

Med-DANet V2: A Flexible Dynamic Architecture for Efficient Medical Volumetric Segmentation

Haoran Shen^{1,*} Yifu Zhang^{1,*} Wenxuan Wang^{1,*} Chen Chen²
 Jing Liu³ Shanshan Song¹ Jiangyun Li^{1†}

¹School of Automation and Electrical Engineering, University of Science and Technology Beijing

²Center for Research in Computer Vision, University of Central Florida

³National Lab of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences

m202220738@xs.ustb.edu.cn, chen.chen@crcv.ucf.edu, lee.jy@ustb.edu.cn

Abstract

Recent works have shown that the computational efficiency of 3D medical image (e.g. CT and MRI) segmentation can be impressively improved by dynamic inference based on slice-wise complexity. As a pioneering work, a dynamic architecture network for medical volumetric segmentation (i.e. Med-DANet [44]) has achieved a favorable accuracy and efficiency trade-off by dynamically selecting a suitable 2D candidate model from the pre-defined model bank for different slices. However, the issues of incomplete data analysis, high training costs, and the two-stage pipeline in Med-DANet require further improvement. To this end, this paper further explores a unified formulation of the dynamic inference framework from the perspective of both the data itself and the model structure. For each slice of the input volume, our proposed method dynamically selects an important foreground region for segmentation based on the policy generated by our Decision Network and Crop Position Network. Besides, we propose to insert a stage-wise quantization selector to the employed segmentation model (e.g. U-Net) for dynamic architecture adapting. Extensive experiments on BraTS 2019 and 2020 show that our method achieves comparable or better performance than previous state-of-the-art methods with much less model complexity. Compared with previous methods Med-DANet and TransBTS with dynamic and static architecture respectively, our framework improves the model efficiency by up to nearly 4.1 and 17.3 times with comparable segmentation results on BraTS 2019. Code will be available at <https://github.com/Rubics-Xuan/Med-DANet>.

1. Introduction

As one of the most prevalent diseases, cancer results in numerous fatalities annually. The precise measurements of medical images play a vital role in accurate diagnosis and appropriate therapy planning. Traditionally, these image analysis approaches rely heavily on the doctors' clinical experience. However, it is labor-intensive and time-consuming, since a 3D volume produced by Magnetic Resonance Imaging (MRI) [22] or Computerized Tomography (CT) [18] typically contains hundreds of 2D slices. Therefore, to improve the accuracy and efficiency of clinical diagnosis, automated and accurate segmentation of tumors and organs' sub-regions is a fundamental requirement for medical image analysis.

Thanks to the rapid development of deep neural networks, they have been extensively applied in medical image segmentation. The mainstream segmentation methods of medical images comprise two categories: (1) applying 2D networks for slice-by-slice predictions and (2) utilizing 3D models to process image volumes consisting of multiple slices. 2D U-Net [41] and its variants such as [30, 51] are the representatives of the former category, while 3D networks like 3D U-Net [13] and V-Net [35] can achieve better results owing to the associations modeling capability between different slices. Besides, as a transition from 2D methods to 3D methods, some 2.5D approaches [21, 25, 48] have combined information from neighboring slices to achieve better segmentation results in the current slice. Benefiting from the ability to capture long-range dependencies, many Transformer-based networks [5, 16, 17, 37] have begun to spring up. Nevertheless, due to the growing model scale and unique U-shaped encoder-decoder design, the high computational cost is difficult to bear in practice, especially for those 3D models.

Therefore, lightweight medical image segmentation models [6, 10, 27, 32] have become one of the research hot spots. Generally, those lightweight methods only focus on efficient structural design, ignoring the unique slice heterogene-

*Equal Contribution. †Corresponding author.




Figure 1. The comparison between the previous dynamic network Med-DANet and our proposed Med-DANet V2 (Ours).

ity property of the medical images. As a pioneering work of applying dynamic inference to achieve efficient medical volumetric segmentation, the recently proposed dynamic architecture network namely Med-DANet [44] reveals its potential for a promising trade-off between accuracy and efficiency. Specifically, Med-DANet employs a decision network to distinguish the specific segmentation difficulty of the current input slice and accordingly selects a suitable candidate network from the pre-defined model bank for inference. In particular, the training of the decision network is supervised by a comprehensive choice metric, which is computed according to the accuracy and complexity of the incorporated candidate models. Nevertheless, as illustrated in Fig. 1, Med-DANet has mainly three drawbacks: (1) Med-DANet lacks comprehensive dynamic and efficient design considering the medical data properties. Although Med-DANet chooses to directly skip to generate the segmentation mask that is filled with only background category, it ignores the spatial redundancy in those simple images (i.e., reaching precise foreground segmentation is feasible utilizing only a portion of the image); (2) Med-DANet needs to train multiple models as candidate networks in the pre-defined model bank, which results in high training cost and limited scalability, and it also doesn't consider further lightweight design of each candidate network. (3) As the decision network requires a pre-calculated comprehensive choice metric for supervision, Med-DANet must be trained with a complicated two-stage pipeline, which is unfriendly for practical deployment.

In this paper, we further explore the potential of dynamic inference in medical volumetric segmentation tasks. To the best of our knowledge, we are the first to unify the spatial-wise dynamic adaption and dynamic model quantization to handle the slice heterogeneity problem in MRI/CT data. Compared to Med-DANet, a more efficient end-to-end framework named Med-DANet V2 is proposed to get closer to clinical applications. To achieve more fine-grained spatial-wise dynamic adaption, we minimize input image redundancy through dynamic resolution selection and apposite crop position determination, thereby reducing the computational complexity of the entire architecture. For more flexible

dynamic architecture adaption, a stage-wise bit selection module is inserted into the segmentation model for dense model quantization, pursuing an extremely efficient model structure. Besides, our proposed Med-DANet V2 has strong compatibility and scalability. The basic segmentation model can be replaced with any 2D network for various accuracy and efficiency requirements.

