

LP-3DCNN: Unveiling Local Phase in 3D Convolutional Neural Networks

Sudhakar Kumawat and Shanmuganathan Raman
Indian Institute of Technology Gandhinagar
Gandhinagar, Gujarat, India

{sudhakar.kumawat, shanmuga}@iitgn.ac.in

Abstract

Traditional 3D Convolutional Neural Networks (CNNs) are computationally expensive, memory intensive, prone to overfit, and most importantly, there is a need to improve their feature learning capabilities. To address these issues, we propose Rectified Local Phase Volume (ReLPV) block, an efficient alternative to the standard 3D convolutional layer. The ReLPV block extracts the phase in a 3D local neighborhood (e.g., $3 \times 3 \times 3$) of each position of the input map to obtain the feature maps. The phase is extracted by computing 3D Short Term Fourier Transform (STFT) at multiple fixed low frequency points in the 3D local neighborhood of each position. These feature maps at different frequency points are then linearly combined after passing them through an activation function. The ReLPV block provides significant parameter savings of at least, 3^3 to 13^3 times compared to the standard 3D convolutional layer with the filter sizes $3 \times 3 \times 3$ to $13 \times 13 \times 13$, respectively. We show that the feature learning capabilities of the ReLPV block are significantly better than the standard 3D convolutional layer. Furthermore, it produces consistently better results across different 3D data representations. We achieve state-of-the-art accuracy on the volumetric ModelNet10 and ModelNet40 datasets while utilizing only 11% parameters of the current state-of-the-art. We also improve the state-of-the-art on the UCF-101 split-1 action recognition dataset by 5.68% (when trained from scratch) while using only 15% of the parameters of the state-of-the-art.

1. Introduction

Over the past few years, research in the area of 2D CNNs has led to unprecedented advances in a number of computer vision tasks such as image classification, semantic segmentation, and image super-resolution. Apart from performance results, 2D CNNs have also made good progress in other complementary areas such as network compression, binarization, quantization, regularization, etc. Unfortunately, unlike their 2D counterparts, 3D CNNs have not enjoyed

the same level of performance jumps on the problems in their domain e.g., video classification and progress in the above mentioned complementary areas. Recent works such as [44] and [11], list down some of the fundamental barriers in modeling and training of deep 3D CNNs such as (1) they are computationally very expensive, (2) they result in large model size, both in terms of memory usage and disk space, (3) they are prone to overfitting, due to a large number of parameters, (4) and there is a need to improve their feature learning capabilities which may require fundamental changes to their network architecture or the standard 3D convolutional layer [44, 25, 39]. Despite the above challenges, the current trend in the literature of deep 3D CNNs is to train computationally expensive, memory intensive, and very deep networks in order to achieve state-of-the-art results [2, 8, 11].

In this work, we take a detour from this trend by proposing an alternative to the fundamental building block of the 3D CNNs, the 3D convolutional layer, which is the primary source of high space-time complexity in 3D CNNs. More precisely, we propose Rectified Local Phase Volume (ReLPV) block, an efficient alternative to the standard 3D convolutional layer in 3D CNNs. The ReLPV block comprises of a local phase module, the ReLU activation function and a set of trainable linear weights. The local phase module extracts the local phase information by computing 3D Short Term Fourier Transform (STFT) [15] (at multiple low frequency points) in a local $n \times n \times n$ (e.g., $3 \times 3 \times 3$) neighborhood/volume of each position of the input feature map. The output of the local phase module is then passed through the ReLU activation function in order to obtain the activated response maps of the local phase information at the fixed low frequency points. Finally, a set of trainable linear weights computes the weighted combinations of these activated response maps. The ReLPV block provides significant parameter savings along with computational and memory savings. The ReLPV block based 3D CNNs have much lower model complexity and are less prone to overfitting. Most importantly, its feature learning capabilities are significantly better than the standard 3D convolutional layer.

Our major contributions in this work are as follows.

- We propose ReLPV block, an efficient alternative to the standard 3D convolutional layer. The ReLPV block significantly reduces the number of trainable parameters, at least 3^3 to 13^3 times compared to the standard 3D convolutional layer with the filter sizes $3 \times 3 \times 3$ to $13 \times 13 \times 13$, respectively.
- We show that the ReLPV block achieves consistently better results on different 3D data representations. We show this on the volumetric ModelNet10 and ModelNet40 datasets by achieving state-of-the-art accuracy using just 11% parameters of the current state-of-the-art. Moreover, we provide results on the spatiotemporal image sequences. In particular, on the UCF-101 split-1 action recognition dataset, improving the current state-of-the-art by 5.68% while using just 15% parameters of the state-of-the-art.
- We present detailed ablation and performance studies of the proposed ReLPV block by varying its various hyperparameters. The analysis will be beneficial for designing ReLPV block based 3D CNNs in future.

2. Related Work

Recently, 2D CNNs have achieved state-of-the-art results in most of the computer vision problems [9]. Moreover, they have also made significant progress in other complementary areas such as network compression [16, 48], binarization [7, 6, 31, 21], quantization [50, 18], regularization [5, 17, 45, 32], etc. Therefore, not surprisingly, there have been many recent attempts to extend this success to the problems in the domain of 3D CNNs e.g., video classification [1], 3D object recognition [26, 2] and MRI volume segmentation [27, 3]. Unfortunately, 3D CNNs are computationally expensive and require large memory and disk space. Furthermore, they overfit very easily owing to the large number of parameters involved. Therefore, there has been recent interest in more efficient variants of 3D CNNs.

Inspired from the progress of network binarization techniques in 2D CNNs such as BinaryConnect [7], BinaryNet [6], and XNORNet [31], Ma *et al.* in [25] introduced BV-CNNs, where they fully binarized some of the state-of-the-art 3D CNN models introduced for recognizing voxelized 3D CAD models from the ModelNet datasets [42]. The binarized version of the 3D CNNs saves significant computation and memory requirements when compared to the floating point baselines. However, this comes at the cost of reduced performance. Furthermore, the binarized network takes binarized inputs only which restricts its application for other 3D data representations such as video classification.

