

HyperSegNAS: Bridging One-Shot Neural Architecture Search with 3D Medical Image Segmentation using HyperNet

Cheng Peng¹ Andriy Myronenko² Ali Hatamizadeh² Vishwesh Nath²
Md Mahfuzur Rahman Siddiquee³ Yufan He² Daguang Xu² Rama Chellappa¹ Dong Yang²
¹Johns Hopkins University ²NVIDIA ³Arizona State University

Abstract

Semantic segmentation of 3D medical images is a challenging task due to the high variability of the shape and pattern of objects (such as organs or tumors). Given the recent success of deep learning in medical image segmentation, Neural Architecture Search (NAS) has been introduced to find high-performance 3D segmentation network architectures. However, because of the massive computational requirements of 3D data and the discrete optimization nature of architecture search, previous NAS methods require a long search time or necessary continuous relaxation, and commonly lead to sub-optimal network architectures. While one-shot NAS can potentially address these disadvantages, its application in the segmentation domain has not been well studied in the expansive multi-scale multi-path search space. To enable one-shot NAS for medical image segmentation, our method, named HyperSegNAS, introduces a HyperNet to assist super-net training by incorporating architecture topology information. Such a HyperNet can be removed once the super-net is trained and introduces no overhead during architecture search. We show that HyperSegNAS yields better performing and more intuitive architectures compared to the previous state-of-the-art (SOTA) segmentation networks; furthermore, it can quickly and accurately find good architecture candidates under different computing constraints. Our method is evaluated on public datasets from the Medical Segmentation Decathlon (MSD) challenge, and achieves SOTA performances.

1. Introduction

Automated medical image segmentation is an active research area with many clinical applications including anatomical analysis [29] and disease diagnosis [14]. Medical image segmentation remains challenging despite advances in deep learning methods. To address the large variability of target objects, the limited number of data samples, and the higher dimensional and high resolution nature of

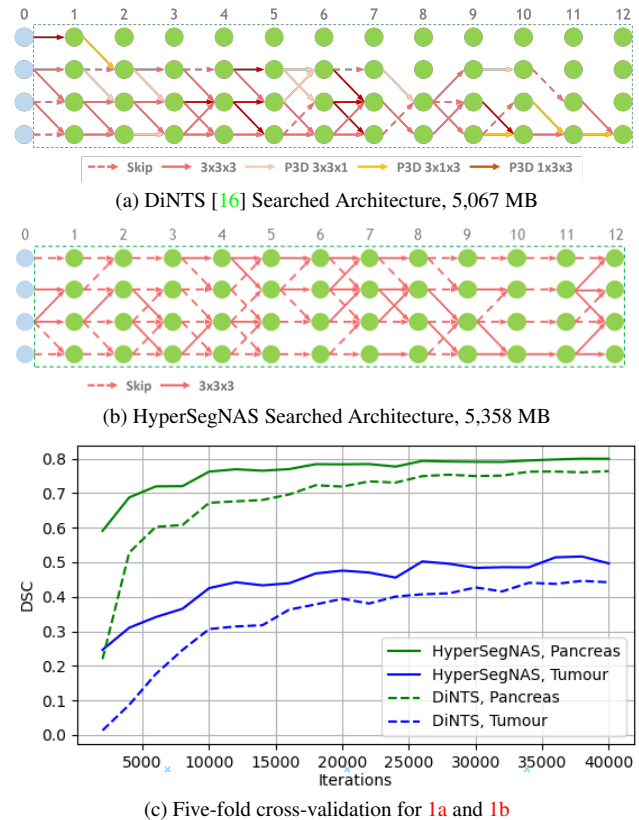


Figure 1. Architectures found from DiNTS [16] (1a) and our proposed HyperSegNAS (1b) based on the Pancreas dataset in MSD [1]. Under similar compute costs, HyperSegNAS’ architecture uses more skip connections and ensure features of multiple scales are propagated for predictions. Our architecture performs significantly better than DiNTS, as shown in 1c.

3D images, researchers have spent a great effort in designing efficient neural architectures under various computing settings [9, 10, 27, 28, 32].

Neural Architecture Search (NAS) has surfaced in response to handcrafted approaches and promises to find well performing architecture through automated algorithms. While initially focusing on classification [5, 26, 39, 43, 48], NAS has been also introduced to the segmentation domain in recent years [16, 21, 25, 37, 38, 40, 44, 47]. Segmentation

networks usually preserve features from multiple scales and aggregate them together to accurately segment objects with different sizes. Accordingly, the search space for multi-scale, multi-path architectures can be very complex, as the feasible network configurations increase exponentially with scale. C2FNAS [44] needs close to *one GPU year* to search for a single 3D segmentation architecture based on the Evolutionary Algorithm (EA), despite building on a U-shape network architecture. On the other hand, DiNTS [16] relaxes the architecture search problem from a discrete formulation to a continuous and differentiable one, thus it greatly speeds up the searching process. Such a relaxation, however, may lead to (1) an optimization gap when the continuous architecture/edge weights are discretized for deployment [8, 35] and (2) possibly infeasible architectures, which require ad-hoc logic to be handled.

Beyond finding the best performing architecture, balancing other computing constraints to the search process is another important aspect of NAS and has led to popularity in methods like one-shot NAS [2, 3, 5–7, 12, 18, 22, 30, 33, 41, 43], where all sub-nets are trained with shared parameters under a large super-net. This aspect is especially important for medical imaging segmentation task for which memory usage is a major challenge when searching for the optimal architecture across various devices and resource constraints. Once training is completed, the shared parameters can be used to evaluate the possible sub-nets and select the best performers given the specified constraints. So far, one-shot approaches have not been applied to the 3D segmentation’s search space yet. As there are no clear training strategies, such as the unilateral augmented (UA) principle [34] used in progressive shrinking [5] and the Sandwich Rule [42, 43], in a multi-scale, multi-path search space [25], we find that the one-shot training scheme with randomly sampled architecture topology frequently leads to sub-optimal performance.

We propose HyperSegNAS, which follows the one-shot NAS approach and seeks to address its issues in the segmentation space. Within HyperSegNAS, we propose a novel Meta-Assistant Network (MAN) to improve super-net training. Specifically, both the sampled architecture and the input image are fed to MAN. MAN then dynamically modifies the shared weights during the training process based on deploying architectures. We show that such an *architecture topology-aware* training method significantly improves sub-net performances. When training is completed, an annealing process is used to gradually remove MAN from the main network. Even after removing MAN, the sub-net performances remain high and can be efficiently evaluated.