In summary, the main contributions are as follows:

- Aiming at the slice heterogeneity in medical volumes, this paper explores the potential of dynamic inference from the perspective of data properties and model structure, resulting in a promising trade-off between volumetric segmentation accuracy and efficiency.
- By introducing the proposed decision network and bit-width selector, we unify the spatial dynamic adaption and dynamic model quantization to realize a highly efficient medical image segmentation task.
- Our proposed Med-DANet V2 exhibits strong compatibility and scalability. The framework can be compatible with any 2D segmentation network to meet various accuracy and efficiency requirements.
- Extensive experiments on the MRI benchmark datasets (BraTS 2019 and BraTS 2020 for brain tumor segmentation) demonstrate that our method reaches competitive or better performance than previous state-of-the-art methods with much less model complexity.

2. Related Work

Static and Lightweight Methods for Medical Image Segmentation. Recently, U-Net [41] and its variants [5, 30, 43, 51] have shown promising performance in medical image segmentation task. However, the high computational complexities of these models pose a great challenge for timely clinical diagnosis. Consequently, great efforts have been dedicated to designing lightweight networks for more efficient segmentation. For instance, S3D-UNet [10] leverages separable 3D convolution to improve model efficiency. DMFNet [6] takes advantage of a 3D dilated multi-fiber network to achieve the trade-off between model efficiency and

accuracy in 3D MRI brain tumor segmentation tasks. HDC-Net [32] replaces 3D convolutions with a novel hierarchical decoupled convolution module to achieve a lightweight yet efficient pseudo-3D model.

Spatial-wise Dynamic Networks. Most spatial-wise dynamic methods are designed for classification tasks [8, 28, 45, 47], and only a few works pursue to implement the dynamic inspiration in fine-grained segmentation tasks. Huang et al. [19] employ a two-stage method. Input images are quickly processed through a lightweight model, and then uncertainty regions are re-evaluated using a more complex network. Li et al. [26] introduce data-dependent routes to adapt to the scale distribution of each image. Verelst et al. [42] propose a dynamic block-based framework, where image blocks are downsampled based on their complexity. Although these spatial-wise dynamic methods on natural images adopt diverse strategies for different regions, the entire input is essentially segmented. However, the characteristics of having a large area of background regions (even pure background) in 3D medical data determine that there is great potential in segmenting mainly the foreground regions, which remains unexplored by previous researchers.

Network Quantization. As an effective method for model compression, quantization reduces the model size and memory requirements, thereby accelerating inference. Although the quantization technique in the natural image field has attracted widespread attention in the community, its potential in medical image analysis has rarely been explored. Additionally, the utilization of low bit-width methods for network quantization leads to significant performance degradation [1, 12, 20, 49]. Thus, some mixed-precision bit-width methods [11, 15, 31, 46] have become popular to improve quantization for a better trade-off. Although applying this method to medical image segmentation models can reduce the varying redundancy of different layers, it cannot adaptively allocate appropriate quantization bit widths for specific medical slices. Therefore, dynamic quantization in the medical field remains challenging and to be exploited, especially for medical image segmentation tasks that require separating tissues and organs with high precision.

3. Methodology

3.1. Preliminary: Designing Details of Med-DANet

Med-DANet [44] is a dynamic architecture network that aims to handle 3D medical images where segmentation targets are sparsely distributed among slices. Given a 3D volume, a slice-specific decision is learned by the Decision Network to dynamically select a suitable model from the pre-defined model bank for subsequent segmentation. The Decision Network undertakes a $n + 1$ -class classification task, where the $n + 1$ categories represent n candidate networks and a skip process. The Decision Network and model

bank are respectively formulated as follows

$$\mathcal{D}(x) = \{\hat{D}|x; \theta\}, \quad (1)$$

$$\mathcal{B}(x) = [\emptyset, M_1(x), M_2(x), \dots, M_n(x)], \quad (2)$$

where θ denotes the parameters of Decision Network and \hat{D} is the prediction of $\mathcal{D}(x)$. $M_1 \sim M_n$ indicate the candidate models and (\emptyset) represents the skip operation.

During the training process, Med-DANet proposes a comprehensive choice metric to supervise the framework, which is computed by the accuracy and complexity of the incorporated models. The metric is calculated as follows.

$$D = \begin{cases} 0, & P_f < 1 \\ \text{argmax}((1 - \alpha) * S_i + \alpha * \text{softmax}(\frac{1}{F_i})) + 1, & P_f \geq 1 \end{cases}, \quad (3)$$

where S_i and F_i are respectively the Dice Score and FLOPs of each candidate model M_i . P_f denotes the number of foreground target pixels.

Limitations of Med-DANet. (1) Although Med-DANet selects models of corresponding scales based on the difficulty of slices, it ignores handling simple images with various background regions from the input perspective. (2) The model bank with multiple models must be pre-established, which requires a high training cost. (3) The pre-calculated comprehensive choice metric requires Med-DANet to be trained using a complex two-stage pipeline, which may pose challenges for users.

Given these limitations, a more flexible and unified dynamic one-stage model architecture named Med-DANet V2 is developed. In terms of the input data, more fine-grained spatial-wise dynamic adaption will be determined by two decision modules. As for the model capacity, the bit-width quantization will be considered to achieve more flexible dynamic architecture adaption.

