

DCNAS: Densely Connected Neural Architecture Search for Semantic Image Segmentation

Xiong Zhang^{1*}, Hongmin Xu², Hong Mo³, Jianchao Tan², Cheng Yang¹, Lei Wang¹, Wenqi Ren⁴

¹Joyy Inc., ² AI Platform, Kwai Inc., ³State Key Lab of VR, Beihang University, ⁴SKLOIS, IIE,CAS

Abstract

Existing NAS methods for dense image prediction tasks usually compromise on restricted search space or search on proxy task to meet the achievable computational demands. To allow as wide as possible network architectures and avoid the gap between realistic and proxy setting, we propose a novel Densely Connected NAS (DCNAS) framework, which directly searches the optimal network structures for the multi-scale representations of visual information, over a large-scale target dataset without proxy. Specifically, by connecting cells with each other using learnable weights, we introduce a densely connected search space to cover an abundance of mainstream network designs. Moreover, by combining both path-level and channel-level sampling strategies, we design a fusion module and mixture layer to reduce the memory consumption of ample search space, hence favoring the proxyless searching. Compared with contemporary works, experiments reveal that the proxyless searching scheme is capable of bridging the gap between searching and training environments. Further, DCNAS achieves new state-of-the-art performances on public semantic image segmentation benchmarks, including 84.3% on Cityscapes, and 86.9% on PASCAL VOC 2012. We also retain leading performances when evaluating the architecture on the more challenging ADE20K and PASCAL-Context dataset.

1. Introduction

The heavy computation overheads of early Neural Architecture Search (NAS) methods [81] hinder their applications in real-world problems. Recently, limited search space strategies [82, 37, 38] significantly shorten the searching time of NAS algorithms, which make NAS approaches achieve superhuman performance in image classification task. However, to meet the low consumption of searching time, NAS methods with a constrained search space throw down the multi-scale representation of high-resolution im-

Method	Proxyless	GPU Days	FLOPs(G)	$\rho \uparrow$	$\tau \uparrow$	mIoU(%) \uparrow
DPC [3]	\times	2600	684.0	0.46	0.37	82.7
Auto-DeepLab [36]	\times	3	695.0	0.31	0.21	82.1
CAS [73]	-	-	-	-	-	72.3
GAS [35]	-	6.7	-	-	-	73.5
FasterSeg [8]	\times	2	28.2	0.35	0.25	71.5
Fast-NAS [44]	\times	8	435.7	0.42	0.33	78.9
SparseMask [59]	\times	4.2	36.4	0.49	0.38	68.6
Ours (DCNAS)	\checkmark	5.6	294.6	0.73	0.55	84.3

Table 1. **Comparisons with state of the arts.** The table presents the comparison results of [3, 36, 73, 35, 8, 66, 59] and our DCNAS comprehensively. In which, ρ and τ represent the Pearson Correlation Coefficient and the Kendall Rank Correlation Coefficient, respectively, both of which are measured according to the performances of the super-net and the fine-tuned stand-alone model, **Proxyless** applies a proxyless searching paradigm, **GPU Days** presents the searching cost, **mIoU(%)** and **FLOPs(G)** report the accuracy and the flops of the best model.

age. As a result, those methods are not suitable for dense image prediction tasks (e.g. semantic image segmentation, object detection, and monocular depth estimation).

To efficiently search appropriate network structure that can combine both the local and global clues of the semantic features, researches recently focus on improving NAS frameworks by designing new search spaces to handle multi-scale features. For instance, DPC [3] introduces a recursive search space, and Auto-DeepLab [36] proposes a hierarchical search space. However, designing new search space for dense image prediction tasks proves challenging: one has to delicately compose a flexible search space that covers as much as possible optimal network architectures. Meanwhile, efficiently address memory consumption and the heavy computation problems accompanied with ample search space for the high-resolution imagery.

In this work, we propose an efficient and proxyless NAS framework to search the optimal model structure for semantic image segmentation. Our approach is based on two principal considerations. Firstly, the search space should be comprehensive enough to handle most of the mainstream architecture designs, even some undiscovered high-quality model structures. As shown in Figure 1, we design a reticular-like and fully densely connected search space, which contains various paths in the search space. Con-

*corresponding author, zhangxiong@yy.com




Figure 1. **Framework of DCNAS.** *Top:* The *densely connected* search space (DCSS) with layer L and max downsampling rate 32. To stable the searching procedure, we keep the beginning *STEM* and final *Upsample* block unchanged. *Dashed lines* represent candidate connections in DCSS, to keep clarity, we only demonstrate several connections among all the candidates. *Bottom Left:* *Fusion module* targets at aggregating feature-maps derived by previous layers, and solving the intensive GPU memory requirement problem by sampling a portion of connections from all possible ones. *Bottom Right:* *Mixture layer* may further save GPU memory consumption and accelerating the searching process by sampling and operating on a portion of features while bypassing the others.

sequently, DCNAS may derive whichever model architecture designs by selecting appropriate paths among the whole set of these connections. Additionally, DCNAS aggregates contextual semantics encoded in multi-scale imageries to derive long-range context information, which has been recently found to be vital in the state-of-the-art manually designed image segmentation models [40, 74, 5, 7, 70]. Secondly, we observe that previous NAS approaches [3, 36] heavily rely on the proxy task or proxy dataset to reduce the cost of GPU hours and to alleviate the high GPU memory consumption problem. However, architectures optimized over the proxy are not guaranteed to be suitable in realistic setting [2], because of the gap between the proxy and the target configuration. Regarding this concern, we relax the discrete architectures into continuous representation to save GPU hours and design a fusion module which applies both path-level and channel-level sampling strategies during the searching procedure to reduce memory demand. Based on that, one may employ the stochastic gradient descent (SGD) to perform the proxyless searching procedure to select the optimal architecture from all the candidate models without the help of proxy datasets or proxy tasks. The searching procedure takes about 5.6 GPU days on Cityscapes [11].

We apply our approach to semantic image segmentation tasks on several public benchmarks, our model achieves the best performance compared with state-of-the-art hand-craft models [51, 74, 80, 7, 24, 12, 17] and other contemporary NAS approaches [36, 3, 59, 44, 8]. We also evaluate the optimal model identified by DCNAS on PASCAL VOC 2012 [15], the model outperforms other leading approaches

[26, 17, 74, 68, 69, 27, 70, 21] and advances the state-of-the-art performance. Transferring the model to ADE20K [78] and PASCAL-Context [43] datasets, our model obtains the best result compared with state-of-the-art approaches [12, 17, 22, 24, 71, 69, 31], according to the corresponding evaluation metrics.

