Dynamic Convolutions: Exploiting Spatial Sparsity for Faster Inference

Thomas Verelst Tinne Tuytelaars
 ESAT-PSI, KU Leuven

{thomas.verelst, tinne.tuytelaars}@esat.kuleuven.be

Abstract

Modern convolutional neural networks apply the same operations on every pixel in an image. However, not all image regions are equally important. To address this inefficiency, we propose a method to dynamically apply convolutions conditioned on the input image. We introduce a residual block where a small gating branch learns which spatial positions should be evaluated. These discrete gating decisions are trained end-to-end using the Gumbel-Softmax trick, in combination with a sparsity criterion. Our experiments on CIFAR, ImageNet, Food-101 and MPII show that our method has better focus on the region of interest and better accuracy than existing methods, at a lower computational complexity. Moreover, we provide an efficient CUDA implementation of our dynamic convolutions using a gather-scatter approach, achieving a significant improvement in inference speed on MobileNetV2 and ShuffleNetV2. On human pose estimation, a task that is inherently spatially sparse, the processing speed is increased by 60% with no loss in accuracy.

1. Introduction

Most research on deep neural networks focuses on improving accuracy without taking into account the model complexity. As the community moves to more difficult problems – e.g. from classification to detection or pose estimation – architectures tend to grow in capacity and computational complexity. Nevertheless, for real-time applications running on consumer devices such as mobile phones, notebooks or surveillance cameras, what matters most is a good trade-off between performance (i.e., frames processed per second) and accuracy [9, 24]. Attempts to improve this trade-off have focused mostly on designing more efficient architectures [21, 39, 44, 53] or compressing existing ones [17, 19, 28, 33, 38, 50].

Interestingly, most neural networks, including the more efficient or compressed ones mentioned above, execute the same calculations for each image, independent of its content. This seems suboptimal: only the complex images

Figure 1: In each residual block, a small gating network (mask unit) predicts pixel-wise masks determining the locations where dynamic convolutions are evaluated.

require such deep and wide networks. Therefore, the domain of conditional execution gained momentum [4, 5, 6]. Compared to static compression methods, the architecture of the network is adapted based on the input image. For instance, the network depth can vary per image since easy and clear images require fewer convolutional layers than ambiguous ones [46, 48, 51]. The neural network chooses which operations to execute. Such practice is often called *gating* [11, 22], and can be applied at the level of convolutional layers [46, 48, 51], channels [3, 13, 30] or other elements in the network.

In this work, we focus on reducing computations by *executing conditionally in the spatial domain*. Typical convolutional networks apply all convolutional filters on every location in the image. In many images, the subject we want to classify or detect is surrounded by background pixels, where the necessary features can be extracted using only few operations. For example, flat regions such as a blue sky can easily be identified. We call such images spatially sparse.

We propose a method, trained end-to-end without explicit spatial supervision, to execute convolutional filters on important image locations only. For each residual block, a small gating network chooses the locations to apply dynamic convolutions on (Fig. 1). Gating decisions are trained end-to-end using the Gumbel-Softmax trick [23, 31]. Those decisions progress throughout the network: the first stages extract features from complex regions in the image, while the last layers use higher-level information to focus on the

region of interest only. Note that the input of a dynamic convolution is a dense matrix, making it fundamentally different from methods operating on sparse data [14, 15].

Many works on conditional execution only report a reduction in the theoretical complexity [12, 46]. When implemented naively, merely applying masks does not save computations. For certain methods, it is actually not clear whether they could be implemented efficiently at all. For instance, Wu *et al.* [51] report an *increase* in execution time when conditionally executing individual layers using a separate policy network. Parallel execution on GPU or FPGA relies on the regularity of standard convolutions to pipeline operations [26, 42] and adding element-wise conditional statements might strongly slow down inference. Moreover, Ma *et al.* [53] show that the number of floating point operations (FLOPS) is not a sufficient metric to estimate inference speed: networks with a similar amount of operations can have largely different execution speeds. Simple element-wise operations such as activation functions, summations and pooling can have a significant impact, but are not included in many FLOPS-estimations.

In contrast, we demonstrate an actual improvement of wall-clock time with our CUDA implementation of dynamic convolutions. Our method is designed with practical feasibility in mind and requires minimal changes to existing libraries: efficient spatially sparse execution is made possible by rearranging tensor elements in memory, which has similar overhead as a simple copy operation. Our code is available at <https://github.com/thomasverelst/dynconv>.