3.2. Overall Architecture

The overall architecture of the proposed Med-DANet V2 is depicted in Fig. 2. In general, our framework consists of the data decision part (i.e. Policy Network (P) and Crop Position Network (C)) and the Dynamic Quantization Network (Q). To deal with the medical datasets where segmentation targets are sparsely distributed among inter and intra-slices, a slice-specific decision dynamically selects a suitable crop size (i.e. obtained by P) and a crop-center (i.e. obtained by C) to locate the cropping image. Then the cropped image will be segmented by the Dynamic Quantization Network, which can automatically allocate the bit-width inside the model to achieve different model scales. Dynamic choice of architecture and the subsequent segmentation task are formulated by Eq. 4:

$$y = Q[P(x) \circ C(x)], \quad (4)$$

where x denotes the input image and y is the corresponding prediction. $Q[P \circ C]$ indicates feeding the slice of a dy-

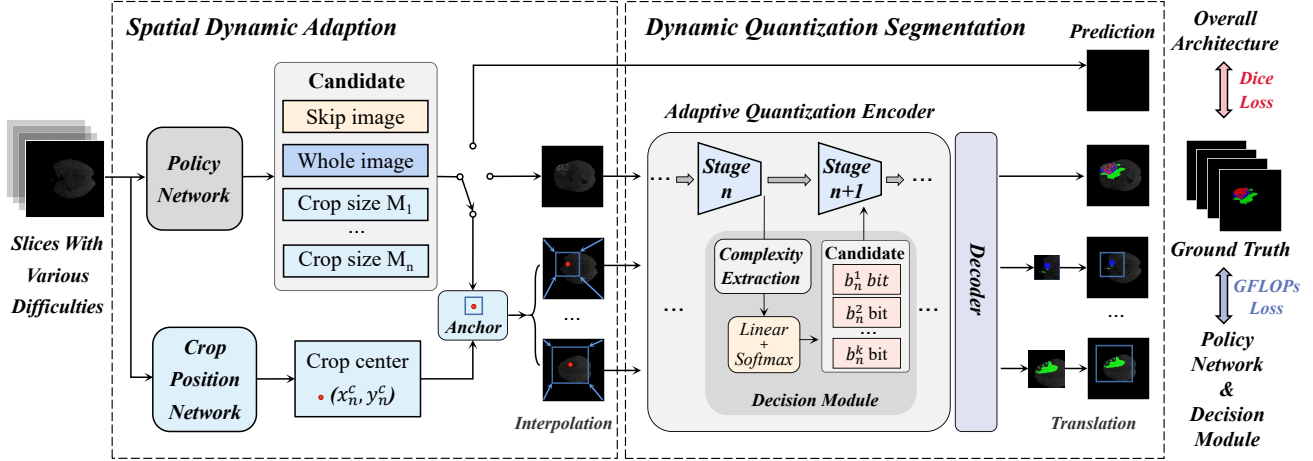


Figure 2. The illustration of our Med-DANet V2. The Policy Network takes a 2D image slice as input and generates a choice depending on the segmentation difficulty of the current slice. Based on the optimal choice made by the Policy Network, our method can adaptively decide whether to skip the current slice (i.e. directly generating the result with all zero – “background” class) or send the input with a suitable resolution to the Dynamic Quantization Model (with the cropping region determined by the Crop Position Network). A stage-wise Decision Module is inserted into the Dynamic Quantization Model for suitable model capacity selection. By unifying data-architecture dynamic inference, our method can achieve accurate and efficient segmentation.

dynamic crop into the dynamic quantization model. In the rest of this section, the designing details about Spatial Dynamic Adaption (Sec. 3.3), Dynamic quantization (Sec. 3.4), and Training/Inference Strategy (Sec. 3.5) will be explained. As a whole, the Policy Network combined with the Crop Position Network will comprehensively consider the foreground and background regional distribution inter and intra slice, which can make the most appropriate crop choice (i.e. containing skip). Regarding the Dynamic Quantization Model, any 2D network can be tailored as the foundation to flexibly accommodate diverse accuracy and efficiency demands. More discussions on the model choices and ablation study are presented in Sec. 4.3.

3.3. Fine-grained Spatial-wise Dynamic Adaption

Due to the large amount of redundant background regions in some simple medical image slices, feeding only the foreground part into the network is sufficient for obtaining good results. Besides, minimizing the proportion of the meaningless regions can make the segmentation model more focused on significant feature extraction (e.g. features of crucial tumor regions), accordingly achieving better performance. To achieve this, two decision networks are introduced to determine the resolution size and center point of the cropped image tile for the following segmentation. Then the selected slice region based on the decisions is sent to the segmentation network to obtain the final result.

Following Med-DANet, the high-efficiency ShuffleNet V2 [33] is selected as our Policy Network and Crop Position Network. Hence, only negligible computation cost is added to the entire framework. The Policy Network acts as a classifier to distinguish $n + 1$ categories, which refer

to n candidate resolutions and a direct skip operation. Besides, we utilize Gumbel-Softmax Sampling [23] strategy to enable the network optimization process through standard back-propagation. Specifically, when choosing the skip operation, the whole architecture will skip the subsequent segmentation process and produce the segmentation maps with the background class. If the choice of full-resolution image is made, the input slice will be straightly fed into the Dynamic Quantization Model for final prediction. Otherwise, the Crop Position Network comes into effect and undertakes a regression task to determine the center point for cropping, which is used before and after the segmentation model for resolution downsampling and restoration.

For the process of altering resolution, we follow AdaFocus V2 [45] to use differentiable bilinear interpolation for downsampling. Assuming the sizes of the whole input and the cropped image are $H \times W$ and $P_n \times P_n$. Each pixel (x_n, y_n) in the cropped patch can be expressed by the center position (x_n^c, y_n^c) and a fixed offset o_{ij} :

$$(x_n, y_n) = (x_n^c, y_n^c) + o_{ij}, o_{ij} \in \left\{ -\frac{P_n}{2}, -\frac{P_n}{2} + 1, \dots, \frac{P_n}{2} \right\}. \quad (5)$$

Then the actual values of the cropped pixels are calculated by the continuous (x_n^c, y_n^c) via interpolation. To effectively adapt the dynamic resolution cropping strategy to segmentation tasks, we propose a reverse process with little modification in the segmentation network for the generation of full-resolution segmentation masks. Concretely, feature maps of the cropped input image are initially positioned at the upper left corner of a full-sized map with all elements set to zero. Afterward, an affine translation is performed on the full-sized map according to the center point position. Intuitively, directly interpolating the cropped segmentation result

to full resolution may not provide precise results. Instead of designing a brand-new module to refine, we interpolate the feature maps before the last decoder stage in the Dynamic Quantization Model, and then use size-independent last decoder stage to obtain the final segmentation results.