Another way to reduce the model complexity of 3D CNNs is to replace the 3D convolutions with separable con-

volution. This technique has been explored recently in a number of 3D CNN architectures proposed for the task of video classification. The idea of separable convolutions is to first convolve spatially in 2D and then convolve temporally in 1D. This factorization is similar in spirit to the depth-wise separable convolutions used in [43], except that here the idea is to apply it to the temporal dimension instead of the feature dimension. The idea has been used in a variety of recent works, including R(2+1)D networks[39], separable-3D CNNs [44], Pseudo-3D networks [30], and factorized spatio-temporal CNNs [36]. The 3D CNNs based on the idea of separable convolutions achieve competitive results compared to the state-of-the-art on the task of video classification at a reduced space-time complexity.

3. Method

Notation. We denote the feature map output by a layer in a 3D CNN network with the tensor $\mathcal{I} \in \mathbb{R}^{c \times d \times h \times w}$ where h, w, d , and c are the height, width, depth, and number of channels of the feature map, respectively.

The ReLPV Block Architecture. The ReLPV block is a four-layer alternative representation of the standard 3D convolutional layer. Fig. 1 illustrates the architecture of the ReLPV block.

Layer 1. This layer is the standard 3D convolutional layer with a single filter of size $1 \times 1 \times 1$. It takes a feature map of size $c \times d \times h \times w$ as input from the previous layer, and converts it into a single channel feature map of size $1 \times d \times h \times w$. This layer prepares the input for the 3D STFT operation which is computed in *Layer 2*. Let $f(\mathbf{x})$ be the feature map output of *Layer 1* with size $1 \times d \times h \times w$. Here, \mathbf{x} is a variable denoting positions on the feature map $f(\mathbf{x})$.

Layer 2. Local phase has been successfully used in images to detect edges and contours for feature extraction [23]. Phase represents the local coherence of different spatial frequencies. Edges and skeletons in image are expressed by their coherence and play a significant role in image understanding [47]. Same property holds true for 3D data representations too. e.g., videos [29]. There are many methods for extracting local phase in multiple dimensions [14]. Our method is inspired from [29]. *Layer 2* extracts the local phase spectra of $f(\mathbf{x})$ by computing the 3D Short Term Fourier Transform (STFT) in a local $n \times n \times n$ neighborhood $\mathcal{N}_{\mathbf{x}}$ at each position \mathbf{x} of $f(\mathbf{x})$ using Equation 1.

$$F(\mathbf{v}, \mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{N}_{\mathbf{x}}} f(\mathbf{x} - \mathbf{y}) \exp^{-j2\pi\mathbf{v}^T\mathbf{y}} \quad (1)$$

Here, $\mathbf{v} \in \mathbb{R}^3$ is a frequency variable and $j = \sqrt{-1}$. Using vector notation [20], we can rewrite Equation 1 as shown in Equation 2.

$$F(\mathbf{v}, \mathbf{x}) = \mathbf{w}_{\mathbf{v}}^T \mathbf{f}_{\mathbf{x}} \quad (2)$$

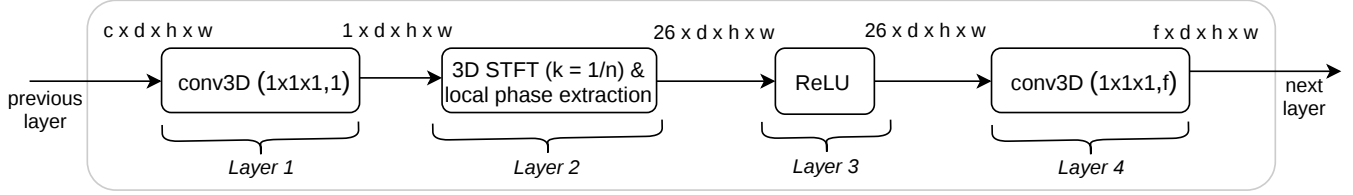


Figure 1: The ReLPV block architecture.

Here, \mathbf{w}_v is the basis vector of the 3D STFT at frequency variable \mathbf{v} and \mathbf{f}_x is a vector containing all the positions from the neighborhood \mathcal{N}_x . Note that, due to the separability of the basis functions, 3D STFT can be computed efficiently for all the positions \mathbf{x} in $f(\mathbf{x})$ by using simple 1D convolutions for each dimension. In this work, we consider 13 lowest non-zero frequency variables which are defined as below. The selected frequency variables are shown as red

$$\begin{array}{lll}
 \mathbf{v}_1 = [k, 0, 0]^T, & \mathbf{v}_2 = [k, 0, k]^T, & \mathbf{v}_3 = [k, 0, -k]^T, \\
 \mathbf{v}_4 = [0, k, 0]^T, & \mathbf{v}_5 = [0, k, k]^T, & \mathbf{v}_6 = [0, k, -k]^T, \\
 \mathbf{v}_7 = [k, k, 0]^T, & \mathbf{v}_8 = [k, k, k]^T, & \mathbf{v}_9 = [k, k, -k]^T, \\
 \mathbf{v}_{10} = [k, -k, 0]^T, & \mathbf{v}_{11} = [k, -k, k]^T, & \mathbf{v}_{12} = [k, -k, -k]^T, \\
 \mathbf{v}_{13} = [0, 0, k]^T, & \text{where } k = 1/n &
 \end{array}$$

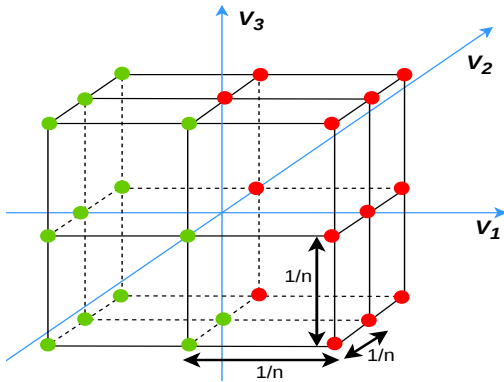


Figure 2: Frequency points used to compute the 3D STFT. The selected frequency points are marked as red dots. The other frequency points in the green dots are ignored, as they are the complex conjugates of the selected ones.