The main contributions of our paper are threefold:

1. We present an approach to train pixel-wise gating masks end-to-end using the Gumbel-Softmax trick, with a focus on efficiency.
2. Our method achieves state-of-the-art results on classification tasks with ResNet [18] and MobileNetV2 [39], and we show strong results on human pose estimation, improving the performance-accuracy trade-off over non-conditional networks.
3. We provide a CUDA implementation of residual blocks with dynamic convolutions on GPU, not just reducing the theoretical amount of floating-point operations but also offering practical speedup with MobileNetV2 and ShuffleNetV2.

2. Related work

Static compression methods have been extensively studied to reduce storage and computational cost of existing neural networks, *e.g.* pruning [27, 28, 33], knowledge distillation [19, 37], structured matrices [41, 49] or quantization [17, 50]. Recent methods vary computations based on the input image. So-called conditional execution can be ap-

plied on several aspects of a network: we make a distinction between layer-based, channel-based and spatial methods.

Layer-based methods conditionally execute certain network layers or blocks depending on the input. Easy images require a less deep network than complex examples. One of the first methods, Adaptive Computation Time [16], interprets residual blocks as a refinement of features. Execution of layers is halted when features are ‘good enough’ for the classifier. Another approach is to use early-exit branches in the network [7, 35, 45]. Later methods improve the flexibility by conditionally executing individual layers. Those methods are based on the observation that residual architectures are robust against layer dropout [20, 47]. SkipNet [48] learns gating decisions using reinforcement learning. ConvNet-AIG [46] uses the Gumbel-Softmax trick and BlockDrop [51] trains a separate policy network using reinforcement learning.

Channel-based methods prune channels dynamically and on-the-fly during inference. The main motivation is that advanced features are only needed for a subset of the images: features of animals might differ from those of airplanes. Gao *et al.* [13] rank channels and only execute the top-k ones. Lin *et al.* [30] propose a method to train an agent for channel-wise pruning using reinforcement learning, while Bejnordi *et al.* [3] use the Gumbel-Softmax trick to gate channels conditionally on the input.

Spatial methods exploit the fact that not all regions in the image are equally important. A first set of methods [1, 32, 40] uses glimpses to only process regions of interest. Such a two-stage approach is limited to applications where the object of interest is clearly separated, since all information outside the crop is lost. The glimpse idea has been extended for pixel-wise labeling tasks, such as semantic segmentation, using cascades [29]. Later methods offer a finer granularity and more control over the conditional execution. The closest work to ours is probably Spatially Adaptive Computation Time (SACT) [12]. It is a spatial extension of a work by Graves [16] and varies the amount of residual blocks executed per spatial location. Features are processed until a halting score determines that the features are good enough. Since the method relies on refinement of features, it is only applicable to residual networks with many consecutive layers. Our method is more general and flexible as it makes skipping decisions per residual block and per spatial location. In addition, they do not show any inference speedup.

One of the only works showing practical speedup with spatially conditional execution is SBNet [36]. Images are divided into smaller *tiles*, which can be processed with dense convolutions. A low-resolution network first extracts a mask, and the second main network processes tiles according to this mask. Tile edges overlap to avoid discontinuities in the output, causing significant overhead when tiles

are small. Therefore the tile size typically is 16×16 pixels, which makes the method only suitable for large images. They demonstrate their method on 3D object detection tasks only. In contrast, our approach integrates mask generation and sparse execution into a single network, while providing finer pixel-wise control and efficient inference.

3. Method

For each individual residual block, a small gating network generates execution masks based on the input of that block (see Fig. 1). We first describe how pixel-wise masks are learned using the Gumbel-Softmax trick. Afterwards, we elaborate on the implementation of dynamic convolutions, used to reduce inference time. Finally, we propose a sparsity criterion that is added to the task loss in order to achieve the desired reduction in computations.

3.1. Trainable masks

Pixel-wise masks define the spatial positions to be processed by convolutions. These discrete decisions, for every spatial location and every residual block independently, are trained end-to-end using the Gumbel-Softmax trick [23].

3.1.1 Block architecture

Denote the input of the residual block b as $X_b \in \mathbb{R}^{c_b \times w_b \times h_b}$. Then the operation of the residual block is described by

$$X_{b+1} = r(\mathcal{F}(X_b) + X_b) \quad (1)$$

with $\mathcal{F}$ the residual function, typically two or three convolutions with batchnorm (BN), and r an activation function. Our work makes $\mathcal{F}$ conditional in the spatial domain. A small mask unit $\mathcal{M}$ outputs soft gating decisions $M_b \in \mathbb{R}^{w_{b+1} \times h_{b+1}}$, based on input X_b . We compare the mask unit of SACT [12] (referred to as *squeeze unit*), incorporating a squeeze operation over the spatial dimensions, to a 1×1 convolution. We use the squeeze unit in classification, and the 1×1 convolution for pose estimation.