3.4. Flexible Dynamic Architecture Adaption

Previous quantization works [1, 12, 39] have normally quantized the network with a fixed bit-width. Nevertheless, different network stages exhibit different degrees of quantization sensitivity (performance drop from a fixed-bit quantization). Therefore, we aim to dynamically assign bit widths for the layers at different stages based on the feature complexity of the slices, avoiding the significant accuracy reduction and computing resources waste of uniformly quantizing a fixed bit-width network. To achieve high average flops-reduction with less accuracy loss, we propose a novel stage-distribution dynamic quantization method that allocates optimal bits to quantize each operator at different stages of the model.

Generally, to replace the majority of floating-point operations with lower-bit operations in a universal model, the input feature and weight of each layer can be respectively quantized [12, 24]. Given the operator weight of the j -th layer, it will be quantized to $Q_b(w^j)$:

$$Q_b(w^j) = \left[\text{clip}(w^j, a) \cdot \frac{r(b)}{a} \right] \cdot \frac{a}{r(b)}, \quad (6)$$

where w^j is first truncated with the clip function and scaled to $[-1, 1]$ with the scale parameter a . Then, it is scaled to the integer range $[-r(b), r(b)]$ of the given bit-width, where $r(b) = 2^{b-1}$. We use the simplest quantization strategy in implementation, the weight scale parameter a is determined simply by $a = \max(w^j)$.

In our framework, each convolution and linear layer will be quantized, one of which is in turn selected by a bit selector consisting of K bit-width quantization candidates, as illustrated in Fig. 2. To realize the bit-width selection in the model's different stages, we use a lightweight bit selection module that assigns quantization bits adaptively according to the feature complexity extracted from the slice (i.e. gradient of spatial feature and standard deviation of channel feature). The bit-width selector is applied before each stage and the highest probability bit-width will be selected.

The selector outputs a quantization function of the max probability, which is a discrete non-differentiable process and cannot be optimized end-to-end. The straight-through estimator trick [4] substitutes the discrete process to make the process differentiable, as shown in Eq. 7, during forward-propagation and back-propagation:

$$b_n^{k*} = \begin{cases} \text{argmax}_{b_n^k} P_{b_n^k}(\mathbf{x}^n) & \text{forward} \\ \sum_{k=1}^K b_n^k \cdot P_{b_n^k}(\mathbf{x}^n) & \text{backward} \end{cases}, \quad (7)$$

where $P_{b_n^k}$ is the probability assigned to the bit-width b_n^k of the n -th stage among K numbers of candidate.

Previous works [36, 40] have mainly focused on optimizing a fixed bit-width quantized network in a static manner, by taking the original network as the teacher model with a knowledge distillation loss and a conventional pixel-wise supervision loss. To realize a more efficient quantization strategy to adaptively adjust the model capacity in a dynamic way, we directly employ the obtained computational complexity of the dynamic quantized model as the loss value. To achieve a better trade-off between the computational cost and restoration performance, the bit-widths of quantization modules with a larger impact on precision should be allocated a large number of quantization bits, resulting in the following GFLOPs loss:

$$GFLOPs\ Loss = \sum_0^N GFLOPs(Q_b(w^j), b_j^{k*}), \quad (8)$$

where N represents the stage numbers.

3.5. Training and Inference Strategy

Training. The two decision networks and the Dynamic Quantization Model are trained at three different periods. Firstly, we train the full-precision segmentation model to ensure benchmark accuracy. Next, we train the dynamic quantization model, which adaptively assigns bit-widths based on feature complexity and adjusts the weights of the entire model accordingly. Finally, we train the entire architecture, including the decision networks and quantization model together, pursuing an optimal balance between accuracy and computational complexity. To address the problem of inter-slice image and intra-slice region distribution (i.e. background occupies a significant portion or the entirety of a slice, which is a substantial part of the dataset), the combination of Dice Loss (Eq. 9) and GFLOPs Loss (Eq. 8) is employed to meet well trade-off between segmentation accuracy and model complexity. Specifically, the Dice Loss is used to supervise the entire network, while GFLOPs Loss is utilized to supervise the selection of the Policy Network and Dynamic Quantization Model.

$$Dice\ Loss = \sum_{i=1}^C \left(1 - \frac{2|pred_i \cap truth_i|}{|pred_i| + |truth_i|} \right), \quad (9)$$

$$Loss = Dice\ Loss + \lambda GFLOPs\ Loss, \quad (10)$$

where $Loss$ represents the overall loss, C denotes class numbers and λ is a weight parameter.

Inference. After the one-stage training phase mentioned above, the well-trained decision networks and Dynamic Quantization Model are cascaded sequentially to achieve efficient inference. Given a 2D slice as input, our lightweight

Policy Network will decide to skip the current slice or generate a specific crop decision according to the segmentation difficulty of the current slice. A segmentation map with all zeros will be directly adopted as the final result if the skip operation is determined by the Policy Network. Otherwise, based on the current slice’s feature complexity, an appropriate combination of crop size and crop center point is determined by the Policy Network and Crop Position Network. Subsequently, the cropped or whole image will be segmented by our Dynamic Quantization Model, which can allocate bit-width quantization for different model stages according to the complexity of image features. In this way, a dynamic slice-dependent architecture with greatly improved efficiency is realized by our method in terms of both input information and model scale. On one hand, compared with the previously proposed lightweight static structure, our Med-DANet V2 simultaneously takes advantage of spatial-wise dynamic selection and dynamic architecture to adjust the inter-slice image and intra-slice regions instead of treating all inputs and regions equally. On the other hand, compared with the previously proposed dynamic methods that train multiple models for constructing a pre-defined model bank or utilize early exit to construct the cascaded dynamic architecture, our highly efficient Med-DANet V2 not only requires low training costs (i.e. only trains a single model) but also implements an extremely lightweight dynamic architecture network by adopting dynamic resolution adjustment and adaptive quantization simultaneously.

