Latency-aware Spatial-wise Dynamic Networks

Yizeng Han[1*] **Zhihang Yuan**[2*] **Yifan Pu**[1*] **Chenhao Xue**[2]
Shiji Song[1] **Guangyu Sun**[2] **Gao Huang**[1†]
[1] Department of Automation, BNRist, Tsinghua University, Beijing, China
[2] School of Electronics Engineering and Computer Science, Peking University, Beijing, China
{hanyz18, pyf20}@mails.tsinghua.edu.cn, {shijis, gaohuang}@tsinghua.edu.cn
{yuanzhihang, xch927027, gsun}@pku.edu.cn

Abstract

Spatial-wise dynamic convolution has become a promising approach to improving the inference efficiency of deep networks. By allocating more computation to the most informative pixels, such an *adaptive* inference paradigm reduces the spatial redundancy in image features and saves a considerable amount of unnecessary computation. However, the *theoretical* efficiency achieved by previous methods can hardly translate into a *realistic* speedup, especially on the multi-core processors (*e.g.* GPUs). The key challenge is that the existing literature has only focused on designing algorithms with minimal *computation*, ignoring the fact that the practical latency can also be influenced by *scheduling strategies* and *hardware properties*. To bridge the gap between theoretical computation and practical efficiency, we propose a *latency-aware* spatial-wise dynamic network (LASNet), which performs *coarse-grained* spatially adaptive inference under the guidance of a novel *latency prediction model*. The latency prediction model can efficiently estimate the inference latency of dynamic networks by simultaneously considering algorithms, scheduling strategies, and hardware properties. We use the latency predictor to guide both the algorithm design and the scheduling optimization on various hardware platforms. Experiments on image classification, object detection and instance segmentation demonstrate that the proposed framework significantly improves the practical inference efficiency of deep networks. For example, the average latency of a ResNet-101 on the ImageNet validation set could be reduced by 36% and 46% on a server GPU (Nvidia Tesla-V100) and an edge device (Nvidia Jetson TX2 GPU) respectively without sacrificing the accuracy. Code is available at `https://github.com/LeapLabTHU/LASNet`.

1 Introduction

Dynamic neural networks [7] have attracted great research interests in recent years. Compared to static models [11, 17, 13, 23] which treat different inputs equally during inference, dynamic networks can allocate the computation in a *data-dependent* manner. For example, they can conditionally skip the computation of network layers [15, 9, 32, 30] or convolutional channels [19, 1], or perform *spatially* adaptive inference on the most informative image regions (*e.g.* the foreground areas) [6, 5, 31, 35, 33, 8]. Spatial-wise dynamic networks, which typically decide whether to compute each feature pixel with *masker* modules [5, 31, 35, 8] (Figure 1 (a)), have shown promising results in improving the inference efficiency of convolution neural networks (CNNs).

Despite the remarkable *theoretical* efficiency achieved by spatial-wise dynamic networks [5, 31, 35], researchers have found it challenging to translate the theoretical results into *realistic* speedup,

[*] Equal contribution.
[†] Corresponding author.

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

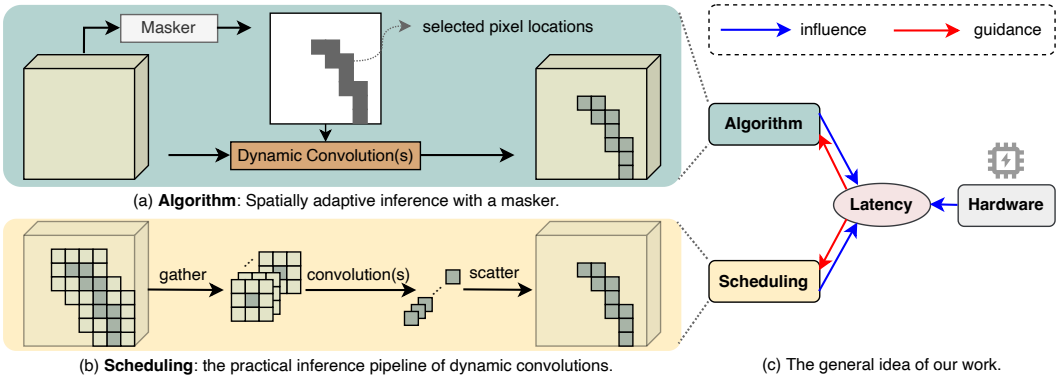

Figure 1: An overview of our method. (a) illustrates the spatially adaptive inference *algorithm*; (b) is the *scheduling* strategy; and (c) presents the three key factors to the practical latency. For a given hardware, the latency is used to *guide* our algorithm design and scheduling optimization.

especially on some multi-core processors, *e.g.,* GPUs [35, 3, 8]. The challenges are two-fold: 1) most previous approaches [5, 31, 35] perform spatially adaptive inference at the finest *granularity*: every *pixel* is flexibly decided whether to be computed or not. Such flexibility induces non-contiguous *memory access* [35] and requires specialized *scheduling strategies* (Figure 1 (b)); 2) the existing literature has only adopted the *hardware-agnostic* FLOPs (floating-point operations) as an inaccurate proxy for the efficiency, lacking latency-aware guidance on the algorithm design. For dynamic networks, the adaptive computation with sub-optimal scheduling strategies further enlarges the discrepancy between the theoretical FLOPs and the practical latency. Note that it has been validated by previous works that the latency on CPUs has a strong correlation with FLOPs [8, 35]. Therefore, we mainly focus on the GPU platform in this paper, which is more challenging and less explored.

We address the above challenges by proposing a *latency-aware* spatial-wise dynamic network (LASNet). Three key factors to the inference latency are considered: the *algorithm*, the *scheduling strategy*, and the *hardware properties*. Given a target hardware device, we directly use the latency, rather than the FLOPs, to *guide* our algorithm design and scheduling optimization (Figure 1 (c)).

Because the memory access pattern and the scheduling strategies in our dynamic operators differ from those in static networks, the libraries developed for static models (*e.g.* cuDNN) are sub-optimal for the acceleration of dynamic models. Without the support of libraries, each dynamic operator requires scheduling optimization, code optimization, compiling, and deployment for each device. Therefore, it is laborious to evaluate the network latency on different hardware platforms. To this end, we propose a novel *latency prediction model* to efficiently estimate the realistic latency of a network by simultaneously considering the aforementioned three factors. Compared to the hardware-agnostic FLOPs, our predicted latency can better reflect the practical efficiency of dynamic models.

Guided by this latency prediction model, we establish our latency-aware spatial-wise dynamic network (LASNet), which adaptively decides whether to allocate computation on feature *patches* instead of *pixels* [5, 31, 35] (Figure 2 top). We name this paradigm as spatially adaptive inference at a *coarse granularity*. While less flexible than the pixel-level adaptive computation in previous works [5, 31, 35], it facilitates more contiguous memory access, benefiting the realistic speedup on hardware. The scheduling strategy and the implementation are further ameliorated for faster inference.

It is worth noting that LASNet is designed as a general framework in two aspects: 1) the coarse-grained spatially adaptive inference paradigm can be conveniently implemented in various CNN backbones, *e.g.,* ResNets [11], DenseNets [17] and RegNets [25]; and 2) the latency predictor is an off-the-shell tool which can be directly used for various computing platforms (*e.g.* server GPUs and edge devices).

We evaluate the performance of LASNet on multiple CNN architectures on image classification, object detection, and instance segmentation tasks. Experiment results show that our LASNet improves the efficiency of deep CNNs both theoretically and practically. For example, the inference latency of ResNet-101 on ImageNet [4] is reduced by 36% and 46% on an Nvidia Tesla V100 GPU and an Nvidia Jetson TX2 GPU, respectively, without sacrificing the accuracy. Moreover, the proposed method outperforms various lightweight networks in a low-FLOPs regime.

