21, no. 3, pp. 1352–1368, 2012.
- [27] S. Hussain, T. Napoleon, and F. Jurie, "Face recognition using local quantized patterns," in *BMVC*, 2012.
- [28] S. Xie, S. Shan, X. Chen, and J. Chen, "Fusing local patterns of gabor magnitude and phase for face recognition," *IEEE TIP*, vol. 19, no. 5, pp. 1349–1361, 2010.
- [29] N.-S. Vu, "Exploring patterns of gradient orientations and magnitudes for face recognition," *IEEE Trans. Information Forensics and Security*, vol. 8, no. 2, pp. 295–304, 2013.
- [30] Z. Chai, Z. Sun, H. Mndez-Vzquez, R. He, and T. Tan, "Gabor ordinal measures for face recognition," *IEEE Trans. Information Forensics and Security*, vol. 9, no. 1, pp. 14–26, Jan. 2014.
- [31] G. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: a database for studying face recognition in

unconstrained environments," in *Technical Report 07-49, University of Massachusetts, Amherst*, 2007.

- [32] L. Wolf, T. Hassner, and Y. Taigman, "Effective face recognition by combining multiple descriptors and learned background statistics," *IEEE TPAMI*, vol. 33, no. 10, 2011.
- [33] O. Barkan, J. Weill, L. Wolf, and H. Aronowitz, "Fast high dimensional vector multiplication face recognition," in *IEEE ICCV*, 2013.
- [34] Y. Taigman, M. Yang, M. A. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *IEEE CVPR*, 2014.
- [35] H. Fan, Z. Cao, Y. Jiang, and Q. Yin, "Learning deep face representation," in *arXiv: 1403.2802v1*, 2014.
- [36] D. Chen, X. Cao, F. Wen, and J. Sun, "Blessing of dimensionality: high-dimensional feature and its efficient compression for face verification," in *CVPR*, 2013.
- [37] Z. Cui, W. Li, D. Xu, S. Shan, and X. Chen, "Fusing robust face region descriptors via multiple metric learning for face recognition in the wild," in *CVPR*, 2013.
- [38] H. Larochelle, D. Erhan, A. Courville, J. Bergstra, and Y. Bengio, "An empirical evaluation of deep architectures on problems with many factors of variation," in *ICML*, 2007.
- [39] D. Ciresan, U. Meier, and J. Schmidhuber, "Multi-column deep neural networks for image classification," in *CVPR*, 2012.
- [40] K. Sohn, G. Zhou, C. Lee, and H. Lee, "Learning and selecting features jointly with point-wise gated Boltzmann machine," in *ICML*, 2013.
- [41] K. Yu, Y. Lin, and J. Lafferty, "Learning image representations from the pixel level via hierarchical sparse coding," in *CVPR*, 2011.
- [42] S. Belongie, J. Malik, and J. Puzicha, "Shape matching and object recognition using shape contexts," *IEEE TPAMI*, vol. 24, no. 4, pp. 509–522, 2002.
- [43] D. Keysers, T. Deselaers, C. Gollan, and H. Ney, "Deformation models for image recognition," *IEEE TPAMI*, vol. 29, no. 8, pp. 1422–1435, 2007.
- [44] M. D. Zeiler and R. Fergus, "Stochastic pooling for regularization of deep convolutional neural networks," in *ICLR*, 2013.
- [45] S. Rifai, P. Vincent, X. Muller, X. Glorot, and Y. Bengio, "Contractive auto-encoders: explicit invariance during feature extraction," in *ICML*, 2011.
- [46] K. Sohn and H. Lee, "Learning invariant representations with local transformations," in *ICML*, 2012.
- [47] M. Varma and A. Zisserman, "A statistical approach to material classification using image patch exemplars," *IEEE TPAMI*, vol. 31, no. 11, pp. 2032–2047, 2009.
- [48] E. Hayman, B. Caputo, M. Fritz, and J. O. Eklundh, "On the significance of real-world conditions for material classification," in *ECCV*, 2004.
- [49] M. Crosier and L. Griffin, "Using basic image features for texture classification," *IJCV*, pp. 447–460, 2010.
- [50] R. E. Broadhurst, "Statistical estimation of histogram variation for texture classification," in *Proc. Workshop on Texture Analysis and Synthesis*, 2006.
- [51] M. D. Zeiler, D. Krishnan, G. Taylor, and R. Fergus, "Deconvolutional networks," in *CVPR*, 2010.
- [52] B. Chen, G. Polatkin, G. Sapiro, D. B. Dunson, and L. Carin, "The hierarchical beta process for convolutional factor analysis and deep learning," in *ICML*, 2011.
- [53] W. Y. Zou, S. Zhu, A. Y. Ng, and K. Yu, "Deep learning of invariant features via simulated fixations in video," in *NIPS*, 2012.



**Tsung-Han Chan** received the B.S. degree from the Department of Electrical Engineering, Yuan Ze University, Taiwan, in 2004 and the Ph.D. degree from the Institute of Communications Engineering, National Tsing Hua University, Taiwan, in 2009. He is currently working as a Project Lead R&D Engineer with Sunplus Technology Co., Hsinchu, Taiwan. His research interests are in image processing and convex optimization, with a recent emphasis on computer vision and hyperspectral remote sensing.



**Kui Jia** received the B.Eng. degree in marine engineering from Northwestern Polytechnical University, China, in 2001, the M.Eng. degree in electrical and computer engineering from National University of Singapore in 2003, and the Ph.D. degree in computer science from Queen Mary, University of London, London, U.K., in 2007. He is currently a Research Scientist at Advanced Digital Sciences Center. His research interests are in computer vision, machine learning, and image processing.




**Shenghua Gao** received the B.E. degree from the University of Science and Technology of China in 2008, and received the Ph.D. degree from the Nanyang Technological University in 2013. He is currently a postdoctoral fellow in Advanced Digital Sciences Center, Singapore. He was awarded the Microsoft Research Fellowship in 2010. His research interests include computer vision and machine learning.



**Jiwen Lu** is currently a research scientist at the Advanced Digital Sciences Center (ADSC), Singapore. His research interests include computer vision, pattern recognition, machine learning, and biometrics. He has authored/co-authored more than 70 scientific papers in peer-reviewed journals and conferences including some top venues such as the TPAMI, TIP, CVPR and ICCV. He is a member of the IEEE.



**Zinan Zeng** received the master degree and B.E. degree with First Honour from the School of Computer Engineering, Nanyang Technological University, Singapore. He is now a senior software engineer in Advanced Digital Sciences Center, Singapore. His research interests include statistical learning, optimization with application in computer vision.



**Yi Ma** (F'13) is a professor of the School of Information Science and Technology of ShanghaiTech University. He received his Bachelors' degree in Automation and Applied Mathematics from Tsinghua University, China in 1995. He received his M.S. degree in EECS in 1997, M.A. degree in Mathematics in 2000, and his PhD degree in EECS in 2000 all from UC Berkeley. From 2000 to 2011, he was an associate professor of the ECE Department of the University of Illinois at Urbana-Champaign, where he now holds an adjunct position. From 2009 to early 2014, he was a principal researcher and manager of the visual computing group of Microsoft Research Asia. His main research areas are in computer vision and high-dimensional data analysis. Yi Ma was the recipient of the David Marr Best Paper Prize from ICCV 1999 and Honorable Mention for the Longuet-Higgins Best Paper Award from ECCV 2004. He received the CAREER Award from the National Science Foundation in 2004 and the Young Investigator Program Award from the Office of Naval Research in 2005. He has been an associate editor for IJCV, SIIMS, IEEE Trans. PAMI and Information Theory.