4. Experimental Evaluation

4.1. Experimental Setup

Data and Evaluation Metric. The first 3D MRI dataset used for experiments is provided by the Brain Tumor Segmentation Challenge (BraTS) 2019 [2, 3, 34]. It comprises a training set with 335 cases and a validation set with 125 cases, each with four modalities (T1, T1c, T2, and FLAIR) rigidly aligned. The size of each modality is resampled to $240 \times 240 \times 155$. All cases are labeled by four categories: background (label 0), necrotic and non-enhancing tumor (label 1), peritumoral edema (label 2), and GD-enhancing tumor (label 4). The segmentation performance on three classes: enhancing tumor region (ET, label 4), regions of the tumor core (TC, labels 1 and 4), and the whole tumor region (WT, labels 1, 2, and 4) is evaluated by the Dice score and the Hausdorff distance (95%) metrics, while FLOPs is utilized for computational complexity measurement. The second 3D MRI dataset employed in our study is obtained from the Brain Tumor Segmentation Challenge (BraTS) 2020. The dataset consists of 369 cases for training and 125 cases for validation. Apart from the difference in sample quantity, the remaining information for these two MRI datasets is identical. Due to space limitations, we place the visualization

results in the *supplementary material*.

Implementation Details. Our Med-DANet V2 is implemented based on PyTorch [38] and trained with 4 Geforce RTX 3090 GPUs. The candidates for Policy Network are set to skip, whole image, and crop size 96, while the quantization bit selection space is set to 8 bits and 16 bits. We trained our method for 350 epochs with a batch size of 64. The training process consists of three phases: the initial 200 epochs when the Policy Network and full precision segmentation module are trained, then the middle 100 epochs to train the employed Crop Position Network and Dynamic Quantization Model, and the final 50 epochs to train the entire framework. The initial learning rate is set to $2e^{-4}$ and $1e^{-4}$ for BraTS 2019 and 2020 respectively. The Adam optimizer and poly learning rate strategy with warm-up are employed for model training. We follow the data augmentation techniques and model regularization in [44]. A combination of the softmax Dice Loss and GFLOPs Loss is employed, while the weight factor λ for balancing the two losses is set to 0.06.

Baseline Selection. As the representation of CNNs, U-Net [41] achieves state-of-the-art performance in medical segmentation owing to its powerful information extraction ability. Compared to the original U-Net, the modified version makes an improvement in both accuracy and complexity. Thus we choose the modified 2D U-Net with 16 base channels as the baseline in our framework.

4.2. Results and Analysis

BraTS 2019. We conduct experiments on the BraTS 2019 validation set and compare our Med-DANet V2 with previous state-of-the-art (SOTA) approaches. The quantitative results are as reported in Table 1. Obviously, our method achieves comparable or higher performance than previous SOTA methods with significantly less computational complexity. Specifically, our method reaches the Dice scores of 80.08%, 90.27%, 81.28% on ET, WT, TC, respectively. In addition, Hausdorff Distance metrics of 3.494%, 5.871%, 6.170% on ET, WT, TC prove the credibility of our segmentation results from another perspective. Besides, Med-DANet V2 has impressively lower complexity compared to other SOTA methods. For instance, Med-DANet [44] has a computational complexity **4.1** times of ours per slice, while the model complexity of TransBTS [43] is incredibly **17.3** times of ours. These findings strongly support the effectiveness of unifying data and model structure for dynamic inference.

BraTS 2020. The comparisons between our Med-DANet V2 with previous SOTA approaches on BraTS 2020 are as reported in Table 2. Our method reaches 80.38%, 90.14%, 81.16% on Dice scores and 10.172%, 6.153%, 8.221% on Hausdorff Distance with a GFLOPs of 33.85 per case. It shows that our framework can significantly improve inference efficiency while maintaining comparable segmentation accuracy. For example, Med-DANet [44] has a computa-

Table 1. Performance Comparison on BraTS 2019 Validation Set.

Backbone	Dice Score (%) \uparrow			Hausdorff Dist. (mm) \downarrow			FLOPs (G) \downarrow	
	ET	WT	TC	ET	WT	TC	Per Case	Per Slice
3D U-Net [13]	70.86	87.38	72.48	5.062	9.432	8.719	1,669.53	13.04
V-Net [35]	73.89	88.73	76.56	6.131	6.256	8.705	749.29	5.85
Attention U-Net [37]	75.96	88.81	77.20	5.202	7.756	8.258	132.67	1.04
Wang et al. [13]	73.70	89.40	80.70	5.994	5.677	7.357	-	-
Chen et al. [9]	74.16	90.26	79.25	4.575	4.378	7.954	-	-
Li et al. [26]	77.10	88.60	81.30	6.033	6.232	7.409	-	-
TransUNet [7]	78.17	89.48	78.91	4.832	6.667	7.365	1205.76	9.42
Swin-UNet [5]	78.49	89.38	78.75	6.925	7.505	9.260	250.88	1.96
TransBTS [43]	78.36	88.89	81.41	5.908	7.599	7.584	333.09	2.60
Med-DANet [44]	79.99	90.13	80.83	4.086	5.826	6.886	77.78	0.61
Ours	80.08	90.27	81.28	3.494	5.871	6.170	19.36	0.15

Table 2. Performance Comparison on BraTS 2020 Validation Set.

