

iNAS: Integral NAS for Device-Aware Salient Object Detection

Yu-Chao Gu¹ Shang-Hua Gao¹ Xu-Sheng Cao¹ Peng Du²
Shao-Ping Lu¹ Ming-Ming Cheng^{1*}

¹TKLNDST, CS, Nankai University ²Huawei Technologies

<https://mmcheng.net/inas/>

Abstract

Existing salient object detection (SOD) models usually focus on either backbone feature extractors or saliency heads, ignoring their relations. A powerful backbone could still achieve sub-optimal performance with a weak saliency head and vice versa. Moreover, the balance between model performance and inference latency poses a great challenge to model design, especially when considering different deployment scenarios. Considering all components in an integral neural architecture search (iNAS) space, we propose a flexible device-aware search scheme that only trains the SOD model once and quickly finds high-performance but low-latency models on multiple devices. An evolution search with latency-group sampling (LGS) is proposed to explore the entire latency area of our enlarged search space. Models searched by iNAS achieve similar performance with SOTA methods but reduce the 3.8 \times , 3.3 \times , 2.6 \times , 1.9 \times latency on Huawei Nova6 SE, Intel Core CPU, the Jetson Nano, and Nvidia Titan Xp. The code is released at <https://mmcheng.net/inas/>.

1. Introduction

Salient object detection (SOD) aims to segment the most attractive objects in images [1, 59]. Served as a pre-processing step, SOD is required by many downstream applications, *i.e.*, image editing [8], image retrieval [22], visual tracking [24], and video object segmentation [20]. These applications often require the SOD model to be deployed with low inference latency on multiple devices, *i.e.*, GPUs, CPUs, mobile phones, and embedded devices. Each device has unique properties. For instance, GPUs are good at massively parallel computing [43] while the embedded devices are energy-friendly at the cost of a low computing budget [27]. Thus, different deployment scenarios require quite different designs of SOD models.

State-of-the-art (SOTA) SOD methods mostly design handcraft saliency heads [36, 44, 47, 78, 81] to aggregate multi-level features from the pre-trained backbone,

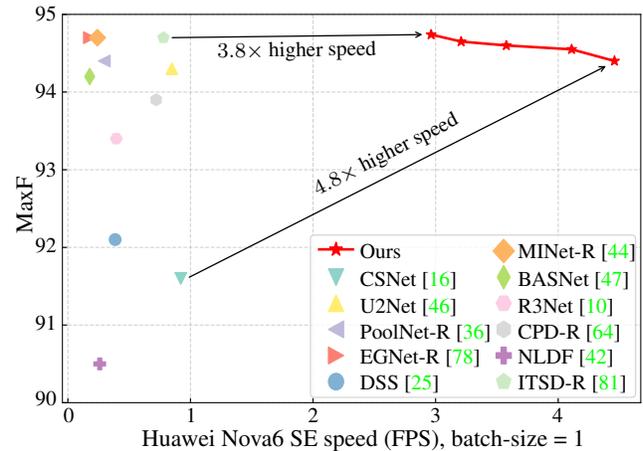


Figure 1. Mobile latency and performance comparison between our iNAS and recent state-of-the-art SOD models.

e.g., VGG [51], ResNet [23] and Res2Net [13]. The prohibitive inference latency often prevents them from been applied on other devices except for GPUs. On the other hand, handcraft low-latency SOD models designed for resource-constrained scenarios [16, 46] suffer from large performance drop. It causes heavy workloads to manually design SOD models for different devices because of the dilemma between model performance and inference latency. Therefore, we aim at a device-aware search scheme to quickly find suitable low-latency SOD models on multiple devices.

There are several obstacles to achieve low-latency SOD models on different devices, as shown in Fig. 2. Firstly, the relative latency of operators varies among different devices due to different parallel computation abilities, IO bottlenecks, and implementations. Transfer the SOD model designed for one device to another would result in sub-optimal latency and performance. Secondly, conventional handcraft SOD models design more powerful saliency heads [36, 44, 47, 81] or more efficient backbones [16, 46], while ignoring their relations. Similarly, most neural architecture search (NAS) methods focus on the backbone for the classification task [35, 53] or incorporate a fixed segmentation head [33, 34] while ignoring the backbone and head relationship. We observe that a powerful backbone achieves

*M.M. Cheng (cmm@nankai.edu.cn) is the corresponding author.

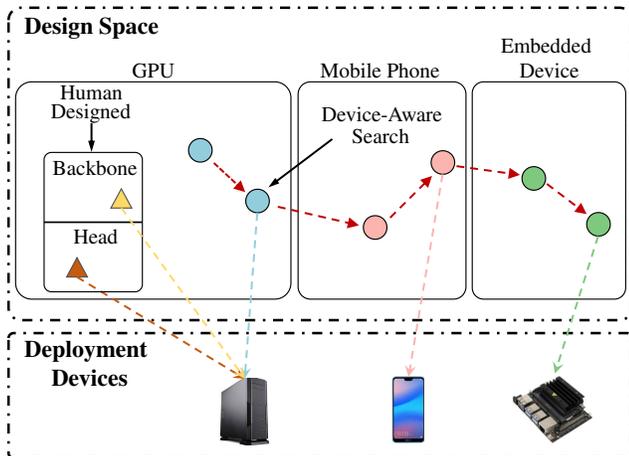


Figure 2. iNAS unifies backbone and head design into an integral design space and specializes low-latency SOD models for different devices.

sub-optimal efficiency with a weak saliency head and vice versa. These obstacles prevent the community from designing device-aware low-latency SOD models either with handcraft or NAS schemes.

To deal with these problems, we propose a device-aware search scheme with an integral search space to train the model once and quickly find high-performance but low-latency SOD models on multiple devices. Specifically, we propose an integral search space for SOD models that holistically consider the backbone and saliency head. To meet multi-scale requirements of SOD models while avoiding the latency increased by multi-branch structures, we construct a searchable multi-scale unit (SMSU). The SMSU supports searchable parallel convolutions with different kernel sizes, and reparameterizes searched multi-branch convolutions to one branch for low inference latency. We also generalize handcraft saliency heads [25, 36, 41, 44, 75] into searchable transport and decoder parts, resulting in a rich saliency head search space for cooperating with the backbone space.

With multi-scale architectures, the proposed integral SOD search space is significantly larger than NAS spaces for the classification task [2, 72]. After training the one-shot supernet, previous methods adopt evolution search with uniform sampling [2, 21, 72] to explore the search space. Uniform sampling can ensure different architecture choices within one layer have equal sampling probability. However, the overall latency of sampled models obeys a multinomial distribution, which causes extremely low-latency or extremely high-latency areas to be under-sampled. This imbalance sampling problem prevents uniform sampling from exploring the entire latency area of our enlarged search space. To overcome this imbalance sampling problem, we propose a latency-group sampling (LGS) that introduces the device latency to guide sampling. Dividing the layer-wise search space into several latency groups, and aggregating samples in specific latency groups, LGS preserves the off-

spring in the under-sampled area but controls the samples of the over-sampled area. Compared with uniform sampling, the evolution search with LGS can explore the entire integral search space and finds a group of models on a higher and wider Pareto frontier.

The main contributions of this paper are:

- An integral SOD search space that considers the backbone-head relation and covers existing SOTA handcraft SOD designs.
- A device-aware evolution search with latency-group sampling for exploring the entire latency area of the proposed search space.
- A thorough evaluation of the iNAS on five popular SOD datasets. Our method can reach a similar performance with handcraft SOTA methods but largely reduces inference latency on different devices, which helps to scale up the application of SOD to different deployment scenarios.