dots in Fig. 2. Low frequency variables are used because they usually contain most of the information, and therefore they have better signal-to-noise ratio than the high frequency components [14]. Let

$$\mathbf{W} = [\Re\{\mathbf{w}_{\mathbf{v}_1}, \mathbf{w}_{\mathbf{v}_2}, \dots, \mathbf{w}_{\mathbf{v}_{13}}\}, \Im\{\mathbf{w}_{\mathbf{v}_1}, \mathbf{w}_{\mathbf{v}_2}, \dots, \mathbf{w}_{\mathbf{v}_{13}}\}]^T \quad (3)$$

Here, \mathbf{W} is a $26 \times n^3$ transformation matrix corresponding to the 13 frequency variables. $\Re\{\cdot\}$ and $\Im\{\cdot\}$ return the real and the imaginary parts of a complex number, respectively. Hence, from Equation 2 & 3, the vector form of 3D STFT for all the 13 frequency points $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{13}$ can be written as shown in Equation 4.

$$\mathbf{F}_x = \mathbf{W}\mathbf{f}_x \quad (4)$$

Since, \mathbf{F}_x is computed for all positions \mathbf{x} of the input $f(\mathbf{x})$, it results in an output feature map with size $26 \times d \times h \times w$. A more detailed mathematical formulation of Layer 2 is provided in the supplementary document.

Layer 3. Applying non-linearity to the local phase information enables the network to learn complex representations. This layer creates activated response maps of the feature maps obtained from Layer 2 by using an activation function. We use the ReLU activation function for better efficiency and faster convergence [28].

Layer 4. This layer is the standard 3D convolutional layer with f filters each of size $1 \times 1 \times 1$ which takes a feature map of size $26 \times d \times h \times w$ as input from Layer 3 and outputs a feature map of size $f \times d \times h \times w$. Note that, Layer 1 and 4 get learned during the training phase of the 3D CNN.

We shall use the notation $\text{ReLPV}(n, f)$ for the ReLPV block, where n and f are its hyperparameters. Here n denotes the size of the local 3D neighborhood from Layer 2 and f is the number of $1 \times 1 \times 1$ filters used in Layer 4.

Importance of using STFT and Local Phase. STFT in multidimensional space was first studied by Hinman *et al.* in [15] as an efficient tool for image encoding. It has two important properties which make it useful for our purpose: (1) Natural images are often composed of objects with sharp edge features. It has been observed that the Fourier phase information accurately represents these edge features. Since STFT in 3D space is simply a windowed Fourier transform, the same property applies [15]. Thus, the local phase has the ability to accurately capture the local features in the same way as done by the convolutional filters. (2) STFT decorrelates the input signal [15]. Regularization is key for deep learning since it allows training of more complex models while keeping lower levels of overfitting and achieves better generalization. Decorrelation of features, representations,

and hidden activations has been an active area of research for better regularization of deep neural nets, with a variety of novel regularizers proposed such as DeCov [5], Decorrelated Batch Normalization (DBN) [17], Structured Decorrelation Constraint (SDC) [45] and OrthoReg [32]. As STFT decorrelates the input representations and due to the reduced number of learnable parameters, the ReLPV block based 3D CNNs are less prone to overfitting and generalize better (for results see Section 5.2).

Forward-Backward Propagation in the ReLPV Block.

The end-to-end training of a 3D CNN network with the ReLPV blocks instead of the standard 3D convolutional layers is straightforward. The steps of forward and backward propagation through the *Layers 1, 3 and 4* of the ReLPV block are standard operations in all deep learning libraries. Back propagation in the *Layer 2* is similar to propagating gradients through layers without learnable parameters (e.g. Add, Multiply etc.) as it involves applying the fixed basis matrix \mathbf{W} to the input. Note that, during training, only the $1 \times 1 \times 1$ filters in *Layers 1 and 4* are updated while the weights in the matrix \mathbf{W} remain unaffected.

Parameter analysis of the ReLPV Block. The ReLPV block uses significantly less trainable parameters when compared to the standard 3D convolutional layer with the same filter size/volume and number of input-output channels. Consider a standard 3D convolutional layer with c input and f output channels. Let $n \times n \times n$ be the size/volume of the filters. Thus, the total number of trainable parameters in a standard 3D convolutional layer is $c \cdot n^3 \cdot f$. An ReLPV block with c input channels and f output channels consists of just $c \cdot 1 + f \cdot 26$ trainable parameters. Thus, the ratio of the number of trainable parameters in a standard 3D convolutional layer and the proposed ReLPV block is calculated as below.

$$\frac{\# \text{ params. in 3D conv. layer}}{\# \text{ params. in ReLPV block}} = \frac{c \cdot n^3 \cdot f}{c \cdot 1 + f \cdot 26} \quad (5)$$

For simplicity, let us assume $f = c$, i.e., the number of input and output channels are same. Furthermore, in practice, in most deep 3D CNNs $f \geq 27$. Therefore, let $f = 27$. This reduces the above ratio to n^3 . Thus, for a filter of size $3 \times 3 \times 3$ in the standard 3D convolutional layer, the ReLPV block uses 27 times less trainable parameters. Therefore, numerically, ReLPV block saves atleast $27\times$, $125\times$, $343\times$, $729\times$, $1331\times$, and $2197\times$ parameters during learning for $3 \times 3 \times 3$, $5 \times 5 \times 5$, $7 \times 7 \times 7$, $9 \times 9 \times 9$, $11 \times 11 \times 11$, and $13 \times 13 \times 13$ 3D convolutional filters, respectively.

4. Experiments

In this section, we show that the proposed ReLPV block produces consistently better results on different 3D data representations compared to the standard 3D convolutional

layer. We demonstrate this on voxelized 3D CAD models and on spatiotemporal image sequences.