Further, considering NAS methods depend on agent metrics to explore promising networks efficiently. If the agent metric provides a strong relative ranking of models, this enables the discovery of high performing models when trained to convergence. However, studies about the relative ranking of models in the searching period are less explored in literature [44, 59, 8, 3, 36]. In this work, we enrich this field and investigate most related methods on how well the accuracy in searching phase for specific architectures correlates with that of the fine-tuned stand-alone model.

To summarize, our main contributions are as follows:

- We design a novel *entirely densely connected search space*, allowing to explore various existing designs and to cover arbitrary model architecture patterns.
- A novel *proxyless searching paradigm* is implemented to efficiently and directly explore the most promising model among all the candidates encoded in DCNAS, on large-scale segmentation datasets (e.g., Cityscapes [11]).
- The DCNAS outperforms contemporary NAS methods [44, 59, 8, 3, 36] and demonstrates new state-of-the-art performance on Cityscapes [11], PASCAL-VOC 2012 [15], PASCAL-Context [43] and ADE20K [78] datasets.

- We conduct a large scale experiment to understand the correlation of model performance between the searching and training periods for most contemporary works [36, 3, 59, 44, 8] focusing on semantic segmentation task.

2. Related Work

2.1. Neural Architecture Search Space

Early NAS works [81, 49] directly search the whole network architecture. Although those works achieve impressive results, their expensive computation overheads (e.g. thousands of GPU days) hinder their applications in real-world scenarios. To alleviate this situation, researchers proposed restricted search spaces. NASNet [82] first proposed a cell-based search space. Concretely, NASNet searches the fundamental cell structure, which is easier than searching the whole network architectures on a small proxy dataset. After that, many works [38, 47, 2, 56, 58] adopt the cell-based search space design and make improvements in several ways. One line of works [2, 56, 58] attempt to search high-efficiency network architectures for resource-constrained platforms, such as mobile phones. Another line of works [9, 16] try to improve previous cell-based NAS methods. P-DRATS [9] first bridge the depth gap between search and evaluation phases. Fang *et al.* [16] proposes the routing blocks to determine the number of blocks automatically.

Despite the success of repeated cell-based methods for image classification tasks, other dense image prediction tasks, such as semantic image segmentation and object detection, demand more delicate and complicated network architectures to capture multi-scale information of the image. Therefore, researchers begin to explore more flexible search space for dense image prediction tasks. For instance, DPC [3] proposes a recursive search space to build a multi-scale representation of a high-resolution image. Auto-DeepLab [36] introduces a hierarchical search space to exhibit multi-level representations of the image.

2.2. Semantic Image Segmentation

The pioneer work FCN [42] makes great progress in semantic segmentation with a fully convolutional network. Afterwards, researchers devote enormous efforts to explore the tremendous potentials of convolutional neural networks (CNNs) on dense image prediction tasks. Several kinds of research works [74, 34, 6, 14, 4, 50, 46, 33, 48, 7, 18] focus on how to efficiently utilize multi-scale context information to improve the performance of semantic segmentation systems. Some works [34, 6, 14] takes an image pyramid as input to capture larger objects in down-sampled input image. Some other works try to make use of multi-scale context information. PSPNet [74] uses spatial pyramid pooling at different manually set scales. Deeplab [4] performs multiple atrous convolution operations. Also, there are some

works [50, 46, 33, 48, 7, 18] that employ encoder-decoder structure to capture long-range context information.

Recent works [3, 36, 44, 52, 59, 8, 30, 73, 72, 60, 41] shift to automatically design network architectures. Research works, such as [44, 52, 59, 8, 30, 73], target to search network structures that are friendly to resource-constrained devices for semantic segmentation task. Other works, such as DPC [3] and Auto-DeepLab [36], intent to directly search the highest-performance network architecture for semantic segmentation without considering computation overheads. Most of those works [36, 44, 52, 59, 8, 30, 73] suffer from limited search space and search on proxy tasks, because of high GPU memory consumption and expensive computational burden. Yet the optimal network structures for multi-scale feature representations require ample search space.

The most similar works to ours are DPC [3], Auto-DeepLab [36], and DenseNAS [16]. [3] proposes an ample search space but requires thousands of GPU days for searching, whereas we can efficiently search over at least the same search space using only 5.6 GPU days. [36] achieves comparable searching time like ours, yet contains a much smaller search space than DCNAS. [16] places candidate connections between adjacent M cells only, while DCNAS comprises a wholly densely connected search space. Besides, both [3, 36, 16] search on a proxy dataset, whereas DCNAS directly search on the target dataset.

3. Methodology

In this section, we introduce several vital components that constitute DCNAS in detail, including a densely connected search space that is general enough to cover enormous network design, the derived continuously differentiable representation of the search space, and the approach that can save the cost of GPU hours and reduce excessive memory footprint, which favors proxyless searching.

3.1. Densely Connected Search Space

The densely connected search space (**DCSS**) involves two primitives, the *mixture layer*, and the *fusion module*. The *mixture layer* is defined as a mixture of the candidate operators, and the *fusion module* aims at aggregating semantic features from preceding fusion modules. A range of reticular-like and densely connected fusion modules constitute the architecture of DCSS, as shown in Figure 1.

Mixture Layer. The *mixture layer* is the elementary structure in search space represents a collection of available operations. Similar to [2], we construct the operator space \mathcal{O} with various configurations of mobileNet-v3 [23], i.e., kernel sizes $k \in \{3, 5, 7\}$, expansion ratios $r \in \{3, 6\}$. Due to the dense connections contained in the search space, the DCSS inherently supports the zero operation and identity mapping. Hence \mathcal{O} does not contain the above operators

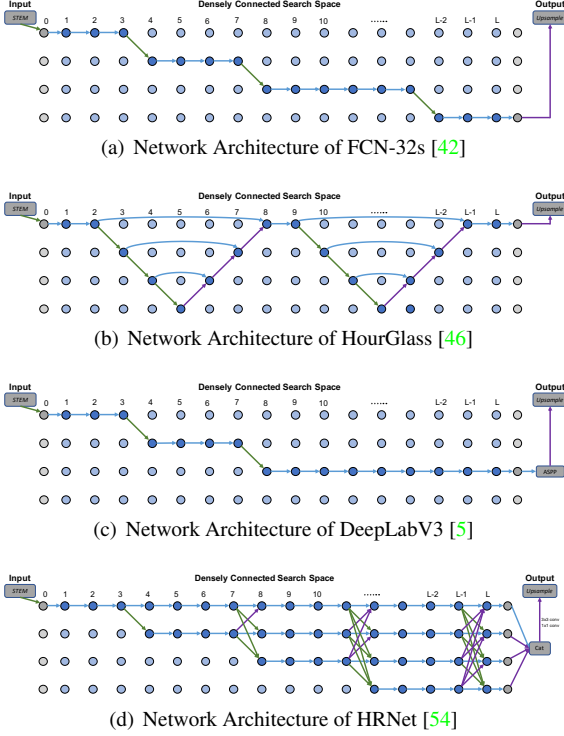


Figure 2. **Mainstream Network Designs in DCNAS.** The figure demonstrates several mainstream network design patterns [42, 46, 5, 54] that embedded in our *Densely Connected Search Space*.

explicitly. Similarly, considering the search space inherently favors the multi-scale semantic features aggregation, which increases the receptive field, therefore, the mixture layer does not support atrous convolution explicitly.