Our main contributions are summarized as follows:

1. We propose LASNet, which performs coarse-grained spatially adaptive inference guided by the practical latency instead of the theoretical FLOPs. To the best of our knowledge, LASNet is the first framework that directly considers the real latency in the design phase of dynamic neural networks;
2. We propose a latency prediction model, which can efficiently and accurately estimate the latency of dynamic operators by simultaneously considering the algorithm, the scheduling strategy, and the hardware properties;
3. Experiments on image classification and downstream tasks verify that our proposed LASNet can effectively improve the practical efficiency of different CNN architectures.

2 Related works

Spatial-wise dynamic network is a common type of dynamic neural networks [7]. Compared to static models which treat different feature locations evenly during inference, these networks perform spatially adaptive inference on the most informative regions (*e.g.*, foregrounds), and reduce the unnecessary computation on less important areas (*e.g.*, backgrounds). Existing works mainly include three levels of dynamic computation: resolution level [36, 37], region level [33] and pixel level [5, 31, 35]. The former two generally manipulate the network inputs [33, 37] or require special architecture design [36]. In contrast, pixel-level dynamic networks can flexibly skip the convolutions on certain feature pixels in arbitrary CNN backbones [5, 31, 35]. Despite its remarkable *theoretical* efficiency, pixel-wise dynamic computation brings considerable difficulty to achieving *realistic* speedup on multi-core processors, *e.g.*, GPUs. Compared to the previous approaches [5, 31, 35] which only focus on reducing the theoretical computation, we propose to directly use the latency to guide our algorithm design and scheduling optimization.

Hardware-aware network design. To bridge the gap between theoretical and practical efficiency of deep models, researchers have started to consider the real latency in the network design phase. There are two lines of works in this direction. One directly performs speed tests on targeted devices, and summarizes some guidelines to facilitate *hand-designing* lightweight models [23]. The other line of work *searches* for fast models using the neural architecture search (NAS) technique [29, 34]. However, all existing works try to build *static* models, which have intrinsic computational redundancy by treating different inputs in the same way. However, speed tests for dynamic operators on different hardware devices can be very laborious and impractical. In contrast, our proposed latency prediction model can efficiently estimate the inference latency on any given computing platforms by simultaneously considering algorithm design, scheduling strategies and hardware properties.

3 Methodology

In this section, we first introduce the preliminaries of spatially adaptive inference, and then demonstrate the architecture design of our LASNet. The latency prediction model is then explained, which guides the granularity settings and the scheduling optimization for LASNet. We further present the implementation improvements for faster inference, followed by the training strategies.

3.1 Preliminaries

Spatially adaptive inference. The existing spatial-wise dynamic networks are usually established by attaching a masker $\mathcal{M}$ in each convolutional block of a CNN backbone (Figure 1 (a)). Specifically, let $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ denote the input of a block, where H and W are the feature height and width, and C is the channel number. The masker $\mathcal{M}$ takes $\mathbf{x}$ as input, and generates a binary-valued spatial mask $\mathbf{M} = \mathcal{M}(\mathbf{x}) \in \{0, 1\}^{H \times W}$. Each element of $\mathbf{M}$ determines whether to perform convolution operations on the corresponding location of the output feature. The unselected regions will be filled with the values from the input [5, 31] or obtained via interpolation [35]. We define the *activation rate* of a block as $r = \frac{\sum_{i,j} \mathbf{M}_{i,j}}{H \times W}$, representing the ratio of the calculated pixels.

Scheduling strategy. During inference, the current scheduling strategy for spatial-wise dynamic convolutions generally involve three steps [26] (Figure 1 (b)): 1) *gathering*, which re-organizes the selected pixels (if the convolution kernel size is greater than 1×1, the neighbors are also required) along the *batch* dimension; 2) *computation*, which performs convolution on the gathered input; and 3) *scattering*, which fills the computed pixels on their corresponding locations of the output feature.

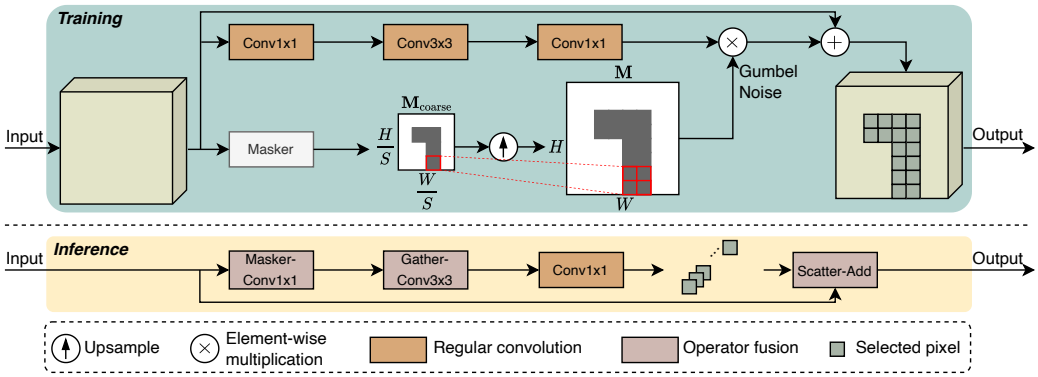

Figure 2: Our proposed LASNet block. Top: we first generate a low-resolution spatial mask $\mathbf{M}_{\text{coarse}}$, which is then upsampled to obtain the mask $\mathbf{M}$ with the same size as the output feature. Gumbel Softmax [18, 24] is used for end-to-end training (Sec. 3.5). Bottom: the scheduling optimization is performed to decrease the memory access for faster inference (Sec. 3.4).

Limitations. Compared to performing convolutions on the entire feature map, the aforementioned scheduling strategy reduces the computation while bringing considerable overhead to *memory access* due to the mask generation and the non-contiguous memory access. Such overhead would increase the overall latency, especially when the *granularity* of dynamic convolution is at the finest pixel level.

3.2 Architecture design

Spatial granularity. As mentioned above, *pixel*-level dynamic convolutions [5, 31, 35] raise substantial challenges to achieving realistic speedup on multi-core processors due to the non-contiguous memory access. To this end, we propose to optimize the *granularity* of spatially adaptive inference. Specifically, take the commonly used bottleneck structure in [11] as an example, our coarse-grained spatial-wise dynamic convolutional block is illustrated in Figure 2. Instead of directly producing a mask with the shape of $H \times W$, we first generate a low-resolution mask $\mathbf{M}_{\text{coarse}} \in \{0, 1\}^{\frac{H}{S} \times \frac{W}{S}}$, where S is named as the *spatial granularity*. Each element in $\mathbf{M}_{\text{coarse}}$ determines the computation of an $S \times S$-sized feature patch. For instance, the feature size in the first ResNet stage[3] is 56×56. Then the possible choices for S are $\{1, 2, 4, 7, 8, 14, 28, 56\}$. The mask $\mathbf{M}_{\text{coarse}}$ is then upsampled to the size of $H \times W$. Notably, $S = 1$ means that the granularity is still at the pixel level as previous methods [5, 31, 35]. Note that the other extreme situation ($S = 56$) is not considered in this paper, when the masker directly determines whether to skip the whole block (*i.e.* layer skipping [30, 32]). Such an overly aggressive approach will lead to considerable drop of accuracy, as we presented in Appendix C.2. The masker is composed of a pooling layer followed by a 1×1 convolution.

Differences to existing works. Without using the interpolation operation [35] or the carefully designed two-branch structure [8], the proposed block architecture is simple and sufficiently general to be plugged into most backbones with minimal modification. Our formulation is mostly similar to that in [31], which could be viewed as a variant of our method with the spatial granularity S=1 for all blocks. Instead of performing spatially adaptive inference at the finest pixel level, our granularity S is optimized under the guidance of our *latency prediction model* (details are presented in the following Sec. 4.2) to achieve *realistic speedup* on target computing platforms.