Benefiting from the one-shot paradigm, HyperSegNAS has many advantages. Firstly, HyperSegNAS is much more efficient at evaluating over *many architectures* under different computing constraints, as the shared parameters only need to be trained once. In comparison, approaches in

[16, 44] need to search from scratch to obtain each single architecture. Secondly, HyperSegNAS evaluates on discrete architectures, which eliminates the discretization gap and infeasible configurations that arise in differentiable NAS. Finally, HyperSegNAS can accurately fit to the given constraints due to the decoupling between training and searching. As demonstrated in Figures 1a and 1b, our search algorithm results in a very different architecture as compared to DiNTS’s under similar computing constraints. Not only is our architecture significantly better performing, as shown in Fig. 1c, it also follows conventional intuitions in propagating multi-scale features through the use of skip connections. Finally, HyperSegNAS achieves new SOTA performance on multiple tasks in the Medical Segmentation Decathlon (MSD) challenge [1] using similar computing budget as DiNTS.

In summary, our main contributions are listed below:

- We propose HyperSegNAS, a one-shot NAS approach at finding optimal 3D segmentation architectures. HyperSegNAS can search for efficient architectures that accurately adapt to different computing constraints.
- We propose a HyperNet structure called Meta-Assistant Network (MAN) to address the large search space of segmentation networks by incorporating relevant meta information to NAS; MAN can be removed after training, and does not add computing overhead for searching and re-training.
- We achieve better results on both low- and high-compute architectures compared to DiNTS and new SOTA performances in multiple tasks in MSD.

2. Related Work

2.1. Medical Image Segmentation

Medical image segmentation has been studied and developed for decades in many medical imaging applications. Particularly, deep learning based-algorithms have been widely adopted in segmentation applications in recent years. Among the deep neural network models, U-Net [9, 32] is a seminal work in this area due to its effectiveness and structural simplicity for processing multi-scale image features. Various improvements have been introduced subsequently [19, 23, 28, 46], which include adding components like residual blocks or dense blocks and combining 2D and 3D U-Nets. On the other hand, nn-UNet [19] ensembles predictions from multiple existing 2D/3D U-Net models, and focuses on finding the correct configurations/hyper-parameters to improve the performance. There are also other existing architectures without U-shape structure, e.g. AH-Net [27], HDenseNet [24], HyperDenseNet [10], etc. A common theme across these

works is to ensure features of various dimension and resolution are preserved while maintaining manageable computational cost; therefore, the use of deep learning modules (e.g., convolution, skip connections) has been a major area of exploration for manually designed architectures.

2.2. Neural Architecture Search (NAS)

Manual architecture designs have resulted in steady performance improvements across vision tasks; however, they iterate gradually and with great recurring cost. NAS proposes to automate architecture design in a performance-driven way, and has received considerable attention. NAS has been frequently applied on image recognition tasks [5, 26, 39, 43, 48]. On a high level, evolutionary algorithm (EA) [31] and reinforcement learning (RL) [48] are used to perform discrete search/optimization, differentiable relaxation [26] on connections and operations allow search algorithms to leverage gradient descent optimization. Finally, training an architecture performance estimator, e.g. one-shot NAS [2, 3, 5–7, 12, 18, 22, 30, 33, 41, 43], which our method is based on, is another popular approach as it disentangles the training and searching processes. While previous methods [3, 45] have used HyperNets to directly estimate sub-net performances, HyperSegNAS focuses on accelerating feature learning through an assistant HyperNet, which does not introduce an overhead in the search process.

More recently, NAS was explored in medical image segmentation [16, 38, 40]. As 3D segmentation architectures need to process and preserve multi-resolution features, they are much more computationally intense than 2D classification architectures. Early attempts [21, 37, 44, 47] build on the prior architecture of U-Net, and modify edge operations and connections while preserving the overall U-shape structure. To go beyond a U-Net based structure while remaining computationally feasible, DiNTS [16] applies a differentiable NAS formulation similar to [26, 39] on a segmentation search space first proposed by Auto-DeepLab [25]. Due to the efficiency introduces in differentiable NAS, DiNTS [16] can feasibly search for architectures that fit under different computing budgets, and is the current best performer in the MSD [1] challenge.

3. Method

3.1. Network Search Space

Topology. We follow previous works [16, 25] in defining the search space for 3D segmentation architectures. As shown in Fig. 2, we down-sample the input image I to generate multi-scale $I_{\downarrow r}$, where r denotes the down-sampling factor. Overall, the search space has $L = 12$ layers and contains features of $|r| = 4$ scales, denoted as $F_{\downarrow r}^l$, at every layer, where $l \in \{1, 2, \dots, L\}$. There are $E = 10$ possible edges between layers; each edge contains a basic opera-

tion and an up-sampling/down-sampling/identity operation depends on the output feature scale. When multiple edges point to the same feature scale, their outputs are summed and normalized. As shown in [16, 25], this search space is very general in topology and contains many major segmentation architectures (like U-shape networks).

Edge Operations. For each edge in the search space, a basic operation is chosen from a pre-defined operation set. As the complexity in topology is already high, we take a minimalist approach at choosing the basic operation set, which includes skip connection and $3 \times 3 \times 3$ 3D convolution. Compared to [16], we omit three pseudo 3D operators (P3D). Such a design choice (1) reduces the search space complexity and allows for better clarity on network topology, (2) still contains the majority of popular medical segmentation architectures, e.g. 3D U-Net and its variants.

3.2. HyperSegNAS

Background. The goal of NAS is to find the best performing sub-network a from the search space \mathcal{A} under a given computing constraint \mathcal{C} , e.g., peak GPU memory at training, as defined by its minimum validation loss

$$\arg \min_a \mathcal{L}_{\text{val}}(a, \omega_a^*), \forall a \in \mathcal{A} \mid \text{Cost}(a) < \mathcal{C}, \quad (1)$$

where ω_a^* is the optimal weights of a after training. Since ω_a^* is generally very costly to obtain from a full training process, the **one-shot NAS** defines a super-net with weights ω , which covers the entire search space. After ω is trained, it is used as a performance estimator of a based on the sub-set weights $\omega_a \in \omega$. This approach means that all sub-nets are evaluated based on shared weights; the training objective is formulated as

$$\arg \min_{\omega} \frac{\sum \mathcal{L}_{\text{train}}(a, \omega_a)}{|\mathcal{A}|}, \forall a \in \mathcal{A}. \quad (2)$$