Fusion Module. To explore various paths in DCSS, we introduce the *fusion module* with the ability of aggregating semantic features from preceding fusion modules besides the feature pyramids and attaching transformed semantic features to succeeding ones. As illustrated by the bottom-left part in Figure 1, the fusion module consists of the shape-alignment layer besides the mixture layer. The shape-alignment layer is in a multi-branch parallel form. Given an array of feature-maps with different shapes (e.g., spatial resolutions and channel widths), the shape-alignment layer dispatches each feature-map to the corresponding branches to align them to the target shape. Semantic features are well-aligned and fully aggregated, then feed into the mixture layer to perform efficient multi-scale features fusion.

Search Space. Benefiting from the sophisticated design of the fusion module, we introduce a set of reticular-arranged fusion modules and place dense connections between them to form a search space that supports more flexible architecture configurations than similar works [36, 8]. The search algorithm is admitted to select a subset of the fusion models together with the appropriate connections be-

tween them to derive whatever model architecture (see Figure 2 for instance), hence enables a general architecture search. Specifically, we construct the super network with an array of fusion modules $\{\mathcal{M}_{(s,l)}\}$, where s indicates the spatial configuration and l refers to the layer index, which is demonstrated in Figure 1, in which, each $\mathcal{M}_{(s,l)}$ aggregates the semantic features coming from $\{\mathcal{M}_{(s',l')}\}_{l' < l}$.

3.2. Differentiable Representation

Reinforcement learning representations in [1, 81] and evolutionary representations in [53, 49] both tend to be computationally intensive, hence probably not suitable for semantic image segmentation task. Regarding this concern, we draw on the experience of continuous relaxation in [38, 62] and relax the discrete architectures into continuous representation. Notably, we derive a continuous representation for both the mixture layer and the fusion module, hence leading to a fully differentiable search space, based on which one can apply the stochastic gradient descent method to search promising architectures.

Mixture Layer. We assign an architecture parameter $\alpha_{(s,l)}^o$ to each operator o in mixture layer $\ell_{(s,l)}$ contained in fusion module $\mathcal{M}_{(s,l)}$, and derive the continuously representation of the mixture layer by defining it as a weighted sum of outputs from all candidate operations. The architecture weight of the operation is computed as a soft-max over all architecture parameters in mixture layer $\ell_{(s,l)}$,

$$w_{(s,l)}^o = \frac{\exp\{\alpha_{(s,l)}^o\}}{\sum_{o' \in \mathcal{O}} \exp\{\alpha_{(s,l)}^{o'}\}}, \quad (1)$$

and the output of mixture layer $\ell_{(s,l)}$ can be estimated as,

$$O_{(s,l)} = \sum_{o \in \mathcal{O}} w_{(s,l)}^o o(\mathcal{I}_{(s,l)}), \quad (2)$$

where $\mathcal{I}_{(s,l)}$ refers to the input feature-maps of $\ell_{(s,l)}$.

Fusion Module. Similar to the relaxation paradigm adopted in the mixture layer, to relax the connections as a continuous representation, we assign an architecture parameter $\beta_{(s',l') \rightarrow (s,l)}$ for the path from $\mathcal{M}_{(s',l')}$ to $\mathcal{M}_{(s,l)}$. Since the $\mathcal{M}_{(s,l)}$ aggregates all semantic features $\{\mathcal{I}_{(s',l')}\}_{l' < l}$, we estimate the transmission probability of each path using a soft-max function over all available connections,

$$p_{(s',l') \rightarrow (s,l)} = \frac{\exp\{\beta_{(s',l') \rightarrow (s,l)}\}}{\sum_{s'' \in \mathbb{S}, l'' < l} \exp\{\beta_{(s'',l'') \rightarrow (s,l)}\}}, \quad (3)$$

and the aggregating process can be formalized as,

$$\mathcal{I}_{(s,l)} = \sum_{s' \in \mathbb{S}, l' < l} p_{(s',l') \rightarrow (s,l)} H_{(s',l') \rightarrow (s,l)}(O_{(s',l')}), \quad (4)$$

where $H_{(s',l') \rightarrow (s,l)}$ denotes the corresponding transformation branch in shape-alignment layer of fusion module $\mathcal{M}_{(s,l)}$, which transforms the semantic features produced by mixture layer $\ell_{(s',l')}$ to the shape specified by $\mathcal{M}_{(s,l)}$.

3.3. Reducing Memory Footprint

Continuous representation of network architectures can largely reduce the cost of GPU hours, while the memory footprint grows linearly w.r.t. the size of candidate operation set and the number of connections, hence suffering from the high GPU memory consumption problem. While image segmentation intrinsically requires high-resolution semantic features, which may lead to even more GPU memory consumption. To resolve the excessive memory consumption problem, we apply the sampling strategy both in the fusion module and the mixture layer to solve the excessive memory footprint problem and further reduce the cost of GPU hours. Taking the above tactics, one can proxylessly search the model over large-scale dataset, hence avoiding the gap between target and proxy dataset.