4.1. Experiments and Results on 3D CAD models

Datasets. ModelNet [42] is a large 3D repository of clean CAD models (shapes). The ModelNet10 with 4,899 shapes (train: 3991, test: 908) and ModelNet40 with 12,311 shapes (train: 9843, test: 2468) are commonly used as benchmarking datasets and consist of 10 and 40 categories, respectively. Each model is aligned to a canonical frame and then rotated at 12 and 24 evenly-sampled orientations about the z-axis ($Az \times 12$ and $Az \times 24$ augmentation). These rotated models are then voxelized to a $32 \times 32 \times 32$ grid. We use the voxelized versions of [26]. The task here is to classify a given voxelized 3D model into its corresponding class.

4.1.1 ModelNet: Comparison with the baselines

Baselines. We start our experiments by replacing the standard 3D convolutional layer with the proposed ReLPV block (with skip connections) in the baseline networks VoxNet [26], VoxNetPlus [25] and LightNet [49], and call these new networks as LP-VoxNet, LP-VoxNetPlus, and LP-LightNet, respectively. Here LP stands for Local Phase. The standard 3D convolutional layer is replaced with the ReLPV block in a straightforward manner. For example the VoxNet network [26] has the following architecture: $\text{conv3D}(5, 32, 2) - \text{conv3D}(3, 32, 1) - \text{MP}(2) - \text{FC}(128) - \text{FC}(K)$. Here, $\text{conv3D}(n, f, s)$ is the standard 3D convolutional layer with f filters each of size $n \times n \times n$ applied with stride s . MP denotes Max Pooling. FC stands for fully connected layer. K is the number of classes. The equivalent local phase version of VoxNet is: $\text{ReLPV}(5, 32, 2) - \text{ReLPV}(3, 32, 1) - \text{MP}(2) - \text{FC}(128) - \text{FC}(K)$. In our earlier discussion on the architecture of the ReLPV block, we focused only on the important hyperparameters and did not discuss other hyperparameters that are commonly used in the standard 3D conv layer, such as the stride information. Such information can easily be incorporated in the ReLPV architecture. Similar procedure is followed while preparing LP-VoxNetPlus and LP-LightNet networks.

Training. We train these new networks using SGD as optimizer with momentum 0.9 and categorical crossentropy as loss. During training, we start with a learning rate of 0.008 and reduce it by a factor of 2 if the validation loss plateaus. For LP-VoxNet and LP-VoxNetPlus networks, following [26, 25], we first train them on ModelNet40 and then fine-tune on ModelNet10. The opposite is done on LP-LightNet network as done in [49]. Following [26, 25, 49], all networks were trained on 12 evenly-sampled rotations of each instance about the z-axis ($Az \times 12$ augmentation). No data augmentation was done on the test data.

Results. Table 1 presents the comparison of the new net-

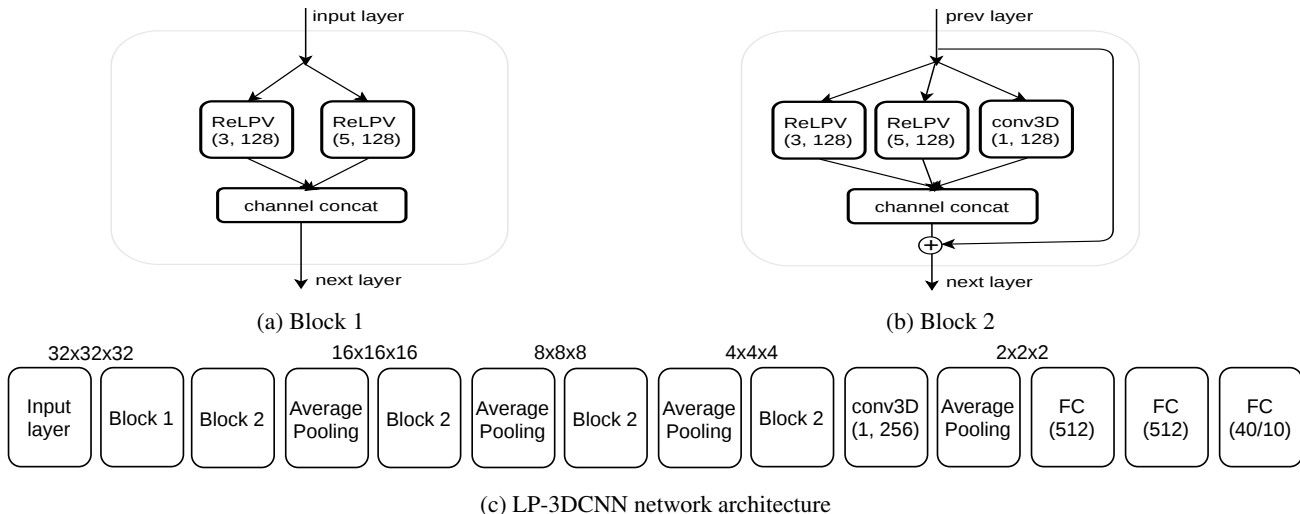


Figure 3: **Experiments and comparison with the state-of-the-art.** LP-3DCNN network building blocks and architecture.

Network	ModelNet40 (%)	ModelNet10 (%)
VoxNet [26] (baseline)	83	92
Binary VoxNet [25]	81.63	90.69
LP-VoxNet (ours)	86.26	92.24
VoxNetPlus [25] (baseline)	83.91	93.36
Binary VoxNetPlus [25]	85.47	92.32
LP-VoxNetPlus (ours)	88.1	93.4
LightNet [49] (baseline)	86.90	93.39
Binary LightNet [25]	84.24	92.36
LP-LightNet (ours)	87.5	92.95

Table 1: **Comparison of the baseline networks with the Local Phase and binarized versions.** The Local Phase version outperforms the baseline and their binarized versions.

works with their corresponding baselines. We also compare the new networks with the binarized version of the baselines [25] (as described in section 2). The local phase version clearly outperforms the corresponding baselines and their binarized versions on both the ModelNet10 and the ModelNet40 datasets.

