

RF-Next: Efficient Receptive Field Search for Convolutional Neural Networks

Shanghua Gao Zhong-Yu Li Qi Han Ming-Ming Cheng Liang Wang

Abstract—Temporal/spatial receptive fields of models play an important role in sequential/spatial tasks. Large receptive fields facilitate long-term relations, while small receptive fields help to capture the local details. Existing methods construct models with hand-designed receptive fields in layers. Can we effectively search for receptive field combinations to replace hand-designed patterns? To answer this question, we propose to find better receptive field combinations through a global-to-local search scheme. Our search scheme exploits both global search to find the coarse combinations and local search to get the refined receptive field combinations further. The global search finds possible coarse combinations other than human-designed patterns. On top of the global search, we propose an expectation-guided iterative local search scheme to refine combinations effectively. Our RF-Next models, plugging receptive field search to various models, boost the performance on many tasks, *e.g.*, temporal action segmentation, object detection, instance segmentation, and speech synthesis. The source code is publicly available on <http://mmcheng.net/rfnext>.

Index Terms—dilation, receptive field, spatial convolutional network, temporal convolutional network, temporal action segmentation



1 INTRODUCTION

Due to the strong representation ability, convolutional neural networks (CNN) have been widely used in spatially visual recognition tasks, *e.g.*, object detection [2], saliency detection [3], instance segmentation [4], and semantic segmentation [5], [6], as well as sequential perception tasks, *e.g.*, temporal action segmentation [7], [8], and speech synthesis [9], [10]. CNN processes short/long-term features by stacking convolutional filters with different receptive fields. Spatial convolutional networks (SCN) for visual recognition tasks process local and global features to represent the texture and semantic information. Temporal convolutional networks (TCN) [7], [11]–[14] are widely adopted in sequential tasks with their ability to capture both long-term and short-term information. Appropriate receptive fields in layers are crucial for both SCN and TCN as large receptive fields contribute to long-term dependencies while small receptive fields benefit the local details. State-of-the-art (SOTA) SCN [2], [4], [15]–[17] and TCN [7], [8], [13], [18], [19] methods rely on human-designed receptive field combinations, *i.e.*, dilation rate or pooling size in each layer, to make the trade-off between capturing long and short term dependencies. Questions have been raised: Are there other effective receptive field combinations that perform comparable or better than hand-designed patterns? Will the receptive field combinations vary among different datasets? To answer those questions, we propose to find the possible receptive field combinations in a coarse-to-fine scheme through the global-to-local search.

As shown in Fig. 1, unlike the existing network architecture search spaces [20]–[22] that only contain several operation options within a layer, the available search space of receptive field combinations could be huge. Suppose a TCN/SCN has L convolutional layers and D possible receptive fields in each layer,




Fig. 1. Search space comparison between searching for network architecture and receptive field combinations. Left: Network architecture search mostly searches for several operations with different functions. Right: The search space of receptive field combinations is huge. The white, green, blue nodes, and orange shade represent the dilation rate candidates, the sparse search space in global search, one of the global searched results, and the local search space, respectively.

there are D^L possible combinations, *i.e.*, the MS-TCN [12] for the long-sequence temporal action segmentation task, consisting of 40 layers and 1024 possible receptive fields in each layer, has 1024^{40} possible receptive field combinations. Directly applying network architecture searching algorithms [20], [22]–[24] to such a huge search space is impractical. For example, conventional reward-based searching methods [24]–[26] are unsuitable for CNN-based models with a huge search space. The model training and performance evaluation of each possible combination are too costly. Differentiable architecture searching methods (DARTS) [20], [21], [23] rely on shared big networks to save training time, thus only supporting several operators within a layer due to the constraint of model size. Moreover, they are heavily dependent on the initial combination and fail to find new combinations with a huge difference from the initial one. While our goal is to explore effective receptive field combinations other than human-designed patterns in the huge search space, those algorithms are either too costly or cannot support the large search space.

To explore the effective receptive fields with a low cost, we exploit both a genetic-based global search to find the coarse

- S. Gao, Z.Y. Li, Q. Han, and M.M. Cheng are with the TMCC, CS, Nankai University, Tianjin 300350, China.
- L. Wang is with National Laboratory of Pattern Recognition.
- M.M. Cheng (cmm@nankai.edu.cn) is the corresponding author.
- A preliminary version of this work has been presented in the CVPR 2021 [1].

receptive field combinations and an expectation-guided iterative (EGI) local search to get the refined combinations. Specifically, we follow common settings in many existing methods [5], [12], [27], [28] to use dilation rates to determine layers' receptive fields. A genetic-based global search scheme is proposed to find coarse combinations within a sparsely sampled search space at an affordable cost. The global search discovers various combinations that achieve even better performance than human designings but have completely different patterns. Based on the global-searched coarse combinations, we propose the local search to determine fine-grained dilation rates. In local search, a convolutional weight-sharing scheme enforces learned dilation coefficients to approximate the probability mass distribution for calculating the expectation of dilation rates. The expectation-guided searching transfers the discrete dilation rates into a distribution, allowing fine-grained dilation rates search. With an iterative searching process, the local search gradually finds more effective fine-grained receptive field combinations with a low cost. Models enhanced by our proposed global-to-local search scheme, namely RF-Next models, surpass human-designed structures with impressive performance gain on many tasks. In summary, we make two major contributions:

- The expectation-guided iterative local search scheme enables searching fine-grained receptive field combinations in the dense search space.
- The global-to-local search discovers effective receptive field combinations with better performance than hand-designed patterns.

The conference version [1] mainly explores the proposed receptive field searching scheme on the temporal action segmentation task. In this work, we improve and give more analysis of our proposed searching scheme, *i.e.*, parallel receptive fields for multi-scale enhancement, searching cost analysis, generalization to different tasks, observations of receptive fields. We generalize our proposed receptive field searching method to multiple tasks on both temporal and spatial dimensions, *i.e.*, speech synthesis, sequence modeling, instance segmentation, object detection, and semantic segmentation. To meet the multi-scale requirements of some tasks, we take advantage of the proposed expectation-guide search scheme to enable the searched structure with parallel multiple receptive fields and shared convolutional weights. We give observations on the receptive field requirements of multiple tasks based on the searched results: 1) Proper receptive fields of CNN are beneficial to many tasks. 2) The receptive field requirements of different parts of the network are quite different.

2 RELATED WORK

2.1 Receptive Field in Networks

The effect of receptive fields has been widely studied [29]–[34]. Although the theoretical receptive field could be huge, [29] shows that the effective receptive field occupies a fraction of the theoretical receptive field. [30] assumes that useful predictive information for a pixel comes from nearby locations rather than far pixels and proposes to gradually suppress the distant pixel values with the scale-sensitive regularization. [35] enlarges the receptive field to improve the performance of the image super-resolution task. [31] observes that model depth must be congruent with the receptive field size for the image super-resolution task. [34] forms continuous receptive fields with the help of Gaussian scale-space representation. Parallel receptive fields are proposed

to enable a more flexible receptive field within a layer [32], [33], [36]. Inception net series [36]–[38] explores the parallel asymmetric convolutions with different receptive fields to enhance the model representation ability. Atrous spatial pyramid modules [5], [32] prove the effectiveness of parallel receptive fields in semantic segmentation. OctConv [33] decomposes the convolution to process two feature scales simultaneously. Some works model the connections of all positions to form arbitrary receptive fields theoretically [39]–[41]. Multi-head attention is presented by [39] to model the relation between every two pixels with attention mechanisms. Similarly, [40] utilizes the non-local operation to aggregate features at all positions. A graph-based global reasoning module [41] is proposed to capture relations between arbitrary regions. Despite these methods exploring receptive fields' great potential, choosing effective receptive fields for different tasks is still an open question.

2.2 Sequential Tasks

Sequential tasks process data in the form of sequences, *e.g.*, video stream and audio stream. As the sequence length could have a large variance for sequential tasks, models with a proper range of effective receptive fields are needed. In this work, we mainly tackle two sequential tasks with long sequences of data, *i.e.*, temporal action segmentation and speech synthesis, which represent video and audio dimensions, respectively.

2.2.1 Temporal Action Segmentation

Temporal action recognition segments the action of each video frame, playing an important role in computer vision applications such as clips tagging [42], video surveillance [43], [44], and anomaly detection [45]. While conventional works [46]–[49] have continuously refreshed the recognition performance of short trimmed videos containing a single activity, segmenting each frame densely in long untrimmed videos remains challenging as those videos contain many activities with different temporal lengths. In this work, we do our major experiments on the temporal action segmentation task. Therefore, we give a thorough introduction to related works for the temporal action segmentation task.

Many approaches have been proposed for modeling dependencies for temporal action segmentation. Early works [50]–[52] mostly model the changing state of appearance and actions with sliding windows [53]–[55]. Thus they mainly focus on short-term dependencies. Capturing both short-term and long-term dependencies then gradually becomes the focus of temporal action segmentation.