Fusion Module. The DCSS comprises $8L(L + 1)$ connections in total, where L is layer number, thus making it impossible to optimize the whole set of candidate connections in each iteration because of the extravagant GPU memory demand. Therefore, in each iteration during the search, for each fusion module, we sample several connections among the corresponding potential connections. Concretely, taking fusion module $\mathcal{M}_{(s,l)}$ for instance, all the preceding fusion modules $\{\mathcal{M}_{(s',l')}\}_{l' < l}$ are associated with it, in each search iteration, we perform sampling without replacement to activate several transmission paths, and the probability distribution is defined as:

$$p_{(s',l') \rightarrow (s,l)} = \frac{\exp \{ \beta_{(s',l') \rightarrow (s,l)} / \tau \}}{\sum_{s'' \in \mathbb{S}, l'' < l} \exp \{ \beta_{(s'',l'') \rightarrow (s,l)} / \tau \}}, \quad (5)$$

where $\beta_{(s'',l'') \rightarrow (s,l)}$ share the same definition in Equation 3, and the temperature variable τ starts from a high temperature then anneal to a small but non-zero value. Assume there are n fusion modules $\{\mathcal{M}_{(s_i,l_i)}\}_{i < n}$ selected from $\{\mathcal{M}_{(s',l')}\}_{l' < l}$, then we may approximate Equation 4 with:

$$\mathcal{I}_{(s,l)} = \sum_{k < n} w_{(s_k,l_k) \rightarrow (s,l)} H_{(s_k,l_k) \rightarrow (s,l)} (O_{(s_k,l_k)}), \quad (6)$$

where $w_{(s_k,l_k) \rightarrow (s,l)}$ refers to the normalized blending weight, H and O shares the same definition in Equation 4.

Mixture Layer. A typical way to apply the sampling trick is to sample the operator o from the candidate operator space \mathcal{O} in each mixture layer. While in practice, we find that the naive sampling strategy makes the searching process unstable and sometime does not convergence. Inspired by [64], a portion of channels are sampled among the input features, then the lucky ones are sent into the mixed transformation of $|\mathcal{O}|$ operators while bypasses the others. Specifically, for each $\ell_{(s,l)}$, we assign a random variable $S_{(s,l)}$ that control the ratio of sampled channels, which trades off the search accuracy and efficiency, and we may

estimate $O_{(s,l)}$ in Equation 2 as:

$$\sum_{o \in \mathcal{O}} w_{(s,l)}^o o (S_{(s,l)} \mathcal{I}_{(s,l)}) + (1 - S_{(s,l)}) \mathcal{I}_{(s,l)}, \quad (7)$$

where $w_{(s,l)}^o$ indicates the normalized blending weight defined by Equation 1. While different to [64], the edge normalization layers are removed and the feature masks $S_{(s,l)}$ remain unchanged in the searching procedure, and experimental result demonstrates that this simplified variant proves more suitable for semantic image segmentation. Regarding the sampling ratio, we do not tune the sampling ratio and fix it to be $1/4$, empirically, our experiments demonstrate that it works well. With this scheme, one shall gain $4 \times$ acceleration of the searching process, and save about 75% GPU memory consumption.

3.4. Search Procedure

Benefiting from the fully differentiable representation of the search space and the sampling strategies mentioned above, one may apply the stochastic gradient descent algorithm to search for appropriate models by optimizing architecture parameters $\{\alpha, \beta\}$ on the target dataset without proxy. Similar to previous works [38, 3, 36, 8], we partition the training data into two parts *trainA* and *trainB*, which are used for updating the convolutional weights w and architecture parameters $\{\alpha, \beta\}$ respectively. We solve this optimization problem alternatively:

- update weights w by $\nabla_w \mathcal{L}_{trainA}(w, \alpha, \beta)$,
- update parameters $\{\alpha, \beta\}$ by $\nabla_{\alpha, \beta} \mathcal{L}_{trainB}(w, \alpha, \beta)$,

in which the loss function \mathcal{L} mainly consists of the classical cross entropy calculated on each mini-batch.

We shall point out that, in practice, \mathcal{L} comprises of several practical regularization terms in addition to cross-entropy, experiments reveal the regularization terms yield a faster convergence rate of the searching procedure and lead to a better searching accuracy (see supplementary material).

3.5. Decoding Network Structure

Once the search procedure terminates, one may derive the suitable operator for each mixture layer and the optimal architecture based on the architecture parameters α and β . For mixture layer $\ell_{(s,l)}$, we select the candidate operation that has maximum operation weight, i.e., $\arg \max_{o \in \mathcal{O}} \alpha_{(s,l)}^o$. In terms of the network architecture, we use Breadth-First Search algorithm (see supplemental material) to derive the network architecture.

4. Experiments

To evaluate DCNAS comprehensively, we first verify the effect of the proxyless searching paradigm by measuring the

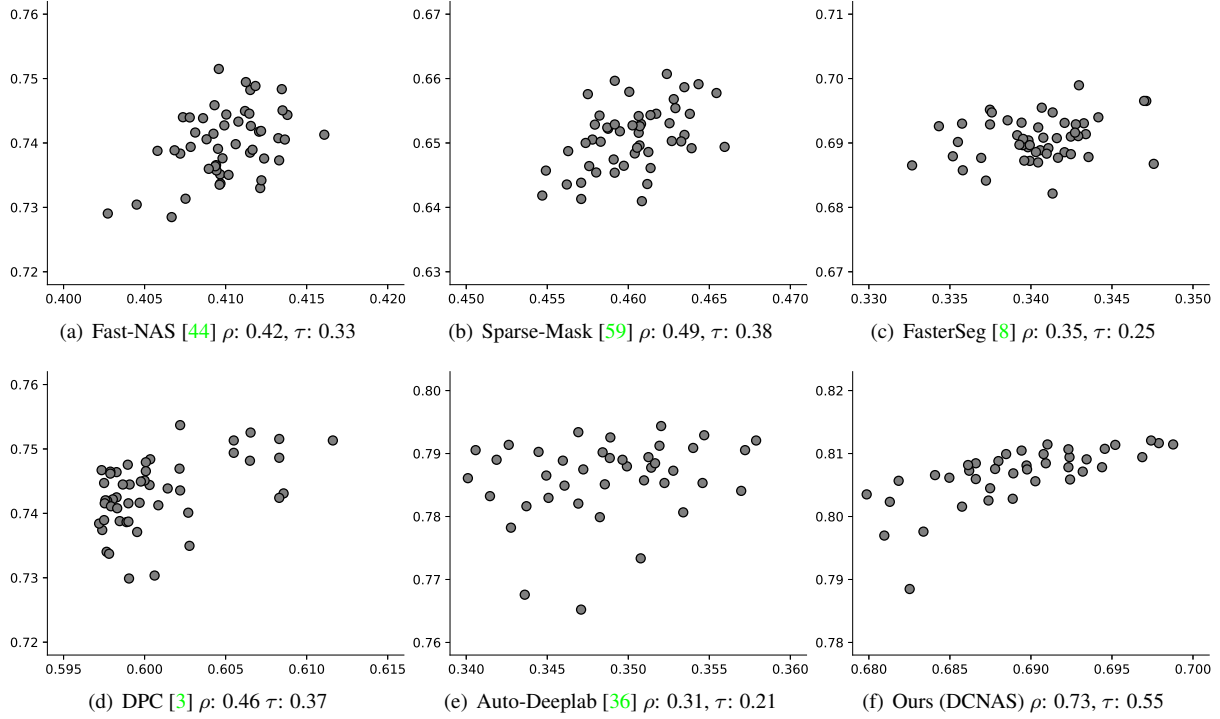


Figure 3. **Correlation of Performance.** The figures present the performance of the state of the art methods [44, 59, 8, 3, 36] and our DCNAS. For each method, we plot the validation performance for both the searching and training stages. Further, we report the Pearson Correlation Coefficient (ρ) and the Kendall Rank Correlation Coefficient (τ), which measured according to pairs of the performances.

correlation of the performance between the searching and training environments. Secondly, we quantitatively demonstrate the superiority of DCNAS on several broadly used benchmarks according to corresponding evaluation metrics. Finally, we perform ablation studies to better understand the impact of different designs on the segmentation tasks.