3.3 Latency prediction model

As stated before, it is laborious to evaluate the latency of dynamic operators on different hardware platforms. To efficiently seek preferable granularity settings on arbitrary hardware devices, we propose a latency prediction model $\mathcal{G}$, which can directly *predict* the delay of executing dynamic operators on any target devices. For a spatial-wise dynamic convolutional block, the latency predictor $\mathcal{G}$ takes the hardware properties $\mathbf{H}$, the layer parameters $\mathbf{P}$, the spatial granularity S, and the activation rate r as input and predicts the latency ℓ of a dynamic convolutional block: $\ell = \mathcal{G}(\mathbf{H}, \mathbf{P}, S, r)$.

Hardware modeling. We model a hardware device as multiple processing engines (PEs), and parallel computation can be executed on these PEs. As shown in Figure 3, we model the memory system as a

[3]Here we refer to a stage as the cascading of multiple blocks which process features with the same resolution.

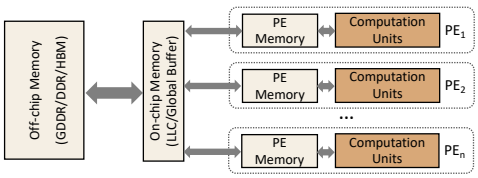

Figure 3: Our hardware model.

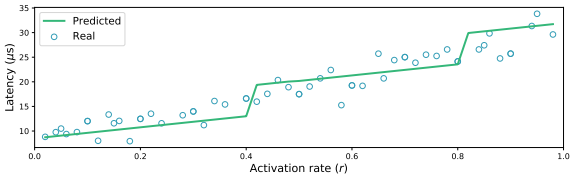

Figure 4: Latency prediction results.

three-level structure [12]: 1) off-chip memory, 2) on-chip global memory, and 3) memory in PE. Such a hardware model enables us to accurately predict the cost on both *data movement* and *computation*.

Latency prediction. When simulating the data movement procedure, the efficiency of non-contiguous memory accesses under different granularity S settings is considered. As for the computation latency, it is important to adopt a proper scheduling strategy to increase the parallelism of computation. Therefore, we search for the optimal scheduling (the configuration of tiling and in-PE parallel computing) of dynamic operations to maximize the utilization of hardware resources. A more detailed description of our latency prediction model is presented in Appendix A.

Empirical validation. We take the first block in ResNet-101 as an example and vary the activation rate r to evaluate the performance of our prediction model. The comparison between our predictions and the real testing latency on the Nvidia V100 GPU is illustrated in Figure 4, from which we can observe that our predictor can accurately estimate the real latency in a wide range of activation rates.

3.4 Implementation details

We use general optimization methods like fusing activation functions and batch normalization layers into convolution layers. We also optimize the specific operators in our spatial-wise dynamic convolutional blocks as follows (see also Figure 2 for an overview).

Fusing the masker and the first convolution. As mentioned in Sec. 3.1, the masker in each block consumes very little computation, but it takes the whole feature map as input. Therefore, it is a *memory-bounded* operation (the inference time is mainly spent on memory access). Since the masker and the first convolution in the block share the same input, there is an opportunity to fuse these two operations to avoid the repeated access of the input data. Note that a spatial-wise dynamic convolution requires the output of the masker. If we fuse the two layers, the first convolution will be changed to a static operation, which may increase the inference latency. There exists a threshold of activation rate r_{th}, when $r > r_{\mathrm{th}}$, the overall latency can be reduced. We decide whether to fuse them according to the average activation rate. See more details in Appendix B.1.

Fusing the gather operation and the dynamic convolution. Traditional approaches first gather the input pixels of the first dynamic convolution in a block (Figure 1 (b)). The gather operation is also a *memory-bounded* operation. Furthermore, when the size of the convolution kernel exceeds 1×1, the area of input patches may overlap, resulting in repeated memory load/store. We fuse the gather operation into the dynamic convolution to reduce the memory access.

Fusing the scatter operation and the add operation. Traditional approaches scatter the output pixels of the last dynamic convolution, and then execute the element-wise addition (Figure 1 (b)). We fuse these two operators to reduce the memory access. The ablation study in Sec. 4.4 validates the effectiveness of the proposed fusing methods.

3.5 Training

Optimization of non-differentiable maskers. The masker modules are required to produce binary-valued spatial masks for making discrete decisions, and cannot be directly optimized with back propagation. Following [35, 31, 8], we adopt straight-through Gumbel Softmax [18, 24] to train the network in an end-to-end fashion. Specifically, let $\tilde{\mathbf{M}} \in \mathbb{R}^{H\times W\times 2}$ denote the output of the mask generator. The decisions are obtained with the argmax function during inference. In the training phase, a differentiable approximation is defined by replacing the argmax operation with a Softmax:

$$\hat{\mathbf{M}} = \frac{\exp\left\{\left(\log\left(\tilde{\mathbf{M}}_{:,:,0}\right) + \mathbf{G}_{:,:,0}\right)/\tau\right\}}{\sum_{k=0}^{1}\exp\left\{\left(\log\left(\tilde{\mathbf{M}}_{:,:,k}\right) + \mathbf{G}_{:,:,k}\right)/\tau\right\}} \in [0,1]^{H\times W}, \tag{1}$$

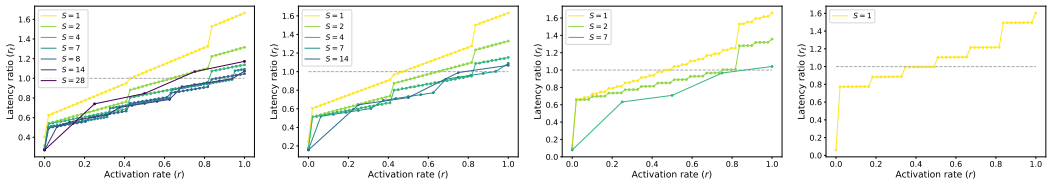

(a) Relationship between r_ℓ and r for LAS-ResNet blocks on the Nvidia Tesla V100 GPU.

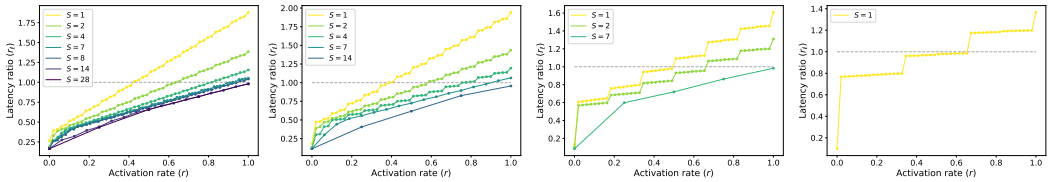

(b) Relationship between r_ℓ and r for LAS-RegNetY-800MF blocks on the Nvidia Jetson TX2 GPU.

Figure 5: Latency prediction results of LAS-ResNet blocks on V100 (a) and LAS-RegNet blocks on TX2 (b). For both networks, we plot the relationship between the latency ratio r_ℓ and the activation rate r for the blocks in 4 stages with the convolutional stride 1. The practical efficiency is only improved when $r_\ell < 1$. Note that $S = 1$ can harm the practical latency even for a small r (reduced computation), while a larger S will alleviate this problem. See detailed analysis in Sec. 4.2.

where τ is the Softmax temperature. Following the common practice [31, 8], we let τ decrease exponentially from 5.0 to 0.1 in training to facilitate the optimization of maskers.

Training objective. The FLOPs of each spatial-wise dynamic convolutional block can be calculated based on our defined activation rate r [31]. Then we can obtain the FLOPs of the overall dynamic network F_{dyn}. Let F_{stat} denotes the FLOPs of its static counterpart. We optimize their ratio to approximate a target $0 < t < 1$:$L_{\text{FLOPs}} = (\frac{F_{\text{dyn}}}{F_{\text{stat}}} - t)^2$. In addition, we define loss item L_{bounds} as in [31] to constrain the upper bound and the lower bound of activation rates in early training epochs.