Clearly, it is infeasible to enumerate all architectures for inference and average their gradients; in practice, one or more architectures are sampled at each training step for local averaging. After training, the best architecture a^* is then selected by evaluating

$$\arg \min_a \mathcal{L}_{\text{val}}(a, \omega_a), \forall a \in \mathcal{A} \mid \text{Cost}(a) < \mathcal{C}. \quad (3)$$

The one-shot formulation reduces the complexity of obtaining ω_a^* in Eq. 1 through weight sharing and can directly evaluate discrete architectures; however, it also has limitations. When the search space \mathcal{A} is large, optimizing Eq. 2 becomes very different from optimizing individual sub-nets. The *multi-scale, multi-path* architectures in our segmentation search space also introduces challenges in training the super-net. Unlike search dimensions like kernel size and channel size [5, 42, 43], a multi-path architecture topology




Figure 2. HyperSegNet introduces an additional HyperNet \mathcal{H} on top of the segmentation search space. When training the super-net, \mathcal{H} introduces additional channel-wise weights, which depend on the input image and the sampled architecture. The circled block in 2a is expanded in 2b for a detailed view of the feature re-weighting process. An annealing process is introduced to gradually replace $\omega_{\mathcal{H}\theta}^{a,I}$ with ω_{fixed} and removes \mathcal{H} from the inference process. As such, efficiency and better estimated performance are achieved.

does not have a self-evident order and cannot be trained in a progressive manner. We find that training the super-net from randomly sampled topology leads to sub-optimal performance evaluations on sub-nets. As such, we hypothesize that due to the lack of good training strategies and a large search space, these super-net features tend to have high redundancy and suffer in accuracy.

Meta-Assistant Network. To address the problem discussed above, we propose a Meta-Assistant Network (MAN), which is architecture topology-aware. MAN relaxes the formulation in Eq. 2 by providing extra degrees of freedom on ω to improve training, i.e. changing ω based on meta information about the problem. In particular, we focus on two pieces of meta information: the architecture a and the input image I . As shown in Fig. 2, MAN consists of an architecture sampler, an image encoder \mathcal{M} , and a HyperNet \mathcal{H}_θ , where θ denotes \mathcal{H} 's parameters. We represent the sampled architecture $a \in \mathbb{N}^{L \times E \times O}$ as a matrix, where $O = 2$ denotes the cardinality of the edge operations. All activated edges and cell operations in the respective layers are labeled as 1, and 0 otherwise.

The conventional implementations of HyperNet [3, 13] directly output network weights ω based on relevant information. However, the total number of parameters $|\omega|$ is exceedingly large and costly to generate. Instead, MAN generates channel-wise scalar weights, as demonstrated in Fig. 2b. More precisely, we denote every convolution kernel in the super-net as $\omega(l, e, c) \in \mathbb{R}^{k \times k}$, where k is the kernel size. Respectively, $l \in \{1, 2, \dots, L\}$, $e \in \{1, 2, \dots, E\}$, and c are the layer, edge, and channel index. For every training iteration, an architecture a and a training image I is sampled. We obtain an image vector l_{image} from $\mathcal{M}(I)$, and an architecture vector l_{arch} by flattening the architecture matrix representation. These vectors are then concatenated as the

input to \mathcal{H} to obtain $\omega_{\mathcal{H}\theta}^{a,I} = \mathcal{H}_\theta(l_{\text{arch}} \oplus l_{\text{image}})$. Finally, the channel-wise weights $\omega_{\mathcal{H}\theta}^{a,I}$ are multiplied with the kernels in the super-net and transform training of Eq. 2 into

$$\arg \min_{\omega, \theta} \frac{\sum \mathcal{L}_{\text{train}}(a, \omega_{\mathcal{H}\theta}^{a,I}(l, e, c) * \omega_a(l, e, c))}{|\mathcal{A}|}. \quad (4)$$

Under this formulation, individual sub-networks gain additional capacity to optimize their own configurations based on a given image. We observe that this setting significantly improves performance.

HyperNet Annealing. While MAN improves sub-network learning, it also leads to potentially biased evaluations due to the introduction of a new network component and the overall increased representational capacity. This can affect the correlation between sub-network's estimated performance from shared weights and true performance in deployment, especially for smaller architectures. To obtain shared weights that are independent of MAN and ensure fairness in searching, HyperSegNAS consists of a second stage to remove MAN for evaluation. Specifically, we introduce a constant channel-wise scalar weight $\omega_{\text{fixed}} = 0.5$ and an annealing process to gradually replace $\omega_{\mathcal{H}\theta}^{a,I}$ with ω_{fixed} . Such a process can be expressed as:

$$\arg \min_{\omega, \theta} \frac{\sum \mathcal{L}_{\text{train}}(a, (\lambda \omega_{\text{fixed}} + (1 - \lambda) \omega_{\mathcal{H}\theta}^{a,I}) * \omega_a)}{|\mathcal{A}|}, \quad (5)$$

where λ is the annealing temperature that gradually moves from 0 to 1. We find that a linear annealing schedule works well and allows ω to maintain reasonable performance after MAN is removed, thus achieving an overall improvement in estimating performances based on ω alone.

Architecture Search. After ω is properly trained, we evaluate different architectures and select the best perform-

ing candidates that fit under the desired computing constraints. While HyperSegNAS yields observably better estimated sub-network performances, these estimations are still slightly worse in comparison to that of ω_a^* from training the architecture from scratch, as is generally observed in one-shot literature [4, 5]. More competitive performances can be reached by applying a few architecture-specific finetuning steps on ω .

As doing so across all architectures is very expensive, we deploy a coarse-to-fine search algorithm, which can be described in three parts.

1. We sample a sufficient number of candidate architectures, i.e. architectures that fit under a given computing constraint σ , to form a valid search space \mathcal{A}_σ .
2. We evaluate all candidates in \mathcal{A}_σ based on the super-net and Eq. 3, and select the top N performers.
3. We fine-tune the N architectures for a few iterations, evaluate their fine-tuned performances, and select the best performer as the final architecture.

While any form of computing constraint, e.g., latency, is applicable to HyperSegNAS, we follow [16] in using training time GPU memory cost as the computing constraint. Due to the decoupling of training and searching in Eq. 2 and Eq. 3, HyperSegNAS can precisely calculate the computing budgets of the evaluated model. Specifically, during the sampling stage in our search algorithm, HyperSegNAS can directly measure training memory cost by performing training for few iterations. This is a distinct advantage of HyperSegNAS compared to differentiable NAS methods like DiNTS [16], as training memory cost is hard to be analytically formulated and predicted due to various low-level optimization and overheads at runtime [11]. In fact, incorporating a differentiable analytical formulation of memory usage in DiNTS [16] leads to not only imprecise cost estimations, but also potentially biased architectures, e.g., as shown in Fig. 1, due to the discretization gap.