4.1. Architecture Search Implementation Details

We first search suitable model structure on Cityscapes [11] dataset and then evaluate this derived model on other benchmarks. In our experiments, for the shape of feature-maps, we set the spatial resolution space \mathbb{S} to be $\{1/4, 1/8, 1/16, 1/32\}$, and the corresponding widths are set to $F, 2F, 4F, 8F$, where we set F to be 64 for our best model. Regarding the DCSS, we set the hyper-parameter L to be 14 and the sampling ratio r to be $1/4$. To deliver end-to-end searching, we hand-craft a *stem module* that aims to extract feature-pyramids for DCSS, and we also design a simple *head block* for aggregating all of the feature-maps from DCSS to make the final prediction (We refer the supplemental materials for more implementation details).

4.2. Correlation of Performance

In this section, we quantify the association of the model’s performance during searching and training configurations. In practice, we take the Pearson Correlation Coefficient (ρ)

Method	Backbone	Validation	Coarse	ImageNet	mIoU(%)
PSPNet [74]	Dilated-ResNet-101	✓	✗	✓	78.4
PSANet [75]	Dilated-ResNet-101	✓	✗	✓	80.1
PADNet [63]	Dilated-ResNet-101	✓	✗	✓	80.3
Auto-Deeplab [36]	-	✓	✗	✗	80.4
DenseASPP [65]	WDenseNet-161	✓	✗	✓	80.6
SVCNet [12]	ResNet-101	✓	✗	✓	81.0
PSPNet [74]	Dilated-ResNet-101	✓	✓	✓	81.2
CPNet101 [67]	ResNet-101	✓	✗	✓	81.3
DeepLabv3 [5]	ResNet-101	✓	✓	✓	81.3
CCNet [24]	ResNet-101	✓	✗	✓	81.4
DANet [17]	Dilated-ResNet-101	✓	✗	✓	81.5
SpyGR [29]	ResNet-101	✓	✓	✓	81.6
RPCNet [76]	ResNet-101	✓	✗	✓	81.8
Mapillary [†] [51]	ResNeXt-101	✓	✓	✓	82.0
HANet [†] [10]	ResNet-101	✓	✗	✓	83.2
DeepLabv3+ [7]	Dilated-ResNet-101	✓	✓	✓	82.1
Auto-Deeplab [36]	-	✓	✓	✗	82.1
HRNetV2 + OCR [†] [57]	HRNetV2	✓	✗	✓	83.9
DPC [3]	Dilated-Xception-71	✓	✗	✓	82.7
DRN [80]	ResNet-101	✓	✓	✓	82.8
GSCNN [†] [55]	Wide-ResNet	✓	✗	✓	82.8
DecoupleSegNets [†] [28]	Wide-ResNet	✓	✗	✓	83.7
DCNAS [†]	-	✓	✗	✗	82.8
DCNAS [†]	-	✓	✗	✗	83.1
DCNAS [†]	-	✓	✓	✗	83.6
DCNAS + ASPP [†]	-	✓	✓	✗	84.3

Table 2. **Performance on Cityscapes.** The table summarizes the performance on Cityscapes testing dataset. **Validation:** Models trained with both train-fine and val-fine parts. **ImageNet:** The backbones of model trained on ImageNet. **Coarse:** Models that exploit extra datas in Cityscapes with coarse annotation. [†]: Adopts the well labeled Mapillary Vistas [45] dataset.

and Kendall Rank Correlation (τ) Coefficient as the evaluation metric and compare our approach with contemporary works [44, 59, 8, 3, 36].

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mIoU(%)
AAF [26]	91.2	72.9	90.7	68.2	77.7	95.6	90.7	94.7	40.9	89.5	72.6	91.6	94.1	88.3	88.8	67.3	92.9	62.6	85.2	74.0	82.2
ResNet38 [61]	94.4	72.9	94.9	68.8	78.4	90.6	90.0	92.1	40.1	90.4	71.7	89.9	93.7	91.0	89.1	71.3	90.7	61.3	87.7	78.1	82.5
DANet [17]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	82.6
PSPNet [74]	91.8	71.9	94.7	71.2	75.8	95.2	89.9	95.9	39.3	90.7	71.7	90.5	94.5	88.8	89.6	72.8	89.6	64.0	85.1	76.3	82.6
DFN [68]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	82.7
EncNet [69]	94.1	69.2	96.3	76.7	86.2	96.3	90.7	94.2	38.8	90.7	73.3	90.0	92.5	88.8	87.9	68.7	92.6	59.0	86.4	73.4	82.9
PAN [27]	95.7	75.2	94.0	73.8	79.6	96.5	93.7	94.1	40.5	93.3	72.4	89.1	94.1	91.6	89.5	73.6	93.2	62.8	87.3	78.6	84.0
CFNet [70]	95.7	71.9	95.0	76.3	82.8	94.8	90.0	95.9	37.1	92.6	73.0	93.4	94.6	89.6	88.4	74.9	95.2	63.2	89.7	78.2	84.2
APCNet [22]	95.8	75.8	84.5	76.0	80.6	96.9	90.0	96.0	42.0	93.7	75.4	91.6	95.0	90.5	89.3	75.8	92.8	61.9	88.9	79.6	84.2
DMNet [21]	96.1	77.3	94.1	72.8	78.1	97.1	92.7	96.4	39.8	91.4	75.5	92.7	95.8	91.0	90.3	76.6	94.1	62.1	85.5	77.6	84.3
SAANet [77]	95.1	65.9	95.4	72.0	80.5	93.5	86.8	94.5	40.5	93.3	74.6	94.1	-	-	-	-	-	-	-	-	83.2
SpyGR [29]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	84.2
CaC-Net[39]	96.3	76.2	95.3	78.1	80.8	96.5	91.8	96.9	40.7	96.3	76.4	94.3	95.8	91.3	89.1	73.1	93.3	62.2	86.7	80.2	85.1
DCNAS (Ours)	96.5	75.2	96.1	80.7	85.2	97.0	93.8	96.6	49.5	94.0	77.6	95.1	95.7	93.9	89.7	76.1	94.7	70.9	89.7	79.4	86.9

Table 3. **Performance on PASCAL VOC 2012.** The table presents per-class semantic segmentation results on the PASCAL VOC 2012 test dataset. Our method advances the new state-of-the-art with mIoU 86.9%.