Sequential Model. Sequential models capture long-short term dependencies in an iterative form. Vo and Bobick [56] apply the Bayes network to segment actions represented by a stochastic context-free grammar. Tang *et al.* [57] use a hidden Markov model to model transitions between states and durations. Later, hidden Markov models are combined with context-free grammar [58], Gaussian mixture model [59], and recurrent networks [60], [61] to model long-term action dependencies. Cheng *et al.* [62] apply the sequence memorizer to capture long-range dependencies in visual words learned from the video. However, these sequential models are inflexible in parallelly modeling long-term dependencies and usually suffer from information forgetting [7], [12].

Multi-stream Architecture. Some researchers [63]–[66] utilize multi-stream models to model dependencies from long and short term. Richard and Gall employ [63] dynamic programming to inference models composed of length model, language model, and action classifier. Singh *et al.* [64] learn short video chunks representation with a two-stream network and pass these chunks to a bi-directional network to predict temporal action segmentation results sequentially. A three-stream architecture is proposed in [65], which contains egocentric cues, spatial and temporal streams. Tricorner [66] utilizes a hybrid temporal convolutional and recurrent network to capture local motion and memorize long-term action dependencies. CoupledGAN [67] uses a GAN model to utilize multi-modal data to better model human actions' evolution. Capturing long-short term information with multiple streams increases computational redundancy.

Temporal Convolutional Network. Recently, temporal convolutional networks (TCN) have been introduced to model dependencies of different ranges within a unified structure by adjusting receptive fields and can process long videos in parallel. Lea *et al.* [11] propose the encoder-decoder style TCN for the temporal action segmentation to capture long-range temporal patterns and apply the dilated convolution to enlarge the receptive field. TDRN [68] further introduces the deformable convolution to process the full-resolution residual stream and low-resolution pooled stream. MS-TCN [7], [12] utilizes multi-stage dilated TCNs with hand-designed dilation rate combinations to capture information from various temporal receptive fields. However, the adjustment of receptive fields still relies on human design, which may not be appropriate. Our proposed efficient receptive field combinations searching scheme can automatically discover more efficient structures, improving these TCN based methods.

Complementary Techniques. Instead of capturing long-term and short-term information, some works [13], [69] further improve the temporal action segmentation performance with boundary refinement. Li *et al.* [69] utilize an iterative training procedure with transcript refinement and soft boundary assignment. Wang *et al.* [13] leverage semantic boundary information to refine the prediction results. Other researchers focus on temporal action segmentation under the weakly supervised [59], [60], [69], or unsupervised [70] settings. These works still rely on the efficient TCN to model the action dependencies, thus complementing the proposed method.

2.2.2 Speech Synthesis

Speech synthesis, also known as text to speech (TTS), aims to synthesize human-like natural speech from text [71], [72], which has thrived due to the strong feature representation ability of the neural networks [9], [10], [27], [73]. The large difference in modality and feature-length between text and speech makes it hard to implement TTS in an end-to-end style [74]. Therefore, common approaches for speech synthesis decompose this process into the following steps: 1) text-to-linguistic features transformation [75], [76]. 2) transfer linguistic features or text to acoustic features [10], [77]–[79]. 3) generate the final speech waveform using the acoustic features [9], [27], [28], [73]. During the process, features are transformed from the short-length text sequence to the long-length acoustic features and speech waveform, *i.e.*, text with about 20 words results in a 5-second speech sequence of 80k

sampling points. We tackle the challenge of transferring acoustic features to speech waveform, as this procedure requires a suitable receptive field to model the short/long-term dependencies in the waveform [27], [74]. Since we only focus on the receptive field of TTS models, we refer readers to the survey [74] for more details of speech synthesis.

2.3 Spatial Tasks

Unlike the sequential tasks that mostly process one-dimension sequences, spatial tasks process images with two dimensions, *i.e.*, height and width dimensions. To extract features of objects of various sizes in the scene, the model needs small receptive fields to detect small objects and large receptive fields to cover large objects or capture surrounding context information [16]. We mainly tackle two popular vision tasks with the proposed receptive field search, *i.e.*, object detection and instance segmentation.

2.3.1 Object Detection

Object detection aims to localize objects with bounding boxes and assign categories accordingly [80]–[84]. Common object detection methods can be divided into single-stage [81], [82], [85]–[87] and two-stage [2], [88], [89] pipelines. More details about object detection can be referred to related surveys [90]–[92]. Single-stage detectors, *e.g.*, SSD [85], YOLO [86], and CornerNet [87], require one inference to end-to-end localize and categorize objects, which are effective in latency but hard to cover all objects. Two-stage methods, *e.g.*, R-CNN [93], Faster-RCNN [2], and Cascade R-CNN [88], decompose object detection into region proposals generation and objects detection from the proposal, enhancing the detection quality at the cost of slow inference speed. Despite the difference, object detectors tend to enhance the multi-scale ability to handle objects of various sizes [2], [85], [88], [89]. SSD [85] merges features from multiple stages to detect objects. Faster-RCNN [2] utilizes a feature pyramid network to aggregate features with multiple scales. Cascade R-CNN [88] and HTC [89] perform cascaded multi-stage feature fusion and refinement. We show that the proper receptive field settings for these methods can further improve the detection ability.

2.3.2 Instance Segmentation

Instance segmentation aims to assign category labels to each pixel of instances [94]–[96], which is similar to object detection as they both require localizing objects. Therefore, common instance segmentation methods add the segmentation branch on object detectors to segment instances from the bounding boxes, *i.e.*, Mask-RCNN [4] extends Faster-RCNN by adding an object mask predicting branch. The multi-scale ability in object detectors, *e.g.*, Cascade R-CNN [88] and HTC [89], is naturally inherited to the instance segmentation. Some works focus on refining the boundary of segmentation masks [97]–[101], which still rely on feature extractors with proper receptive fields. We observe that the receptive field search benefits the performance of instance segmentation.

2.4 Network Architecture Search

The genetic algorithm [102] has achieved remarkable performance on a wide range of applications. Many genetic-based methods have been recently introduced for the neural networks architecture search of vision tasks [24]–[26], [103], [104]. An evolutionary

coding scheme is proposed in Genetic CNN [24] to encode the network architecture to a binary string. A hierarchical representation is presented by Liu *et al.* [26] to constrain the search space. Real *et al.* [25] regularize the evolution by an age property selection operation. Sun *et al.* [103] introduce a variable-length encoding method for effective architecture designing. However, the genetic algorithm requires the training of each candidate, consuming too much computational cost when faced with a huge search space.

Differentiable architecture search [20], [23], [105]–[108] saves the training time by introducing a large network containing subnetworks with different searching options. The importance of searched blocks is determined by gradient backpropagation [109]. However, these network architecture search methods are designed for finding a limited number of operations such as convolution, ReLU, batch normalization, short connection, *etc.* Thus, these methods cannot be directly used for receptive field search due to the different searching targets, *e.g.*, they cannot handle the huge receptive field combinations search space. Fair DARTS [105] solves the problem of performance collapse due to the unfair advantage in exclusive competition between different operators. The receptive field searching has no such problem as it contains the same operation. [106] reduces the memory cost of supernet by randomly sampling a proportion of channels for operation search and bypassing the held out part in a shortcut. This approach is unsuitable for receptive field search as the shortcut does not belong to the receptive field search space. [107] conducts channel number and feature resolution searching with a masking mechanism for feature map reuse. The masking mechanism for feature map reuse cannot be applied to the receptive field search represented by dilation rate because different dilation rates do not belong to each other and cannot be chosen with different masks. [108] narrows the performance gap between models searched from a small dataset and evaluated on the large dataset. Theoretically, [108] is orthogonal to receptive field search space. However, the high efficiency of the proposed local search allows to directly search on the large dataset instead of on a small dataset. In this paper, we propose a global search to handle the huge search space with sparse sampling. The expectation-guided iterative local search then transfers the sparse search space of receptive fields into the dense one for fine-level searching.

This differentiable search idea is further extended [110] to deal with semantic segmentation [23], [111], and other tasks beyond image classification [21]. Auto-deeplab [23] and DCNAS [111] focus on searching the feature resolution for different stages of the semantic segmentation network. We show that our proposed receptive field searching scheme can find better receptive fields on these searched segmentation networks.

3 METHOD

The pipeline of our proposed global-to-local search method has two components: (i) a genetic-based global search algorithm that produces coarse but competitive combinations of the receptive fields; (ii) an expectation-guided iterative local search scheme that locally refines the global-searched coarse structures.

3.1 Description

Our objective is to efficiently search for optimal receptive field combinations for the given dataset. The receptive field can be represented in multiple forms: the dilation rate, kernel size, pooling size, stride, and the stack number of layers. Our method is

initially designed for temporal action segmentation. We mainly follow the MS-TCN [12] to formulate the receptive fields using the combinations of dilation rates in layers and propose to evolve these combinations during the searching process. Other receptive field representations can also be applied to the proposed global-to-local search with minor adjustments. Though we conduct major experiments on the temporal action segmentation task, our receptive field searching method can easily be generalized to new tasks, as introduced in Sec. 3.4.