4.1.2 ModelNet: Comparison with the state-of-the-art

Network Architecture. We follow ideas from the Voxception-ResNet (VRN) architecture of [2] which adopts a simple inception-style architecture with ResNet-style skip connections. The intuition behind this design is to have a maximum number of possible pathways for information to flow through the network. For the first non-downsampling block that follows the input layer (Fig. 3a), we concatenate an equal number (128) of feature maps from two ReLPV blocks with different local phase volume sizes ($3 \times 3 \times 3$ and $5 \times 5 \times 5$). For other non-downsampling blocks, we augment the above structure with an additional $1 \times 1 \times 1$ convolutional layer that outputs the same number (128) of feature maps as

the ReLPV blocks and concatenate it with the other feature maps as shown in Fig. 3b. This architecture allows the network to choose between taking a weighted average of the feature maps in the previous layer (i.e. by heavily weighting the $1 \times 1 \times 1$ convolutions) or focusing on local phase information (i.e., by heavily weighting the ReLPV blocks). Along with this, skip connections are added as shown in Fig. 3b for smoother flow of the gradients to the previous layers. For downsampling, we use average pooling with pool size 2 and stride 2. Our final model is shown in Fig. 3c with five non-downsampling blocks, followed by two fully connected layers each of size 512, and a final softmax layer for classification. All non-downsampling layers (after batch normalization) and fully connected layers are followed by the ReLU activation function. The layer conv3D(1, 256) is used after the final non-downsampling layer to reduce the number of parameters in the fully connected layers.

Training and Testing. The input to our network are voxels of size $32 \times 32 \times 32$ from the ModelNet datasets. Following [2], we change the binary voxel range from $\{0, 1\}$ to $\{-1, 5\}$ to encourage the network to pay more attention to positive entries. Network is trained using SGD as optimizer with momentum 0.9 and categorical crossentropy as loss. During training, we start with a learning rate of 0.008 and reduce it by a factor of 5 if the validation loss plateaus. All weights are initialized using orthogonal initialization. The network is first trained on the $Az \times 12$ augmented data, then it is fine-tuned on the $Az \times 24$ augmented data at low learning rate. No data augmentation was done on the test data. Apart from rotations, the data is augmented by adding noise, random translations and horizontal flips to each training example, as done in [26, 2].

Results. Table 2 compares our results with other methods that use voxelized/volumetric ModelNet datasets as input.

Network	Framework	Augmentation	Parameters (Millions)	ModelNet40 (%)	ModelNet10 (%)
3D ShapeNets [42]	Single, Volumetric	Az×12	≈ 38	77	83.5
Beam Search [46]	Single, Volumetric	-	≈ 0.08	81.26	88
3D-GAN [41]	Single, Volumetric	-	≈ 11	83.3	91
VoxNet [26]	Single, Volumetric	Az×12	≈ 0.92	83	92
LightNet [49]	Single, Volumetric	Az×12	≈ 0.30	86.90	93.39
ORION [33]	Single, Volumetric	Az×12	≈ 0.91	-	93.8
VRN [2]	Single, Volumetric	Az×24	≈ 18	91.33	93.61
LP-3DCNN (ours)	Single, Volumetric	Az×12	≈ 2	89.4	93.76
LP-3DCNN (ours)	Single, Volumetric	Az×24	≈ 2	92.1	94.4
FusionNet [13]	Ensemble, Vol.+ Mul.	(Az, El)×60	≈ 118	90.8	93.11
VRN Ensemble [2]	Ensemble, Volumetric	Az×24	≈ 108	95.54	97.14

Table 2: **Performance results on the ModelNet datasets.** Az stands for azimuth rotation and El stands for elevation rotation. “-” means that information is not provided for the item in the paper. Vol. stands for volumetric, Mul. stands for multi-view.

In order to make a fair comparison, we only consider volumetric network frameworks in this work. We do not include multi-view networks or point cloud-based networks. In single network framework, our proposed network outperforms all the previous networks on both the ModelNet10 and the ModelNet40 datasets. Furthermore, it uses just 2 million parameters compared to the current state-of-the-art, the VRN network, that uses 18 million parameters. In the ensemble framework, the VRN achieves the best performance on both the ModelNet10 and ModelNet40 datasets. However, it has the most complex network architecture with up to 45 layers and 108 million parameters, taking almost 6 days to train. In ensemble framework, our network outperforms FusionNet [13] while using almost 59 times less parameters and significantly less data augmentation.

4.2. Experiments and Results on Spatiotemporal Image Sequences

Dataset. We use the UCF-101 split-1 action recognition dataset [34]. The dataset has been used as a benchmark dataset in [37, 38, 8] for the performance studies and for searching 3D CNN network architectures and hyperparameters for action recognition tasks.

Baseline. We use the experimental 3D CNN network proposed by [37] for action recognition as baseline which is a smaller version of the C3D network [37]. For simplicity, we call this network as mini C3D network or mC3D. The mC3D network with filter size $n \times n \times n$ denoted as mC3D $_n$ has the following architecture: conv3D($n, 64$) – MP(2) – conv3D($n, 128$) – MP(2) – conv3D($n, 256$) – MP(2) – conv3D($n, 256$) – MP(2) – conv3D($n, 256$) – MP(2) – FC(2048) – FC(2048) – FC(101). Each 3D convolutional and fully connected layer is followed by a ReLU activation function. All the convolution layers are applied with appropriate padding and stride 1 such that there is no change in size of the tensor from the input to the output of these layers. Following [37], the input to the network are videos of dimension $3 \times 16 \times 112 \times 112$.

The equivalent local phase version of the above network,

denoted as LP-mC3D $_n$, is prepared by replacing the standard 3D convolutional layers with the ReLPV blocks as done in Section 4.1.1. Here, n denotes the size of the local 3D neighborhood in which STFT is computed.