To make a fair comparison, we conduct the experiment about 40 times for each method. For each experiment, we run the exploration algorithm on the training dataset of Cityscapes for 80 epochs. As the searching procedure convergences, we train the derived model for 120 epochs on the training dataset. We report the accuracy of both stages on the validation part in each experiment for quantitative comparisons. We shall point out that we reproduce [44, 59, 8, 36] based on the officially released code. As to [3], we refer the announced results from the published paper rather than reproduce the method from scratch due to the intensive computation resource requirements.

As shown in Figure 3, one can observe that our method obtains a higher validation accuracy in the searching phase compared with [44, 59, 8, 3, 36]. This is reasonable because our method does not depend on a proxy dataset, which alleviates the gap between searching and validation data, hence leading to a better performance. Regarding the performance when fine-tuning the derived model, our approach still achieves the best result. This is consistent with our expectation because our search space (DCSS) is more general than [36, 59, 8] and our gradient-based framework is easier to train compared to [3, 44] that rely on reinforcement learning. In terms of the correlation coefficient, our approach outperforms other contemporary methods. It is not astonishing because [44, 59, 8, 3, 36] perform the searching with proxy, yet architectures optimized with proxy are not guaranteed to be optimal on the target task [2]. In our case, searching and training environments share a similar protocol. Consequently, the paired validation performances shall be approximately linear to a certain extent.

4.3. Semantic Segmentation Results

In this section, we evaluate our optimal model structure on Cityscapes [11], PASCAL VOC 2012 [15], ADE20K [78], and PASCAL-Context [43] datasets.

Cityscapes [11] is a large-scale and challenging dataset, following the conventional evaluation protocol [11], we evaluate our model on 19 semantic labels without considering the void label. Table 2 presents the comparison results of our method and several state-of-the-art methods. Without

Methods	ADE20K		Pascal Context	
	mIoU	Pix-Acc	mIoU 59-cls	mIoU 60-cls
DeepLabv2 [4]	-	-	-	45.7
GCPNet [25]	38.37	77.76	-	46.5
RefineNet [33]	40.70	-	-	47.3
MSCI [32]	-	-	-	50.3
PSANet [75]	43.77	81.51	-	-
PSPNet [74]	44.94	81.69	47.8	-
SAC [71]	44.30	-	-	-
CCL [13]	-	-	-	51.6
EncNet [69]	44.65	81.69	52.6	51.7
DSSPN [31]	43.68	-	-	-
CFNet [70]	44.89	-	54.0	-
CCNet [24]	45.22	-	-	-
DeepLabv3+ [7]	45.65	82.52	-	-
Auto-DeepLab [36]	43.98	81.72	-	-
APCNet [22]	45.38	-	55.6	54.7
DANet [17]	-	-	-	52.6
SVCNet [12]	-	-	53.2	-
DRN [80]	-	-	49.0	-
SAANet [77]	-	-	55.3	54.4
SpyGR [29]	-	-	-	52.8
CPNet [67]	46.27	81.85	-	53.9
HRNetV2 + OCR [57]	45.66	-	-	56.2
CaC-Net [39]	46.12	-	-	55.4
DCNAS (Ours)	47.12	84.31	57.1	55.6

Table 4. **Performance on ADE20K and PASCAL-Context.** The table presents the semantic segmentation performance on the validation part of ADE20K and PASCAL-Context according to various evaluation metrics.

any pre-training, our base model achieves the performance with 83.1% mIoU on test dataset, which outperforms most state-of-the-art methods. In addition, we can further improve the test mIoU to 83.6% when pre-training our model with the coarse data. Exploiting ASPP [5] as the prediction head, our method outperforms all state of the art methods with 84.3% mIoU, including [36, 28] that also utilize ASPP.

PASCAL VOC 2012 [15] is another gold-standard benchmark for object segmentation, which contains 20 foreground object classes and one background class. Following conventional [4, 74, 69, 70, 36], we exploit the augmented set [20] by pre-training our model on the train + val parts on the augmented set, then fine-tune the model on the original PASCAL VOC 2012 benchmark. As a result, our model advances the new state-of-the-art result to 86.9% mIoU, which largely outperforms the state-of-the-art (absolute 2.5% mIoU improvement). We also report the per-class comparisons with state-of-the-art methods in

Search Scheme	$L = 8$	$L = 10$	$L = 12$	$L = 14$
ProxlessNAS [2]	64.3	68.1	73.7	77.5
PC-DARTs [64]	70.8	72.4	76.2	80.6
P-DARTs [9]	68.5	70.4	76.9	79.4
SPOS [19]	65.2	68.5	73.2	77.4
DCNAS (Ours)	72.8	75.4	78.9	81.2

Table 5. **Comparisons with State of the Art Search Methods.** The table presents the comparison results (by T-mIoU on validation set of Cityscapes) of our proxyless search paradigm with contemporary approaches [2, 64, 9, 19] under various configurations.

Table 3, one can observe that our model achieves superior performance on many categories.

PASCAL-Context [43] and **ADE20K** [78] are substantial benchmarks toward scene parsing tasks. Being consistent with [70, 24, 79, 36], we employ not only the classical mIoU but also the pixel accuracy as the evaluation metrics for ADE20K. As to PASCAL-Context, we report the mIoUs both with (60-cls) and without (59-cls) considering the background. As shown in Table 4, we present the performances of our model on the validation parts of Pascal Context and ADE20K according to corresponding evaluation metrics. Our DCNAS outperforms other state-of-the-art methods on both Pascal Context and ADE20K datasets.

It is consistent with our expectation that our approach retained leading performance on the above mentioned extensively used benchmarks with various evaluation metrics. Specifically: (1) the proxyless searching paradigm can bridge the gap between searching and training environments, which is conducive to the discovery of promising model structures; (2) our DCNAS includes an abundance of possible model structures, accompanied with the proxyless searching method, so that one may easily obtain an outstanding model; (3) the multi-scale semantic features aggregation has been proved to be instrumental for visual recognition [42, 4, 74, 7, 55], and our DCNAS inherently and repeatedly applying top-down and bottom-up multi-scale features fusion, hence results in leading performance.