Suppose a TCN has L convolutional layers and $D = \{d_1, d_2, \dots, d_N\}$ is the possible dilation-rates/receptive-fields in each layer. The combination of receptive fields is represented with $C = \{c_1, \dots, c_l, \dots, c_L\}$, where $l \in [1, L]$ is the index of layers with dilated convolutions, and $c_l \in D$ is the receptive field of each layer. There are $|D|^L$ possible combinations of receptive fields, *i.e.*, the possible receptive field combinations in MS-TCN [12] is 1024^{40} when dilation rates range from 1 to 1024. Directly searching for effective combinations in such a large search space is impractical. We thus decompose the searching process into the global and local search to find the combination in a coarse-to-fine manner.

3.2 Global Search

The global search aims to find possible coarse receptive field combinations in the huge search space that are very different from human-designed structures, focusing more on discovering new structures with large diversity to human-designings instead of performance. To guarantee the diversity of new structures, we utilize the random sparse sampling strategy and apply a generic algorithm with the random crossover and random mutation specifically designed for receptive fields.

Population initialization with gradually sparse sampling. The objective of the global search is to find the coarse receptive field combinations at an affordable cost. Therefore, we reduce the search space by sparsely sampling the dilation rates within layers. Multiple sparse discrete sampling strategies such as uniform sampling, gradually sparse sampling, and gradually dense sampling can be applied to sparse the search space. Because small receptive fields benefit the extraction of precise local details, and large receptive fields contribute to coarse long-term dependencies. A gradually sparse sampling scheme from small to large dilation rates is appropriate for common tasks, *e.g.*, temporal action segmentation. Therefore, we formulate the receptive field space in global search as:

$$D_g = \{d_i = k^i, i \in [0, 1, \dots, T]\}, \quad (1)$$

where k is the controller of the search space sparsity, and T determines the largest receptive field. With the same maximum receptive field, $|D_g| \ll |D|$, the search space is greatly reduced. *i.e.*, when set $k = 2$, and set the maximum receptive field to 1024 as in MS-TCN, the search space is reduced from 1024^{40} to 11^{40} . The population of receptive field combinations can be described as a group of candidate structures $P = \{C_i, i \in [1, M]\}$, where C_i is the candidate structure in the global search space, and M is the number of individuals in the population.

However, the reduced space of receptive field combinations can still be huge, and unaffordable for a brute force search. We propose a genetic algorithm [102] based method to find coarse combinations that are competitive or even better than human




Fig. 2. Illustration of one iteration in our genetic-based global search algorithm. Step1. Gradually sparse random sampling initial receptive field combinations; Step2. Crossover between segments of the receptive field combinations; Step3. Randomly mutating the receptive fields to generate new individuals; Step4. Selecting individuals for the next iteration based on estimated performance of models trained with early stopping strategy.

designing. We now detail the selection, crossover, and mutation process within our proposed global search method.

Selection according to early stopped training. We need to select samples from the population of receptive field combinations P for each iteration. The selection operation selects individuals to be kept in P based on the estimated performance of each structure C_i , denoted by $E(C_i)$:

$$E(C_i) = f(V|C_i, \theta_n), \quad (2)$$

where $f(\cdot)$ is the task-specific evaluation metrics on the validation set V , e.g., frame-wise accuracy for temporal action segmentation, θ_n is a model trained with n epochs. The major cost of the global search is the performance evaluation of candidate structures. The global search aims to find coarse structures and reasonable performance, allowing searched structures to have sub-optimal performance. Also, we observe the receptive field combinations play a key role in the model convergence, i.e., a model with good receptive fields converges much faster than a model equipped with bad receptive fields. To reduce evaluation cost, we choose to early stop the training of candidate structures when the trained models can roughly show the relative performance gap of different structures, e.g., training MS-TCN for 5 epochs can reflect the structure performance. The early stop training strategy substantially reduces the structure evaluation cost.

Crossover between segments of the receptive field combination. This operation generates new samples of receptive field combinations. Every two combinations in the population are exchanged to bear new patterns of the combination while maintaining the local structures. Each C_i will be selected for the crossover operation with probability $p(C_i)$:

$$p(C_i) = \frac{E(C_i)}{\sum_i^M E(C_i)}. \quad (3)$$

Since the representation ability lies in the combination patterns, we want to reserve the local combination patterns during the crossover. Instead of randomly exchanging individual points, we choose to exchange random segments of the receptive field combination. Specifically, we randomly choose two anchors and exchange receptive field combination segments within two anchors to generate new samples.

Random receptive field mutation. The mutation operation avoids getting stuck in local optimal results by choosing an individual with pre-defined probability $p_m \in [0, 1]$ and randomly changes each value within the selected combination with pre-defined probability $p_s \in [0, 1]$. To reduce the searching cost, we also apply the

gradually sparse sampling strategy when choosing a new receptive field value.

Algorithm 1 Global Search.

Input: Iterations N , training epoch n , mutation probability p_m , and population size M ;
 Gradually sparse sampling initial receptive field combinations P ;
for iter in $[1, N]$ **do**
 Selecting individuals for the crossover with the probability obtained with estimated performance in Equ. (3);
 Crossover between random segments of every two selected receptive field combinations;
 Randomly choose combinations with probability p_m , and mutate the receptive fields with probability p_s within sparse sampling search space to generate new individuals;
 Training each individual with early stopped n epochs to save evaluation cost;
 Selecting the top M individuals based on estimated performance in Equ. (2) as the new population P ;
end for
return P .

The global search process can be summarised as Algorithm (1), and the illustration of one iteration in global search is given in Fig. 2. With the coarse search space and the global search method, we can find receptive field combinations with different patterns than human-designed structures while having similar or even better performance. We further propose the local search to locally find the more efficient combinations on top of the global-searched structures. We show in Tab. 5(b) that local search heavily relies on the initial structure, revealing the importance of global search.

3.3 Expectation-Guided Iterative Local Search

The local search aims to find more efficient receptive field combinations at a fine-grained level at a low cost. A naive approach is to sample finer-grained dilation rates near the initial dilation rate searched by the global search and apply existing DARTS algorithms [20], [21] to choose the proper one. However, even with the good initial structure provided by the global search, the available range of fine-grained dilation rates is still large. Existing search algorithms are designed to search sparse operators with several choices in each layer, thus can not handle dilation rates with hundreds of choices. While too sparsely sampling conflicts with our goal of searching for the finer-grained receptive fields. Also, DARTS methods search operators with different functionality [20], while the searching on receptive fields only contains one functional dimension. Different subsets in the dataset sometimes




Fig. 3. The approximated probability mass function of dilation rates is determined by the multi-dilated convolutional layer with shared convolutional weights. d_i is the dilation rate, and α_i is the PMF in Equ. (4).

prefer different searching options. Searching within a functional dimension enables us to determine dilation rates with the expectation of all subsets instead of choosing the option required by one majority subset. Therefore, we propose an expectation-guided iterative (EGI) local search scheme to determine the finer-level dilation rates on top of the global-searched structures.

Suppose that the receptive field of a layer l is D_l . For a dataset, once we get the probability mass distribution of dilation rates around D_l , we can obtain the expected dilation rate with the weighted average of the dilation rates required by all subsets. However, the probability mass of dilation rates for the dataset is inaccessible. Therefore, we utilize a convolutional weight-sharing scheme to enforce the learned importance coefficients of dilation rates to approximate the probability mass. To get the approximated probability mass function of dilation rates, we first evenly sample S dilation rates near the initial dilation rate D_l within the range of $[D_l \pm \Delta D_l]$. The set of available dilation rates within this layer is $T_l = \{d_i | i \in [1, S]\}$, where $d_i = D_l - \Delta D_l + (i-1) \cdot 2\Delta D_l / (S-1)$. ΔD_l is the finer controller of the search space that results in a more dense sampling than that of the global search.

Algorithm 2 Expectation-Guided Iterative Local Search.

Input: Iterations N , initial receptive fields D ;

Initialize model using given D ;

for iter in $[1, N]$ **do**

Construct T_l for each layer based on D and initialize W with the same value;

Train model to get the PMF in Equ. (4);

Obtain new dilation rates through Equ. (6);

Update D ;

end for

return local-searched D .