Training. Following [37], we use SGD as optimizer with Nesterov momentum with value 0.9 and categorical crossentropy as loss. We train the networks for 16 epochs starting with a learning rate of 0.003 and decreasing it by a factor of 10 after every 4 epochs. Note that all the networks are trained from scratch. No data augmentation such as frame translation, rotation, or scaling is used. We re-trained all the baseline networks (for $n = 3, 5, 7$). The results were found to be consistent with Fig.2 in [37].

Results. Early works such as [37, 22] showed that training relatively shallow 3D CNNs from scratch on the UCF-101 split-1 dataset achieve performance between 41 – 44%. Recent works such as [38, 8] use deep 3D Residual ConvNet architectures to achieve better results. Table 3 reports our results on the UCF-101 split-1 dataset. We improve the state-of-the-art by 5.68% while using just five ReLPV

Network	Parameters (Millions)	Model Size (Mb)	FLOP (Millions)	Acc. (%)
2D-ResNet 18 [12, 38]	≈ 11.2	-	-	42.2
2D-ResNet 34 [12, 38]	≈ 21.5	-	-	42.2
3D-ResNet 18 [38]	≈ 33.2	254	-	45.6
3D-ResNet 34 [38]	≈ 63.5	485	-	45.9
3D-ResNet 101 [8]	≈ 86.06	657	-	46.7
3D STC-ResNet 101 [8]	-	-	-	47.9
mC3D $_3$ [37] (baseline)	≈ 18	139.6	34.88	44
LP-mC3D $_3$ (ours)	≈ 13	106.2	26.072	53.58
mC3D $_5$ [37] (baseline)	≈ 34.32	274.6	68.64	42.5
LP-mC3D $_5$ (ours)	≈ 13	106.2	26.077	51.44
mC3D $_7$ [37] (baseline)	≈ 71.88	575	143.72	42.3
LP-mC3D $_7$ (ours)	≈ 13	106.2	26.08	50.54
mC3D $_9$ (baseline)	≈ 138.34	1100	276.68	36.17
LP-mC3D $_9$ (ours)	≈ 13	106.2	26.083	48.99

Table 3: **Performance results on the UCF-101 split-1 action recognition dataset.** Comparison of the ReLPV block based 3D CNNs with their corresponding baselines and other state-of-the-art networks. All the networks are trained from scratch.

blocks. Our network uses 13 million parameters compared to the 3D STC-ResNet 101 network [8], which is built on the top of 3D ResNet 101 network and uses more than 86 million parameters. Furthermore, all the local phase versions with different local phase volumes significantly outperform the corresponding baseline networks.

5. Discussion and Analysis

In this section, we present detailed ablation and performance studies of the ReLPV block. Furthermore, we discuss some statistical advantages afforded by the ReLPV block over the standard 3D convolutional layer.

5.1. Space-time Complexity of the ReLPV block

Model size. Table 3 shows that the ReLPV block based 3D CNNs use less parameters and occupy less disk space when compared to the corresponding baselines. Furthermore, with an increase in the local phase volume (while keeping other hyperparameters constant) from 3 to 9, there is no change in the number of trainable parameters or model size in the ReLPV block based networks. In contrast, there is a significant rise in the number of parameters and model size in baseline networks with an increase in filter size. We believe this feature of the ReLPV block can be of huge benefit for 3D CNNs in resource constraint environment.

Computational cost. We discussed in Section 3 that due to the separability of the basis functions, STFT can be computed efficiently by using simple 1D convolutions for each dimension. This technique of computing 3D STFT using separable convolutions saves huge computational costs and has been of recent interest in 3D CNNs as discussed in Section 2. Table 3 reports the computation cost in terms of the number of Floating Point Operations (FLOP) of the models. The FLOP values of the ReLPV block based 3D CNN are less when compared to the corresponding baselines. Furthermore, they vary very little with an increase in the local phase volume. However, for the baseline networks, the FLOP values increase by almost 8 times with an increase of filter size from 3 to 9.

5.2. Statistical advantages of the ReLPV block

As discussed earlier, one of the major challenges in training deep 3D CNNs is to avoid overfitting [37, 38, 11]. A recent study by Hara *et al.* in [11] shows that even a relatively shallow 3D CNN such as 3D ResNet-18 tends to overfit significantly on action recognition datasets such as UCF-101 [34] and HMDB-51 [24]. This is partly due to the large number of trainable parameters in 3D CNNs in comparison to their 2D counterparts and partly due to the unavailability of large scale 3D datasets [38, 11]. These pose a major bottleneck in training deep 3D CNNs. In order to curb overfitting, various training methods such as data augmentation, training shallow networks, and novel regularizers such

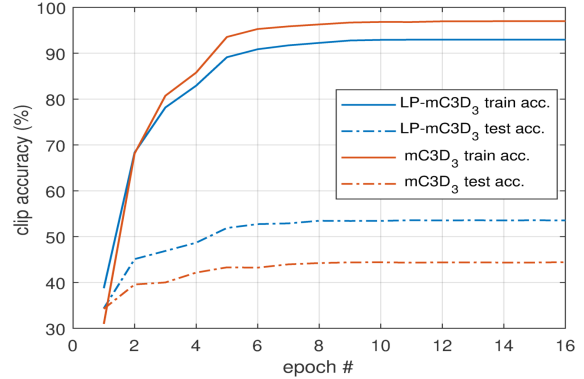


Figure 4: **Overfitting results on the UCF-101 split-1 dataset.** The LP-mC3D₃ network overfits less and generalizes significantly better compared to the baseline mC3D₃ network.