We further propose to leverage the static counterparts of our dynamic networks as "teachers" to guide the optimization procedure. Let $\mathbf{y}$ and $\mathbf{y}'$ denote the output logits of a dynamic model ("student") and its "teacher", respectively. Our final loss can be written as

$$L = L_{\text{task}} + \alpha(L_{\text{FLOPs}} + L_{\text{bounds}}) + \beta T^2 \cdot \text{KL}(\sigma(\mathbf{y}/T)||\sigma(\mathbf{y}'/T)), \tag{2}$$

where L_{task} represents the task-related loss, *e.g.*, cross-entropy loss in image classification. $\text{KL}(\cdot||\cdot)$ denotes the Kullback–Leibler divergence, and α, β are the coefficients balancing these items. We use σ to denote the log-Softmax function, and T is the temperature for computing KL-divergence.

4 Experiments

In this section, we first introduce the experiment settings in Sec. 4.1. Then the latency of different granularity settings are analyzed in Sec. 4.2. The performance of our LASNet on ImageNet is further evaluated in Sec. 4.3, followed by the ablation studies in Sec. 4.4. Visualization results are illustrated in Sec. 4.5, and we finally validate our method on the object detection task (Sec. 4.6). The results on the instance segmentation task are presented in . For simplicity, we add "LAS-" as a prefix before model names to denote our LASNet, *e.g.*, LAS-ResNet-50.

4.1 Experiment settings

Latency prediction. Various types of hardware platforms are tested, including a server GPU (Tesla V100), a desktop GPU (GTX1080) and edge devices (*e.g.*, Nvidia Nano and Jetson TX2). The major properties considered by our latency prediction model include the number of processing engines (#PE), the floating-point computation in a processing engine (#FP32), the frequency and the bandwidth. It can be observed that the server GPUs generally have a larger #PE than the IoT devices. The batch size is set as 1 for all dynamic models and computing platforms.

Image classification. The image classification experiments are conducted on the ImageNet [4] dataset. Following [31], we initialize the backbone parameter from a pre-trained checkpoint[4], and

[4]We use the torchvision pre-trained models at `https://pytorch.org/vision/stable/models.html`.

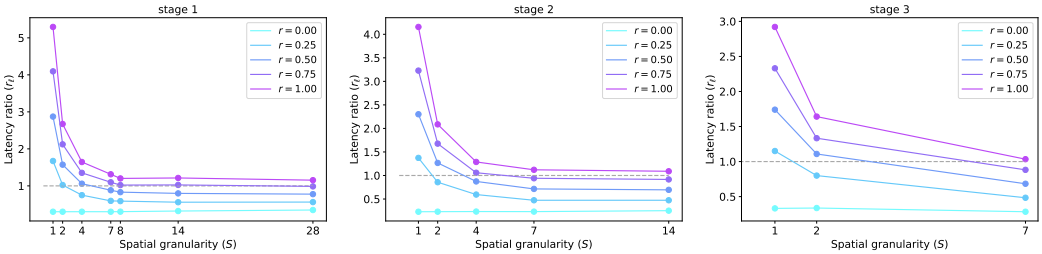

(a) Relationship between r_ℓ and S for LAS-ResNet blocks on V100.

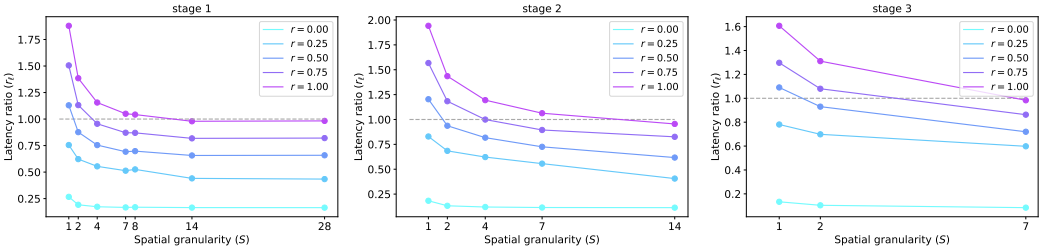

(b) Relationship between r_ℓ and S for LAS-RegNetY-800MF blocks on Nvidia Jetson TX2 GPU.

Figure 6: The relationship between the latency ratio r_ℓ and the spatial granularity S.

finetune the whole network for 100 epochs with the loss function in Eq. (2). We fix $\alpha = 10, \beta = 0.5$ and $T = 4.0$ for all dynamic models. More details are provided in Appendix B.

4.2 Latency prediction results

In this subsection, we present the latency prediction results of the spatial-wise dynamic convolutional blocks in two different models: LAS-ResNet-101 [11] (on V100) and LAS-RegNetY-800MF [25] (on TX2). All the blocks have the bottleneck structure with different channel numbers and convolution groups, and the RegNetY is equipped with Squeeze-and-Excitation (SE) [14] modules.

We first define ℓ_{dyn} as the latency of a spatial-wise dynamic convolutional block, and ℓ_{stat} as that of a static block without a masker. Their ratio is denoted as $r_\ell = \frac{\ell_{\text{dyn}}}{\ell_{\text{stat}}}$. We investigate the relationship between r_ℓ and the activation rate r (cf. Sec. 3.5) for different *granularity* settings. The results in Figure 5 demonstrate that: 1) even equipped with our special optimization on the scheduling strategies, pixel-level spatially adaptive inference (S=1) *cannot always* improve the practical efficiency. Such fine-grained adaptive inference is adopted by most previous works [31, 35], and our result can explain the reason why they can only achieve realistic speedup on less powerful CPUs [35] or specialized devices [3]; 2) a proper granularity $S > 1$ effectively alleviates this problem on both hardware devices. By setting $S > 1$, realistic speedup could be achieved with larger activation rates.

The latency prediction results are further used to seek for a preferable granularity setting for the first 3 stages (we fix $S = 1$ for the last stage, where the feature resolution is 7×7). Therefore, we plot the relationship between r_ℓ and S in Figure 10. It can be observed that: 1) r_ℓ generally decreases with S increasing for a given r; 2) an overly large S (less flexible adaptive inference) brings insignificant improvement on both devices. Especially, enlarging S from 8 to 28 in the first stage of a LAS-ResNet brings very little improvement on V100. Based on the results in Figure 10, we can trade off between flexibility and efficiency by selecting appropriate S for different models and hardware devices. For example, we can simply set S_{net}=8-4-7-1[5] in a LAS-ResNet-101 to achive realistic speedup. The accuracy-latency plots in Figure 7 also validate this observation. More results of our latency prediction model on the desktop-level GPU, GTX1080, are presented in Appendix C.1.

4.3 ImageNet classification results

We now empirically evaluate our proposed LASNet on the ImageNet dataset. The network performance is measured in terms of the trade-off between classification accuracy and inference efficiency. Both theoretical (*i.e.* FLOPs) and practical efficiency (*i.e.* latency) are tested in our experiments.

[5] We use this form to represnet the S settings for the 4 stages of a network.

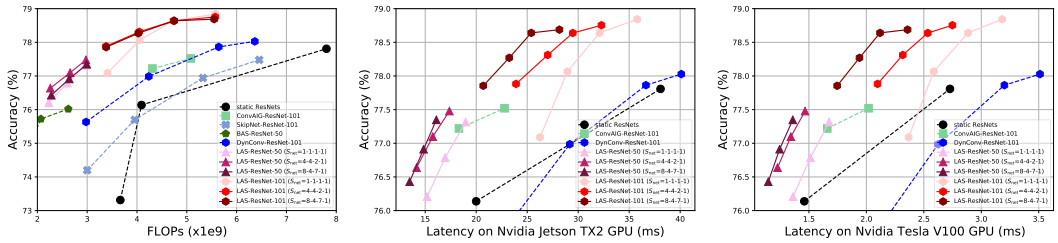

(a) LAS-ResNet results.