2. Related Work

2.1. Salient Object Detection.

Traditional SOD methods [1, 6, 55, 83] mainly rely on handcraft features and heuristic priors. [28, 29, 79] make an early attempt to use convolution neural networks (CNNs) to extract patch-level features. Inspired by FCN [41], the recent SOD methods [39, 57, 60] formulate SOD as a pixel-wise prediction task, which achieves large improvement over traditional or CNN-based methods. We refer readers to comprehensive surveys [1, 59, 82].

Most of the SOD methods handcraft the saliency head to effectively fuse the multi-scale information of the multi-level feature extracted by the pre-trained backbone [14, 23, 51], *e.g.*, ResNet [23]. These methods [4, 17, 25, 36, 38, 58, 67] inherit an encoder-decoder structure, in which the decoder is responsible for the bottom-up feature fusion. Transport layers [12, 44, 74, 75, 80] are included inside the saliency head, enabling both the bottom-up and top-down feature fusion. Methods that introduce edge cues into the saliency head for precise boundary refinement [30, 62, 78] are orthogonal to our search space.

The gradually complicated SOD models bring improvements in performance steadily while increasing prohibitive inference latency. Recent works [16, 20, 63, 64, 81] try to design lightweight models to eliminate the large inference latency. Among them, CPD [64] and ITSD [81] design lightweight saliency heads, achieving fast speed on CPUs and GPUs, respectively. CSNet [16] designs a light SOD backbone to achieve the low-latency on the mobile phone and embedded device. However, separating the design and the deployment devices causes sub-optimal latency when the hardware characteristics are quite different.

In this work, we introduce an integral search space that

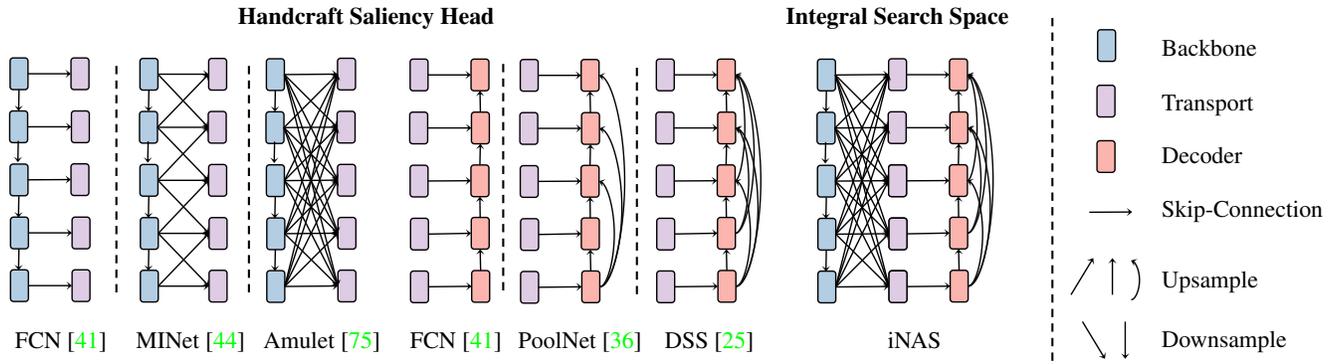


Figure 3. The designs of recent handcraft SOD models and the proposed integral search space.

Stage	Backbone					Transport			Decoder		
	Operator	Resolutions	Channels	Layers	Kernel	Level	Kernel	Fusions	Level	Kernel	Fusions
stem	Conv	256x256-384x384	32-40	1	3	1	3,5,7,9	1-5	1	3,5,7,9	2-5
1	MBconv1	128x128-192x192	16-24	1-2	3	2	3,5,7,9	1-5	2	3,5,7,9	2-4
2	MBconv6	128x128-192x192	24-32	2-3	3	3	3,5,7,9	1-5	3	3,5,7,9	2-3
3	MBconv6	64x64-96x96	32-48	2-3	3,5,7,9	4	3,5,7,9	1-5	4	3,5,7,9	2
4	MBconv6	32x32-48x48	64-88	2-4	3,5,7,9	5	3,5,7,9	1-5	5	3,5,7,9	1
5	MBconv6	32x32-48x48	96-128	2-6	3,5,7,9						
6	MBconv6	16x16-24x24	160-216	2-6	3,5,7,9						
7	MBconv6	16x16-24x24	320-352	1-2	3,5,7,9						

Table 1. Detailed configurations of the proposed integral search space.

covers most of the handcraft SOD designs. Based on our integral search space, we propose a device-aware search scheme, which achieves similar performance to SOTA methods but largely reduces latency on different devices.

2.2. Neural Architecture Search.

Neural architecture search (NAS) demonstrates its potential to design efficient networks for various tasks automatically [15, 18, 32, 34, 49, 70, 73, 76]. Early methods based on reinforcement learning [84, 85] and evolutionary algorithm [48, 65] train thousands of candidate architectures to learn a meta-controller, cost hundreds of GPU days to search. Later, differentiable NAS [19, 35] and one-shot NAS [2, 21, 72] exploit the idea of weight-sharing [45] to reduce the search cost, where the one-shot NAS decouples the supernet training and architecture search. Most one-shot NAS methods [2, 21, 72] target improving the supernet training and adopt evolution search with uniform sampling to explore the search space. However, we find uniform sampling causes an imbalance sampling problem when taking model latency into account.

Apart from the search method, the search space plays a vital role in NAS. Early methods [35, 45, 48, 65] utilize cell-based search space, where the cell is composed of multiple searchable operations. Based on cell-based search space, Auto-deeplab [34] additionally supports searching for the macro-structure of scale transformation. In order to adapt the segmentation task, Auto-deeplab incorporates fixed parallel ASPP [3] decoders. However, the searched structures of cell-based search space have complicated branch con-

nections, which is hard to be parallelized in current deep learning frameworks [52], limiting its potential to low-latency applications. Exploiting the human expert knowledge, MnasNet [53] and following works [9, 54, 61] develop a MobileNet-based [50] search space, which supports more hardware-friendly architectures than cell-based search space. However, since these methods are designed for classification tasks, it has less multi-scale representation capability and can not be directly applied to SOD.

Two design principles make iNAS different from Auto-deeplab and MnasNet: 1) The integral search for all components reduces the overall inference latency; 2) The searchable multi-scale unit supports searching for multi-branch structures without additional inference latency cost. To fully explore the proposed integral search space, we propose a latency-group sampling to address the imbalance sampling problem of previous one-shot NAS methods [2, 21, 72]. Different from FairNAS [9], which aims to improve the fairness of optimizing different components in the supernet training stage, our proposed latency-group sampling hopes to explore the search space in a balanced way in the search stage.