4. Experiments

Dataset. We follow the settings in [16, 44]. Specifically, our search algorithm is performed on the MSD Pancreas dataset, which has been shown by previous work to lead to generally performant architectures. We resample 281 provided scans to a voxel resolution of $1.0 \times 1.0 \times 1.0 mm^3$, use 225 scans for training, and 56 for validation.

Implementation Details. We perform architecture search on the same topological search space as in [16]. For details, the channel size c for features of scales $r = \{1, 2, 3, 4\}$ are $\{32, 64, 128, 256\}$; overall, there are a total number of $c_{tot} = 27,648$ channels in the search space.

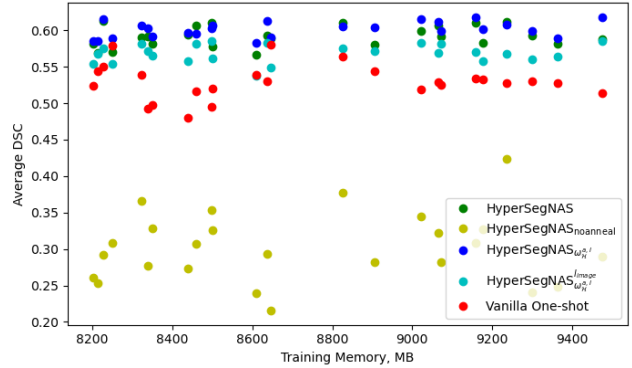


Figure 3. Visualization of the average Dice-Sørensen (DSC) scores on twenty-five evaluated architectures in Table 1 sorted based on its training memory cost.

Method	Pancreas Tumor Avg.	Inf. Speed
Vanilla One-Shot	71.93 33.96 52.92	44.60s
HyperSegNAS $_{\omega_{\mathcal{H}}^{a,I}}$	75.64 37.42 56.53	67.91s
HyperSegNAS $_{\omega_{\mathcal{H}}^{a,I}}$	75.90 44.36 60.13	69.41s
HyperSegNAS $_{noanneal}$	32.15 28.17 30.16	46.89s
HyperSegNAS	76.31 41.99 59.15	47.27s

Table 1. Ablation study on different implementations of HyperSegNAS. Twenty-five architectures are randomly sampled and evaluated on the Pancreas dataset. Dice-Sørensen (DSC) Scores on respective labels and segmentation inference speed per 3D volume (in seconds) are reported.

During training, patches of size $96 \times 96 \times 96$ are sampled and augmented with rotation and flip. The image encoder \mathcal{M} consists of six 3D convolutional layer with Instance Normalization [36] and ReLU activation, and outputs $l_{image} \in \mathbb{R}^{256}$. The HyperNet \mathcal{H} consists of five fully connected layers with ReLU activations in between and a Sigmoid activation at the end; it outputs $\omega_{\mathcal{H}_\theta}^{a,I} \in \mathbb{R}^{c_{tot}}$. For every sampled architecture a , the inactivated portion of $\omega_{\mathcal{H}_\theta}^{a,I}$ is masked off from back-propagation. We use a combination of dice and cross-entropy loss as the training objective function \mathcal{L}_{train} . To speed up the evaluation during the search stage, we sample the foreground regions with a maximum size of $192 \times 192 \times 192 mm^3$ and evaluate them patch-wise. The architecture sampling algorithm is implemented with several heuristics to eliminate undesirable configurations. We ensure that at least one feasible edge is sampled per layer so that the architecture always has L layers. This prevents early termination and an excessive number of very small architectures. Feasible edges and operations are sampled with uniform likelihood. Super-net training is performed on eight NVIDIA V100 GPUs with a batch size of of eight and 160,000 iterations; HyperNet is then annealed over 20,000 iterations.

Since HyperSegNAS can fit to cost constraint with accuracy, we constraint our searched architectures to have similar training memory costs as the three architectures found in [16]. Specifically, we sample two thousand architectures that cost within 300MB of DiNTS’s respective architecture to form a \mathcal{A}_σ , and select the top $N = 10$ architectures for quick fine-tuning with 5,000 iterations. Training memory is directly measured by monitoring the NVIDIA System Management Interface (*nvidia-smi* command). Overall, HyperSegNAS takes 32 GPU days to train the super-net, and 4 GPU days to search under one computing constraint. This is much faster than C2FNAS [44], which takes 333 GPU days to search for one architecture. Compared to DiNTS [16], which takes 5.8 GPU days to search for one architecture, our method has a front-loaded cost in super-net training, but is faster at adapting to different computing constraints once super-net weights are obtained. Please refer to the Supplemental Material for more algorithm/implementation details.




Figure 4. Improved ranking of architectures based on fine-tuned weights. The Kendall-Tau correlation between estimations from fine-tuned super-net weights and fully trained performances is $\tau = 0.55$, and $\tau = 0.22$ based on using non fine-tuned super-net weights. Spearman’s rank correlations are $r_s = 0.7$ and $r_s = 0.4$, respectively.

4.1. Ablation Study

We evaluate the effectiveness of HyperSegNAS against alternative implementations. Specifically, we compare the proposed pipeline with the following implementations using the same training hyper-parameters:

- A) Vanilla one-shot: The training scheme used in most one-shot methods, e.g. [5, 43], where the architecture is sampled and trained with shared weights.
- B) HyperSegNAS $_{\omega_{\mathcal{H}_\theta}^{l_I}}$: A HyperSegNAS implementation where the image vector l_{image} is the only input to \mathcal{H} and \mathcal{H} is not removed.
- C) HyperSegNAS $_{\omega_{\mathcal{H}}^{a,I}}$: A HyperSegNAS implementation using both l_{image} and l_{arch} as input to \mathcal{H} ; \mathcal{H} is not removed.
- D) HyperSegNAS $_{\text{noanneal}}$: A HyperSegNAS implementation where \mathcal{H} is removed without annealing.

As shown in Figure 3 and Table 1, the traditional one-shot training scheme leads to sub-optimal performance in our segmentation search space. The shared weights need to implicitly accommodate all randomly selected architectures; as the features are generally applicable to random architectures, their capacity is limited.