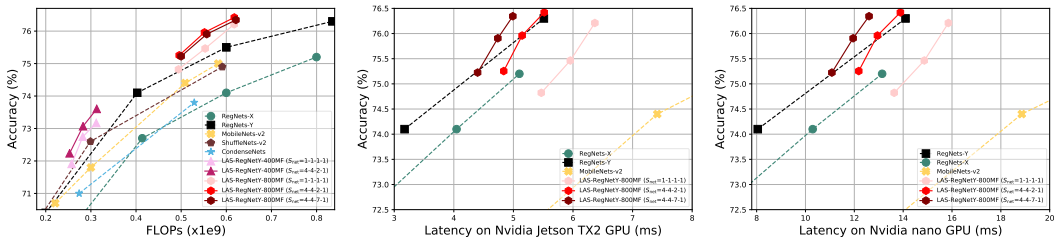

(b) LAS-RegNetY results.

Figure 7: Experimental results on ImageNet. The proposed coarse-grained spatially adaptive inference is tested on standard ResNets (a) and lightweight RegNets (b).

4.3.1 Standard baseline comparison: ResNets

We first establish our LASNet based on the standard ResNets [11]. Specifically, we build LAS-ResNet-50 and LAS-ResNet-101 by plugging our maskers in the two common ResNet structures.

The baselines include various types of dynamic inference approaches: 1) layer skipping (SkipNet [32] and Conv-AIG [30]); 2) channel skipping (BAS [1]); and 3) pixel-level spatial-wise dynamic network (DynConv [31]). For our LASNet, we compare various settings of the spatial granularity S_{net}. We set training targets (cf. Sec. 3.5) $t \in \{0, 4, 0.5, 0.6, 0.7\}$ for our dynamic models to evaluate their performance in different sparsity regimes. We apply the same operator fusion (Sec. 3.4) for both our models and the compared baselines [30, 31] for fair comparison.

Results are presented in Figure 7 (a). On the left we plot the relationship of accuracy *v.s.* FLOPs. It can be observed that our LAS-ResNets with different granularity settings significantly outperform the competing dynamic neural networks. Surprisingly, coarse-grained spatially adaptive inference (S_{net}=4-4-2-1 and S_{net}=8-4-7-1 for the 4 stages) can achieve even higher accuracy when consuming similar FLOPs on ResNets, despite the sacrificed flexibility compared to S_{net}=1-1-1-1. We conjecture that a larger S is also beneficial to the optimization of maskers.

We compare the practical latency of three granularity settings in Figure 7 (a) predicted by our latency prediction model (middle on TX2 and right on V100). We can witness that although they achieve comparable theoretical efficiency (Figure 7 (a) left), larger S is more hardware-friendly compared to the finest granularity. For example, the inference latency of LAS-ResNet-101 (S_{net}=1-1-1-1) is significantly higher than the ResNet-101 baseline on V100 (Figure 7 (a) right), even though its theoretical computation is much smaller than that of the static model. However, larger granularities (S_{net}=4-4-2-1 and S_{net}=8-4-7-1) can effectively improve the inference latency due to its lower burden on the memory access. Remarkably, the latency of ResNet-101 could be reduced by 36% and 46% on V100 and TX2 respectively without sacrificing the accuracy when t=0.4. The classification accuracy is increased by 1.9% with similar inference efficiency. It can be observed that the realistic speedup ratio r_ℓ is more close to the theoretical FLOPs ratio target t on the less powerful TX2, because the latency is *computation-bounded* (*i.e.* the latency is mainly spent on computation) on such IoT devices. In contrast, there is a larger gap between practical and theoretical efficiency on the more powerful V100, as the latency is bounded by the memory access cost.

4.3.2 Lightweight baseline comparison: RegNets

We further evaluate our LASNet in lightweight CNN architectures, *i.e.* RegNets-Y [25]. Two different sized models are tested: RegNetY-400MF and RegNetY-800MF. Compared baselines include other types of efficient models, *e.g.*, MobileNets-v2 [28], ShuffletNets-v2 [23] and CondenseNets [16].

The results are presented in Figure 7 (b). The x-axis for the three sub-figures are the FLOPs, the latency on TX2, and the latency on the Nvidia nano GPU, respectively. We can observe that our

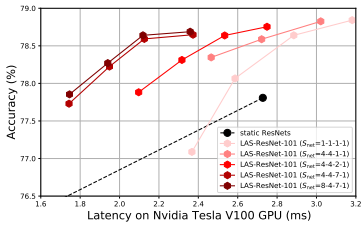

Figure 8: Ablation studies on S.

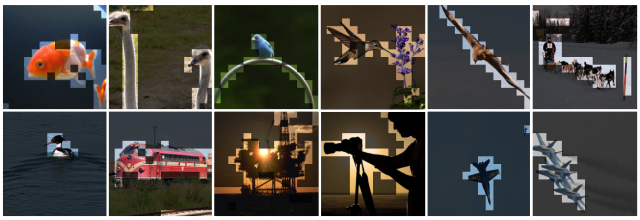

Figure 9: Visualization results.

method outperforms various types of static models in terms of the trade-off between accuracy and efficiency. More results on image classification are provided in Appendix C.2.

4.4 Ablation studies

We conduct ablation studies to validate the effectiveness of our *coarse-grained* spatially adaptive inference (Sec. 3.2) and operator fusion operations (Sec. 3.4).

More granularities settings. We test various granularity settings on LAS-ResNet-101 to examine the effects of S in different stages. The results on the Tesla-V100 GPU are presented in Figure 8. It can be found that the finest granularity (S_{net}=1-1-1-1) leads to substantial inefficiency despite the reduced FLOPs (cf. Figure 7 (a) left). Coarse-grained spatially adaptive inference in the first two stages (S_{net}=4-4-1-1) effectively reduces the inference latency. We further increase S in the third stage to 2 and 7, and this procedure consistently improves the realistic efficiency on the V100 GPU. This trend also holds on the GTX 1080 GPU (see the results in Appendix C.2.

Operator fusion. We investigate the effect of our operator fusion introduced in Sec. 3.4. One convolutional block in the first stage of a LAS-ResNet-101 (S=4, r=0.6) is tested. The results in Table 1 validate that every step of operator fusion benefits the practical latency of a block, as the overhead on memory access is effectively reduced. Especially, the fusion of the masker operation and the first convolution is crucial to reducing the latency.

Table 1: Ablation studies on operator fusion.

Masker-Conv1x1	Gather-Conv3x3	Scatter-Add	Latency (µs)
✗	✗	✗	163.2
✓	✗	✗	90.1
✓	✓	✗	86.7
✓	✓	✓	**71.4**

4.5 Visualization

We visualize the masks generated by our masker in the third block of a LAS-ResNet-101 (S_{net}=4-4-2-1) in Figure 12. The brilliant areas correspond to the locations of 1 elements in a mask, and the computation on the dimmed regions is skipped by our dynamic model. We can witness that the masker accurately locate the most task-related regions (even the tiny aircraft at the corner), which helps reduce the unnecessary computation on background areas. These resultsalso suggest that for the first stage, the granularity S=4 is sufficiently flexible to recognize the important regions, and a *win-win* can be achieved between accuracy and efficiency. Interestingly, the masker could select some objects that are *not labeled* for the sample, *e.g.*, the flower beside the hummingbird and the human holding the camera. This indicates that our spatial-wise dynamic networks can automatically recognize the regions with semantics, and their capability is not limited by the classification labels. This property is helpful in some downstream tasks, such as object detection (Sec. 4.6) and instance segmentation (Sec. 4.7), which require detecting multiple classes and objects in an image. More visualization results could be found in Appendix C.3.