Backbone	Dice Score (%) \uparrow			Hausdorff Dist. (mm) \downarrow			FLOPs (G) \downarrow	
	ET	WT	TC	ET	WT	TC	Per Case	Per Slice
3D U-Net [13]	68.76	84.11	79.06	50.983	13.366	13.607	1,669.53	13.04
V-Net [35]	61.79	84.63	75.26	47.702	20.407	12.175	749.29	5.85
Deeper V-Net [35]	68.97	86.11	77.90	43.518	14.499	16.153	-	-
3D Residual U-Net [50]	71.63	82.46	76.47	37.422	12.337	13.105	407.37	3.18
Liu et al. [29]	76.37	88.23	80.12	21.390	6.680	6.490	-	-
Ghaffari et al. [14]	78.00	90.00	82.00	-	-	-	-	-
TransUNet [7]	78.42	89.46	78.37	12.851	5.968	12.840	1205.76	9.42
Swin-UNet [5]	78.95	89.34	77.60	11.005	7.855	14.594	250.88	1.96
TransBTS [43]	78.50	89.00	81.36	16.716	6.469	10.468	333.09	2.60
Med-DANet [44]	80.57	90.28	81.34	6.474	6.718	7.416	77.71	0.61
Ours	80.38	90.14	81.16	10.172	6.153	8.221	33.85	0.26

tional complexity **2.3** times of our method, while the complexity of TransBTS [43] is incredibly **10.0** times of ours.

4.3. Ablation Study

To better understand the proposed Med-DANet V2, we conduct ablation experiments on components and settings. All studies are conducted with a U-Net baseline on the BraTS 2019 training set under the five-fold cross-validation setting unless specified otherwise.

Ablation Study on the Dynamic Architecture. We first conduct effectiveness validation of Med-DANet V2 framework for the dynamic adaption in the spatial domain (S) and quantization architecture (Q). As shown in Table 3a, our method achieves significant improvements in both segmentation accuracy and model efficiency compared to baseline. The fine-grained spatial-wise dynamic adaption reduces the computational complexity by nearly 2.2 times and improves all the Dice Scores. The flexible dynamic architecture adaption reduces computational complexity by about 1.9 times without sacrificing too much precision. By integrating the advantages of both aspects, the results of accuracy

and efficiency prove the significance of taking spatial and architecture-wise adaption into consideration for dynamic medical inference.

Ablation Study on the Weight Factor λ of the GFLOPs Loss. In this section, we explore the optimal balance between model complexity and performance. As described in Sec. 3.5, we introduce λ to control the computational complexity of the model. Ablation results are listed in Table 3b, which indicates that $\lambda = 0.06$ is the optimal weight for achieving the best trade-off. Continuing to increase the constraint on computational complexity will make the model select some slices to use higher resolution while they perform better in low resolution. And also the high-resolution slices with small foregrounds may weaken the model’s ability to focus on large targets.

Ablation Study on the Policy Space. We explore the effect of policy space by using different combinations of resolution and skipping. The resolution candidates are set to 2^n , namely 32, 64, 96, and the whole image. As shown in Table 3d, the policy space of {Skip, 96, Whole} achieves the best trade-off in accuracy and efficiency compared to all the alternatives.

Dynamic		Dice Score (%) \uparrow				FLOPs (G) \downarrow	λ	Dice Score (%) \uparrow				FLOPs (G) \downarrow	Quantization Space	Dice Score (%) \uparrow				FLOPs (G) \downarrow
S	Q	ET	WT	TC	Per Slice	ET		WT	TC	Per Slice	ET	WT		TC	Per Slice			
-	-	78.31	90.61	82.59	17.07	0.1	78.5	90.86	82.88	2.17	{8, 12}	77.07	90.41	82.63	2.74			
✓		78.40	90.95	83.39	7.87	0.08	78.43	90.93	82.80	2.23	{8, 16}	78.55	91.01	82.85	2.75			
	✓	77.10	90.76	81.86	9.01	0.06	78.55	91.01	82.85	2.75	{12, 16}	76.81	90.27	81.58	2.69			
✓	✓	78.55	91.01	82.85	2.75	0.04	78.32	90.43	82.36	4.12	{8, 12, 16}	77.69	90.59	82.73	3.25			

(a) Ablation Study on the Components of Med-DANet V2 Architecture. (b) Ablation Study on the Weight Factor λ of the GFLOPs Loss. (c) Ablation Study on the Quantization Space.

Policy Space	Dice Score (%) \uparrow			FLOPs (G) \downarrow	Backbone	Method	Dice Score (%) \uparrow			FLOPs (G) \downarrow
	ET	WT	TC	Per Slice			ET	WT	TC	Per Slice
{Skip, 32, Whole}	78.24	89.41	81.37	2.67	U-Net [41]	Baseline	77.04	90.27	82.53	17.07
{Skip, 64, Whole}	78.89	90.60	82.81	3.63		Ours	78.55	91.01	82.85	2.75
{Skip, 96, Whole}	78.55	91.01	82.85	2.75	TransBTS [43]	Baseline	78.55	91.25	82.80	37.73
{Skip, 64, 96, Whole}	78.64	90.52	82.26	3.67		Ours	79.01	91.03	82.18	3.55

(d) Ablation Study on the Policy Space. (e) Ablation Study on Different Baselines.

Table 3. Ablation study of various aspects of our proposed Med-DANet V2 approach.

And the further increase of candidates in policy space may not achieve better performance.