The Gumbel-Softmax module $\mathcal{G}$ turns soft decisions M_b into hard decisions $G_b \in \{0, 1\}^{w_{b+1} \times h_{b+1}}$ by applying a binary Gumbel-Softmax trick (see Section 3.1.2) on each element of M_b :

$$G_b = \mathcal{G}(\mathcal{M}(X_b)). \quad (2)$$

Gating decisions G_b serve as *execution masks*, indicating active spatial positions where the residual block should be evaluated. The residual block with spatially sparse inference is then described by

$$X_{b+1} = r(\mathcal{F}(X_b) \circ G_b + X_b) \quad (3)$$

with $\circ$ an element-wise multiplication over the spatial dimensions ($w_{b+1} \times h_{b+1}$), broadcasted over all channels.

Figure 2: Training spatial execution masks using the Gumbel-Softmax trick. The mask unit generates a floating-point mask, after which the Gumbel-Softmax trick converts soft-decisions into hard-decisions and enables backpropagation for end-to-end learning.

During training, this is an actual multiplication with the mask elements in order to learn gating decisions end-to-end (see Fig. 2). During inference, the residual function is only evaluated on locations indicated by the execution mask G_b .

3.1.2 Binary Gumbel-Softmax

The Gumbel-Softmax trick turns soft decisions into hard decisions while enabling backpropagation, needed to optimize the weights of the mask unit. Take a categorical distribution with class probabilities $\pi = \pi_1, \pi_2, \dots, \pi_n$, then discrete samples z can be drawn using

$$z = \text{one_hot}(\arg \max_i [\log(\pi_i) + g_i]) \quad (4)$$

with g_i being noise samples drawn from a Gumbel distribution. The Gumbel-Softmax trick defines a continuous, differentiable approximation by replacing the argmax operation with a softmax:

$$y_i = \frac{\exp((\log(\pi_i) + g_i)/\tau)}{\sum_{j=1}^k \exp((\log(\pi_j) + g_j)/\tau)}. \quad (5)$$

Gating decisions are binary, which makes it possible to strongly simplify the Gumbel-Softmax formulation. A soft-decision $m \in (-\infty, \infty)$, outputted by a neural network, is converted to a probability π_1 indicating the probability that a pixel should be executed, using a sigmoid σ .

$$\pi_1 = \sigma(m). \quad (6)$$

Then, the probability that a pixel is not executed is

$$\pi_2 = 1 - \sigma(m). \quad (7)$$

Substituting π_1 and π_2 in Equation 5, for the binary case of $k = 2$ and $i = 1$, makes it possible (see supplementary) to reduce this to

$$y_1 = \sigma\left(\frac{m + g_1 - g_2}{\tau}\right). \quad (8)$$

Our experiments use a fixed temperature $\tau = 1$, unless mentioned otherwise. We use a straight-through estimator, where hard samples are used during the forward pass and gradients are obtained from soft samples during the backward pass:

$$z = \begin{cases} y_1 > 0.5 \equiv \frac{m+g_1-g_2}{\tau} > 0 & \text{(forward),} \\ y_1 & \text{(backward).} \end{cases} \quad (9)$$

Note that this formulation has no logarithms or exponentials in the forward pass, typically expensive computations on hardware platforms. During inference, we do not add Gumbel noise and therefore models are finetuned during the last 20 percent of epochs with g_1 and g_2 fixed to 0, making it similar to the straight-through estimator of Bengio *et al.* [6].

3.2. Efficient inference implementation

The residual function should be evaluated on active spatial positions only. Efficiently executing conditional operations can be challenging: hardware strongly relies on regularity to pipeline operations. Especially spatial operations, *e.g.* 3×3 convolutions, require careful optimization and data caching [26].

Our method copies elements at selected spatial positions to an intermediate, dense tensor using a *gather* operation. Non-spatial operations, such as pointwise 1×1 convolutions and activation functions, can be executed efficiently by applying existing implementations on the intermediate tensor. The result is copied back to its original position afterwards using a *scatter* operation. More specifically, let the input I of a residual block be a 4D tensor with dimensions $N \times C \times H \times W$, being the batch size, number of channels, height and width respectively. The gather operation copies the active spatial positions to a new intermediate tensor T with dimensions $P \times C \times 1 \times 1$, where P is the number of active spatial positions spread over the N inputs of the batch. The intermediate tensor can be seen as P images of size 1×1 with C channels, and non-spatial operations can be applied as usual. The inverted residual block of MobileNetV2 relies heavily on non-spatial operations and was designed for low computational complexity, making it a good fit for conditional execution. It consists of a pointwise convolution expanding the feature space, followed by a lightweight depthwise (DW) convolution and linear pointwise bottleneck. The gather operation is applied before the first pointwise convolution, which then operates on the intermediate tensor T . The depthwise convolution is the only spatial operation in the block and should be adapted to operate on the atypical dimensions of T . The second pointwise convolution is followed by the scatter operation, after which the residual summation is made. The architecture of the residual block with dynamic convolutions is shown in Figure 3 and next we describe the role of each component:

Figure 3: Architecture of a residual block for efficient sparse inference. The mask unit $\mathcal{M}$ generates a mask based on the block’s input. The gather operation uses the mask to copy selected spatial positions (yellow) to a new intermediate tensor. Non-spatial operations use standard implementations, while the 3×3 convolution is modified to operate on the intermediate tensor.

- **Mask dilation:** Gating decisions G_b indicate positions where the spatial 3×3 convolution should be applied. The first 1×1 convolution should also be applied on neighboring spatial positions to avoid gaps in the input of the 3×3 convolution. Therefore, the mask G_b is morphologically dilated, resulting in a new mask $G_{b,dilated}$.
- **Masked gather operation:** The gather operation copies active spatial positions from input I with dimension $N \times C \times H \times W$ to an intermediate tensor T with dimension $P \times C \times 1 \times 1$. The index mapping from I to T depends on the execution mask G : an element $I_{n,c,h,w}$ being the p th active position in a flattened version of G , is copied to $T_{p,c,1,1}$.
- **Modified 3x3 Depthwise Convolution:** The depthwise convolution applies a 3×3 convolutional kernel to each channel separately. We implement a custom CUDA kernel that applies the channelwise filtering efficiently on T . The spatial relation between elements of T is lost due to its dimensions being $P \times C \times 1 \times 1$. When processing elements t in T , our implementation retrieves the memory locations of spatial neighbors using an index mapping from T to I .

Traditional libraries for sparse matrix operations have considerable overhead due to indexing and bookkeeping. Our solution minimizes this overhead by gathering elements in smaller, dense matrix. The extra mapping step in the modified 3×3 DW convolution has minimal impact since its computational cost is small compared to 1×1 convolutions. Note that the gather-scatter strategy combined with a modified depthwise convolution can be applied on other networks, such as ShuffleNetV2 [53] and MnasNet [44].

3.3. Sparsity loss

Without additional constraints, the most optimal gating state is to execute every spatial location. We define a *computational budget* hyperparameter $\theta \in [0, 1]$, indicating the relative amount of desired operations. For instance, 0.7 means that on average 70% of the FLOPS in the conditional layers should be executed. The total number of floating point operations for convolutions in a MobileNetV2 residual block b with stride 1 is

$$\mathbb{F}_b = H \cdot W \cdot (9C_{b,e} + 2C_b C_{b,e}), \quad (10)$$

with C_b the number of base channels and $C_{b,e}$ the amount of channels for the depthwise convolution (typically $6C_b$) and $H \times W$ the spatial dimensions. For sparse inference with dynamic convolutions, this becomes

$$\mathbb{F}_{b,sp} = N_{b,dilated} C_b C_{b,e} + N_b (9C_{b,e} + C_{b,e} C_b), \quad (11)$$

with $N_{b,dilated}$ and N_b the number of active spatial positions in the dilated mask and mask respectively. The value N_b is made differentiable by calculating it as the sum of all elements in the execution mask of that block (Eq. 2):

$$N_b = \sum G_b. \quad (12)$$

The following loss then minimizes the difference between the given computational budget θ and the budget used by a network consisting of B residual blocks:

$$\mathcal{L}_{sp,net} = \left(\frac{\sum_b^B \mathbb{F}_{b,sp}}{\sum_b^B \mathbb{F}_b} - \theta \right)^2. \quad (13)$$

In practice, we average the FLOPS over the batch size, and the network is free to choose the distribution of computations over the residual blocks and batch elements. However, without proper initialization this could lead to a suboptimal state where the network executes all positions in a certain block or none. Minimizing the sparsity loss is trivial compared to the task loss and the mask units never recover from this state. This problem occurs often in conditional execution and existing solutions consist of dense pretraining with careful initialization [12, 48], curriculum learning [51] or setting a computational budget for each residual block individually [46]. The latter can be formulated as

$$\mathcal{L}_{sp,per\,layer} = \sum_b^B \left(\frac{\mathbb{F}_{b,sp}}{\mathbb{F}_b} - \theta \right)^2. \quad (14)$$

Such a constraint per layer is effective but limits the flexibility of computation allocations. We propose a solution to ensure proper initialization of each block, by adding an extra constraint that keeps the percentage of executed operations $\mathbb{F}_{b,sp}/\mathbb{F}_b$ in each residual block between an upper

and lower bound. Those bounds are annealed from the target budget θ to 0 and 1 respectively.