4.6 COCO Object detection

We further evaluate our LASNet on the COCO [22] object detection task. The mean average precision (mAP), the average backbone FLOPs, and the average backbone latency on the validation set are used to measure the network performance. We test two commonly used detection frameworks: Faster R-CNN [27] with Feature Pyramid Network [20] and RetinaNet [21]. Thanks to the generality of our method, we can conveniently replace the backbones with ours pre-trained on ImageNet, and the whole models are finetuned on COCO with the standard setting for 12 epochs (see detailed

Table 2: Object detection results on the COCO dataset.

Detection Framework	Backbone	Backbone FLOPs (G)	Backbone Latency (ms)			mAP (%)
			V100	GTX1080	TX2	
Faster R-CNN	ResNet-101 (Baseline)	141.2	39.5	119.3	729.4	39.4
	LAS-ResNet-101 (S_{net}=4-4-2-1, t=0.6)	90.7	33.8	90.6	524.9	**40.3**
	LAS-ResNet-101 (S_{net}=4-4-2-1, t=0.5)	79.3	30.7	82.5	477.1	39.8
	LAS-ResNet-101 (S_{net}=4-4-7-1, t=0.5)	79.5	29.0	79.9	464.8	40.0
	LAS-ResNet-101 (S_{net}=4-4-7-1, t=0.4)	**67.9**	**25.3**	**69.1**	**401.3**	39.5
RetinaNet	ResNet-101 (Baseline)	141.2	39.5	119.3	729.4	38.5
	LAS-ResNet-101 (S_{net}=4-4-2-1, t=0.5)	77.8	30.4	81.8	472.6	**39.3**
	LAS-ResNet-101 (S_{net}=4-4-7-1, t=0.5)	79.4	28.9	79.9	464.8	**39.3**
	LAS-ResNet-101 (S_{net}=4-4-7-1, t=0.4)	**66.4**	**25.3**	**69.1**	**401.3**	38.9

Table 3: Instance Segmentation results on the COCO dataset.

Segmentation Framework	Backbone	Backbone FLOPs (G)	Backbone Latency (ms)			AP^{mask} (%)	AP^{box} (%)
			V100	GTX1080	TX2		
Mask R-CNN	ResNet-101 (Baseline)	141.2	39.5	119.3	729.4	36.1	40.0
	LAS-ResNet-101 (S_{net}=4-4-2-1, t=0.5)	80.5	31.1	83.3	481.9	**37.0**	**41.0**
	LAS-ResNet-101 (S_{net}=4-4-2-1, t=0.4)	69.2	27.9	74.8	431.6	36.1	40.0
	LAS-ResNet-101 (S_{net}=4-4-7-1, t=0.4)	**68.8**	**25.8**	**70.9**	**411.8**	36.2	40.0

setup in Appendix B.3). The input images are resized to a short side of 800 and a long side not exceeding 1333. The results of our LAS-ResNet-101 with different S_{net} settings are presented in Table 2. We can observe from the results that when setting the training target as 0.4, the latency of our LAS-ResNet-101 with S_{net}=4-4-7-1 is significantly lower than the static baseline on all devices without sacrificing mAP in both detection frameworks. With a larger training targets, our LASNet can increase the mAP by 0.9% and 0.8% in Faster R-CNN and RetinaNet respectively, while still being faster than the baseline method.

4.7 COCO instance segmentation

We also present the results of instance segmentation on COCO, which demonstrate the effectiveness of our LASNet on the dense prediction task. From the results in Table 3, we can observe that when setting the training target as 0.4, the Mask R-CNN [10] models (S_{net}=4-4-2-1 and S_{net}=4-4-7-1) runs faster on all tested hardware devices without sacrificing the performance. With a training target of 0.5, the AP^{mask} and AP^{box} of the Mask R-CNN model could be increased by 0.9% and 1.0% respectively while still running faster than the baseline method.

5 Conclusion

In this paper, we propose to build *latency-aware* spatial-wise dynamic networks (LASNet) under the guidance of a *latency prediction model*. By simultaneously considering the algorithm, the scheduling strategy and the hardware properties, we can efficiently estimate the practical latency of spatial-wise dynamic operators on arbitrary computing platforms. Based on the empirical analysis on the relationship between the latency and the *granularity* of spatially adaptive inference, we optimize both the algorithm and the scheduling strategies to achieve realistic speedup on many multi-core processors, *e.g.*, the Tesla V100 GPU and the Jetson TX2 GPU. Experiments on image classification, object detection and instance segmentation tasks validate that the proposed method significantly improves the practical efficiency of deep CNNs, and outperforms various competing approaches.

Acknowledgement

This work is supported in part by the National Key R&D Program of China under Grant 2020AAA0105200, the National Natural Science Foundation of China under Grants 62022048 and the Tsinghua University-China Mobile Communications Group Co.,Ltd. Joint Institute..

References

[1] Babak Ehteshami Bejnordi, Tijmen Blankevoort, and Max Welling. Batch-shaping for learning conditional channel gated networks. In *ICLR*, 2020.

[2] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. Tvm: An automated end-to-end optimizing compiler for deep learning. In *OSDI*, 2018.

[3] Steven Colleman, Thomas Verelst, Linyan Mei, Tinne Tuytelaars, and Marian Verhelst. Processor architecture optimization for spatially dynamic neural networks. In *VLSI-SoC*, 2021.

[4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.

[5] Xuanyi Dong, Junshi Huang, Yi Yang, and Shuicheng Yan. More is less: A more complicated network with less inference complexity. In *CVPR*, 2017.

[6] Michael Figurnov, Maxwell D Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, and Ruslan Salakhutdinov. Spatially adaptive computation time for residual networks. In *CVPR*, 2017.

[7] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. Dynamic neural networks: A survey. *TPAMI*, 2021.

[8] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Yitian Zhang, and Haojun Jiang. Spatially adaptive feature refinement for efficient inference. *TIP*, 2021.

[9] Yizeng Han, Yifan Pu, Zihang Lai, Chaofei Wang, Shiji Song, Junfen Cao, Wenhui Huang, Chao Deng, and Gao Huang. Learning to weight samples for dynamic early-exiting networks. In *ECCV*, 2022.

[10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

[11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[12] John L Hennessy and David A Patterson. *Computer architecture: a quantitative approach*. Elsevier, 2011.

[13] 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.

[14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.

[15] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Weinberger. Multi-scale dense networks for resource efficient image classification. In *ICLR*, 2018.

[16] Gao Huang, Shichen Liu, Laurens Van der Maaten, and Kilian Q Weinberger. Condensenet: An efficient densenet using learned group convolutions. In *CVPR*, 2018.

[17] Gao Huang, Zhuang Liu, Geoff Pleiss, Laurens Van Der Maaten, and Kilian Weinberger. Densely connected convolutional networks. *TPAMI*, 2019.

[18] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017.

[19] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In *NeurIPS*, 2017.

[20] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.

[21] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.

[22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.

[23] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, 2018.

[24] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*, 2017.

[25] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, 2020.

[26] Mengye Ren, Andrei Pokrovsky, Bin Yang, and Raquel Urtasun. SBNet: Sparse Blocks Network for Fast Inference. *CVPR*, 2018.

[27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.

[28] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018.

[29] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019.

[30] Andreas Veit and Serge Belongie. Convolutional networks with adaptive inference graphs. In *ECCV*, 2018.

[31] Thomas Verelst and Tinne Tuytelaars. Dynamic convolutions: Exploiting Spatial Sparsity for Faster Inference. In *CVPR*, 2020.

[32] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *ECCV*, 2018.

[33] Yulin Wang, Kangchen Lv, Rui Huang, Shiji Song, Le Yang, and Gao Huang. Glance and focus: a dynamic approach to reducing spatial redundancy in image classification. In *NeurIPS*, 2020.

[34] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10734–10742, 2019.

[35] Zhenda Xie, Zheng Zhang, Xizhou Zhu, Gao Huang, and Stephen Lin. Spatially adaptive inference with stochastic feature sampling and interpolation. In *ECCV*, 2020.