With the dilation rates set T_l , we propose a multi-dilated layer composed of a shared convolutional weight and multiple branches with different dilation rates, as shown in Fig. 3. Each branch has a unique coefficient to determine the importance of the dilation rate. During the searching process, the coefficients are updated with the gradient backpropagation to reflect the receptive field requirements of the dataset. Existing DARTS schemes [20], [110] have separated operator weights in each branch. In contrast, our convolutional weight-sharing strategy forces the model to

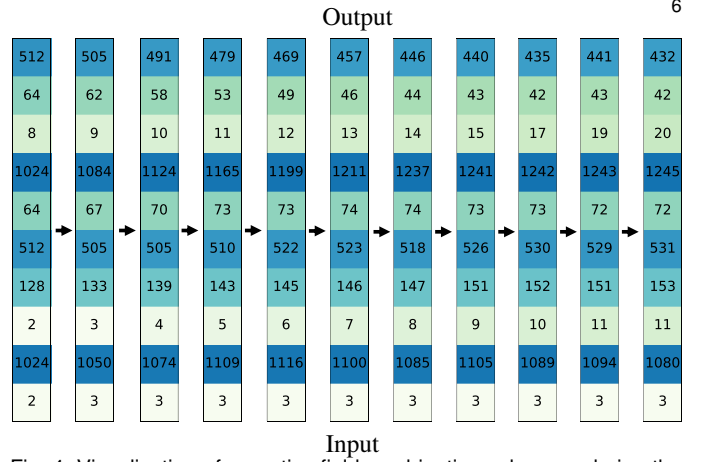


Fig. 4. Visualization of receptive field combinations changes during the EGI local searching process.

learn the approximated probability of receptive fields and ease the model convergence. Specifically, the dilation rates in the multi-dilated convolutional layer are set to T_l . Apart from the shared convolutional θ , the multi-dilated layer contains coefficient $W = \{w_1, w_2, \dots, w_i, i \in [1, S]\}$ to determine the importance of the dilation rates. Both θ and W are learnable parameters and can be trained with gradient backpropagation. For each iteration, each value in W is reinitialized with the same initial value.

W is unbounded, thus cannot be directly used to determine the dilation rates probability. Therefore, we propose a normalization function to get the approximated probability mass function $PMF(d_i)$ of dilation rates through normalizing w_i :

$$PMF(d_i) = \alpha_i = \frac{|w_i|}{\sum_i^S |w_i|}. \quad (4)$$

With the probability mass function, given the input feature x , the output y of the multi-dilated convolutional layer can be written as follows:

$$y = \sum_i^S \alpha_i \Psi(x, d_i, \theta), \quad (5)$$

where $\Psi(x, d_i, \theta)$ is the convolutional operation with the shared convolutional weight θ and dilation rate d_i . α_i is updated with gradient optimization. Once we get the probability mass function, the newly searched dilation rate D'_l is obtained with the expectation:

$$D'_l = \lfloor \sum_{d_i \in T_l} PMF(d_i) \cdot d_i \rfloor. \quad (6)$$

To reduce the computational cost during the local search process, we reduce the number of dilation rates in T_l to 3 by default and apply the iterative search scheme to find the more suitable dilation rate based on the D'_l from the last iteration. The local search process can be summarised as Algorithm (2). Furthermore, Fig. 4 visualizes the dilation rates changes during the local searching process.

Parallel receptive fields for multi-scale enhancement. The local search results in one dilation rate for each convolution. However, we observe that some spatial tasks, e.g., instance segmentation and object detection, require the parallel multi-scale ability to process small and large objects in the senses. Our expectation-guided local search scheme can provide the parallel multi-scale ability with different dilation rates and shared convolutional weights. Therefore, we extend the local-searched structure to the parallel

TABLE 1

Performance of the global and local searching stages of our global-to-local searching method using MS-TCN [12] as the baseline. The global search finds new receptive field combinations that are better than baseline. Local search further refines the global searched structures to achieve better performance.

	BreakFast					50Salads					GTEA				
	F@0.1	F@0.25	F@0.5	Edit	Acc	F@0.1	F@0.25	F@0.5	Edit	Acc	F@0.1	F@0.25	F@0.5	Edit	Acc
MS-TCN [12]	52.6	48.1	37.9	61.7	66.3	76.3	74.0	64.5	67.9	80.7	87.5	85.4	74.6	81.4	79.2
Reproduce	69.1	63.7	50.1	69.9	67.3	78.8	75.3	64.4	71.4	77.8	87.1	83.6	70.4	81.1	75.5
Global	72.2	66.0	51.5	71.0	69.2	79.3	76.5	68.1	71.9	81.2	89.1	87.1	74.4	84.2	78.6
Global+Local	74.9	69.0	55.2	73.3	70.7	80.3	78.0	69.8	73.4	82.2	89.9	87.3	75.8	84.6	78.5

TABLE 2

Details of three temporal action segmentation datasets. #Cls and #Vid are the numbers of classes and videos, respectively. #Frame is the average frame of videos.

	#Cls	#Vid	#Frame	Scene
GTEA [52]	11	28	1115	daily activities
50Salads [112]	17	50	11552	preparing salads
BreakFast [113]	48	1712	2097	cooking breakfast

version by keeping the dilation rates in T_l instead of merging them after the last iteration of searching. The parallel version only has $|T_l|$ extra parameters compared with the single branch version. The parallel structures have significant improvement over the single branch structures in instance segmentation and object detection tasks.

3.4 RF-Next: Next Generation Receptive Field Models

Our global-to-local receptive field searching scheme is suitable for various models that utilize convolutions. Given an initial network structure, we apply the searching scheme to convolutions with kernel sizes larger than one. For easy implementation, we utilize dilation rates to represent the receptive field. The global search aims to find receptive field combinations beyond human knowledge, which is optional as many models have been manually tuned. The local search finds suitable fine-grained receptive fields with a small extra cost, and thus it can be easily applied to human-designed models of various tasks. Enhanced by our receptive field search, these **Next** generation **Receptive Field** models, namely RF-Next models, show advantages on many tasks, *e.g.*, object detection, instance segmentation, semantic segmentation, speech synthesis, and sequence modeling.

4 EXPERIMENTS ON TEMPORAL ACTION SEGMENTATION

Temporal action segmentation requires a relatively large range of receptive fields, which is suitable for verifying the effectiveness of our proposed global-to-local search. Therefore, we do our major experiments on the temporal action segmentation task. This section introduces the implementation details of our proposed global-to-local search scheme and shows the superiority of searched receptive field combinations over the human-designed patterns on the temporal action segmentation task. We also give an analysis of the search scheme and the property of searched structures.

4.1 Implementation Details

Structure Searching and Training. Our proposed method is implemented with the PyTorch [114], and Jittor [115] frameworks. Following existing works [7], [12], features are first extracted from videos using the I3D network [48] and then passed to temporal action segmentation models to get the temporal segmentation. Since our proposed global-to-local search scheme is model-agnostic, the training settings for model evaluation, *i.e.*, training epochs, optimizer, learning rate, batch size, keep the same with the cooperation methods [7], [8], [13]. In the global search stage, we set the total iterations $N = 100$, $k = 2$ in Equ. (1), the initialized population size $M = 50$, and mutation probability $p_m=p_s=0.2$. The T in Equ. (1) is set to 10, indicating the maximum dilation rate of the global search space is 1024. We observe that 5 epochs of training can reflect the structure performance, and therefore models are trained with 5 epochs for evaluation. In the EGI local search stage, ΔD_l and S are set to be $0.1D_l$ and 3, respectively. We train the model for 30 epochs during local search, and each iteration contains 3 epochs.

Datasets. Following [7], [8], [12], [13], we evaluate our proposed method on three popular temporal action segmentation datasets: Breakfast [113], 50Salads [112], and GTEA [52]. The details of the three datasets are summarized in Tab. 2. As far as we know, the Breakfast dataset is the largest public dataset for the temporal action segmentation task, which has a larger number of categories and samples compared with the other two datasets. So we perform our ablations mainly on the Breakfast dataset if not otherwise stated. Following common settings [7], [8], [12], [13], we perform 4-fold cross-validation for the Breakfast and GTEA datasets and 5-fold cross-validation for the 50Salads dataset.

Evaluation Metrics. We follow previous works [7], [8], [12], [13] to use the frame-wise accuracy (Acc), segmental edit score (Edit) [11], and segmental F1 score [116] at the temporal intersection over union with thresholds 0.1, 0.25, 0.5 (F@0.1, F@0.25, F@0.5) as our evaluation metrics.

4.2 Performance Evaluation

Global2Local Search. Our proposed global-to-local search aims to find new combinations of receptive fields better than human designings. We mainly take MS-TCN [12] as our baseline architecture to perform the global-to-local search. When testing the MS-TCN on the Breakfast dataset, we train all models with batch size 8 to save training time. The reproduced results shown in Tab. 1 indicate that a large batch size achieves much better




Fig. 5. Visualization of the global-to-local searched structures of three datasets with the MS-TCN baseline. Each row represents the dilations of one structure, which contains four stages.

TABLE 3

Cooperation with existing temporal action segmentation methods. We perform the whole search pipeline based on MS-TCN [12]. Because of the limited computing resources, we only perform the EGI local search on MS-TCN++ [7] and BCN [13], denoted by †. SSTDA [8] uses MS-TCN [12] as a backbone, so we directly add our searched structure to SSTDA, denoted by ‡.