4.4. Ablation Study

In this section, we evaluate the impact of the key designs in DCNAS. We take the mIoU(%) on validation part of Cityscapes both in searching (S-mIoU) and training (T-mIoU) periods, the Pearson Correlation Coefficient (ρ), the Kendall Rank Correlation Coefficient (τ), and the searching cost (GPU Days) as the evaluation metrics.

Since the proxyless searching is instrumental in DCNAS, we compare our searching strategy with searching methods [2, 64, 9, 19]. Table 5 reports the comparing results under several configurations, as a result, our proxyless searching method is more suitable for semantic image segmentation. Besides, our searching approach obtains significant performance improvement comparing with PC-DARTs [64].

Methods	Proxyless	Dense Space	ρ	τ	S-mIoU	T-mIoU
Auto-DeepLab	\times	\times	0.31	0.21	34.9	80.3
Auto-DeepLab	\checkmark	\times	0.61	0.45	37.6	80.6
Auto-DeepLab	\times	\checkmark	0.27	0.19	49.7	80.7
Auto-DeepLab	\checkmark	\checkmark	0.63	0.43	60.9	80.9
Auto-DeepLab \dagger	\times	\times	0.42	0.33	40.8	56.9
Auto-DeepLab \dagger	\checkmark	\times	0.64	0.48	43.1	67.5
Auto-DeepLab \dagger	\times	\checkmark	0.48	0.41	46.2	70.2
Auto-DeepLab \dagger	\checkmark	\checkmark	0.69	0.54	50.5	72.4
DCNAS (Ours)	\times	\times	0.37	0.25	45.2	60.1
DCNAS (Ours)	\checkmark	\times	0.69	0.56	51.7	73.3
DCNAS (Ours)	\times	\checkmark	0.39	0.32	61.5	78.9
DCNAS (Ours)	\checkmark	\checkmark	0.73	0.55	69.9	81.2

Table 6. **Ablation of the Search Strategy and Search Space.** The table presents the effects of the proxyless search paradigm and the densely connected search space in both DCNAS and Auto-DeepLab frameworks, mIoUs are evaluated on validation set. \dagger replace the ASPP with a simpler head used in DCNAS. For DCNAS, **Proxyless** denotes that setting the sampling ratio $r = 1$ in the mixture layer and uses a much lower resolution image for searching, while **Dense Space** indicates constraining connections to the previous layers only. Regrading Auto-DeepLab, **Proxyless** transfer the mixture-layer and exploit higher resolution image for searching, **Dense Space** transfer the fusion module and optimize whole candidate connections rather than the adjacent cells only.

Besides, to deeply study the impact of the proxyless search paradigm and the densely connected search space, we transfer the searching space and searching method to the Auto-DeepLab [36] framework. As Table 6 presents the ablation results, one may observe that, (1) the densely connected search space remains indispensable for semantic image segmentation; (2) proxyless searching scheme is capable of bridging the gap between searching and training situations, which can further improve the performance.

5. Discussion

We proposed a novel NAS framework, dubbed DCNAS, to directly and proxylessly search the optimal multi-scale network architectures for dense image prediction tasks. We introduce a densely connected search space (DCSS), which contains most of the widespread human-designed network structures. With DCSS, we attempt to fully automatically search the optimal multi-scale representations of high-resolution imagery and to fuse those multi-scale features. Meanwhile, to enable the efficient searching process, we propose a novel fusion module to deal with the high GPU memory consumption and expensive computation overhead issues. Consequently, we implement a novel proxyless NAS framework for dense image prediction tasks. Experiment results demonstrate that the architectures obtained from our DCNAS can surpass not only the human-invented architectures but also the automatically designed architectures from previous NAS methods the focusing on semantic segmentation task. For future work, one promising direction is investigating the lightweight neural architecture under the DCSS for the resource-constrained devices.

References

- [1] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2016. 4
- [2] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations (ICLR)*, 2018. 2, 3, 7, 8
- [3] Liang-Chieh Chen, Maxwell Collins, Yukun Zhu, George Papandreou, Barret Zoph, Florian Schroff, Hartwig Adam, and Jon Shlens. Searching for efficient multi-scale architectures for dense image prediction. In *Advances in neural information processing systems (NIPS)*, 2018. 1, 2, 3, 5, 6, 7
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. 2017. 3, 7, 8
- [5] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. 2017. 2, 4, 6, 7
- [6] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2, 3, 6, 7, 8
- [8] Wuyang Chen, Xinyu Gong, Xianming Liu, Qian Zhang, Yuan Li, and Zhangyang Wang. Fasterseg: Searching for faster real-time semantic segmentation. In *International Conference on Learning Representations (ICLR)*, 2019. 1, 2, 3, 4, 5, 6, 7
- [9] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3, 8
- [10] Sungha Choi, Joanne T Kim, and Jaegul Choo. Cars can't fly up in the sky: Improving urban-scene segmentation via height-driven attention networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 6, 7
- [12] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. Semantic correlation promoted shape-variant context for segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 6, 7
- [13] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. Context contrasted feature and gated multi-scale aggregation for scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 7
- [14] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2015. 3
- [15] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. 2010. 2, 7
- [16] Jiemin Fang, Yuzhu Sun, Qian Zhang, Yuan Li, Wenyu Liu, and Xinggang Wang. Densely connected search space for more flexible neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [17] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 6, 7
- [18] Jun Fu, Jing Liu, Yuhang Wang, Jin Zhou, Changyong Wang, and Hanqing Lu. Stacked deconvolutional network for semantic segmentation. 2019. 3
- [19] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 8
- [20] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2011. 7
- [21] Junjun He, Zhongying Deng, and Yu Qiao. Dynamic multi-scale filters for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2, 7
- [22] Junjun He, Zhongying Deng, Lei Zhou, Yali Wang, and Yu Qiao. Adaptive pyramid context network for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 7
- [23] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3
- [24] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2, 6, 7, 8
- [25] Wei-Chih Hung, Yi-Hsuan Tsai, Xiaohui Shen, Zhe Lin, Kalyan Sunkavalli, Xin Lu, and Ming-Hsuan Yang. Scene parsing with global context embedding. In *Proceedings of*