3. Methodology

3.1. Integral SOD Design Space.

The previous handcraft SOD models [1, 39, 59] are mainly based on the fixed pre-trained backbone (e.g., VGG [51] and ResNet [23]) and design saliency head to fuse the multi-level feature from the backbone. Some

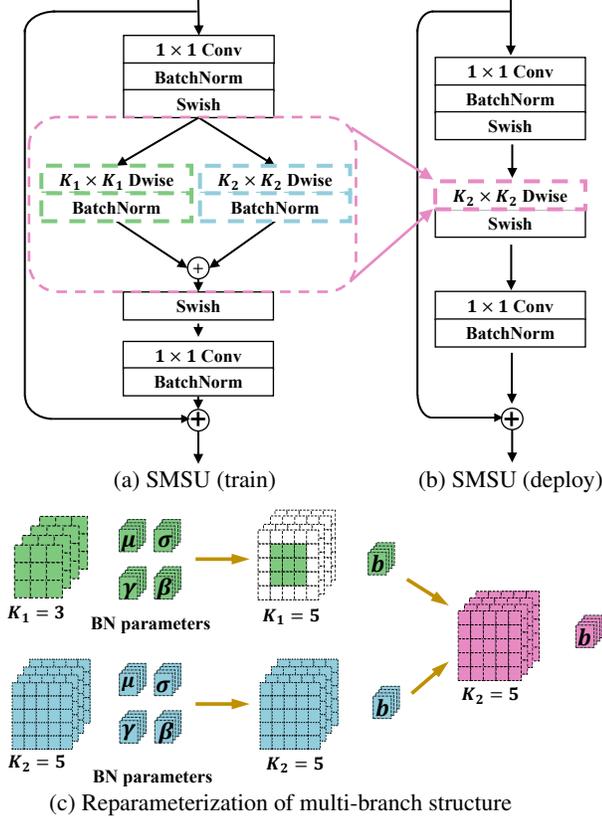


Figure 4. Illustration of the searchable multi-scale unit (SMSU).

recent works have noticed that the pre-trained backbone accounts for most of the latency cost [16]. Instead of adopting a heavy backbone, they design lightweight backbones for SOD. However, both design strategies separate the backbone and decoder design, which hinder finding the low-latency high-performance SOD model in the integral design space. This section introduces an integral SOD design space, composed of the basic search unit (*i.e.*, searchable multi-scale unit) in Sec. 3.1.1 and the searchable saliency head in Sec. 3.1.2.

3.1.1 Searchable Multi-Scale Unit.

Since previous general backbones account for most of latency cost, recent designs [16, 46] of SOD backbones replace vanilla convolution with group convolution [66] or separable convolution [50] for reducing latency. To capture multi-scale representations in images, they design several branches to encode features with different receptive fields and fuse multi-scale features. However, multi-branch structures are not hardware-friendly [50, 61, 77], which will slow down inference speed. For example, the CSNet [16] has reduced $13.4\times$ flops of the ITSD-R [81] but only achieves similar inference latency on the GPU. We thus propose a searchable multi-scale unit (SMSU), which automatically supports finding suitable multi-scale fusions. The SMSU

enables multi-branch structures to capture multi-scale feature representation in training and adopt the reparameterization strategy [11] to fuse multiple branches into a single branch for fast inference.

We show a two-branches setting of SMSU in Fig. 4 (a,b). The SMSU can extract multi-scale feature representation with different kernel sizes. Specifically, assume there are 3×3 Conv and 5×5 Conv, we denote the depthwise convolution parameters $W_1 \in \mathcal{R}^{C\times 1\times 3\times 3}$ and $W_2 \in \mathcal{R}^{C\times 1\times 5\times 5}$. The batchnorm (BN) parameters following 3×3 Conv and 5×5 Conv are denoted as $\mu_1, \sigma_1, \gamma_1, \beta_1$ and $\mu_2, \sigma_2, \gamma_2, \beta_2$, respectively. Given an input feature $F_{in} \in \mathcal{R}^{C\times H\times W}$, we denote the output feature as $M = F_{in} * W$, where $*$ is the convolution. The fusion of two branches can be denoted as

$$F_{out}^{(i)} = (M_1^{(i)} - \mu_1^{(i)}) \frac{\sigma_1^{(i)}}{\gamma_1^{(i)}} - \beta_1^{(i)} + (M_2^{(i)} - \mu_2^{(i)}) \frac{\sigma_2^{(i)}}{\gamma_2^{(i)}} - \beta_2^{(i)}, \quad (1)$$

where i represents i -th channel. Eqn. (1) describes multi-scale fusions of SMSU in the training time. In deployment, we merge the convolution weight and its following BN parameters into a single convolution, defined as

$$V^{(i)} = \frac{\gamma^{(i)}}{\sigma^{(i)}} W^{(i)}, \quad b^{(i)} = -\frac{\mu^{(i)} \gamma^{(i)}}{\sigma^{(i)}} + \beta^{(i)}, \quad (2)$$

where V is the merged convolution weight and b is the bias. Then we zero-pad the small kernel in given branches to match the size of the largest kernel. Finally, we average these two branches to get a single convolution weight and bias.

The introduced two-branches fusion can be easily extended to any branches. Thus, we enable searching for fusion kernel combinations in the SMSU. We replace the inverted bottleneck of MobileNet search space with SMSU and summarize the search space in Tab. 1.

3.1.2 Searchable Saliency Head.

Previous handcraft saliency head incorporates transports or decoders to fuse multi-level features from the backbone. The high-level feature provides a rough location of the salient object, and the low-level feature provides the detailed information for recovering the edge and boundary. As shown in Fig. 3, typical transport designs [44, 75] enable both bottom-up and top-down fusions of multi-level features. Our searchable transport connects to all resolution levels of the backbone. In our largest child transport, each level can aggregate features from all five resolution levels like Amulet [75], while our smallest child transport only keeps identity branches, like FCN [41]. The downsample and upsample branches are composed of 1×1 Conv-BN

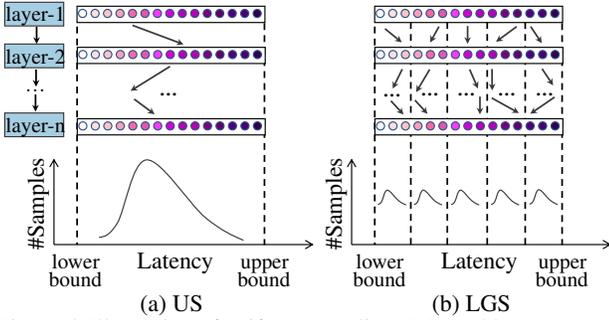


Figure 5. Illustration of uniform sampling (US) and latency-group sampling (LGS).

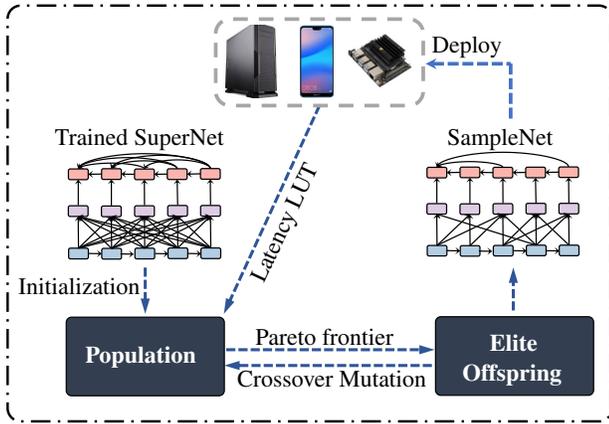


Figure 6. Illustration of iNAS search and deployment.

and maxpool operation/bilinear interpolation. Our searchable transport covers many SOTA SOD transport designs [12, 37, 44, 80].

Unlike the transport, the decoders [25, 36] only support a bottom-up prediction refinement and gradually add in low-level features to recover the boundary. Thus, we do not support top-down fusion branches in the decoder. The identity and upsample branches from the adjacent resolution level are fixed, while other branches are searchable. The largest child decoder has a similar structure to DSS [25], while the smallest child decoder is similar to FCN [41]. The searchable decoder covers many handcraft SOD decoder designs [4, 38, 58, 64, 67].