BreakFast	F@0.1	F@0.25	F@0.5	Edit	Acc
ED-TCN [11]	-	-	-	-	43.3
HTK (64) [58]	-	-	-	-	52.0
TCFPN [69]	-	-	-	-	56.3
GRU [60]	-	-	-	-	60.6
GTRM [19]	57.5	54.0	43.3	58.7	65.0
MS-TCN [12]	52.6	48.1	37.9	61.7	66.3
RF-MS-TCN	74.9	69.0	55.2	73.3	70.7
MS-TCN++ [7]	64.1	58.6	45.9	65.6	67.6
RF†-MS-TCN++	72.4	66.8	53.5	70.2	69.6
BCN [13]	68.7	65.5	55.0	66.2	70.4
RF†-BCN	72.5	69.9	60.2	69.0	72.9
SSTDA [8]	75.0	69.1	55.2	73.7	70.2
RF‡-SSTDA	76.3	69.9	54.6	74.5	70.8

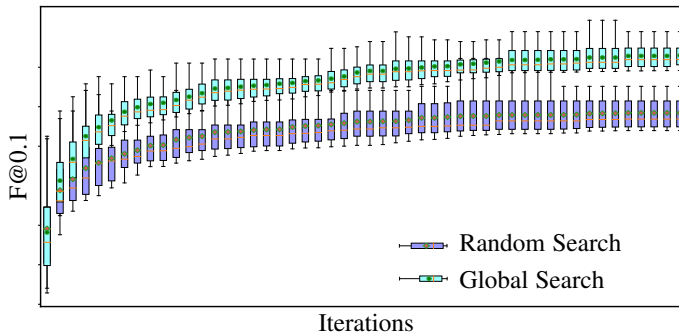


Fig. 6. Performance comparison between our proposed genetic-based search and random search during the global search stage.

performance. Tab. 1 shows that global-to-local searched structures achieve considerable performance improvements than human-designed baselines, *i.e.*, the searched structure surpasses the reproduced baseline with 5.8% in terms of F@0.1. The global-to-local search focuses on receptive field combinations, thus cooperating with existing SOTA temporal action segmentation methods to improve their performance. As shown in Tab. 3, on the large-scale BreakFast dataset, global-to-local search consistently boosts the performance of MS-TCN++ [7], BCN [13], and SSTDA [8]. Also, we give comparisons on two small-scale datasets, 50Salads and GTEA datasets in Tab. 4, proving the effectiveness of our proposed global-to-local search.

Global Search. Global search reduces the computational cost with the sparse search space and our proposed genetic-based searching scheme. Fig. 6 shows the performance change of models during

TABLE 4

Comparison with existing temporal action segmentation methods on the 50Salads and GTEA datasets.

50Salads	F@0.1	F@0.25	F@0.5	Edit	Acc
Spatial CNN [117]	32.3	27.1	18.9	24.8	54.9
Bi-LSTM [64]	62.6	58.3	47.0	55.6	55.7
Dilated TCN [11]	52.2	47.6	37.4	43.1	59.3
ST-CNN [117]	55.9	49.6	37.1	45.9	59.4
TUNet [118]	59.3	55.6	44.8	50.6	60.6
ED-TCN [11]	68.0	63.9	52.6	59.8	64.7
TResNet [15]	69.2	65.0	54.4	60.5	66.0
TricorNet [66]	70.1	67.2	56.6	62.8	67.5
TRN [68]	70.2	65.4	56.3	63.7	66.9
TDRN [68]	72.9	68.5	57.2	66.0	68.1
MS-TCN++ [7]	80.7	78.5	70.1	74.3	83.7
MS-TCN [12]	76.3	74.0	64.5	67.9	80.7
RF-MS-TCN	80.3	78.0	69.8	73.4	82.2
BCN [13]	82.3	81.3	74.0	74.3	84.4
RF-BCN	85.8	83.6	76.5	78.1	85.5

GTEA	F@0.1	F@0.25	F@0.5	Edit	Acc
Spatial CNN [117]	41.8	36.0	25.1	-	54.1
Bi-LSTM [64]	66.5	59.0	43.6	-	55.5
Dilated TCN [11]	58.8	52.2	42.2	-	58.3
ST-CNN [117]	58.7	54.4	41.9	-	60.6
TUNet [118]	67.1	63.7	51.9	60.3	59.9
ED-TCN [11]	72.2	69.3	56.0	-	64.0
TResNet [15]	74.1	69.9	57.6	64.4	65.8
TricorNet [66]	76.0	71.1	59.2	-	64.8
TRN [68]	77.4	71.3	59.1	72.2	67.8
TDRN [68]	79.2	74.4	62.7	74.1	70.1
MS-TCN++ [7]	88.7	87.4	73.5	83.0	78.2
MS-TCN [12]	87.5	85.4	74.6	81.4	79.2
Reproduce	87.1	83.6	70.4	81.1	75.5
RF-MS-TCN	89.9	87.3	75.8	84.6	78.5
BCN [13]	88.5	87.1	77.3	84.4	79.8
RF-BCN	92.1	90.2	79.2	87.2	80.6

the global searching process. Compared with the random search, the genetic-based global search converges faster. The standard deviation of model performance searched by genetic-based search is smaller than the random search, showing the stability of our proposed search scheme. The visualized well-performed global-searched structures in Fig. 5 prove that the global search discovers various structures completely different from human-designed patterns. Tab. 5(b) also shows that the local search heavily relies on global-searched structures to achieve better performance.

Local Search. Based on the global-searched structures, our proposed EGI local search aims to fine-tune the receptive field in a finer search space. We compare the DARTS [20] method and EGI local search based on the global searched structures, as shown in Tab. 5(a). Compared with the DARTS method that only supports several search options, the EGI local search iteratively finds the accurate dilations in a dense space, obtaining structures

TABLE 5
Ablation about the proposed EGI local search.

(a) Performance of EGI local search and DARTS-related methods.					
BreakFast	F@0.1	F@0.25	F@0.5	Edit	Acc
DARTS [20]	73.4	67.3	53.1	72.7	69.5
+Early Stop [119]	73.8	67.6	52.8	72.8	69.3
DARTS+Early Stop	73.8	67.6	52.8	72.8	69.3
+ Fair DARTS [120]	73.3	67.5	52.9	71.9	69.9
Ours	74.9	69.0	55.2	73.3	70.7
(b) Performance of EGI local search initialized by different structures.					
BreakFast	F@0.1	F@0.25	F@0.5	Edit	Acc
random	67.7	61.8	48.3	68.4	67.0
random + local	73.6	67.8	53.7	72.3	69.9
baseline [12]	69.1	63.7	50.1	71.0	69.2
baseline + local	74.1	68.5	55.3	72.3	70.2
global	72.2	66.0	51.8	71.5	69.4
global + local	74.9	69.0	55.2	73.3	70.7
(c) Performance of EGI local search initialized by different structures.					
BreakFast	F@0.1	F@0.25	F@0.5	Edit	Acc
$S = 2$	74.8	68.9	55.0	73.4	70.4
$S = 3$	74.9	69.0	55.2	73.3	70.7
$S = 4$	74.9	68.8	55.1	73.3	70.9
(d) Ablation of possible probability mass functions in EGI local search.					
BreakFast	F@0.1	F@0.25	F@0.5	Edit	Acc
sigmoid	72.7	66.9	52.7	71.8	69.4
softmax	73.2	67.2	52.0	71.6	69.7
Equ. (4)	74.9	69.0	55.2	73.3	70.7

TABLE 6

GPU hours of the global and local search on different temporal action segmentation datasets using the RTX 2080Ti GPU based on the MS-TCN method.

GPU Hours	BreakFast	50Salads	GTEA
Global Search	144h	9h	1h
Local Search	2.2h	0.15h	0.05h
MS-TCN Training	2.0h	0.14h	0.05h

with better performance. We also compare the EGI local search with some variants of DARTS, *i.e.*, the early stop scheme [119] and Fair DARTS [120], and our method outperform them with a clear margin. The early stop scheme improves the performance of DARTS that searched with more epochs, while Fair DARTS [120] has no advantage on the receptive field search task. The early stop scheme [119] stops the searching before the overfitting of DARTS. Fair DARTS [105] solves the problem of performance collapse caused by the unfair advantage in exclusive competition between different operator paths, while the dilation search has no such problem because each path has the same operation.

As shown in Tab. 5(c), EGI local search is insensitive to the number of sampling dilation rates S , as it searches dilation rates with the expectation. Tab. 5(b) shows that the EGI local search can boost the performance of randomly generated, human-designed, and global-searched structures. Still, the performance of the local-searched structures is related to the initial structures, as local search focuses on searching for receptive fields within

TABLE 7
Cross-validation performance (F@0.1) of searched structures among the fold 1 of different datasets. Arch-dataset indicates the structure is searched on which dataset.