Ablation Study on the Quantization Space. In this section, we investigate the impact of using different combinations of Quantization Space, as displayed in Table 3c. Intermediate quantization bits such as 8, 12, and 16 are selected as candidates for the quantization space. Among them, {8,16} achieves the best balance between accuracy and computational complexity. Setting the quantization space as {8,12,16} and {12,16} would increase the model’s capacity, but the fixed constraints of λ introduce biases towards smaller values in both the decision space and quantization space, resulting in inadequate segmentation performance. Although {8,12} has the similar computational complexity as {8,16}, its upper limit of 12 bits hinders the segmentation for some complex slices.

Ablation Study on Different Baselines. We choose two baselines (the modified 2D UNet and the 2D version of lightweight TransBTS) to verify the compatibility and scalability of our framework because they represent two popular architectures (i.e. convolutional neural networks and vision Transformers). As shown in Table 3e, by unifying spatial adaption and dynamic quantization architecture, our framework significantly surpasses both baselines in computational complexity with comparable or even better performance.

5. Discussion and Conclusion

In this paper, we present a study to explore the potential of dynamic inference in medical volumetric segmentation. We propose to unify the spatial dynamic adaption and dynamic model quantization to handle the slice heterogeneity in medical volume data. From the perspective of data

properties and model structure, we focus on the 3D MRI brain tumor segmentation and propose a novel framework named Med-DANet V2 with adaptive input selection and dynamic architectures to pursue the trade-off between segmentation accuracy and efficiency. In comparison with the previous work Med-DANet in this research direction, more fine-grained dynamic spatial adapting, and more flexible dynamic structural adapting are jointly incorporated to greatly promote the model efficiency. Extensive experiments on two benchmark datasets for multimodal 3D MRI brain tumor segmentation demonstrate that our Med-DANet V2 reaches comparable or better performance than previous state-of-the-art methods with significantly less model complexity.

Broader Impact and Limitation. Our study offers a novel perspective and solution to realize efficient medical volumetric segmentation for clinical applications by holistically considering the unique characteristics of medical volume data along with adaptive model capacity adjustments. Our approach could be especially beneficial for time-sensitive medical diagnoses and treatment planning, where our framework’s efficiency could lead to quicker and potentially more accurate clinical decisions. The proposed framework opens up a new avenue for achieving efficient volumetric segmentation in clinical settings, thereby serving as a catalyst for future research in this domain.

However, one potential limitation could be that our current method for fine-grained, spatially-aware dynamic adaptation is restricted to isolating important foreground regions with regular geometries. This shortcoming identifies an avenue for future research, specifically the development of more nuanced dynamic input adaptation techniques capable of accurately identifying and segmenting foreground regions with irregular shapes.

References

- [1] MohammadHossein AskariHemmat, Sina Honari, Lucas Rouhier, Christian S Perone, Julien Cohen-Adad, Yvon Savaria, and Jean-Pierre David. U-net fixed-point quantization for medical image segmentation. In *Large-Scale Annotation of Biomedical Data and Expert Label Synthesis and Hardware Aware Learning for Medical Imaging and Computer Assisted Intervention: International Workshops, LABELS 2019, HAL-MICCAI 2019, and CuRIOUS 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13 and 17, 2019, Proceedings*, pages 115–124. Springer, 2019. [3](#), [5](#)
- [2] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4:170117, 2017. [6](#)
- [3] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018. [6](#)
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. [5](#)
- [5] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation. In *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*, pages 205–218. Springer, 2023. [1](#), [2](#), [7](#)
- [6] Chen Chen, Xiaopeng Liu, Meng Ding, Junfeng Zheng, and Jiangyun Li. 3d dilated multi-fiber network for real-time brain tumor segmentation in mri. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III 22*, pages 184–192. Springer, 2019. [1](#), [2](#)
- [7] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021. [7](#)
- [8] Jin Chen, Xijun Wang, Zichao Guo, Xiangyu Zhang, and Jian Sun. Dynamic region-aware convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8064–8073, 2021. [3](#)
- [9] Minglin Chen, Yaozu Wu, and Jianhuang Wu. Aggregating multi-scale prediction based on 3d u-net in brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 142–152. Springer, 2019. [7](#)
- [10] Wei Chen, Boqiang Liu, Suting Peng, Jiawei Sun, and Xu Qiao. S3d-unet: separable 3d u-net for brain tumor segmentation. In *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 4th International Workshop, BrainLes 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers, Part II 4*, pages 358–368. Springer, 2019. [1](#), [2](#)
- [11] Weihang Chen, Peisong Wang, and Jian Cheng. Towards mixed-precision quantization of neural networks via constrained optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5350–5359, 2021. [3](#)
- [12] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*, 2018. [3](#), [5](#)
- [13] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*, pages 424–432. Springer, 2016. [1](#), [7](#)
- [14] Mina Ghaffari, Arcot Sowmya, and Ruth Oliver. Brain tumour segmentation using cascaded 3d densely-connected unet. *arXiv preprint arXiv:2009.07563*, 2020. [7](#)
- [15] Liming Guo, Wen Fei, Wenrui Dai, Chenglin Li, Junni Zou, and Hongkai Xiong. Mixed-precision quantization of u-net for medical image segmentation. In *2022 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 2871–2875. IEEE, 2022. [3](#)
- [16] Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger R Roth, and Daguang Xu. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 7th International Workshop, BrainLes 2021, Held in Conjunction with MICCAI 2021, Virtual Event, September 27, 2021, Revised Selected Papers, Part I*, pages 272–284. Springer, 2022. [1](#)
- [17] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 574–584, 2022. [1](#)
- [18] Nicholas Heller, Fabian Isensee, Klaus H Maier-Hein, Xiaoshuai Hou, Chunmei Xie, Fengyi Li, Yang Nan, Guangrui Mu, Zhiyong Lin, Miao Han, et al. The state of the art in kidney and kidney tumor segmentation in contrast-enhanced ct imaging: Results of the kits19 challenge. *Medical image analysis*, 67:101821, 2021. [1](#)
- [19] Yu-Hui Huang, Marc Proesmans, Stamatios Georgoulis, and Luc Van Gool. Uncertainty based model selection for fast semantic segmentation. In *2019 16th International Conference*