The upper and lower bound are imposed by the following loss functions, where we use cosine annealing to vary p from 1 to 0 between the first and last epoch of training:

$$\mathcal{L}_{sp,low} = \frac{1}{B} \sum_b^B \max(0, p \cdot \theta - \frac{\mathbb{F}_{b,sp}}{\mathbb{F}_b})^2, \quad (15)$$

$$\mathcal{L}_{sp,up} = \frac{1}{B} \sum_b^B \max(0, \frac{\mathbb{F}_{b,sp}}{\mathbb{F}_b} - (1 - p(1 - \theta)))^2. \quad (16)$$

The final loss to be optimized is then given by

$$\mathcal{L} = \mathcal{L}_{task} + \alpha(\mathcal{L}_{sp,net} + \mathcal{L}_{sp,lower} + \mathcal{L}_{sp,upper}) \quad (17)$$

where α is a hyperparameter, chosen so that the task and sparsity loss have the same order of magnitude. We choose $\alpha = 10$ for classification and $\alpha = 0.01$ for pose estimation.

4. Experiments and results

We first compare our masking method with other conditional execution approaches on CIFAR and ResNet, and show that our method improves the accuracy-complexity trade-off. Then we demonstrate inference speedup on Food-101 with MobileNetV2 and ShuffleNetV2. Finally we apply our method on pose estimation, a task that is typically spatially sparse. We study the impact of the proposed sparsity criterion on this task.

4.1. Classification

4.1.1 Comparison with state-of-the-art

The ResNet [17] CNN for classification is typically used to compare the performance of conditional execution methods. We compare the theoretical number of floating point operations and accuracy. The main work related to ours is SACT [12], also exploiting spatial sparsity. ConvNet-AIG [46], conditionally executing complete residual blocks, can be seen as a non-spatial variant of our method.

CIFAR-10 We perform experiments with ResNet-32 on the standard train/validation split of CIFAR-10 [25]. We use the same hyperparameters and data augmentation as ConvNet-AIG, being an SGD optimizer with momentum 0.9, weight decay $5e-4$, learning rate 0.1 decayed by 0.1 at epoch 150 and 250 with a total of 350 epochs. Results for SACT and ConvNet-AIG are obtained with their provided code. We evaluate our method for different budget targets $\theta \in \{0.1, 0.2, \dots, 0.9\}$. The mask unit is a squeeze unit with the same architecture as in SACT. Non-adaptive baseline ResNet models have 8, 14, 20, 26 and 32 layers. Figure 4a shows that our method (DynConv) outperforms the other conditional execution methods for all

Figure 4: Comparison with state-of-the-art models

computational costs (MACs) while improving the accuracy-complexity trade-off. Moreover, there is a smaller drop in accuracy when reducing the computational budget.

ImageNet We use ResNet-101 [18] and ImageNet [10] to compare DynConv against results reported in SACT [12], ConvNet-AIG [46], SkipNet [48] and standard ResNet [18]. Just like SACT, we initialize convolutional layers with weights from a pretrained ResNet-101 since the large amount of conditional layers makes the network prone to dead residuals, where some layers are not used at all. We use the standard training procedure of ResNet [18] with InceptionV3 [43] data augmentation. Models are trained on a single GPU with batch size 64 and learning rate 0.025 for 100 epochs. Learning rate is decayed by 0.1 at epoch 30 and 60. The Gumbel temperature is gradually annealed from 5 to 1, for more stable training of this deep model. The mask unit is a squeeze unit with the same architecture as in SACT. Results in Figure 4b show that DynConv outperforms SACT by a large margin and matches the performances of the best layer-based methods. Those methods strongly benefit from the large amount of residual blocks in ResNet-101 and therefore perform relatively better than they did in the CIFAR-10 experiment.