	MS-TCN	Arch-50Salads	Arch-GTEA	Arch-BF
50Salads	67.1	75.4	68.8	72.6
GTEA	83.8	82.4	88.9	85.6
BF	69.9	75.1	72.5	76.4

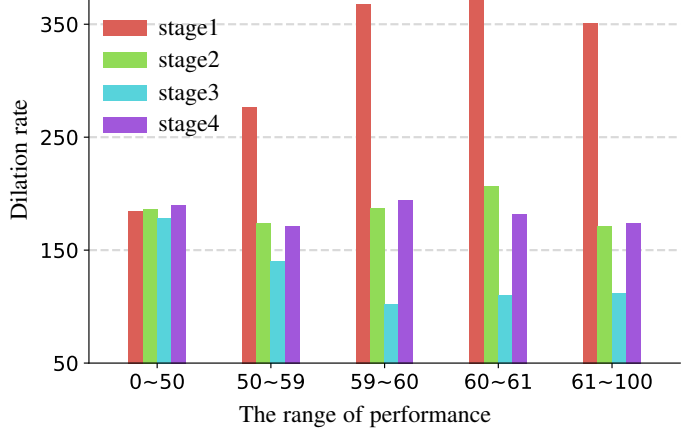


Fig. 7. Visualization of average dilation rates in each stage and the range of performance of global-searched structures.

a finer local search space. We visualize the searching process of the iterative local search in Fig. 4. The dilation rates for each layer gradually converge to a suitable state during the iterative searching process. Tab. 5(d) verifies different ways to get the approximated probability mass function $PMF(d_i)$ from coefficient w . Equ. (4) is superior to the sigmoid and softmax functions because it maintains the probability distribution while the other two functions change the distribution non-linearly.

Searching Cost. We report the cost of our proposed global-to-local search method. When cooperating with MS-TCN, the size of the receptive field combination search space is 1024^{40} . The cost of searching on such a huge space is unaffordable when using existing search methods. Our proposed global-to-local search decomposes the searching process into the global and local search to find the combination in a coarse-to-fine manner. Since the main bottleneck of the search method is the GPU resources, we report the GPU hours of the proposed global-to-local search in Tab. 6. The global search requires more computational cost to find multiple new well-performed structures with different patterns than human-designed structures. The local search needs a small training cost to fine-tune the global-searched/human-designed structures in the dense but local search space.

4.3 Observations

In this section, we try to exploit the common knowledge contained in the global-to-local searched structures.

Connections between Receptive Fields and Data. We want to know if receptive field combinations vary among data. Therefore, we evaluate the generalization ability of the searched structures on the subsets of the same dataset and different datasets, respectively. Within the BreakFast dataset, we perform the global-to-local

TABLE 8

Performance of local search on object detection with COCO [121] dataset using Faster-RCNN as the baseline method. Local-P means the local-searched structure with parallel receptive fields as described in Sec. 3.3. -R50 and -R101 denote using ResNet-50 and ResNet-101 as backbones, respectively. S indicates using S branches in local search as shown in Equ. (4).

	P	validation set						test set					
		mAP	mAP ₅₀	mAP ₇₅	mAP _s	mAP _m	mAP _l	mAP	mAP ₅₀	mAP ₇₅	mAP _s	mAP _m	mAP _l
Faster-RCNN-R50 [2]		37.4	58.3	40.6	22.0	41.1	48.1	37.8	59.0	41.0	22.1	40.8	46.4
+RF ($S = 3$)		39.3	60.6	42.9	23.5	43.0	51.0	39.2	60.9	42.6	22.6	41.8	48.9
+RF ($S = 3$)	✓	40.2	61.7	43.8	23.5	43.9	52.2	40.4	62.1	44.0	23.6	43.0	50.5
+RF ($S = 2$)		39.1	60.5	42.3	23.2	42.8	50.3	39.1	60.8	42.3	22.7	41.7	48.7
+RF ($S = 2$)	✓	40.0	61.4	43.8	23.9	43.6	51.9	40.3	62.1	43.9	23.7	43.0	50.4
Faster-RCNN-R101 [2]		39.4	60.1	43.1	22.4	43.7	51.1	39.7	60.7	43.2	22.5	42.9	49.9
+RF ($S = 3$)		41.1	62.4	44.7	24.5	45.1	53.9	41.2	62.8	44.9	23.6	44.1	52.0
+RF ($S = 3$)	✓	42.0	63.2	45.7	25.0	45.9	55.4	42.1	63.8	45.8	24.3	45.1	53.3




Fig. 8. Visualization of the local-searched structures of Faster-RCNN [2] for object detection (OB) and Mask-RCNN [4] for instance segmentation (IN). S2-S4 denotes stage 2 to stage 4 of the ResNet backbone. FPN, RPN, and MASK mean the feature pyramid network, region proposal network, and the mask segmentation head in Faster-RCNN and Mask-RCNN.

TABLE 9

Cross-validation performance ($F@0.1$) of searched structures among different folds of the Breakfast dataset. Arch-n means the structure is searched on fold n.

BreakFast	Arch-1	Arch-2	Arch-3	Arch-4
fold1	76.4	76.3	76.2	75.7
fold2	74.1	75.3	75.1	74.6
fold3	76.1	76.6	76.1	75.4
fold4	71.7	72.1	72.0	71.8

search on one fold and then evaluate the searched structures on other folds. Tab. 9 shows almost no obvious performance gap on different folds, indicating that receptive field combinations almost have no difference within a dataset. However, when searching and evaluating structures across different datasets, different structures searched on different datasets have a large performance gap, as shown in Tab. 7. We can conclude that different data distributions will result in different receptive field combinations. We visualize the structures searched from different datasets in Fig. 5. The searched structures are based on both global and local search. Since the global search introduces the randomness of each structure, we cannot fairly compare these structures for different datasets. Still, we give a rough explanation based on the searched receptive field of each structure. The structures searched on the Breakfast and 50Salads datasets tend to have larger receptive fields, while the structure searched on the GTEA dataset has smaller receptive fields. The number of video frames shown in Table 2 positively correlates with receptive fields. We assume that the average video length of a dataset might influence the receptive fields of the structure. A video with more frames normally requires larger receptive fields to capture long-term relations. The structures searched on Breakfast and 50Salads datasets have similar average receptive fields, but the average

video length of 50Salads is longer than Breakfast. We assume that understanding the content of cooking breakfast in the Breakfast dataset requires more long-range information than the preparing salads content in the 50Salads dataset. Analyzing different video contents might need features from different ranges of receptive fields. Our work mainly focuses on searching receptive fields, but fully explaining why the receptive field combination of a certain task on a certain dataset looks like the searched one is still an open question.

Receptive Fields for Different Stages. Our global-to-local search is based on MS-TCN. MS-TCN contains four stages, and all stages share the same receptive field combination in human design. The visualized searched structures shown in Fig. 5 demonstrate that different stages have different receptive field combinations, which conflicts with human design. We further count the average receptive fields of each stage among all individuals. The range of performance and the average dilation rates of each stage are shown in Fig. 7. The average dilation rate in the first stage of MS-TCN tends to be large on high-performance structures. In contrast, the average dilation rate in the third stage of MS-TCN is relatively small on high-performance structures. We assume that the first stage of MS-TCN requires large receptive fields to get the long-term context for coarse prediction, while the following stages need small receptive fields to refine the results locally.

5 RF-NEXT FOR MULTIPLE NETWORKS AND TASKS

This section shows that RF-Next models can be applied to multiple networks and tasks. We use prefixes RF to denote RF-Next models, and P means the parallel receptive field version.




Fig. 9. Visualization and probability of each receptive field in the parallel searched structure of ResNet-50 based Faster-RCNN (a) and Mask-RCNN (b) when $S=3$ during searching. S2-S4 denotes stage 2 to stage 4 of the ResNet backbone. FPN, RPN, and MASK mean the feature pyramid network, region proposal network, and the mask segmentation head in Faster-RCNN and Mask-RCNN.

TABLE 10

Performance of local search on instance segmentation with COCO [121] dataset using Mask-RCNN as the baseline method. Local-P means the local-searched structure with parallel receptive fields as described in Sec. 3.3. -R50 and -R101 denote using ResNet-50 and ResNet-101 as backbones, respectively. R50-NonLocal means adding the non-local block [40] to each residual block in stage 4 of the ResNet-50 backbone. S indicates using S branches in local search as shown in Equ. (4).