- on *Machine Vision Applications (MVA)*, pages 1–6. IEEE, 2019. 3
- [20] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. *Advances in neural information processing systems*, 29, 2016. 3
- [21] Alex Ling Yu Hung, Haoxin Zheng, Qi Miao, Steven S Raman, Demetri Terzopoulos, and Kyunghyun Sung. Cat-net: A cross-slice attention transformer model for prostate zonal segmentation in mri. *IEEE transactions on medical imaging*, 42(1):291–303, 2022. 1
- [22] Yuankai Huo, Jiaqi Liu, Zhoubing Xu, Robert L Harrigan, Albert Assad, Richard G Abramson, and Bennett A Landman. Robust multicontrast mri spleen segmentation for splenomegaly using multi-atlas segmentation. *IEEE Transactions on Biomedical Engineering*, 65(2):336–343, 2017. 1
- [23] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. 4
- [24] Sangil Jung, Changyong Son, Seohyung Lee, Jinwoo Son, Jae-Joon Han, Youngjun Kwak, Sung Ju Hwang, and Changkyu Choi. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4350–4359, 2019. 5
- [25] Han Li, Long Chen, Hu Han, and S Kevin Zhou. Satr: Slice attention with transformer for universal lesion detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 163–174. Springer, 2022. 1
- [26] Xiangyu Li, Gongning Luo, and Kuanquan Wang. Multi-step cascaded networks for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 163–173. Springer, 2019. 3, 7
- [27] Zhuoying Li, Junquan Pan, Huisi Wu, Zhenkun Wen, and Jing Qin. Memory-efficient automatic kidney and tumor segmentation based on non-local context guided 3d u-net. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part IV 23*, pages 197–206. Springer, 2020. 1
- [28] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. *arXiv preprint arXiv:2202.07800*, 2022. 3
- [29] Chenyu Liu, Wangbin Ding, Lei Li, Zhen Zhang, Chenhao Pei, Liqin Huang, and Xiahai Zhuang. Brain tumor segmentation network using attention-based fusion and spatial relationship constraint. In *International MICCAI Brainlesion Workshop*, pages 219–229. Springer, 2020. 7
- [30] Ange Lou, Shuyue Guan, and Murray Loew. Dc-unet: rethinking the u-net architecture with dual channel efficient cnn for medical image segmentation. In *Medical Imaging 2021: Image Processing*, volume 11596, pages 758–768. SPIE, 2021. 1, 2
- [31] Qian Lou, Feng Guo, Lantao Liu, Minje Kim, and Lei Jiang. Autoq: Automated kernel-wise neural network quantization. *arXiv preprint arXiv:1902.05690*, 2019. 3
- [32] Zhengrong Luo, Zhongdao Jia, Zhimin Yuan, and Jialin Peng. Hdc-net: Hierarchical decoupled convolution network for brain tumor segmentation. *IEEE Journal of Biomedical and Health Informatics*, 25(3):737–745, 2020. 1, 3
- [33] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*, pages 116–131, 2018. 4
- [34] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014. 6
- [35] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016. 1, 7
- [36] Asit Mishra and Debbie Marr. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. *arXiv preprint arXiv:1711.05852*, 2017. 5
- [37] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018. 1, 7
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [39] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV*, pages 525–542. Springer, 2016. 5
- [40] 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. 5
- [41] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 1, 2, 6, 8
- [42] Thomas Verelst and Tinne Tuytelaars. Segblocks: Block-based dynamic resolution networks for real-time segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 3

- [43] Wenxuan Wang, Chen Chen, Meng Ding, Hong Yu, Sen Zha, and Jiangyun Li. Transbts: Multimodal brain tumor segmentation using transformer. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 109–119. Springer, 2021. [2](#), [6](#), [7](#), [8](#)
- [44] Wenxuan Wang, Chen Chen, Jing Wang, Sen Zha, Yan Zhang, and Jiangyun Li. Med-danet: Dynamic architecture network for efficient medical volumetric segmentation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXI*, pages 506–522. Springer, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [45] Yulin Wang, Yang Yue, Yuanze Lin, Haojun Jiang, Zihang Lai, Victor Kulikov, Nikita Orlov, Humphrey Shi, and Gao Huang. Adafocus v2: End-to-end training of spatial dynamic networks for video recognition. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20030–20040. IEEE, 2022. [3](#), [4](#)
- [46] Bichen Wu, Yanghan Wang, Peizhao Zhang, Yuandong Tian, Peter Vajda, and Kurt Keutzer. Mixed precision quantization of convnets via differentiable neural architecture search. *arXiv preprint arXiv:1812.00090*, 2018. [3](#)
- [47] Zhenda Xie, Zheng Zhang, Xizhou Zhu, Gao Huang, and Stephen Lin. Spatially adaptive inference with stochastic feature sampling and interpolation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 531–548. Springer, 2020. [3](#)
- [48] Xiangyi Yan, Hao Tang, Shanlin Sun, Haoyu Ma, Deying Kong, and Xiaohui Xie. After-unet: Axial fusion transformer unet for medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3971–3981, 2022. [1](#)
- [49] Rongzhao Zhang and Albert CS Chung. Medq: Lossless ultra-low-bit neural network quantization for medical image segmentation. *Medical Image Analysis*, 73:102200, 2021. [3](#)
- [50] Zhengxin Zhang, Qingjie Liu, and Yunhong Wang. Road extraction by deep residual u-net. *IEEE Geoscience and Remote Sensing Letters*, 15(5):749–753, 2018. [7](#)
- [51] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer, 2018. [1](#), [2](#)