Considering that best receptive fields for multi-scale fusions may be different at different resolution levels, we use SMSU as a basic search unit in the transport and decoder. Though the multi-scale fusion is proven to be effective in the SOD, how to prune redundant fusion branches and choose appropriate fusion kernels with latency constraints is a labor-intensive work. Our proposed saliency head makes these key components searchable, automatically designed with backbone to minimize inference latency.

3.2. Latency-group Sampling.

Previous one-shot methods adopt evolution search with uniform sampling, which causes an imbalance sampling problem when considering model latency. As illustrated in

Algorithm 1: Evolution Search with LGS

- Input:** Trained supernet, initial population size N , latency lookup table (LUT), latency groups G , offspring size k , crossover probability p_c , mutation probability p_m , iteration $iter$.
- Output:** Pareto frontier of population P .
- 1 Compute the lower-bound and upper-bound of latency (i.e., LAT_{min}^l and LAT_{max}^l) in each layer l based on LUT;
 - 2 Divide the $(LAT_{min}^l, LAT_{max}^l)$ in each layer l into G groups;
 - 3 Sample $\frac{N}{G}$ child models for each latency group $\{P_i | i = 1 \dots G\}$;
 - 4 Set initial population $P = P_1 \cup \dots \cup P_G$;
 - 5 Evaluate performance for models in P ;
 - 6 **for** $j = 1 \dots iter$ **do**
 - 7 **for each** P_i **do**
 - 8 $S_i \leftarrow$ Select $\frac{k}{G}$ models from the Pareto frontier of P_i ;
 - 9 $S = S_1 \cup \dots \cup S_G$;
 - 10 **for each model in** S **do**
 - 11 Crossover and mutate the model under probability p_c and p_m .
 - 12 Evaluate performance for models in S ;
 - 13 $P = P \cup S$;
 - 14 $P \leftarrow$ Select Pareto frontier of P ;
 - 15 **Return** P

Fig. 5, the whole search space is composed of layer-wise block choices. The block choices within each layer vary in latency. Suppose we uniformly sample the block layer-by-layer, the accumulated latency of overall sampled models will obey a multinomial distribution, i.e., the extremely low-latency or extremely high-latency areas are under-sampled but the middle latency area is over-sampled. To explore the entire latency area of our integral search space, we propose latency-group sampling (LGS). Given a latency lookup table (LUT), we divide the layer-wise search space into several latency groups. To get a model in specific latency group, we sample blocks within this latency group at each layer. Although samples remain imbalance within each local latency group, we can get balanced samples in global latency range if we divide adequate groups. Moreover, when selecting elite offsprings, we also keep a balance number of offsprings in different latency groups.

The general pipeline of device-aware evolution search is depicted in Fig. 6. We first build a latency lookup table (LUT) on target device. Then we perform the evolution search based on LGS. After searching, the searched model inherits the supernet weight and can be directly deployed without retraining. As shown in Algorithm. 1, the evolution search with LGS contains four stages:

Method	FLOPs (G)	Latency (ms)		ECSSD(1000)			DUT-O(5168)			DUTS-TE(5019)			HKU-IS(4447)			PASCAL-S(850)		
		GPU	Embedded	maxF	MAE	S_m	maxF	MAE	S_m	maxF	MAE	S_m	maxF	MAE	S_m	maxF	MAE	S_m
VGG-16/VGG-19																		
NLDF _{CVPR17} [42]	66.68	9.48	505.59	0.905	0.063	0.875	0.753	0.080	0.770	0.813	0.065	0.805	0.902	0.048	0.879	0.822	0.098	0.805
DSS _{CVPR17} [25]	48.75	5.85	N/A	0.921	0.052	0.882	0.781	0.063	0.790	0.825	0.056	0.812	0.916	0.040	0.878	0.831	0.093	0.798
PiCANet _{CVPR18} [39]	59.82	34.21	N/A	0.931	0.046	0.914	0.794	0.068	0.826	0.851	0.054	0.861	0.921	0.042	0.906	0.856	0.078	0.848
CPD-V _{CVPR19} [64]	24.08	3.78	266.40	0.936	0.040	0.910	0.793	0.057	0.818	0.864	0.043	0.866	0.924	0.033	0.904	0.861	0.072	0.845
ITSD-V _{CVPR20} [81]	17.08	9.97	494.93	0.939	0.040	0.914	0.807	0.063	0.829	0.876	0.042	0.877	0.927	0.035	0.906	0.869	0.068	0.856
PoolNet-V _{CVPR19} [36]	48.80	8.81	N/A	0.941	0.042	0.917	0.806	0.056	0.833	0.876	0.042	0.878	-	-	-	0.865	0.072	0.852
EGNet-V _{ICCV19} [78]	120.15	11.58	N/A	0.943	0.041	0.919	0.809	0.057	0.836	0.877	0.044	0.878	0.930	0.034	0.912	0.858	0.077	0.848
MINet-V _{CVPR20} [44]	71.76	14.78	N/A	0.943	0.036	0.919	0.794	0.057	0.822	0.877	0.039	0.875	0.930	0.031	0.912	0.865	0.064	0.854
ResNet-34/ResNet-101/ResNetXt-101																		
R3Net _{IJCAI18} [10]	26.19	6.70	335.14	0.934	0.040	0.910	0.795	0.063	0.817	0.831	0.057	0.835	0.916	0.036	0.895	0.835	0.092	0.807
CPD-R _{CVPR19} [64]	7.19	2.52	124.09	0.939	0.037	0.918	0.797	0.056	0.825	0.865	0.043	0.869	0.925	0.034	0.906	0.859	0.071	0.848
BASNet _{CVPR19} [47]	97.51	16.37	N/A	0.942	0.037	0.916	0.805	0.056	0.836	0.859	0.048	0.865	0.928	0.032	0.909	0.854	0.076	0.838
PoolNet-R _{CVPR19} [36]	38.17	9.13	N/A	0.944	0.039	0.921	0.808	0.056	0.836	0.880	0.040	0.883	0.932	0.033	0.916	0.863	0.075	0.849
EGNet-R _{ICCV19} [78]	120.85	12.01	N/A	0.947	0.037	0.925	0.815	0.053	0.841	0.888	0.039	0.887	0.935	0.031	0.917	0.865	0.074	0.852
MINet-R _{CVPR20} [44]	42.68	7.38	N/A	0.947	0.033	0.925	0.810	0.056	0.833	0.884	0.037	0.884	0.935	0.029	0.919	0.867	0.064	0.856
ITSD-R _{CVPR20} [81]	9.65	3.57	164.76	0.947	0.034	0.925	0.820	0.061	0.840	0.882	0.041	0.884	0.934	0.031	0.917	0.870	0.066	0.859
Handcraft SOD Backbone																		
CSNet _{ECCV20} [16]	0.72	3.63	95.75	0.916	0.065	0.893	0.775	0.081	0.805	0.813	0.075	0.822	0.898	0.059	0.881	0.828	0.103	0.813
U ² -Net _{PR20} [46]	9.77	4.45	173.61	0.943	0.041	0.918	0.813	0.060	0.837	0.852	0.054	0.858	0.928	0.037	0.908	0.847	0.086	0.831
Searched Models on Different Devices																		
iNAS(GPU)-S	0.43	1.32	48.56	0.944	0.037	0.921	0.819	0.055	0.842	0.872	0.043	0.875	0.930	0.033	0.914	0.864	0.071	0.852
iNAS(Embedded)-S	0.41	1.53	40.99	0.944	0.038	0.920	0.816	0.056	0.840	0.871	0.043	0.875	0.931	0.033	0.915	0.865	0.070	0.852
iNAS(GPU)-L	0.70	1.94	71.70	0.947	0.036	0.924	0.824	0.052	0.846	0.879	0.040	0.881	0.935	0.031	0.918	0.867	0.071	0.852
iNAS(Embedded)-L	0.63	2.30	63.39	0.947	0.036	0.924	0.820	0.055	0.842	0.875	0.041	0.879	0.935	0.031	0.919	0.865	0.070	0.852

Table 2. Comparison with existing SOD methods. The FLOPs and latency are measured with 224×224 input images. N/A means that it could not be deployed on the embedded device because of the out-of-memory error.