	P	mAP ^{bb}	mAP ₅₀ ^{bb}	mAP ₇₅ ^{bb}	mAP _s ^{bb}	mAP _m ^{bb}	mAP _l ^{bb}	mAP ^{mk}	mAP ₅₀ ^{mk}	mAP ₇₅ ^{mk}	mAP _s ^{mk}	mAP _m ^{mk}	mAP _l ^{mk}
validation set:													
Mask-RCNN-R50 [4]		38.2	59.0	41.7	22.4	41.6	49.2	34.7	55.9	37.0	16.5	37.4	50.1
+RF ($S = 3$)		39.7	60.7	43.4	23.7	43.3	51.3	36.0	57.8	38.4	17.0	39.1	52.2
+RF ($S = 3$)	✓	41.0	62.3	44.9	24.2	44.7	53.2	37.1	59.1	39.5	18.4	39.9	53.6
+RF ($S = 2$)		39.6	60.8	43.2	23.3	43.1	51.4	35.9	57.6	38.3	16.4	38.9	52.2
+RF ($S = 2$)	✓	40.7	62.1	44.5	24.4	44.3	52.9	36.8	58.9	39.3	18.0	39.7	53.6
Mask-RCNN-R101 [4]		40.0	60.5	44.0	22.6	44.0	52.6	36.1	57.5	38.6	18.8	39.7	49.5
+RF ($S = 3$)		41.7	62.8	45.7	24.8	45.4	55.4	37.4	59.6	40.0	18.5	40.6	55.1
+RF ($S = 3$)	✓	42.7	64.0	46.7	25.6	46.8	55.8	38.4	60.9	41.3	18.5	42.0	55.7
R50-NonLocal [40]		39.4	60.8	43.4	23.2	43.2	50.5	35.5	57.1	37.9	17.2	38.5	51.3
+RF ($S = 2$)		40.2	61.5	43.9	23.9	43.9	52.2	36.4	58.5	38.7	17.6	39.3	52.3
+RF ($S = 2$)	✓	40.8	62.2	44.6	24.4	44.5	53.0	36.9	59.0	39.6	18.0	39.8	53.8
test set:													
Mask-RCNN-R50 [4]		38.5	59.5	41.8	22.3	41.6	47.4	34.9	56.4	37.2	18.9	37.5	44.6
+RF ($S = 3$)		40.0	61.4	43.6	23.2	42.6	49.8	36.2	58.5	38.5	19.9	38.6	46.8
+RF ($S = 3$)	✓	41.0	62.5	45.0	24.0	43.8	51.3	37.1	59.5	39.9	20.5	39.7	48.2
+RF ($S = 2$)		39.8	61.2	43.4	23.2	42.3	49.4	36.1	58.2	38.6	19.8	38.4	46.5
+RF ($S = 2$)	✓	41.0	62.5	44.7	24.1	43.7	51.4	37.1	59.5	39.7	20.7	39.6	48.3
Mask-RCNN-R101 [4]		40.4	61.2	44.1	23.1	43.5	50.8	36.5	58.3	38.9	19.6	39.2	47.8
+RF ($S = 3$)		42.0	63.4	45.9	24.2	45.0	53.1	37.8	60.5	40.4	20.5	40.4	49.7
+RF ($S = 3$)	✓	42.8	64.1	46.9	24.7	46.1	54.2	38.5	61.3	41.3	21.0	41.4	50.7
R50-NonLocal [40]		39.8	61.4	43.3	23.2	42.8	49.1	36.0	58.2	38.3	19.7	38.6	46.3
+RF ($S = 2$)		40.4	61.9	44.0	23.2	42.8	50.4	36.4	58.7	38.9	19.7	38.6	47.2
+RF ($S = 2$)	✓	41.2	62.9	45.0	23.7	43.7	51.7	37.3	59.8	39.9	20.3	39.7	48.3

5.1 Spatial tasks

We apply our proposed searching scheme to find proper receptive fields for spatial tasks, *e.g.*, object detection, instance segmentation, and semantic segmentation. We observe that proper receptive fields significantly improve the performance of these tasks.

5.1.1 Object Detection

Object detection aims to assign bounding boxes and categories to objects of various sizes. We utilize the widely used Faster-RCNN [2] method with a dilation rate of 1 for all convolutions. Faster-RCNN applies the feature pyramid network to aggregate features with multiple scales to handle objects of various sizes. However, the receptive fields for convolutions are ignored. Therefore, we search for dilation rates of convolutions with kernel sizes larger than one in the Faster-RCNN. Due to the extremely large

training cost and the small initial dilation rates, we only use the highly efficient local search scheme.

Training and Searching. We employ the ResNet-50 [15] and ResNet-101 [15] that are pre-trained on the ImageNet [122] dataset as backbones. We verify the effectiveness of local search on the COCO dataset [121] and report the mean average precision (mAP) to evaluate the trained model. Following the official training scheme [2], the images are resized to 1333×800 with a randomly horizontal flip, and the model is trained for 12 epochs, with a batch size of 16 on 4 GPUs. During the local search, we train the model for 12 epochs and update the structure in each epoch for the first 10 epochs. The ΔD_l is set to be $0.5D_l$. Since the weights of the first stage of ResNet in Faster-RCNN are frozen during training, we skip this stage during searching.

Performance and Observation. As shown in Tab. 8, RF-Next model improves the test mAP of Faster-RCNN with ResNet-50 by 1.4%. For the RF-ResNet-101 model, the test mAP is also improved by 1.5%. Theoretically, a network with a larger depth has a larger range of receptive fields. Still, effective receptive field settings have a similar performance gain on both shallow and deep models. We visualize the searched dilation rates of RF-ResNet-50/101 based Faster-RCNN as shown in Fig. 8. The shallow layers require relatively small dilation rates, while some deep layers have large dilation rates. Interestingly, the ResNet-101 based model requires larger dilation rates in stage 4 of the network than ResNet-50 based model.

When utilizing the RF-Next with parallel receptive fields, the performance gain of ResNet-50 and ResNet-101 based models are 2.6% and 2.4% in test mAP, indicating that object detection task needs parallel multi-scale ability. The visualization and probability of each receptive field in the parallel RF-Next are shown in Fig. 9. By default, we utilize the number of sampling dilation rates $S = 3$ in Equ. (4). And we also explore using $S = 2$ for local search, as shown in Tab. 8. Using two/three branches achieve similar performance for parallel RF-Next, showing that using two branches for each layer provides sufficient multi-scale ability. The result is also consistent with the observation in Tab. 5(c) that the proposed expectation-guided search is insensitive to the number of sampling dilation rates.

We analyze the performance gain for objects of different sizes in Tab. 8. For ResNet-50 and ResNet-101 based models, the test mAP improvement for small, medium, and large objects are (1.5%, 2.2%, 4.1%) and (1.8%, 2.2%, 3.4%), respectively. The performance gain gradually increases with the increase of object sizes, showing that the default receptive field settings of Faster-RCNN are not large enough to capture large objects.

5.1.2 Instance Segmentation

Instance segmentation outputs the instances segmentation masks and categories, which is similar to the object detection task. To compare the receptive field requirements of object detection and instance segmentation, we use the widely used Mask-RCNN [4] that extends the Faster-RCNN with a mask segmentation branch. Like the object detection, we apply the local search on convolutions whose kernel is larger than one.

Training and Searching. The ResNet-50 [15] and ResNet-101 [15] with ImageNet [122] pre-training are used as backbones. We use the COCO dataset [121] for object detection and report the mean average precision of bounding box (mAP^{bb}) and instance segmentation (mAP^{mk}) to evaluate the trained model. To fairly compare the searched structure on instance segmentation and object detection, the training scheme of Mask-RCNN is aligned with Faster-RCNN in both searching and re-training stages.

Performance and Observation. We give the performance comparison of searched structure and baseline in Tab. 10. Using the RF-Next with a single branch brings 1.3%/1.3% gain on test mAP^{mk} and 1.5%/1.6% gain on test mAP^{bb} of ResNet-50/101 based models. And the performance gains are further enlarged by the parallel RF-Next, *i.e.*, 2.2%/2.0% gain on test mAP^{mk} and 2.5%/2.4% gain on test mAP^{bb} . We give the visualized dilation rates comparison between Faster-RCNN and Mask-RCNN as shown in Fig. 8. The dilation rates of these two tasks are very

similar because of the large similarity between them. As shown in Fig. 9, the probability of dilation rates in the mask segmentation head of Mask-RCNN shows that the middle two layers require diverse receptive fields while the first and last layer requires a small receptive field.

5.1.3 Semantic Segmentation

Semantic segmentation task requires assigning each pixel of images with category labels. Receptive fields are vital for the dense pixel-level prediction of semantic segmentation. Deeplab series [5], [32] utilize convolution with dilation rates larger than one to enlarge the receptive fields, which becomes the default option for semantic segmentation networks [123], [124]. However, possible better receptive fields than the human-designed ones have not been explored. We apply the Deeplab V3 [32] network as the baseline method and conduct local search to find more effective receptive fields.

Training and Searching. We conduct experiments on both the PASCAL VOC [125] dataset and the ADE20k dataset [126] using the ImageNet pretrained ResNet-50 [15] based Deeplab V3 as the baseline method. Following official training settings [32], models are trained for 20k and 80k iterations for PASCAL VOC and ADE20k, respectively, with a batch size of 16 on 4 GPUs. For both datasets, the images are randomly scaled at a ratio between 0.5 and 2.0 and randomly cropped to 512×512 . An auxiliary loss is applied at the output of the third stage of the ResNet backbone to ease convergence. To avoid the influence of auxiliary loss, we apply the local search to stage 4 and the decoder of the network. During the local search, we update the structure every 2k and 8k iterations for PASCAL VOC and ADE20k, respectively, and conduct 8 iterations of the local search in total. The ΔD_l is set to be $0.15D_l$ for both datasets.

Performance and Observation. We utilize the mean intersection over union (mIoU) and the mean accuracy (mAcc) to evaluate the trained models. As shown in Tab. 11, the RF-Next brings the 1.6% and 0.8% gain in mIoU for PASCAL VOC and ADE20k datasets, respectively. The searched multi-branch structures achieve similar performance to the single branch competitors. We assume the atrous spatial pyramid pooling structure in Deeplab V3 already enhances the parallel multi-scale ability of the network. The visualization of searched receptive fields in Fig. 10 indicates that larger receptive fields are required in stage 4 of the network.