[36] Le Yang, Yizeng Han, Xi Chen, Shiji Song, Jifeng Dai, and Gao Huang. Resolution adaptive networks for efficient inference. In *CVPR*, 2020.

[37] Mingjian Zhu, Kai Han, Enhua Wu, Qiulin Zhang, Ying Nie, Zhenzhong Lan, and Yunhe Wang. Dynamic resolution network. *NeurIPS*, 2021.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section 4.1 and Section 4.6, the public ImageNet and COCO datasets are cited.

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Appendix D.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Appendix D.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See the abstract.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Most important hyperparameters are given in the main paper, and more details are presented in the supplementary material (Appendix B).
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Since the experiment results on the large-scale dataset, ImageNet and COCO, are generally stable, we perform several repeat experiments in the early exploring phase, and dismiss the error bars in our main results as in most previous works.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] The details are presented in the supplementary material (Appendix B).
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes] See Section 4.1. We use the torchvision pre-trained models.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendix

A Latency prediction model.

As the dynamic operators in our method have not been supported by current deep learning libraries, we propose a latency prediction model to efficiently estimate the real latency of these operators on hardware device. The inputs of the latency prediction model include: 1) the structural configuration of a convolutional block, 2) its activation rate r which decides the computation amount, 3) the spatial granularity S, and 4) the hardware properties mentioned in Table 4. The latency of a dynamic block is predicted as follows.

Input/output shape definition. The first step of predicting the latency of an operation is to calculate the shape of input and output. Taking the gather-conv2 operation as an example, the input of this operation is the activation with the shape of $C_{in} \times H \times W$, where C_{in} is the number of input channels, and H and W are the resolution of the feature map. The shape of the output tensor is $P \times C_{out} \times S \times S$, where P is the number of output patches, C_{out} is the number of output channels and S is the spatial granularity. Note that P is obtained based on the output of our maskers.

Operation-to-hardware mapping. Next, we map the operations to hardware. As is mentioned in the paper, we model a hardware device as multiple processing engines (PEs). We assign the computation of each element in the output feature map to a PE. Specifically, we consecutively split the output feature map into multiple *tiles*. The shape of each tile is $T_P \times T_C \times T_{S1} \times T_{S2}$. These split tiles are assigned to multiple PEs. The computation of the elements in each tile is executed in a PE. We can configure different shapes of tiles. In order to determine the optimal shape of the tile, we make a search space of different tile shapes. The tile shape has 4 dimensions. The candidates of each dimension are power-of-2 and do not exceed the corresponding dimension of the feature map.

Latency estimation. Then, we evaluate the latency of each tile shape in the search space and select the optimal tile shape with the lowest latency. The latency includes the *data movement* latency and the *computation* latency: $\ell = \ell_{\text{data}} + \ell_{\text{computation}}$.

1) *Data movement latency* ℓ_{data}. The estimation of the latency for data movement requires us to model the memory system of a hardware device. We model the memory system of hardware as a three-level architecture [12]: off-chip memory, on-chip global memory, and local memory in PE. The input data and weight data are first transferred from the off-chip memory to the on-chip global memory. We assume the hardware can make full use of the off-chip memory bandwidth to simplify the latency prediction model.

After that, the data used to compute the output tiles is transferred from on-chip global memory to the local memory of each PE. The latency of data movement to local memory is estimated by its *bandwidth* and *efficiency*. We assume each PE only moves the corresponding input feature maps and weights once to compute a output tile so as to simplify the prediction model. The input data movement latency ℓ_{in} is calculated by adding the time from off-chip memory to on-chip global memory and the time from on-chip global memory to local-memory together: $\ell_{\text{in}} = \ell_{\text{off2on}} + \ell_{\text{global2local}}$. Contrary to the input data, the output data ℓ_{out} are moved from local memory to on-chip global memory and then to off-chip memory: $\ell_{\text{out}} = \ell_{\text{local2global}} + \ell_{\text{on2off}}$. We calculate the total data movement latency by adding the input and output data movement latency together: $\ell_{\text{data}} = \ell_{\text{in}} + \ell_{\text{out}}$.

The latency of data movement is affected by the granularity S: when the granularity S is small, the same input data has a higher probability of being sent to multiple PEs to compute different output patches, which significantly increases the number of on-chip memory movement. And due to the small amount of data transmitted each time and the data is randomly distributed, the efficiency of data movement will be low. This accounts for our experiment results in the paper that *a larger S will effectively improve the practical efficiency*.

2) *Computation latency* $\ell_{\text{computation}}$. The computation latency of each tile is estimated using the PE's *maximum throughput of FP32 computation* and the *FLOPs* of computing an output tile. The total computation latency can be obtained according to the number of tiles and the number of PEs.

To summarize, our latency prediction model can predict the real latency of dynamic operators by considering both the *data movement* cost and the *computation* cost. Guided by the latency prediction model, we propose our LASNets with coarse-grained spatially adaptive inference ($S > 1$). It is validated in our paper that LASNets achieve better efficiency than previous approaches [35, 31]

($S = 1$), as it effectively reduces the data movement latency, which is rarely considered by other researchers.

B Detailed experimental settings

In this section, we present the detailed experiment settings which are not provided in the main paper due to the page limit.

B.1 Latency prediction

Hardware properties considered by our latency prediction model include the number of processing engines (#PE), the floating-point computation in a processing engine (#FP32), the frequency and the bandwidth. We test four types of hardware devices, and their properties are listed in Table 4.

Table 4: Hardware properties.

Name	#PE	#FP32	frequency (MHz)	bandwidth (G)
Nvidia Tesla V100	80	64	1500	700
Nvidia GTX1080	20	64	1700	320
Nvidia Jetson TX2	2	128	1300	59.7
Nvidia Nano	1	128	921	25.6

It could be found that the server GPU V100 is the most powerful hardware device, especially with the most number of processing engines (#PE). Therefore, spatially adaptive inference could easily fall into a *memory-bounded* operation on V100 due to its high parallelism. Our experiment results in Figure 7 (a) and Figure 8 in the paper can reflect this phenomenon: the more flexibility the computation is, the harder to improve the practical efficiency.

In contrast, on the less powerful computing devices such as the IoT devices, the real acceleration is close to the theoretical effect (compare Figure 7 (a) left with Figure 7 (a) middle).

Operator fusion.

1) *Fusing the masker and the first convolution.* We mentioned in Sec. 3.4 of the paper that the masker operation is fused with the first 1×1 convolution in a block to reduce the cost on memory access. This is feasible because the two operators share the same input feature, and their convolutional kernel sizes are both 1×1.

Note that during the inference stage, we only need to perform $\arg\max$ along the channel dimension of a mask $\mathbf{M} \in \mathbb{R}^{2\times H\times W}$ to obtain the positions of the gathered pixels. Therefore, we can reduce the output channel number of our maskers from 2 to 1 since the convolution is a linear operation:

$$[\mathbf{x} * \mathbf{W}]_{:,:,0} > [\mathbf{x} * \mathbf{W}]_{:,:,1} \Longleftrightarrow \mathbf{x} * (\mathbf{W}_{:,:,0} - \mathbf{W}_{:,:,1}) > 0. \tag{3}$$

Afterwards, we fuse the masker with the first convolution layer by performing once convolution whose output channel number is $C + 1$, where C is the original output width of the first convolution. The output of this step is split into a feature map (for further computation) and a mask (for obtaining the index for gathering). Such operator fusion avoids the repeated reading the input feature, and helps reduce the inference latency (see Table 1 in the paper).

2) *Fusing the gather operation and the dynamic convolution.* To facilitate the scheduling on hardware devices with multiple PEs, the masker generates the indices of activated patches instead of sparse mask at inference time. In this way, it is easy to evenly distribute the computation of output patches to different PEs, thus avoiding unbalanced computation of PEs. Each element in the indices represents the index of an activated patch. PE fetches the input data from the corresponding positions on the feature map according to the index. The output patches could be densely stored in memory. Such operator fusion benefits the contiguous memory access and parallel computation on multiple PEs.