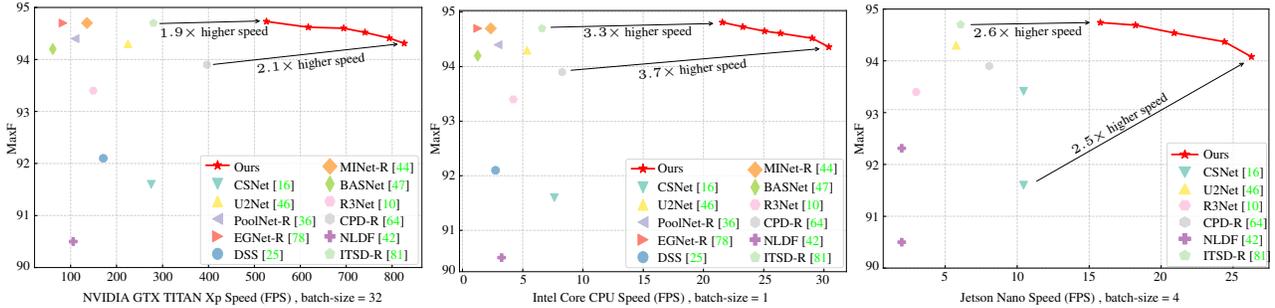


Figure 7. Speed comparison with existing SOD methods on different devices. iNAS achieves SOTA performance and consistent speedup.

- **S1: Initialization.** We divide latency ranges of block choices in each layer into G latency groups. We sample N candidates for an initial population P , where each latency group has $\frac{n}{G}$ samples.
- **S2: Selection.** We select k models from the Pareto frontier of P into a candidate set S , where each latency group contains $\frac{k}{G}$ samples.
- **S3: Crossover.** For each model in S , it has a probability of p_c to crossover with another model in S . We allow swap the stage-wise configuration in the backbone and swap level-wise configuration in the head.
- **S4: Mutation.** For each model in S , each configuration has a probability of p_m to mutate. Then we merge the S into the population P and continue to S2 until target iterations $iter$.

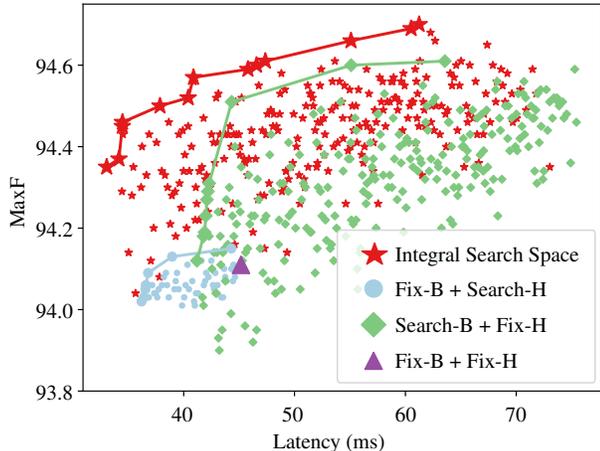
The main difference between LGS and uniform sampling is in the initialization and selection. In the initialization step, LGS balances the samples in different latency area

while uniform sampling over-samples the middle latency area. Then in the selection step, LGS preserves a certain number of elite offsprings in different groups, which enables the evolution search to find better models in different latency areas.

4. Experiments

4.1. Implementation Details.

Details of supernet training. We implement iNAS using Pytorch [52] and Jittor [26] library. We organize the search space as a nested supernet as [2, 72]. Specifically, the weight of smaller convolution kernel copies from the center part of the larger kernel, then transformed by a fully connected layer. And also, the lower-index channels and layers are sharing. The supernet is trained on DUTS-TR for 100 epochs with ImageNet pre-training. The training batch size is set to 40. We use an Adam optimizer with a learning



(a) Search space exploration. B: backbone, H: head.

Searchable		Low Latency Arch		High Performance Arch	
Backbone	Head	Latency (ms)	maxF	Latency (ms)	maxF
✗	✗	45.17	0.941	45.17	0.941
✓	✗	41.20	0.941	63.56	0.946
✗	✓	36.20	0.940	44.30	0.942
✓	✓	33.06	0.944	61.24	0.947

(b) Quantitative analysis of the integral and partial search.

Figure 8. Comparison between the integral and partial search.

rate of $1e-4$ and the poly learning rate schedule [40]. We sample the largest, the smallest, and two middle models for each iteration and fuse their gradients to update the supernet. Following [25], we add deep supervision on the prediction of each decoder level. The supernet training costs 17 hours on four Tesla V100.

Details of search and deployment. We set the initial population size N to 1000, and the latency group G to 10. The evolution iteration $iter$ is set to 20. Each selection step retains $k = 100$ offspring. The crossover and mutation probability (p_c and p_m) are set to 0.2. For evaluating the performance of each child model, we copy their weight from supernet and finetune their BN parameters for 200 iterations [71]. We use the Pytorch-Mobile [52] library to build the LUT on the mobile phone. On other devices, we directly benchmark their speeds with Pytorch toolkit. The search phase costs 0.8 GPU-Days on one Tesla V100 GPU.

Dataset. The supernet is trained with the DUTS-TR dataset [56]. We conduct evaluations on five popular SOD datasets, *i.e.*, ECSSD [68], DUT-O [69], DUTS-TE [56], HKU-IS [28], PASCAL-S [31], containing 1000, 5168, 5019, 4447, and 850 pairs of images and saliency maps.

Evaluation metrics. Following common settings [39, 47], we use MAE [7], Max F-measure (F_β) [6] and S-measure (S_m) [5] as the evaluation metrics to evaluate our results. Since we aim to design low-latency SOD models, the inference latency is also used as the evaluation metric.

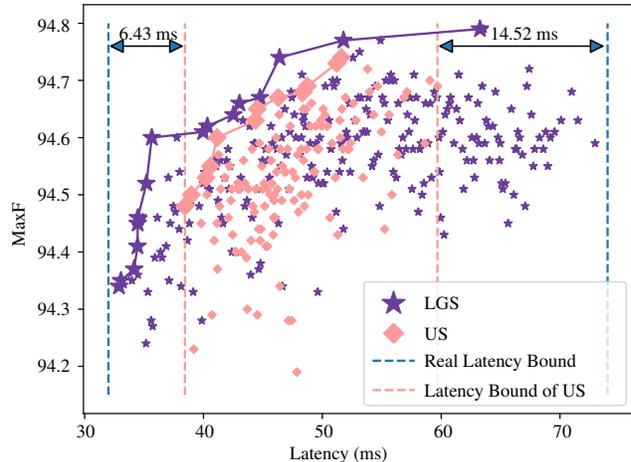


Figure 9. Comparison of the evolution search with uniform sampling (US) and proposed LGS.