Further analysis Figure 5 presents a qualitative comparison between our method and SACT. The amount of computations per spatial location is visualized using *ponder cost maps*, obtained by upscaling the binary execution masks of all residual blocks and subsequently summing them. Our method shows better focus on the region of interest. Analyzing the distribution of computations over ImageNet classes (Fig 6a) shows that the network spends fewer computations on typically sparse images such as birds. When looking at the execution rates per residual block (Fig. 6b), it is clear that downsampling blocks are more important than others: all spatial locations are evaluated in these blocks. The last residual blocks, processing high-level features, are

Table 1: Inference time of dynamic convolutions

Method	Acc.	MACs	Im/Sec
MobileNetV2 x0.75 [39] (our impl.)	82.0	225M	508
MobileNetV2 x0.60 (our impl.)	79.7	150M	638
$\theta = 0.75$	81.2	200M	541
$\theta = 0.50$	80.6	174M	629
$\theta = 0.25$ (G-Binary)	79.8	148M	724
$\theta = 0.25$ (G-Softmax)	79.8	148M	522
ShuffleNetV2 [53] (our impl.)	78.7	149M	710
$\theta = 0.25$	76.5	100M	781
$\theta = 0.25$ with 1×1 conv mask unit	76.3	97M	889

also crucial. This highlights the architectural advantage over SACT, where computation at a spatial location can only be halted.

4.1.2 DynConv inference speedup

We integrate dynamic convolutions in MobileNetV2 [39] and ShuffleNetV2 [53]. Results for different computational budgets θ are shown in Table 1. We use the Food-101 dataset [8], containing 75k/25k train/test images of 101 food classes, with InceptionV3 [43] data augmentation and image size 224×224 . For MobileNetV2, reducing the computational budget θ results in a proportional increase of throughput (images per second). We compared a version with standard Gumbel-Softmax (Eq. 5) to our reformulation (Eq. 9): our reformulation (G-Binary) is significantly faster than the one with softmax and logarithms (G-Softmax). We use ShuffleNetV2 with a residual connection [53]. Our dynamic convolutions are integrated in the convolutional branch of non-strided blocks. This architecture uses narrower residual blocks, and the relative cost of the mask unit with squeeze operation becomes significant. We suggest using a cheaper 1×1 convolution as mask unit for narrow networks.

4.2. Human pose estimation

Human pose estimation is a task that is inherently spatially sparse: many pixels around the person are not relevant for keypoint detection and the output heatmaps are sparse. Most existing dynamic execution methods are not suited for this task: layer-based and channels based methods, such as ConvNet-AIG [46], do not exploit the spatial dimensions. SACT [12] is not directly applicable on branched architectures such as stacked hourglass networks [34], as it can only halt execution.

Experiment setup We base our work on Fast Pose Distillation (FPD) [52], which uses the well-known stacked hourglass network for single-person human pose estima-

Figure 5: Qualitative evaluation and comparison with SACT. Ponder cost maps indicate the amount of residual blocks executed at each spatial location. Both methods have the same average complexity (5.7 GMACs), but ours shows better focus on the region of interest, both in simple examples (top row) as more cluttered ones (bottom row). Ponder cost maps of SACT and input images obtained from [12].

Figure 6: (a) Distribution of MACs over images in the ImageNet validation set. Images of birds, typically sparse, are processed with fewer operations than those of consumer goods. (b) Distribution of computations over residual blocks. Error bars indicate variance. Downsampling blocks (purple) are important and evaluated at all spatial positions.

tion [34]. Unlike their work, we do apply knowledge distillation since such method is complementary. The standard residual blocks are replaced by those of MobileNetV2 with depthwise convolution, in order to achieve efficient inference. Our baseline models has 96 features in the hourglass, expanded to 576 channels for the depthwise convolution. Models of different width are obtained by multiplying the number of channels with a *width multiplier* $\in \{0.5, 0.75, 1.0\}$, while the network depth is varied by using 1, 2 and 4 hourglass stacks (1S, 2S and 4S). For dynamic convolutions, the computational budget of the base model is varied with $\theta \in \{0.125, 0.25, 0.5, 0.75\}$.

We use the MPII dataset [2] with standard test/validation split (22k/3k images). Images are resized to 256×256 and augmented with ± 30 degrees rotation, ± 25 percent scaling and random horizontal flip. No flip augmentation is used during evaluation. The optimizer is Adam with learning rate

Figure 7: Results on pose estimation (MPII validation set) for hourglass models with varying depths and widths. Our conditional execution method (DynConv, in red) always outperforms baseline models (in blue) with the same depth and amount of computations. Our models achieve similar performance as FPD [52] (without knowledge distillation) with fewer computations.

$2e-4$ for a batch size of 6 samples. The mean square error loss for heatmaps is averaged over samples. Training lasts for 100 epochs and the learning rate is reduced by factor 0.1 at epoch 60 and 90. The evaluation metric is the mean Percentage of Correct Keypoints, normalized by a fraction of the head size (PCKh@0.5), as implemented in [52].