In comparison, the addition of \mathcal{H} helps generally improve performances across sampled architectures. Specifically, while HyperSegNAS $_{\omega_{\mathcal{H}}^{l_I}}$ shows some performance improvements, we observe significant improvements from applying both the image and architecture vectors in HyperSegNAS $_{\omega_{\mathcal{H}}^{a,I}}$, demonstrating the utility of each piece of meta-information in accelerating learning. While annealing \mathcal{H} away leads to some performance drop off, as ω is now not architecture or input specific, HyperSegNAS still outperforms Vanilla one-shot by a wide margin. If annealing is not applied and \mathcal{H} is simply removed from inference, performances are severely degraded.

4.2. Method Analysis

Search Correlation. As HyperSegNAS explicitly evaluates the performances of different architectures based on shared fine-tuned weights, we can measure the Kendall-Tau correlation [20] of such performance estimations with the performances of these architectures trained from scratch. As shown Fig. 4, we observe a moderate amount of correlation when shared weights are used for estimations, and a stronger correlation when fine-tuned weights are used.

Searched Architectures. Besides the low-cost architecture shown in Fig. 1, we visualize the other two searched architectures that use comparable training memory to [16] in Fig. 5. The training memory cost and performances based on five-fold cross-validation are presented in Table 2.

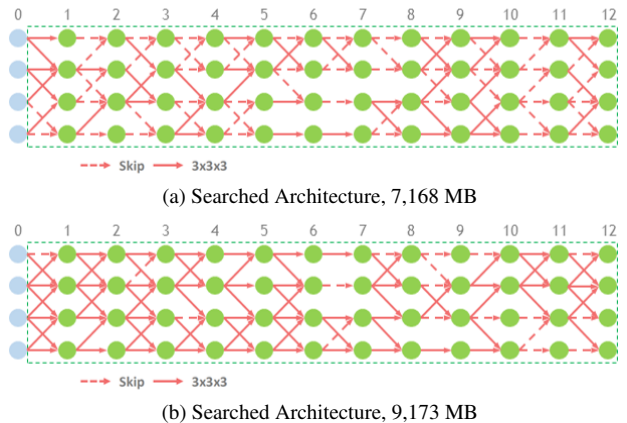


Figure 5. Searched architectures from HyperSegNAS with various compute constraints.

We make several observations. Firstly, under limited computational resources, e.g. at around 5GB of GPU training memory, the network architecture has a significant role

on the final performance. This is reasonable as high compute architectures in the same search space tend towards homogeneity, i.e. the super-net, while low compute architectures have more variety. As such, we can demonstrate the search algorithm’s effectiveness under lower computational resources with better clarity.

Method	Cost (MB)	Panc.	Tumor	Avg.
DiNTS [16]	5,067	77.94	48.07	63.00
	7,239	80.20	52.25	66.23
	8,802	80.06	52.53	66.29
HyperSegNAS	5,358	79.92	52.44	66.18
	7,168	80.10	53.47	66.79
	9,173	79.98	54.88	67.43

Table 2. Performance comparison for architectures found under different compute constraints, based on five-fold cross-validation¹ on the Pancreas dataset.

DiNTS [16] yields a low-cost architecture that down-samples all features to the lowest resolutions. Such an architecture differs from conventional designs, where multi-scale features are desired for segmenting objects of various sizes; as a result, the performance of this architecture is worse than HyperSegNAS’s low-cost architecture. We attribute DiNTS’s counter-intuitive results to differentiable NAS formulation. In DiNTS, an edge can have a near-zero importance (η_j^k in [16]) and estimated memory cost while still propagating its features forward, leading to a significant discrepancy between search and deployment. In comparison, HyperSegNAS evaluates on discrete representations, and reflects its performance in deployment more closely. In fact, *our smallest architecture only performs marginally worse than the largest architecture from DiNTS* in Table 2, demonstrating HyperSegNAS’s ability to truly adapt to different computing constraints.

Secondly, skip connections are used extensively in our architectures. This shows the advantage in occupying as much of the search space as possible, as these connections make the original features more available to other parts of the network with little cost. Such an observation generally agrees with previous works like ResNet [15] and DenseNet [17], where features are shuffled through skip connections to achieve less information loss and improved representations. In particular, we find parallel skip connections to be very efficient, as they introduce very little computing overhead and consistently improve architecture performance according to evaluations based on the super-net. This is in agreement with the design of U-Net [32], where parallel skip connections are also used extensively.

HyperNet Analysis. We provide some insights into the behavior of the HyperNet \mathcal{H} and its output $\omega_{\mathcal{H}_\theta}^{a,I}$. While it is difficult to interpret the channel-wise weights individually,




Figure 6. Visualization of changing $\omega_{\mathcal{H}_\theta}^{a,I}$ depends on different image or architecture inputs. We use our 5GB and 7GB architectures as a_i and a_j in 6b. Results are aggregated based on sliding-window averaging with a stride size of sixteen.

we can observe how $\omega_{\mathcal{H}_\theta}^{a,I}$ changes with different inputs. To do so, we fix the architecture vector l_{arch} and provide \mathcal{H} with l_{image} generated from different patches. Similarly, we can fix l_{image} and change l_{arch} to observe the differences based on architectures. We observe that the majority of the patches yields very similar $\omega_{\mathcal{H}_\theta}^{a,I}$, while a few patches have very different weights. To quantify, we define two metric:

$$\Delta\omega_{\mathcal{H}_\theta}^a(I) = \|\omega_{\mathcal{H}_\theta}^{a,I} - \bar{\omega}_{\mathcal{H}_\theta}^a\|_2 \quad (6)$$

$$\Delta\omega_{\mathcal{H}_\theta}^I(a_i, a_j) = \|\omega_{\mathcal{H}_\theta}^{a_i,I} - \omega_{\mathcal{H}_\theta}^{a_j,I}\|_2, \quad (7)$$

where $\Delta\omega_{\mathcal{H}_\theta}^a(I)$ measures the \mathcal{L}_2 distance between $\omega_{\mathcal{H}_\theta}^{a,I}$, which is from a specific patch I , and $\bar{\omega}_{\mathcal{H}_\theta}^a$, which is the average weights over all patches in a given volume; $\Delta\omega_{\mathcal{H}_\theta}^I(a_i, a_j)$ measures the difference between weights generated from architecture a_i and a_j on the same patch.