Search Dev.	Latency (ms)			
	GPU	CPU	Mobile	Embedded
GPU	1.94	48.90	397.17	71.70
Device-Aware	1.94	42.99	339.61	63.39
Latency Reduction	0%	12.1%	14.5%	10.9%

Table 3. Comparison of searching on GPU and specialized device.

4.2. Performance Evaluation.

Comparison to the state-of-the-art. Tab. 2 shows the comparison between our searched models and previous hand-craft SOD methods. The iNAS(GPU)-L, a large model searched on GPU, requires similar FLOPs to the CSNet, but reduces 47% inference latency and improves 3.1% F_β on ECSSD, which suggests that FLOPs is not highly related to the inference latency. We also show the latency comparison of our searched models on different devices in Fig. 1 and Fig. 7. Our method achieves similar performance to SOTA but reduces $1.9\times$, $3.3\times$, $2.6\times$, $3.8\times$ latency on GPU, CPU, embedded device, and mobile phone, respectively. Compared to the previous fastest methods, the fastest models searched by iNAS speed up $2.1\times$, $3.7\times$, $2.5\times$, and $4.8\times$ on these devices. Current SOD models are mostly designed for GPU while ignoring other devices. Some ResNet-based and VGG-based methods can not even be applied to the embedded device due to the out-of-memory error. In comparison, our device-aware searched models achieve consistent latency reduction on all devices.

Device-aware search. To verify the effectiveness of device-aware search, we compare the models searched on GPU and specialized devices in Tab. 3. We benchmark the latency of the iNAS(GPU)-L on target devices. With aligned performance, models searched on specialized devices achieve 12.1%, 14.5%, 10.9% latency reductions on CPU, mobile phone, and embedded device, respectively. This observation verifies that device-aware search can find suitable mod-

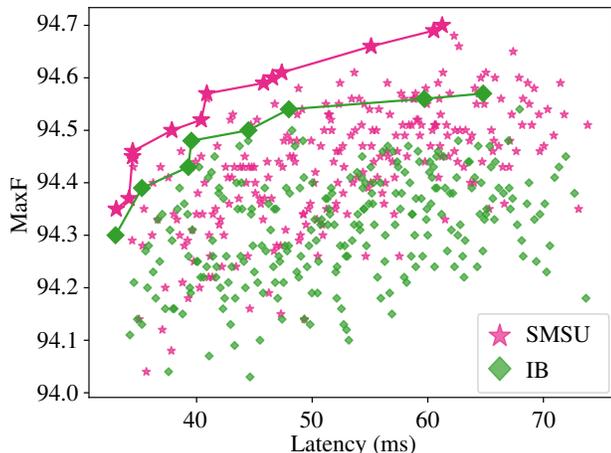


Figure 10. Comparison of the search space constructed by the inverted bottleneck (IB) [50] and our proposed searchable multi-scale unit (SMSU).

els for target devices to reduce latency.

Integral search space. iNAS supports an integral search space for SOD. Fig. 8 verifies the importance of the integral search space. For the baseline network, we use the MobileNetV2 structure [50] as the fixed backbone and combine the Amulet transport [75] and DSS decoder [25] to form the fixed saliency head. As shown in Fig. 8 (b), the fixed baseline network csots 45.17 *ms* inference latency on the CPU and gets 94.1% on ECSSD. Only enabling the searchable backbone or searchable saliency head reduces the lower-bound of latency to 41.20 *ms* (-8.7%) or 36.20 *ms* (-19.8%) with similar performance. While using the integral search space greatly reduces the lower-bound latency to 33.06 *ms* (-26.8%) but improves the performance of the fastest architecture to 94.4%. Similarly, the upper-bound of performance is promoted to 94.7%. Fig. 8 (a) shows the integral search space has a consistently better Pareto frontier over partial searchable space and significantly improves the handcraft structure on both the latency and performance.

Latency-group sampling. Fig. 9 compares the evolution search based on uniform sampling and proposed latency-group sampling (LGS). The lower-bound and upper-bound latency of the search space is 32.12 *ms* and 74.14 *ms*, respectively. As shown in Fig. 9, the lower-bound and upper-bound latency obtained by uniform sampling are 38.55 *ms* and 59.62 *ms*, which only account for 50.2% of whole latency range. While LGS ensures each latency group to have balanced samples and offsprings, thus can explore 99% of the search space. As a result, our proposed LGS obtains a broader Pareto frontier over uniform sampling.

Searchable multi-scale unit. Fig. 10 verifies the effectiveness of the proposed searchable multi-scale unit (SMSU). We compare the search space constructed by SMSU with the inverted bottleneck (IB). Enhancing the IB with multi-scale ability, search space constructed by SMSU shows a

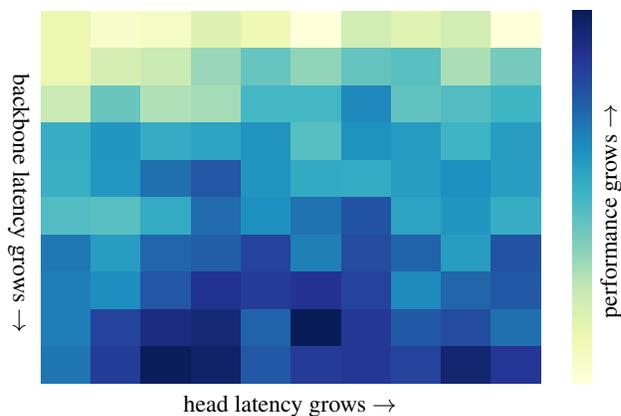


Figure 11. Visualization of the correspondence between the backbone/head latency and the performance.

better latency-performance Pareto frontier. We observe that the improvement of higher latency models is larger, because a relaxed latency constrain enables large kernel, which supports more powerful multi-scale kernel combinations.

4.3. Observation

To explore the relation of performance with the backbone and head latency, we divide the backbone and the head latency into 10 groups and sample 20 models in each grid, resulting in 2000 samples. Observing Fig. 11, we find (1) a more complicated backbone consistently improves the performance; (2) while the complicated saliency head is not always the best choice. These observations show why integral search space can reduce model latency, *i.e.*, iNAS can choose appropriate saliency heads for backbones of specific latency. Because choosing appropriate saliency heads for better latency-performance balance has no apparent pattern, searching may be an efficient solution for designing low-latency SOD models.

5. Conclusion

In this work, we propose an integral search (iNAS) space for SOD, which generalizes the designs of handcraft SOD models. The integral search can automatically find correspondence between backbone and head and get the best performance-latency balance. Then we propose a latency-group sampling to explore our entire integral search space. The experiment demonstrates that iNAS has similar performance to the handcraft SOTA SOD methods but largely reduces their latency in various devices. Our work paves the way for SOD applications on low-power devices.

Acknowledgement. This research was supported by NSFC (61922046), S&T innovation project from Chinese Ministry of Education, BNRist (No. BNR2020KF01001) and the Fundamental Research Funds for the Central Universities (Nankai University, No. 63213090).