Results Figure 7 shows that our method always outperforms non-conditional models with a similar amount of operations. The amount of operations is reduced by more than 45% with almost no loss in accuracy. The performance difference between baselines of similar FLOPS and our method becomes larger for larger reductions of FLOPS.

Ponder cost maps in Figure 8 show that the networks learns to focus on the person. Table 2 demonstrates that our dynamic convolution method can significantly speed up in-

Figure 8: Ponder cost on single-person human pose estimation (4-stack hourglass, $\theta = 0.125$). The network learns to apply the majority of convolutional operations on keypoint locations, even in the presence of clutter.

Table 2: Timings on pose estimation for 2-stack models. Our models achieve significant wall-clock time speedup (measured in images processed per second) on an Nvidia GTX1050 Ti 4GB GPU. Timings of subcomponents (mask, bookkeeping, gather, residual function, scatter) are given as average time per execution, in milliseconds.

Model	PCKh@0.5	# Params	MACs	Images/Sec	Mask	Bookkeeping	Gather	Res. $\mathcal{F}$	Scatter
4-stack baseline	88.1	6.88M	6.90G	30	N.A	N.A.	N.A	31.1	N.A
DynConv ($\theta = 0.75$)	88.1	6.89M	5.39G	33	0.48	0.73	0.59	27.3	0.76
DynConv ($\theta = 0.50$)	88.2	6.89M	3.78G	48	0.48	0.78	0.47	18.4	0.54
DynConv ($\theta = 0.25$)	87.5	6.89M	2.30G	67	0.45	0.70	0.33	10.8	0.36
DynConv ($\theta = 0.125$)	86.7	6.89M	1.71G	85	0.50	0.67	0.30	8.25	0.27
baseline ($\times 0.5$ channels)	85.2	1.83M	1.83G	66	N.A.	N.A.	N.A.	14.1	N.A.

Table 3: Ablation on pose estimation (4-stack, $\theta = 0.125$).

Mask unit	Sparsity criterion	PCKh	Im/Sec
1×1 conv	$\mathcal{L}_{sp,net} + \mathcal{L}_{sp,up} + \mathcal{L}_{sp,low}$	86.7	85
squeeze unit	$\mathcal{L}_{sp,net} + \mathcal{L}_{sp,up} + \mathcal{L}_{sp,low}$	87.0	76
1×1 conv	$\mathcal{L}_{sp,net}$ (Eq. 13)	86.1	85
1×1 conv	$\mathcal{L}_{sp,per,layer}$ (Eq. 14)	86.3	85

ference, giving a 60% speedup without loss in accuracy and a 125% speedup with 0.6% accuracy loss. The table also shows that the time spent on generating the mask, bookkeeping and gather/scatter operations is small compared to the cost of the residual function. Our model also outperforms smaller baseline model with equal FLOPS. Table 3 compares the squeeze masking unit used in SACT [12] to a simple convolution. Using a squeeze function as mask unit performs slightly better at the cost of a lower inference speed. In addition, we compare the sparsity losses proposed in Section 3.3. Our sparsity criterion with upper and lower bound outperforms more simple losses. The effect of our loss during training is shown in Figure 9.

5. Conclusion and future work

In this work we proposed a method to speed up inference using dynamic convolutions. The network learns pixel-wise execution masks in an end-to-end fashion. Dynamic convolutions speed up inference and reduce the number of oper-

Figure 9: Evolution of the percentage of conditional computations executed, per residual block, during training of pose estimation. Early layers are colored blue, while deeper layers are red. The simple network-wise sparsity criterion (Eq. 13) causes early convergence to a suboptimal state, where many layers are not used. Our criterion where a lower and upper bound per block are annealed (Eq. 17) provides better initialization and more stable training.

ations by only operating on these predicted locations. Our method achieves state-of-the-art results on image classification, and our qualitative analysis demonstrates the architectural advantages over existing methods. Our method is especially suitable for processing high-resolution images, e.g. in pose estimation or object detection tasks.

6. Acknowledgement

This work was funded by IMEC through the ICON Lecture+ project and by CELSA through the SfS++ project.