We visualize these two metrics on a sample CT volume in Fig. 6, along with the segmentation labels. For patches that are not relevant to the labels, \mathcal{H} generates very similar $\omega_{\mathcal{H}_\theta}^{a,I}$; on the other hand, patches that contain labels yield significantly different weights. Furthermore, $\Delta\omega_{\mathcal{H}_\theta}^I(a_i, a_j)$ is also significant on foreground patches, and minimum on background patches. This suggests that \mathcal{M} and \mathcal{H} implicitly partitions potential foreground regions, which is a strategy similar to foreground oversampling used in methods like nnU-Net [19]. While our method does not employ explicit foreground oversampling, the HyperNet design appears to automate such a strategy. For more details on the visualization of $\omega_{\mathcal{H}_\theta}^{a,I}$, please refer to the Supplemental Material.

4.3. Quantitative Evaluation on MSD

We test our largest architecture on five tasks in the MSD challenge - the Pancreas and Tumor, Lung Tumor, Hepatic Vessel and Tumor, Brain Tumour, Hippocampus datasets. We select these tasks as a simplified and representative set of MSD to make experiments more efficient, as these tasks cover 80% of all data in MSD and include multiple organs and imaging modalities. The DSC and Normalised Surface Distance (NSD) scores are reported from the MSD leaderboard, and listed in Table 3. HyperSegNAS’ architecture

	Pancreas and Tumor						Lung Tumor		Hepatic Vessel and Tumor					
Method	DSC1	DSC2	Avg.	NSD1	NSD2	Avg.	DSC1	NSD1	DSC1	DSC2	Avg.	NSD1	NSD2	Avg.
Kim et al [21]	80.61	51.75	66.18	95.83	73.09	84.46	63.10	62.51	62.34	68.63	65.49	83.22	78.43	80.83
nnUNet [19]	81.64	52.78	67.21	96.14	71.47	83.81	73.97	76.02	66.46	<u>71.78</u>	69.12	84.43	80.72	<u>82.58</u>
C2FNAS [44]	80.76	54.41	67.59	96.16	75.58	85.87	70.44	72.22	64.30	71.00	67.65	83.78	80.66	82.22
DiNTS [16]	<u>81.02</u>	<u>55.35</u>	<u>68.19</u>	<u>96.26</u>	<u>75.90</u>	<u>86.08</u>	<u>74.75</u>	<u>77.02</u>	<u>64.50</u>	71.76	68.13	83.98	<u>81.03</u>	82.51
HyperSegNAS	80.99	56.16	68.58	96.30	77.68	86.99	76.72	79.67	64.35	72.10	<u>68.23</u>	<u>84.08</u>	81.14	82.61

	Brain Tumor						Hippocampus							
Method	DSC1	DSC2	DSC3	Avg.	NSD1	NSD2	NSD3	Avg.	DSC1	DSC2	Avg.	NSD1	NSD2	Avg.
Kim et al [21]	67.40	45.75	68.26	60.47	86.65	72.03	90.28	82.99	90.11	88.72	89.42	97.77	<u>97.73</u>	<u>97.75</u>
nnUNet [19]	68.04	46.81	68.46	61.10	87.51	72.47	90.78	83.59	90.23	88.69	89.46	<u>97.79</u>	97.53	97.66
C2FNAS [44]	67.62	48.60	69.72	61.98	87.61	72.87	91.16	83.88	89.37	87.96	88.67	97.27	97.35	97.31
DiNTS [16]	<u>69.28</u>	<u>48.65</u>	<u>69.75</u>	<u>62.56</u>	<u>89.33</u>	73.16	<u>91.69</u>	84.73	89.91	88.41	89.16	97.76	97.56	97.66
HyperSegNAS	69.32	49.11	70.01	62.81	89.76	<u>72.58</u>	91.74	<u>84.69</u>	<u>90.21</u>	<u>88.47</u>	<u>89.34</u>	98.11	97.7	97.91

Table 3. Dice-Sørensen score (DSC) and Normalised Surface Distance (NSD) performance on the Brain Tumor (MRI), Hippocampus (MRI), Lung Tumor (CT), Pancreas and Tumor (CT), Hepatic Vessel and Tumor (CT) datasets from MSD. All scores are from the current MSD leaderboard. The best and second best performances are **bold** and underlined.

shows an overall improvement on the Pancreas dataset compared to previous NAS-based methods, particularly on Pancreatic Tumour (DSC2), which is the more challenging and clinically relevant label. For tasks of the same CT modality, our architecture significantly outperforms all baselines in Lung Tumour segmentation, and is 2% better than the previous SOTA from DiNTS in DSC. General improvements can also be found in MRI segmentation tasks against other methods, showing a generality to our architecture.

We note that while nnU-Net [24] outperforms HyperSegNAS on some metrics in Hepatic Vessel and Hippocampus, it achieves so by (1) manually searching for the best hyper-parameters and (2) performing ensemble over different variations of U-Net. Specifically, to generate results for a single task, nnU-Net trains *four* architecture configurations each with five-fold cross validation, which leads to an ensemble of *twenty* models. HyperSegNAS only performs ensemble over five models from cross validation, uses no task-specific hyper-parameter tuning, and obtains similar if not better performances compared to nnU-Net overall.

4.4. Limitation and Future Directions

This work mainly focuses on obtaining general segmentation architectures. Problem specific factors like hyper-parameter selection and data augmentation methods are also important, particularly on the smaller datasets in MSD which we do not explore. The use of 2D models or 2D operators can also be beneficial to tasks where the volume resolution within-slice and between-slice is very different.

As the topology search space for segmentation is expansive, training a super-net requires high compute cost. While the ability to compare architectures yields useful in-

sights, e.g., on the incorporation of skip connections, we plan to investigate ways to partition the segmentation search space into more manageable pieces and improve search efficiency.

5. Conclusions

In this work, we present HyperSegNAS, a novel NAS algorithm for medical image segmentation. HyperSegNAS finds well-performing segmentation architectures by estimating their performances through a trained super-net. To address the complexity in the training process due to a large and unordered segmentation search space, we propose to use a HyperNet called Meta-Assistant Network (MAN) to improve learning. MAN incorporates high level information on the deploying architecture and the input image to output channel-wise weights, which modify the inference process for better performance. To allow fair evaluations during the searching stage, HyperSegNAS then removes MAN through a annealing process. We perform extensive experiments and show that our method can effectively adapt to different compute constraints both in terms of precision in matching the constraints and segmentation performance. Specifically, our low-compute architecture significantly outperforms DiNTS’s architecture on the Pancreas dataset. We also test our large-compute architecture on five representative tasks in MSD, and achieve SOTA performances respectively. Our analysis of architecture performance estimations show that the inclusion of skip connections is crucial to finding well-performing architectures, which agrees with the intuition from a history of manual architecture designs and is not reflected in previous methods.