References

- [1] Ali Borji, Ming-Ming Cheng, Qibin Hou, Huaizu Jiang, and Jia Li. Salient object detection: A survey. *Computational Visual Media*, 5(2):117–150, 2019. [1](#), [2](#), [3](#)
- [2] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *Int. Conf. Learn. Represent.*, 2020. [2](#), [3](#), [6](#)
- [3] 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. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017. [3](#)
- [4] Shuhan Chen, Xiuli Tan, Ben Wang, and Xuelong Hu. Reverse attention for salient object detection. In *Eur. Conf. Comput. Vis.*, pages 234–250, 2018. [2](#), [5](#)
- [5] Ming-Ming Cheng and Deng-Ping Fan. Structure-measure: A new way to evaluate foreground maps. *Int. J. Comput. Vis.*, 129(9):2622–2638, 2021. [7](#)
- [6] Ming-Ming Cheng, Niloy J Mitra, Xiaolei Huang, Philip HS Torr, and Shi-Min Hu. Global contrast based salient region detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(3):569–582, 2015. [2](#), [7](#)
- [7] Ming-Ming Cheng, Jonathan Warrell, Wen-Yan Lin, Shuai Zheng, Vibhav Vineet, and Nigel Crook. Efficient salient region detection with soft image abstraction. In *Int. Conf. Comput. Vis.*, pages 1529–1536, 2013. [7](#)
- [8] Ming-Ming Cheng, Fang-Lue Zhang, Niloy J Mitra, Xiaolei Huang, and Shi-Min Hu. Repfinder: finding approximately repeated scene elements for image editing. *ACM Trans. Graph.*, 29(4):1–8, 2010. [1](#)
- [9] Xiangxiang Chu, Bo Zhang, and Ruijun Xu. Fairnas: Rethinking evaluation fairness of weight sharing neural architecture search. In *Int. Conf. Learn. Represent.*, 2021. [3](#)
- [10] Zijun Deng, Xiaowei Hu, Lei Zhu, Xuemiao Xu, Jing Qin, Guoqiang Han, and Pheng-Ann Heng. R3net: Recurrent residual refinement network for saliency detection. In *IJCAI*, pages 684–690, 2018. [1](#), [6](#)
- [11] Xiaohan Ding, Yuchen Guo, Guiguang Ding, and Jungong Han. Acnet: Strengthening the kernel skeletons for powerful cnn via asymmetric convolution blocks. In *Int. Conf. Comput. Vis.*, pages 1911–1920, 2019. [4](#)
- [12] Mengyang Feng, Huchuan Lu, and Errui Ding. Attentive feedback network for boundary-aware salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1623–1632, 2019. [2](#), [5](#)
- [13] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–1, 2020. [1](#)
- [14] Shang-Hua Gao, Qi Han, Duo Li, Pai Peng, Ming-Ming Cheng, and Pai Peng. Representative batch normalization with feature calibration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. [2](#)
- [15] Shang-Hua Gao, Qi Han, Zhong-Yu Li, Pai Peng, Liang Wang, and Ming-Ming Cheng. Global2local: Efficient structure search for video action segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. [3](#)
- [16] Shang-Hua Gao, Yong-Qiang Tan, Ming-Ming Cheng, Chengze Lu, Yunpeng Chen, and Shuicheng Yan. Highly efficient salient object detection with 100k parameters. In *Eur. Conf. Comput. Vis.*, 2020. [1](#), [2](#), [4](#), [6](#)
- [17] Yanliang Ge, Cong Zhang, Kang Wang, Ziqi Liu, and Hongbo Bi. Wgi-net: A weighted group integration network for rgb-d salient object detection. *Computational Visual Media*, 7(1):115–125, 2021. [2](#)
- [18] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Nas-fpn: Learning scalable feature pyramid architecture for object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7036–7045, 2019. [3](#)
- [19] Yu-Chao Gu, Li-Juan Wang, Yun Liu, Yi Yang, Yu-Huan Wu, Shao-Ping Lu, and Ming-Ming Cheng. Dots: Decoupling operation and topology in differentiable architecture search. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12311–12320, 2021. [3](#)
- [20] Yu-Chao Gu, Li-Juan Wang, Zi-Qin Wang, Yun Liu, Ming-Ming Cheng, and Shao-Ping Lu. Pyramid constrained self-attention network for fast video salient object detection. In *AAAI*, pages 10869–10876, 2020. [1](#), [2](#)
- [21] 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 *Eur. Conf. Comput. Vis.*, pages 544–560, 2020. [2](#), [3](#)
- [22] Junfeng He, Jinyuan Feng, Xianglong Liu, Cheng Tao, and S F Chang. Mobile product search with bag of hash bits and boundary reranking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2012. [1](#)
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. [1](#), [2](#), [3](#)
- [24] Seunghoon Hong, Tackgeun You, Suha Kwak, and Bohyung Han. Online tracking by learning discriminative saliency map with convolutional neural network. In *ICML*, pages 597–606, 2015. [1](#)
- [25] Qibin Hou, Ming-Ming Cheng, Xiaowei Hu, Ali Borji, Zhuowen Tu, and Philip Torr. Deeply supervised salient object detection with short connections. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(4):815–828, 2019. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [26] Shi-Min Hu, Dun Liang, Guo-Ye Yang, Guo-Wei Yang, and Wen-Yang Zhou. Jittor: a novel deep learning framework with meta-operators and unified graph execution. *Science China Information Sciences*, 63(222103):1–21, 2020. [6](#)
- [27] Agus Kurniawan. Introduction to nvidia jetson nano, 2021. [1](#)
- [28] Guanbin Li and Yizhou Yu. Visual saliency based on multiscale deep features. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5455–5463, 2015. [2](#), [7](#)
- [29] Guanbin Li and Yizhou Yu. Deep contrast learning for salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 478–487, 2016. [2](#)
- [30] Xin Li, Fan Yang, Hong Cheng, Wei Liu, and Dinggang Shen. Contour knowledge transfer for salient object detection. In *Eur. Conf. Comput. Vis.*, pages 355–370, 2018. [2](#)