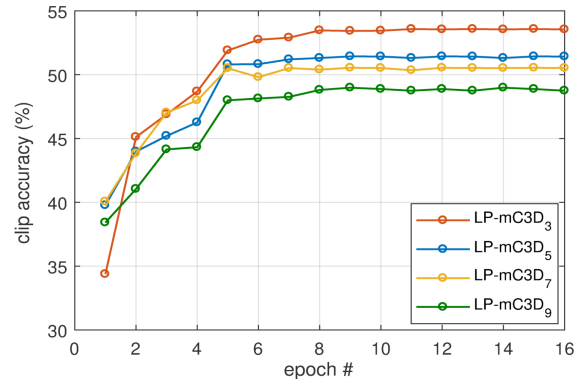


Figure 5: **ReLPV block STFT volume search.** LP-mC3D₃ network with STFT volume of $3 \times 3 \times 3$ performs the best.

as Dropout [35], DropConnect [40], and Maxout [10] have been introduced. While regularizers such as [40, 35, 10] have been proposed to regularize the fully connected layers of the network, recent works such as [4, 19, 35] show that regularizing the convolutional layers of the network is equally important. Our ReLPV block when used in the place of the standard 3D convolutional layer in deep 3D CNNs, naturally regularizes the network due to its use of significantly less trainable parameters and due to the decorrelation property of STFT (see Section 3). Fig. 4 reports our result on the overfitting experiment. The LP-mC3D₃ network clearly overfits less and generalizes significantly better when compared to the baseline mC3D₃ network.

5.3. Exploring the Local Phase Volume of the ReLPV block

As described earlier, the ReLPV block takes two hyperparameters as input, one of which is the size of the local volume in which the STFT is computed (and the local phase is extracted) for each position of the input feature map. In

this section, we explore this hyperparameter. We experiment with different sizes of local volumes, in particular from $3 \times 3 \times 3$ to $9 \times 9 \times 9$. We found that the performance of the ReLPV block decreases with an increase in the STFT volume. Fig. 5 presents the clip accuracy of the LP-mC3D_n network on the UCF-101 test split-1 dataset for various STFT volumes ranging from $3 \times 3 \times 3$ to $9 \times 9 \times 9$ over 16 epochs. The LP-mC3D₃ network with STFT volume of $3 \times 3 \times 3$ performs the best while the LP-mC3D₉ network performs the worst. Note that, an analogous study was carried out in [37] for the standard 3D convolutional layer where it was found that the 3D CNNs with $3 \times 3 \times 3$ convolutional kernels in all the layers perform the best.

5.4. Exploring the number of feature maps output by the ReLPV block

In this section, we explore another hyperparameter, the number of feature maps output by the ReLPV block. In simple words, we explore the effect of varying the number of $1 \times 1 \times 1$ filters in the *Layer 4* of the ReLPV block (see Section 3). For this, we use a modified version of the LP-mC3D_n network and experiment with different pairs of ReLPV block hyperparameters (n, f). Let LP-mC3D_{n,f} be our experimental network with the following architecture: Input layer – ReLPV(n, f) – MP(2) – ReLPV(n, f) – MP(2) – ReLPV(n, f) – MP(2) – ReLPV(n, f) – MP(2) – ReLPV(n, f) – conv3D($1 \times 1 \times 1, 256$) – MP(2) – FC(2048) – FC(2048) – FC(101). The layer Conv3D($1 \times 1 \times 1, 256$) is used after the last ReLPV block so that the number of parameters in the fully-connected layers does not vary across different networks. Table 4 presents our results of the experiment on the UCF-101 split-1 test set. We observe that, for a fixed value of the local STFT volume (the hyperparameter n), performance improves with an increase in the number of $1 \times 1 \times 1$ filters (the hyperparameter f). Another important observation is that the model size and the number of trainable parameters vary by a very small amount with an increase in the value of the hyperparameter f .

5.5. ReLPV Block based Hybrid 3D CNN Models

In this section, we explore the performance effects of using ReLPV blocks and the standard 3D convolutional layers in a single 3D CNN network. We call such networks as hybrid 3D CNNs. We experiment with two types of variations. In the first variation, we replace the top few layers (following the input layer) of a traditional 3D CNN network (baseline mC3D₃) with the ReLPV blocks such that the feature maps learned by the ReLPV blocks are input to the later standard 3D convolutional layers. In the second variation, the bottom layers are replaced with the ReLPV blocks such that the feature maps learned by the standard 3D convolutional layers are input to the later ReLPV blocks. We use the notation mC3D₃(B_l/T_l) to denote that l bot-

Network	Parameters (Millions)	Model Size (Mb)	Acc.
LP-mC3D _{3,64}	≈ 12.84	104.2	50.96
LP-mC3D _{3,128}	≈ 12.93	104.9	51.84
LP-mC3D _{3,256}	≈ 13.20	107.1	53.50
LP-mC3D _{5,64}	≈ 12.84	104.2	50.29
LP-mC3D _{5,128}	≈ 12.93	104.9	51.10
LP-mC3D _{5,256}	≈ 13.20	107.1	53.22
LP-mC3D _{7,64}	≈ 12.84	104.2	47.66
LP-mC3D _{7,128}	≈ 12.93	104.9	50.10
LP-mC3D _{7,256}	≈ 13.20	107.1	51.14

Table 4: **Exploring the number of feature maps output by the ReLPV block.** Performance improves with increase in value of f .

Network	Parameters (Millions)	Model Size (Mb)	Acc.
mC3D ₃ (T_1)	≈ 17.44	139.9	51.51
mC3D ₃ (T_2)	≈ 16.13	138.5	47.67
mC3D ₃ (T_3)	≈ 13.20	132.1	43.95
mC3D ₃ (B_1)	≈ 15.82	126.9	35.1
mC3D ₃ (B_2)	≈ 14.13	113.7	36.47
mC3D ₃ (B_3)	≈ 13.30	107.3	40.84

Table 5: **Results on hybrid 3D CNN architectures.** Performance results on the UCF-101 split-1 test set.

tom/top successive 3D conv layers of mC3D₃ are replaced with the ReLPV block. Table 5 reports our results of the experiments. We observe that replacing standard 3D convolutional layers with the ReLPV blocks at the top of a traditional 3D CNN network improves its performance while the opposite happens when ReLPV blocks are added in the bottom layers. However, the hybrid 3D CNNs do not outperform the LP-mC3D₃ network where all layers are replaced with the ReLPV block (Table 3).