- the *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 7
- [26] Tsung-Wei Ke, Jyh-Jing Hwang, Ziwei Liu, and Stella X Yu. Adaptive affinity fields for semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2, 7
- [27] Hanchao Li, Pengfei Xiong, Jie An, and Lingxue Wang. Pyramid attention network for semantic segmentation. In *The British Machine Vision Conference (BMVC)*, 2018. 2, 7
- [28] Xiangtai Li, Xia Li, Li Zhang, Cheng Guangliang, Jianping Shi, Zhouchen Lin, Yunhai Tong, and Shaohua Tan. Improving semantic segmentation via decoupled body and edge supervision. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 6, 7
- [29] Xia Li, Yibo Yang, Qijie Zhao, Tiancheng Shen, Zhouchen Lin, and Hong Liu. Spatial pyramid based graph reasoning for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6, 7
- [30] Xin Li, Yiming Zhou, Zheng Pan, and Jiashi Feng. Partial order pruning: for best speed/accuracy trade-off in neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [31] Xiaodan Liang, Hongfei Zhou, and Eric Xing. Dynamic-structured semantic propagation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 7
- [32] Di Lin, Yuanfeng Ji, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Multi-scale context intertwining for semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 7
- [33] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3, 7
- [34] Guosheng Lin, Chunhua Shen, Anton Van Den Hengel, and Ian Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [35] Peiwen Lin, Peng Sun, Guangliang Cheng, Sirui Xie, Xi Li, and Jianping Shi. Graph-guided architecture search for real-time semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [36] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [37] Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. Hierarchical representations for efficient architecture search. In *International Conference on Learning Representations (ICLR)*, 2017. 1
- [38] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations (ICLR)*, 2018. 1, 3, 4, 5
- [39] Jianbo Liu, Junjun He, Yu Qiao, Jimmy S Ren, and Hongsheng Li. Learning to predict context-adaptive convolution for semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 7
- [40] Wei Liu, Andrew Rabinovich, and Alexander C Berg. Parsenet: Looking wider to see better. In *International Conference on Learning Representations Workshops (ICLRW)*, 2015. 2
- [41] Zuhao Liu, Huan Wang, Shaoting Zhang, Guotai Wang, and Jin Qi. Nas-scnn: Neural architecture search-based spatial and channel joint attention module for nuclei semantic segmentation and classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (ICMIC)*, 2020. 3
- [42] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3, 4, 8
- [43] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2, 7, 8
- [44] Vladimir Nekrasov, Hao Chen, Chunhua Shen, and Ian Reid. Fast neural architecture search of compact semantic segmentation models via auxiliary cells. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 6, 7
- [45] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 6
- [46] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 3, 4
- [47] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *International Conference on Machine Learning (ICML)*, 2018. 3
- [48] Tobias Pohlen, Alexander Hermans, Markus Mathias, and Bastian Leibe. Full-resolution residual networks for semantic segmentation in street scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [49] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V Le, and Alexey Kurakin. Large-scale evolution of image classifiers. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017. 3, 4
- [50] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmen-

- tation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015. 3
- [51] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. In-place activated batchnorm for memory-optimized training of dnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 6
- [52] Albert Shaw, Daniel Hunter, Forrest Landola, and Sammy Sidhu. Squeezenas: Fast neural architecture search for faster semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (IC-CVW)*, 2019. 3
- [53] Masanori Suganuma, Shinichi Shirakawa, and Tomoharu Nagao. A genetic programming approach to designing convolutional neural network architectures. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2017. 4
- [54] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 4
- [55] Towaki Takikawa, David Acuna, Varun Jampani, and Sanja Fidler. Gated-scnn: Gated shape cnns for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 6, 8
- [56] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [57] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Minghui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. 2020. 6, 7
- [58] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [59] Huikai Wu, Junge Zhang, and Kaiqi Huang. Sparsemask: Differentiable connectivity learning for dense image prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 3, 6, 7
- [60] Yangxin Wu, Gengwei Zhang, Hang Xu, Xiaodan Liang, and Liang Lin. Auto-panoptic: Cooperative multi-component architecture search for panoptic segmentation. 2020. 3
- [61] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. 2019. 7
- [62] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. In *International Conference on Learning Representations (ICLR)*, 2018. 4
- [63] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [64] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. Pc-darts: Partial channel connections for memory-efficient differentiable architecture search. In *International Conference on Learning Representations (ICLR)*, 2019. 5, 8
- [65] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [66] Yukang Chen, Zeming Li, Xiangyu Zhang, Xingang Wang, Jian Sun, Yanwei Li, Lin Song. Learning dynamic routing for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [67] Changqian Yu, Jingbo Wang, Changxin Gao, Gang Yu, Chunhua Shen, and Nong Sang. Context prior for scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6, 7
- [68] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Learning a discriminative feature network for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 7
- [69] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 7
- [70] Hang Zhang, Han Zhang, Chenguang Wang, and Junyuan Xie. Co-occurrent features in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 7, 8
- [71] Rui Zhang, Sheng Tang, Yongdong Zhang, Jintao Li, and Shuicheng Yan. Scale-adaptive convolutions for scene parsing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 2, 7
- [72] Xiong Zhang, Qiang Li, Hong Mo, Wenbo Zhang, and Wen Zheng. End-to-end hand mesh recovery from a monocular rgb image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3
- [73] Yiheng Zhang, Zhaofan Qiu, Jingen Liu, Ting Yao, Dong Liu, and Tao Mei. Customizable architecture search for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 3
- [74] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 3, 6, 7, 8
- [75] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Psanet: Point-wise spatial attention network for scene parsing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 6, 7
- [76] Mingmin Zhen, Jinglu Wang, Lei Zhou, Shiwei Li, Tianwei Shen, Jiaxiang Shang, Tian Fang, and Long Quan. Joint semantic segmentation and boundary detection using iterative

- pyramid contexts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6
- [77] Zilong Zhong, Zhong Qiu Lin, Rene Bidart, Xiaodan Hu, Ibrahim Ben Daya, Zhifeng Li, Wei-Shi Zheng, Jonathan Li, and Alexander Wong. Squeeze-and-attention networks for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7
- [78] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 7, 8
- [79] Zhen Zhu, Mengde Xu, Song Bai, Tengpeng Huang, and Xiang Bai. Asymmetric non-local neural networks for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 8
- [80] Yueqing Zhuang, Fan Yang, Li Tao, Cong Ma, Ziwei Zhang, Yuan Li, Huizhu Jia, Xiaodong Xie, and Wen Gao. Dense relation network: Learning consistent and context-aware representation for semantic image segmentation. In *IEEE international conference on image processing (ICIP)*, 2018. 2, 6, 7
- [81] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2016. 1, 3, 4
- [82] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 3