3) *Fusing the scatter operation and the add operation.*

Similar to the previous operation, each PE fetches a tile of data from the residual feature map according to the index, adds them with the corresponding feature map from previous dynamic

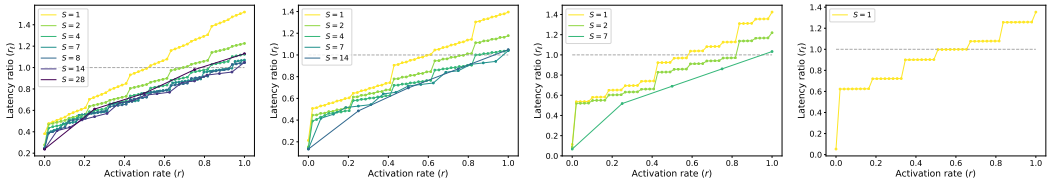

(a) Relationship between r_ℓ and r for LAS-ResNet blocks on Nvidia GeForce GTX1080.

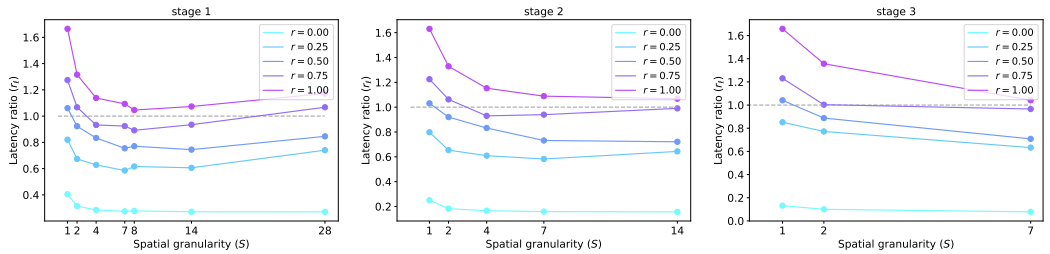

(b) Relationship between r_ℓ and S for LAS-ResNet blocks on Nvidia GeForce GTX1080.

Figure 10: Latency prediction results of LAS-ResNet blocks on the Nvidia GTX1080 GPU.

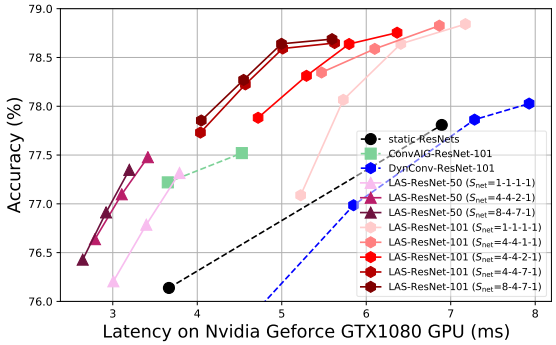

Figure 11: Experimental results on the ImageNet classification task.

convolution, and then stores the results to the corresponding position on the residual feature map according to the index. This optimization can significantly reduce the costs on memory access.

Speed test. We test the latency on real hardware devices to evaluate the accuracy of our latency prediction model. On GPUs, we use TVM [2] and CUDA (version 11.6) for code generation and compilation respectively. The results in Fig. 4 of the paper validate the effectiveness of our model.

B.2 ImageNet classification

As mentioned in the paper, we use pre-trained CNN models in the official torchvision website to initialize our backbone parameters, and finetune the overall models for 100 epochs. The initial learning rate is set as $0.01\times$batch size/128, and decays with a cosine shape. The training batch size is determined on the model size and the GPU memory. For example, we train our LAS-ResNet-101 on 8 RTX 3090 GPUs with the batch size of 512, and the batch size for LAS-ResNet-50 is doubled. We use the same weight decay and the standard data augmentation as in the RegNet paper [25]. For our own hyper-parameter τ in Eq. (1) of the paper, this Gumbel temperature τ exponentially decreases from 5 to 0.1 in the training procedure. For the training hyper-parameter in Eq. (2), we simply fix $\alpha = 10, \beta = 0.5$ and $T = 4.0$ for all dynamic models. We conduct a very simple grid search with a RegNet for $\beta \in \{0.3, 0.5\}$ and $T \in \{1.0, 4.0\}$ to determine their values.

Figure 12: Visualization results.

B.3 COCO object detection & instance segmentation

We use the standard setting suggested in [20, 21, 10], except that we decrease the learning rate for our pre-trained backbone network. We simply set a learning rate multiplier 0.5 for Faster R-CNN [27], 0.2 for RetinaNet [21] and 0.5 for Mask R-CNN [10]. As for the additional loss items, the hyper-parameters are kept the same as training our classification models, except that the temperature is fixed as 0.1 in the 12 training epochs.

C More experimental results

In this section, we report more experimental results which are not presented in the main paper.

C.1 Latency prediction

In Figure 5 and Figure 6 of the paper, we report the latency prediction results of LAS-ResNet on V100 and TX2. Here we present the results on GTX1080 (Figure 10). It can be found that S_{net}=8-4-7-1 (which is the same to the optimal setting on V100 in the paper) will lead to faster inference on GTX1080. This is reasonable since GTX1080 has a large #PE than those IoT devices, and requires more contiguous memory access (a larger S for coarse-grained spatially adaptive inference) to achieve realistic speedup ($r_\ell < 1$). The accuracy-latency curves in Figure 11 further validate this observation.

C.2 ImageNet classification

Results on GTX1080 GPU. In Figure 7 of the paper, we report the ImageNet classification results of LAS-ResNets on V100 and TX2. Here we present the results of LAS-ResNet on GTX1080 (Figure 11). From the results in Figure 11, we can get the same conclusion as in the paper, that a proper large S is more parallel-friendly than the finest granularity. Remarkably, the latency of ResNet-101 could be reduced by 41% without sacrificing the accuracy on GTX1080 under the spatial granularity setting of S_{net}=8-4-7-1. With a similar inference latency, the accuracy of a static ResNet-101 could be increased by 1.0% by our LAS-ResNet-101 (S_{net}=4-4-2-1).

Results of the extreme situation of S_{net}=56-28-14-7. We mentioned in the paper that when we set S_{net}=56-28-14-7, the spatially adaptive inference paradigm will reduce into layer skipping. We experiment on ResNet-50, and find that although being slightly faster (13ms) than our LAS-ResNet-50 (S_{net}=4-4-2-1, 16ms), the accuracy of LAS-ResNet-50 with S_{net}=56-28-14-7 could be significantly degraded from 76.6% (ours, S_{net}=4-4-2-1) to 76.1%. Therefore, we mainly focus on the discussion of spatially adaptive inference in this paper.

C.3 More visualization results

In addition to Figure 9 in the paper, here we present more visualization results of the regions selected by our masker in the 3-rd block of a LAS-ResNet-101 (S_{net}=4-4-2-1) in Figure 12, which demonstrate that our spatially adaptive inference paradigm can effectively locate the most task-related areas in image features, and reduce the unnecessary computation on those background areas.

D Limitations

The current limitations of our work include the following aspects:

1) the latency-ware co-designing framework is only constructed for spatial-wise dynamic networks. Support for more types of dynamic models (*e.g.* channel skipping) will be explored in the future;

2) To achieve faster inference, the spatial masks are defined the same for all input/output channels, which might limit the flexibility of adaptive inference. Future work may explore more flexible forms of dynamic computation;

3) The combination with other acceleration techniques such as Winograd, and the implementation on more CNN backbones may be worth studying in the future.

Social impact. Our work can help reduce the inference cost of deep CNNs, but the training of our models might potentially increase the carbon emissions.