- [31] Yin Li, Xiaodi Hou, Christof Koch, James M Rehg, and Alan L Yuille. The secrets of salient object segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 280–287, 2014. [7](#)
- [32] Yingwei Li, Xiaojie Jin, Jieru Mei, Xiaochen Lian, Linjie Yang, Cihang Xie, Qihang Yu, Yuyin Zhou, Song Bai, and Alan L Yuille. Neural architecture search for lightweight non-local networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10297–10306, 2020. [3](#)
- [33] Yanwei Li, Lin Song, Yukang Chen, Zeming Li, Xiangyu Zhang, Xingang Wang, and Jian Sun. Learning dynamic routing for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8553–8562, 2020. [1](#)
- [34] 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 Conf. Comput. Vis. Pattern Recog.*, pages 82–92, 2019. [1, 3](#)
- [35] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *Int. Conf. Learn. Represent.*, 2019. [1, 3](#)
- [36] Jiang-Jiang Liu, Qibin Hou, Ming-Ming Cheng, Jiashi Feng, and Jianmin Jiang. A simple pooling-based design for real-time salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3917–3926, 2019. [1, 2, 3, 5, 6](#)
- [37] Jiang-Jiang Liu, Zhi-Ang Liu, and Ming-Ming Cheng. Centralized information interaction for salient object detection. *arXiv preprint arXiv:2012.11294*, 2020. [5](#)
- [38] Nian Liu and Junwei Han. DHSNet: Deep hierarchical saliency network for salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 678–686, 2016. [2, 5](#)
- [39] Nian Liu, Junwei Han, and Ming-Hsuan Yang. PiCANet: Learning pixel-wise contextual attention for saliency detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3089–3098, 2018. [2, 3, 6, 7](#)
- [40] Yun Liu, Yu-Chao Gu, Xin-Yu Zhang, Weiwei Wang, and Ming-Ming Cheng. Lightweight salient object detection via hierarchical visual perception learning. *IEEE TCYB*, 2020. [7](#)
- [41] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3431–3440, 2015. [2, 3, 4, 5](#)
- [42] Zhiming Luo, Akshaya Kumar Mishra, Andrew Achkar, Justin A Eichel, Shaozi Li, and Pierre-Marc Jodoin. Non-local deep features for salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6609–6617, 2017. [1, 6](#)
- [43] NVIDIA, Péter Vingelmann, and Frank H.P. Fitzek. Cuda, release: 10.2.89, 2020. [1](#)
- [44] Youwei Pang, Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. Multi-scale interactive network for salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9413–9422, 2020. [1, 2, 3, 4, 5, 6](#)
- [45] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *ICML*, pages 4095–4104, 2018. [3](#)
- [46] Xuebin Qin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Osmar R Zaiane, and Martin Jagersand. U2-net: Going deeper with nested u-structure for salient object detection. *Pattern Recognition*, 106:107404, 2020. [1, 4, 6](#)
- [47] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. BASNet: Boundary-aware salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7479–7489, 2019. [1, 6, 7](#)
- [48] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI*, volume 33, pages 4780–4789, 2019. [3](#)
- [49] Tonmoy Saikia, Yassine Marrakchi, Arber Zela, Frank Hutter, and Thomas Brox. Autodispnet: Improving disparity estimation with automl. In *Int. Conf. Comput. Vis.*, October 2019. [3](#)
- [50] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2018. [3, 4, 8](#)
- [51] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Int. Conf. Learn. Represent.*, 2015. [1, 2, 3](#)
- [52] Benoit Steiner, Zachary DeVito, Soumith Chintala, Sam Gross, Adam Paszke, Francisco Massa, Adam Lerer, Gregory Chanan, Zeming Lin, Edward Yang, et al. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inform. Process. Syst.*, 32:8026–8037, 2019. [3, 6, 7](#)
- [53] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2820–2828, 2019. [1, 3](#)
- [54] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuan-dong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, et al. Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12965–12974, 2020. [3](#)
- [55] Jingdong Wang, Huaizu Jiang, Zejian Yuan, Ming-Ming Cheng, Xiaowei Hu, and Nanning Zheng. Salient object detection: A discriminative regional feature integration approach. *Int. J. Comput. Vis.*, 123(2):251–268, 2017. [2](#)
- [56] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 136–145, 2017. [7](#)
- [57] Linzhao Wang, Lijun Wang, Huchuan Lu, Pingping Zhang, and Xiang Ruan. Saliency detection with recurrent fully convolutional networks. In *Eur. Conf. Comput. Vis.*, pages 825–841, 2016. [2](#)
- [58] Tiantian Wang, Lihe Zhang, Shuo Wang, Huchuan Lu, Gang Yang, Xiang Ruan, and Ali Borji. Detect globally, refine locally: A novel approach to saliency detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3127–3135, 2018. [2, 5](#)
- [59] Wenguan Wang, Qiuxia Lai, Huazhu Fu, Jianbing Shen, Haibin Ling, and Ruigang Yang. Salient object detection in the deep learning era: An in-depth survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2021. [1, 2, 3](#)

- [60] Wenguan Wang, Jianbing Shen, Ming-Ming Cheng, and Ling Shao. An iterative and cooperative top-down and bottom-up inference network for salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5968–5977, 2019. [2](#)
- [61] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10734–10742, 2019. [3](#), [4](#)
- [62] Runmin Wu, Mengyang Feng, Wenlong Guan, Dong Wang, Huchuan Lu, and Errui Ding. A mutual learning method for salient object detection with intertwined multi-supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8150–8159, 2019. [2](#)
- [63] Yu-Huan Wu, Yun Liu, Jun Xu, Jia-Wang Bian, Yuchao Gu, and Ming-Ming Cheng. Mobilesal: Extremely efficient rgb-d salient object detection. *arXiv preprint arXiv:2012.13095*, 2020. [2](#)
- [64] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3907–3916, 2019. [1](#), [2](#), [5](#), [6](#)
- [65] Lingxi Xie and Alan Yuille. Genetic cnn. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1379–1388, 2017. [3](#)
- [66] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1492–1500, 2017. [4](#)
- [67] Yingyue Xu, Dan Xu, Xiaopeng Hong, Wanli Ouyang, Rongrong Ji, Min Xu, and Guoying Zhao. Structured modeling of joint deep feature and prediction refinement for salient object detection. In *Int. Conf. Comput. Vis.*, pages 3789–3798, 2019. [2](#), [5](#)
- [68] Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia. Hierarchical saliency detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1155–1162, 2013. [7](#)
- [69] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3166–3173, 2013. [7](#)
- [70] Lewei Yao, Hang Xu, Wei Zhang, Xiaodan Liang, and Zhen-guo Li. Sm-nas: structural-to-modular neural architecture search for object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12661–12668, 2020. [3](#)
- [71] Jiahui Yu and Thomas S Huang. Universally slimmable networks and improved training techniques. In *Int. Conf. Comput. Vis.*, pages 1803–1811, 2019. [7](#)
- [72] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. In *Eur. Conf. Comput. Vis.*, pages 702–717, 2020. [2](#), [3](#), [6](#)
- [73] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4126–4135, 2020. [3](#)
- [74] Lu Zhang, Ju Dai, Huchuan Lu, You He, and Gang Wang. A bi-directional message passing model for salient object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1741–1750, 2018. [2](#)
- [75] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Xiang Ruan. Amulet: Aggregating multi-level convolutional features for salient object detection. In *Int. Conf. Comput. Vis.*, pages 202–211, 2017. [2](#), [3](#), [4](#), [8](#)
- [76] Wenqiang Zhang, Jiemin Fang, Xinggong Wang, and Wenyu Liu. Efficientpose: Efficient human pose estimation with neural architecture search. *Computational Visual Media*, pages 1–13, 2021. [3](#)
- [77] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6848–6856, 2018. [4](#)
- [78] Jia-Xing Zhao, Jiangjiang Liu, Den-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. EGNNet: Edge guidance network for salient object detection. In *Int. Conf. Comput. Vis.*, pages 8779–8788, 2019. [1](#), [2](#), [6](#)
- [79] Rui Zhao, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Saliency detection by multi-context deep learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1265–1274, 2015. [2](#)
- [80] Xiaoqi Zhao, Youwei Pang, Lihe Zhang, Huchuan Lu, and Lei Zhang. Suppress and balance: A simple gated network for salient object detection. In *Eur. Conf. Comput. Vis.*, pages 35–51, 2020. [2](#), [5](#)
- [81] Huajun Zhou, Xiaohua Xie, Jian-Huang Lai, Zixuan Chen, and Lingxiao Yang. Interactive two-stream decoder for accurate and fast saliency detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2020. [1](#), [2](#), [4](#), [6](#)
- [82] Tao Zhou, Deng-Ping Fan, Ming-Ming Cheng, Jianbing Shen, and Ling Shao. Rgb-d salient object detection: A survey. *Computational Visual Media*, 7(1):37–69, 2021. [2](#)
- [83] Wangjiang Zhu, Shuang Liang, Yichen Wei, and Jian Sun. Saliency optimization from robust background detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2814–2821, 2014. [2](#)
- [84] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *Int. Conf. Learn. Represent.*, 2017. [3](#)
- [85] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8697–8710, 2018. [3](#)