References

- [1] Amjad Almahairi, Nicolas Ballas, Tim Cooijmans, Yin Zheng, Hugo Larochelle, and Aaron Courville. Dynamic capacity networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2549–2558, 2016.
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [3] Babak Ehteshami Bejnordi, Tijmen Blankevoort, and Max Welling. Batch-shaped channel gated networks. *arXiv preprint arXiv:1907.06627*, 2019.
- [4] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*, 2015.
- [5] Yoshua Bengio. Deep learning of representations: Looking forward. In *International Conference on Statistical Language and Speech Processing*, pages 1–37. Springer, 2013.
- [6] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- [7] Tolga Bolukbasi, Joseph Wang, Ofer Dekel, and Venkatesh Saligrama. Adaptive neural networks for efficient inference. In *Proceedings of the 34th International Conference on Machine Learning (ICML) - Volume 70*, pages 527–536. JMLR.org, 2017.
- [8] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 446–461. Springer, 2014.
- [9] Alfredo Canziani, Adam Paszke, and Eugenio Culurciello. An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678*, 2016.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. IEEE, 2009.
- [11] Alain Droniou, Serena Ivaldi, and Olivier Sigaud. Deep unsupervised network for multimodal perception, representation and classification. *Robotics and Autonomous Systems*, 71:83–98, 2015.
- [12] Michael Figurnov, Maxwell D Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, and Ruslan Salakhutdinov. Spatially adaptive computation time for residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1039–1048, 2017.
- [13] Xitong Gao, Yiren Zhao, Lukasz Dudziak, Robert Mullins, and Cheng-zhong Xu. Dynamic channel pruning: Feature boosting and suppression. *arXiv preprint arXiv:1810.05331*, 2018.
- [14] Benjamin Graham. Spatially-sparse convolutional neural networks. *arXiv preprint arXiv:1409.6070*, 2014.
- [15] Benjamin Graham and Laurens van der Maaten. Sub-manifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307*, 2017.
- [16] Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.
- [17] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [19] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [20] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European Conference on Computer Vision (ECCV)*, pages 646–661. Springer, 2016.
- [21] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [22] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, Geoffrey E Hinton, et al. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- [23] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
- [24] Yong-Deok Kim, Eunhyeok Park, Sungjoo Yoo, Taelim Choi, Lu Yang, and Dongjun Shin. Compression of deep convolutional neural networks for fast and low power mobile applications. *arXiv preprint arXiv:1511.06530*, 2015.
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009.
- [26] Andrew Lavin and Scott Gray. Fast algorithms for convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4013–4021, 2016.
- [27] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in Neural Information Processing Systems (NIPS)*, pages 598–605, 1990.
- [28] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.
- [29] Xiaoxiao Li, Ziwei Liu, Ping Luo, Chen Change Loy, and Xiaoou Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3193–3202, 2017.
- [30] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2181–2191, 2017.

- [31] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
- [32] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2204–2212, 2014.
- [33] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. *arXiv preprint arXiv:1611.06440*, 2016.
- [34] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European conference on Computer Vision (ECCV)*, pages 483–499. Springer, 2016.
- [35] Priyadarshini Panda, Abhronil Sengupta, and Kaushik Roy. Conditional deep learning for energy-efficient and enhanced pattern recognition. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 475–480. IEEE, 2016.
- [36] Mengye Ren, Andrei Pokrovsky, Bin Yang, and Raquel Urtasun. Sbnnet: Sparse blocks network for fast inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8711–8720, 2018.
- [37] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.
- [38] Tara N Sainath, Brian Kingsbury, Vikas Sindhwani, Ebru Arisoy, and Bhuvana Ramabhadran. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6655–6659. IEEE, 2013.
- [39] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018.
- [40] Shikhar Sharma, Ryan Kiros, and Ruslan Salakhutdinov. Action recognition using visual attention. *arXiv preprint arXiv:1511.04119*, 2015.
- [41] Vikas Sindhwani, Tara Sainath, and Sanjiv Kumar. Structured transforms for small-footprint deep learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3088–3096, 2015.
- [42] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.
- [43] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016.
- [44] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2820–2828, 2019.
- [45] Surat Teerapittayanon, Bradley McDanel, and Hsiang-Tsung Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *23rd International Conference on Pattern Recognition (ICPR)*, pages 2464–2469. IEEE, 2016.
- [46] Andreas Veit and Serge Belongie. Convolutional networks with adaptive inference graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [47] Andreas Veit, Michael J Wilber, and Serge Belongie. Residual networks behave like ensembles of relatively shallow networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 550–558, 2016.
- [48] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 409–424, 2018.
- [49] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2074–2082, 2016.
- [50] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4820–4828, 2016.
- [51] Zuxuan Wu, Tushar Nagarajan, Abhishek Kumar, Steven Rennie, Larry S Davis, Kristen Grauman, and Rogerio Feris. Blockdrop: Dynamic inference paths in residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8817–8826, 2018.
- [52] Feng Zhang, Xiatian Zhu, and Mao Ye. Fast human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3517–3526, 2019.
- [53] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6848–6856, 2018.