References

- [1] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A. Landman, Geert Litjens, Bjoern H. Menze, Olaf Ronneberger, Ronald M. Summers, Bram van Ginneken, Michel Bilello, Patrick Bilic, Patrick Ferdinand Christ, Richard K. G. Do, Marc Gollub, Stephan Heckers, Henkjan J. Huisman, William R. Jarnagin, Maureen McHugo, Sandy Napel, Jennifer Goli-Pernicka, Kawal S. Rhode, Catalina Tobon-Gomez, Eugene Vorontsov, James A. Meakin, Sébastien Ourselin, Manuel Wiesenfarth, Pablo Arbelaez, Byeonguk Bae, Sihong Chen, Laura Alexandra Daza, Jianjiang Feng, Baochun He, Fabian Isensee, Yuanfeng Ji, Fucang Jia, Namkug Kim, Ildoo Kim, Dorit Merhof, Akshay Pai, BeomHee Park, Mathias Perslev, Ramin Rezaifar, Oliver Rippel, Ignacio Sarasua, Wei Shen, Jaemin Son, Christian Wachinger, Liansheng Wang, Yan Wang, Yingda Xia, Daguang Xu, Zhanwei Xu, Yefeng Zheng, Amber L. Simpson, Lena Maier-Hein, and M. Jorge Cardoso. The medical segmentation decathlon. *CoRR*, abs/2106.05735, 2021. [1](#), [2](#), [3](#)
- [2] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc V. Le. Understanding and simplifying one-shot architecture search. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 549–558. PMLR, 2018. [2](#), [3](#)
- [3] Andrew Brock, Theodore Lim, James M. Ritchie, and Nick Weston. SMASH: one-shot model architecture search through hypernetworks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [2](#), [3](#), [4](#)
- [4] Xingyuan Bu, Junran Peng, Junjie Yan, Tieniu Tan, and Zhaoxiang Zhang. Gaia: A transfer learning system of object detection that fits your needs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 274–283, June 2021. [5](#)
- [5] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [1](#), [2](#), [3](#), [5](#), [6](#)
- [6] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [2](#), [3](#)
- [7] Minghao Chen, Jianlong Fu, and Haibin Ling. One-shot neural ensemble architecture search by diversity-guided search space shrinking. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 16530–16539. Computer Vision Foundation / IEEE, 2021. [2](#), [3](#)
- [8] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 1294–1303. IEEE, 2019. [2](#)
- [9] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: Learning dense volumetric segmentation from sparse annotation. In Sébastien Ourselin, Leo Joskowicz, Mert R. Sabuncu, Gözde B. Ünal, and William Wells, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016 - 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II*, volume 9901 of *Lecture Notes in Computer Science*, pages 424–432, 2016. [1](#), [2](#)
- [10] Jose Dolz, Karthik Gopinath, Jing Yuan, Herve Lombaert, Christian Desrosiers, and Ismail Ben Ayed. Hyperdense-net: A hyper-densely connected CNN for multi-modal image segmentation. *IEEE Trans. Medical Imaging*, 38(5):1116–1126, 2019. [1](#), [2](#)
- [11] Yanjie Gao, Yu Liu, Hongyu Zhang, Zhengxian Li, Yonghao Zhu, Haoxiang Lin, and Mao Yang. Estimating GPU memory consumption of deep learning models. In Prem Devanbu, Myra B. Cohen, and Thomas Zimmermann, editors, *ESEC/FSE '20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*, pages 1342–1352. ACM, 2020. [5](#)
- [12] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XVI*, volume 12361 of *Lecture Notes in Computer Science*, pages 544–560. Springer, 2020. [2](#), [3](#)
- [13] David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. [4](#)
- [14] Stephanie A. Harmon, Thomas H. Sanford, Sheng Xu, Evrim B. Turkbey, Holger Roth, Ziyue Xu, Dong Yang, Andriy Myronenko, Victoria Anderson, Amel Amalou, Maxime Blain, Michael Kassin, Dilara Long, Nicole Varble, Stephanie M. Walker, Ulas Bagci, Anna Maria Ierardi, Elvira Stellato, Guido Giovanni Plensich, Giuseppe Franceschelli, Cristiano Girlando, Giovanni Irmici, Dominic Labella, Dima Hammoud, Ashkan Malayeri, Elizabeth Jones, Ronald M. Summers, Peter L. Choyke, Daguang Xu, Mona Flores, Kaku Tamura, Hirofumi Obinata, Hitoshi Mori, Francesca Patella, Maurizio Cariati, Gianpaolo Carrafiello, Peng An, Bradford J. Wood, and Baris Turkbey. Artificial intelligence for the detection of covid-19 pneumonia on chest ct using multinational datasets. *Nature Communications*, 11(1):4080, Aug 2020. [1](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE*

- Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. [7](#)
- [16] Yufan He, Dong Yang, Holger Roth, Can Zhao, and Daguang Xu. Dints: Differentiable neural network topology search for 3d medical image segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 5841–5850. Computer Vision Foundation / IEEE, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [17] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2261–2269. IEEE Computer Society, 2017. [7](#)
- [18] Sian-Yao Huang and Wei-Ta Chu. Searching by generating: Flexible and efficient one-shot NAS with architecture generator. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 983–992. Computer Vision Foundation / IEEE, 2021. [2](#), [3](#)
- [19] Fabian Isensee, Paul F. Jaeger, Simon A. A. Kohl, Jens Petersen, and Klaus H. Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2):203–211, Feb 2021. [2](#), [7](#), [8](#)
- [20] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938. [6](#)
- [21] Sungwoong Kim, Ildoo Kim, Sungbin Lim, Woonhyuk Baek, Chiheon Kim, Hyungjoo Cho, Boogyeon Yoon, and Taesup Kim. Scalable neural architecture search for 3d medical image segmentation. In Dinggang Shen, Tianming Liu, Terry M. Peters, Lawrence H. Staib, Caroline Essert, Sean Zhou, Pew-Thian Yap, and Ali R. Khan, editors, *Medical Image Computing and Computer Assisted Intervention - MICCAI 2019 - 22nd International Conference, Shenzhen, China, October 13-17, 2019, Proceedings, Part III*, volume 11766 of *Lecture Notes in Computer Science*, pages 220–228. Springer, 2019. [1](#), [3](#), [8](#)
- [22] Changlin Li, Jiefeng Peng, Liuchun Yuan, Guangrun Wang, Xiaodan Liang, Liang Lin, and Xiaojun Chang. Blockwisely supervised neural architecture search with knowledge distillation. *CoRR*, abs/1911.13053, 2019. [2](#), [3](#)
- [23] Xiaomeng Li, Hao Chen, Xiaojuan Qi, Qi Dou, Chi-Wing Fu, and Pheng-Ann Heng. H-denseunet: Hybrid densely connected unet for liver and tumor segmentation from CT volumes. *IEEE Trans. Medical Imaging*, 37(12):2663–2674, 2018. [2](#)
- [24] Xiaomeng Li, Hao Chen, Xiaojuan Qi, Qi Dou, Chi-Wing Fu, and Pheng Ann Heng. H-denseunet: Hybrid densely connected unet for liver and tumor segmentation from ct volumes, 2018. [2](#), [8](#)
- [25] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L. Yuille, and Li Fei-Fei. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 82–92. Computer Vision Foundation / IEEE, 2019. [1](#), [2](#), [3](#)
- [26] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: differentiable architecture search. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [1](#), [3](#)
- [27] Siqi Liu, Daguang Xu, Shaohua Kevin Zhou, Olivier Pauly, Sasa Grbic, Thomas Mertelmeier, Julia Wicklein, Anna K. Jerebko, Weidong Cai, and Dorin Comaniciu. 3d anisotropic hybrid network: Transferring convolutional features from 2d images to 3d anisotropic volumes. In Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger, editors, *Medical Image Computing and Computer Assisted Intervention - MICCAI 2018 - 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II*, volume 11071 of *Lecture Notes in Computer Science*, pages 851–858. Springer, 2018. [1](#), [2](#)
- [28] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Fourth International Conference on 3D Vision, 3DV 2016, Stanford, CA, USA, October 25-28, 2016*, pages 565–571. IEEE Computer Society, 2016. [1](#), [2](#)
- [29] Miguel Monteiro, Virginia FJ Newcombe, Francois Mathieu, Krishma Adatia, Konstantinos Kamnitsas, Enzo Ferrante, Tilak Das, Daniel Whitehouse, Daniel Rueckert, David K Menon, et al. Multiclass semantic segmentation and quantification of traumatic brain injury lesions on head ct using deep learning: an algorithm development and multicenter validation study. *The Lancet Digital Health*, 2(6):e314–e322, 2020. [1](#)
- [30] Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 4092–4101. PMLR, 2018. [2](#), [3](#)
- [31] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. Regularized evolution for image classifier architecture search. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 4780–4789. AAAI Press, 2019. [3](#)
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells III, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015 - 18th International Conference Munich, Germany, October 5 - 9, 2015, Proceedings, Part III*, volume 9351 of *Lecture Notes in Computer Science*, pages 234–241. Springer, 2015. [1](#), [2](#), [7](#)
- [33] Shreyas Saxena and Jakob Verbeek. Convolutional neural fabrics. In Daniel D. Lee, Masashi Sugiyama, Ulrike von

- Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4053–4061, 2016. [2](#), [3](#)
- [34] Xiu Su, Shan You, Fei Wang, Chen Qian, Changshui Zhang, and Chang Xu. Bcnet: Searching for network width with bilaterally coupled network. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 2175–2184. Computer Vision Foundation / IEEE, 2021. [2](#)
- [35] Yunjie Tian, Chang Liu, Lingxi Xie, Jianbin Jiao, and Qixiang Ye. Discretization-aware architecture search. *Pattern Recognit.*, 120:108186, 2021. [2](#)
- [36] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Instance normalization: The missing ingredient for fast stylization. *CoRR*, abs/1607.08022, 2016. [5](#)
- [37] Yu Weng, Tianbao Zhou, Yujie Li, and Xiaoyu Qiu. Nasnet: Neural architecture search for medical image segmentation. *IEEE Access*, 7:44247–44257, 2019. [1](#), [3](#)
- [38] Yingda Xia, Fengze Liu, Dong Yang, Jinzheng Cai, Lequan Yu, Zhuotun Zhu, Daguang Xu, Alan L. Yuille, and Holger Roth. 3d semi-supervised learning with uncertainty-aware multi-view co-training. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2020, Snowmass Village, CO, USA, March 1-5, 2020*, pages 3635–3644. IEEE, 2020. [1](#), [3](#)
- [39] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. PC-DARTS: partial channel connections for memory-efficient architecture search. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [1](#), [3](#)
- [40] Xingang Yan, Weiwen Jiang, Yiyu Shi, and Cheng Zhuo. MS-NAS: multi-scale neural architecture search for medical image segmentation. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz, editors, *Medical Image Computing and Computer Assisted Intervention - MICCAI 2020 - 23rd International Conference, Lima, Peru, October 4-8, 2020, Proceedings, Part I*, volume 12261 of *Lecture Notes in Computer Science*, pages 388–397. Springer, 2020. [1](#), [3](#)
- [41] Shan You, Tao Huang, Mingmin Yang, Fei Wang, Chen Qian, and Changshui Zhang. Greedynas: Towards fast one-shot NAS with greedy supernet. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 1996–2005. Computer Vision Foundation / IEEE, 2020. [2](#), [3](#)
- [42] Jiahui Yu and Thomas S. Huang. Universally slimmable networks and improved training techniques. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 1803–1811. IEEE, 2019. [2](#), [3](#)
- [43] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas S. Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Scaling up neural architecture search with big single-stage models. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VII*, volume 12352 of *Lecture Notes in Computer Science*, pages 702–717. Springer, 2020. [1](#), [2](#), [3](#), [6](#)
- [44] Qihang Yu, Dong Yang, Holger Roth, Yutong Bai, Yixiao Zhang, Alan L. Yuille, and Daguang Xu. C2FNAS: coarse-to-fine neural architecture search for 3d medical image segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 4125–4134. Computer Vision Foundation / IEEE, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)
- [45] Chris Zhang, Mengye Ren, and Raquel Urtasun. Graph hypernetworks for neural architecture search. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [3](#)
- [46] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: Redesigning skip connections to exploit multiscale features in image segmentation. *IEEE Trans. Medical Imaging*, 39(6):1856–1867, 2020. [2](#)
- [47] Zhuotun Zhu, Chenxi Liu, Dong Yang, Alan L. Yuille, and Daguang Xu. V-NAS: neural architecture search for volumetric medical image segmentation. In *2019 International Conference on 3D Vision, 3DV 2019, Québec City, QC, Canada, September 16-19, 2019*, pages 240–248. IEEE, 2019. [1](#), [3](#)
- [48] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. [1](#), [3](#)