6. Conclusion

In this work, we have proposed ReLPV block, an efficient alternative to the standard 3D convolutional layer, in order to reduce the high space-time and model complexity of the traditional 3D CNNs. The ReLPV block when used in place of the standard 3D convolutional layer in traditional 3D CNNs, significantly improves the performance of the baseline architectures. Furthermore, they produce consistently better results across different 3D data representations. Our proposed ReLPV block based 3D CNN architectures achieve state-of-the-art results on the ModelNet and UCF-101 split-1 action recognition datasets. We plan to apply ReLPV block in 3D CNN architectures for other 3D data representations and tasks such as 3D MRI segmentation.

Acknowledgments. The authors gratefully acknowledge the travel grant support from Google Research India. Sudhakar Kumawat was supported by TCS Research Fellowship. Shanmuganathan Raman was supported by SERB Core Research Grant and Imprint 2 Grant. We thank Manisha Verma for early contribution and technical discussion.

References

- [1] Maryam Asadi-Aghbolaghi, Albert Clapes, Marco Bellantoni, Hugo Jair Escalante, Víctor Ponce-López, Xavier Baró, Isabelle Guyon, Shohreh Kasaei, and Sergio Escalera. A survey on deep learning based approaches for action and gesture recognition in image sequences. In *Automatic Face & Gesture Recognition*, pages 476–483. IEEE, 2017.
- [2] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016.
- [3] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, pages 424–432, 2016.
- [4] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *ICLR*, 2016.
- [5] Michael Cogswell, Faruk Ahmed, Ross Girshick, Larry Zitnick, and Dhruv Batra. Reducing overfitting in deep networks by decorrelating representations. In *ICLR*, 2016.
- [6] M Courbariaux and Y Bengio. Binarynet: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv: 1602.02830*, 2017.
- [7] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *NIPS*, pages 3123–3131, 2015.
- [8] Ali Diba, Mohsen Fayyaz, Vivek Sharma, M Mahdi Arzani, Rahman Yousefzadeh, Juergen Gall, and Luc Van Gool. Spatio-temporal channel correlation networks for action classification. In *ECCV*, 2018.
- [9] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.
- [10] Ian J Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron Courville, and Yoshua Bengio. Maxout networks. In *ICML*, pages III–1319, 2013.
- [11] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet. In *CVPR*, pages 18–22, 2018.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [13] Vishakh Hegde and Reza Zadeh. Fusionnet: 3d object classification using multiple data representations. *arXiv preprint arXiv:1607.05695*, 2016.
- [14] Janne Heikkilä and Ville Ojansivu. Methods for local phase quantization in blur-insensitive image analysis. In *LNLA*, pages 104–111, 2009.
- [15] B Hinman, Jared Bernstein, and D Staelin. Short-space fourier transform image processing. In *ICASSP*, volume 9, pages 166–169, 1984.
- [16] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [17] Lei Huang, Dawei Yang, Bo Lang, and Jia Deng. Decorrelated batch normalization. In *CVPR*, 2018.
- [18] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- [19] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456, 2015.
- [20] Anil K Jain. *Fundamentals of digital image processing*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- [21] Felix Juefei-Xu, Vishnu Naresh Boddeti, and Marios Savvides. Local binary convolutional neural networks. In *CVPR*, volume 1, 2017.
- [22] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, pages 1725–1732, 2014.
- [23] P Kovesi. Image features from phase congruency. *Videre: A Journal of Computer Vision Research*, 1(3):1–26, 1999.
- [24] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, pages 2556–2563, 2011.
- [25] Chao Ma, Wei An, Yinjie Lei, and Yulan Guo. BV-CNNs: Binary volumetric convolutional networks for 3d object recognition. In *BMVC*, 2017.
- [26] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *IROS*, 2015.
- [27] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, pages 565–571, 2016.
- [28] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, pages 807–814, 2010.
- [29] Juhani Päivärinta, Esa Rahtu, and Janne Heikkilä. Volume local phase quantization for blur-insensitive dynamic texture classification. In *Scandinavian Conference on Image Analysis*, pages 360–369, 2011.
- [30] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *ICCV*, pages 5534–5542, 2017.
- [31] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, pages 525–542, 2016.
- [32] Pau Rodríguez, Jordi Gonzalez, Guillem Cucurull, Josep M Gonfaus, and Xavier Roca. Regularizing cnns with locally constrained decorrelations. In *ICLR*, 2017.
- [33] N Sedaghat, M Zolfaghari, E Amiri, and T Brox. Orientation-boosted voxel nets for 3d object recognition. In *BMVC*, 2017.

- [34] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [35] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [36] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *ICCV*, pages 4597–4605, 2015.
- [37] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [38] Du Tran, Jamie Ray, Zheng Shou, Shih-Fu Chang, and Manohar Paluri. Convnet architecture search for spatiotemporal feature learning. *arXiv preprint arXiv:1708.05038*, 2017.
- [39] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018.
- [40] Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *ICML*, pages 1058–1066, 2013.
- [41] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *NIPS*, pages 82–90, 2016.
- [42] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, pages 1912–1920, 2015.
- [43] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, pages 5987–5995, 2017.
- [44] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, pages 305–321, 2018.
- [45] Wei Xiong, Bo Du, Lefei Zhang, Ruimin Hu, and Dacheng Tao. Regularizing deep convolutional neural networks with a structured decorrelation constraint. In *ICDM*, pages 519–528, 2016.
- [46] Xu Xu and Sinisa Todorovic. Beam search for learning a deep convolutional neural network of 3d shapes. In *ICPR*, pages 3506–3511, 2016.
- [47] Ido Zachevsky and Yehoshua Y Zeevi. Modelling local phase of images and textures with applications in phase denoising and phase retrieval. *arXiv preprint arXiv:1810.00403*, 2018.
- [48] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018.
- [49] Shuaifeng Zhi, Yongxiang Liu, Xiang Li, and Yulan Guo. Lightnet: A lightweight 3d convolutional neural network for real-time 3d object recognition. In *Eurographics Workshop on 3D Object Retrieval*, 2017.
- [50] Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. Incremental network quantization: Towards lossless cnns with low-precision weights. *ICLR*, 2017.