As described above, we skip the local search in the first three stages of the network to avoid the side effect of auxiliary loss. We show in Tab. 11 that searching all convolutions achieves 76.3% mIoU, similar to the human-designed baseline. We observe that searching receptive fields for all convolutions results in a much smaller auxiliary loss value than human-designed baseline (0.102 VS. 0.191) and searching after stage 3 (0.102 VS. 0.189). Also, compared to searching after stage 3, searching in the early stages achieves the performance gain of 15.4% in mIoU for output in stage 3 of the network. Since the local search relies on the gradient backpropagation to find receptive fields, adding auxiliary loss makes the local search find better receptive fields for output in stage 3 instead of the final output.

5.2 Sequential tasks

Except for the temporal action segmentation task, we also conduct receptive field search on other sequential tasks, *e.g.*, speech

TABLE 11

Performance of local search on semantic segmentation with PASCAL VOC [125] and ADE20k datasets [126]. Local-P means the local-searched structure with parallel receptive fields as described in Sec. 3.3. †means apply local search to all convolutions of the network. Local-S3 indicates the output in stage 3 of the network, where an auxiliary loss is added to ease convergence [32].

P	VOC [125]		ADE20K [126]	
	mIoU	mAcc	mIoU	mAcc
DeepLabV3 [32]	76.2	85.7	42.4	53.6
+RF	77.8	87.4	43.2	54.1
+RF	✓ 77.9	87.6	43.0	53.7
+RF†	76.3	85.8	-	-
+RF-S3	51.7	64.7	-	-
+RF-S3†	67.7	79.8	-	-

Baseline	2	4	4	12	24	36
PASCAL VOC	2	8	6	11	23	43
ADE20K	2	8	6	12	24	40
	S4			Decoder		

Fig. 10. Visualization of the local searched receptive fields of stage 4 and decoder in Deeplab V3 on the semantic segmentation task.

synthesis, P-MNIST digit classification, and polyphonic music modeling. We verify the effectiveness of global-to-local search on the polyphonic music modeling and P-MNIST digit classification task. Due to the high training cost of speech synthesis, we apply the local search on top of the human-designed structure.

5.2.1 Speech Synthesis

As described in Sec. 2.2.2, we focus on the procedure of transferring acoustic features to speech waveform in speech synthesis. We use WaveGlow [127] as the baseline method that combines the advantage of Glow [128] and WaveNet [28]. WaveGlow network has 12 coupling layers, where each contains 8 layers of dilated convolutions with human-designed gradually expanded dilation rates. To save computational cost, we utilize the local search to find more effective dilation rates of these layers based on human-designed structures.

Training and Searching. We conduct experiments on the widely used LJ speech [129] dataset, including 13,100 audio clips with a total length of about 24 hours. Each sample is randomly cropped to 16,000 for training, and the sampling rate is 22,050Hz. The mel-spectrograms [79], [130], generated through a short-time Fourier transform, are fed to the network for speech synthesis. The network is optimized by Adam optimizer with a learning rate of $1e-4$ for 100 epochs. We train the model with 4 GPUs using the batch sizes of 20 and 48 during searching and training, respectively. During the local searching, the ΔD_l and S are set to be $0.6D_l$ and 3, respectively. We train the model for 60 epochs during the local search and update the structure every 3 epochs.

Performance and Observation. We use three metrics, *i.e.*, mel-cepstral distortion (MCD) [131], perceptual evaluation of speech quality (PESQ) [132] and log-likelihood ratio (LLR) [133], to evaluate the speech synthesis quality, as shown in Tab. 12. MCD measures the difference between two sequences of mel-cepstra of

(a)	1	1	2	2	4	4	8	8
(b)	1	2	16	8	8	8	8	2
(c)	1	2	14	7	8	7	8	2

Fig. 11. Visualization of the baseline structure (a), the global searched structure (b) and the global-to-local searched structure (c) of TCN on polyphonic music modeling task.

(a)	1	1	2	2	4	4	8	8	16	16	32	32	64	64	128	128
(b)	4	1	32	64	1	256	16	128	128	4	64	256	128	64	32	64
(c)	4	1	35	64	1	237	14	141	128	4	71	294	141	70	29	70

Fig. 12. Visualization of the baseline structure (a), the global searched structure (b) and the global-to-local searched structure (c) of TCN on P-MNIST Classification task.

speech, and a small MCD indicates the synthesized and natural speeches are close. Similarly, LLR measures the difference between two speeches. The structure with searched receptive field combinations achieves better performance than human-designed receptive fields in MCD and PESQ. PESQ assesses the voice quality, and a higher value between the synthesized and natural speech means better the synthesized speech quality. We calculated the PESQ under the narrowband of 8,000Hz. The speech synthesis results of the local-searched structure also have a better PESQ score, indicating that more proper receptive fields benefit the speech synthesis quality.

We visualize the local searched and baseline receptive fields (dilation rates) of WaveGlow [127] in Fig. 13. We observe that the maximum dilation rate of the human-designed structure is much larger than the searched structure, indicating that too large receptive fields may not be necessary for this task. Unlike the human-designed structure with the same receptive field combination for each coupling layer, the searched structure has small receptive fields on the shallow layers and larger receptive fields on the deeper layers. We assume that the speech synthesis task requires local features in the shallow layers, and the deeper layers are responsible for modeling the long-term dependencies.

5.2.2 Sequence Modeling using TCN

Bai *et al.* [134] have verified the performance of TCN on several sequence modeling tasks. We further show the effectiveness of RF-Next models on two sequential tasks, *i.e.*, polyphonic music modeling and P-MNIST digit classification.

P-MNIST Classification. P-MNIST classification aims to classify the handwritten digit images in which the order of the pixels is disrupted. P-MNIST dataset [135], [136] randomly permutes images in the MNIST dataset [137] to 784 length sequences, which is widely used for long-term relation modeling [134], [138]–[141]. A TCN with 8 layers is used for P-MNIST classification [134], where each layer has a convolution with a kernel size of 7 and a channel number of 25. We apply the global-to-local search on the TCN to find more effective receptive fields. Following the settings in [134], the final model is trained for 20 epochs with an Adam optimizer. The initial learning rate is $2e-3$ and is multiplied by 0.1 at the 10-th epoch. In P-MNIST, the order of the pixels is randomly permuted, and we fix the permutation order for all experiments. During the global search, we set the




Fig. 13. Visualization of the local searched dilation rates of WaveGlow [127] for the speech synthesis task.

TABLE 12

Speech synthesis performance of WaveGlow [127] based on local-searched structure on the LJ speech dataset [129].

	MCD↓ [131]	LLR↓ [133]	PESQ↑ [132]
WaveGlow [127]	5.79	1.29	1.52
RF-WaveGlow	5.59	0.71	1.84

TABLE 13

Performance of global-to-local search on TCN [134]. We evaluate the performance of the polyphonic music modeling and P-MNIST digit classification task.

Task	baseline	global	global+local
Permuted MNIST (accuracy↑)	97.2	97.6	97.8
Music Nottingham (NLL↓)	2.97	2.73	2.69

iterations $N = 50$ and the initial population size $M = 25$. The model is trained with 11 epochs for each sample. During the local search, the ΔD_l is set to $0.1D_l$, and the structure is trained for 15 epochs and updated every 3 epochs. The classification accuracy is used as the evaluation metric. As shown in Tab. 13, the global search improves the accuracy from 97.2% to 97.6%, and the local search further improves the performance to 97.8%. The visualized structure in Fig. 12 shows the searched receptive fields are very different from the human-designed patterns.

Polyphonic Music Modeling. Polyphonic music modeling aims to predict the subsequent musical notes based on the history of the played notes. The polyphonic music modeling is conducted on the widely used Nottingham [142]–[144] dataset composed of 1200 British and American folk tunes. For polyphonic music modeling, we utilize a TCN with 4 layers, where each layer has two convolutions with a kernel size of 5 and a channel number of 150. Following [134], the model is trained for 100 epochs with an Adam optimizer. The initial learning rate is $1e-3$, which is multiplied by 0.1 every 30 epochs. The dropout with the rate of 0.25 and the gradient clipping with the maximum norm of 0.2 is applied. We set the iterations $N = 50$ and the initial population size $M = 25$ for the global search, and each sample is trained with 30 epochs. For local search, the ΔD_l is set to $0.15D_l$, and the model is trained for 60 epochs and updated in every 10 epochs. We evaluate models using the negative-log-likelihood (NLL), as shown in Tab. 13. The NLL is improved from 2.97 to 2.73 by the global search, and the local search improves the performance to 2.69. The global and local-searched structures are shown in Fig. 11.

5.3 Receptive field search on Modern Networks

We apply the receptive field searching method on multiple networks, *e.g.*, SOTA attention/convolution based networks, multi-scale networks, and searched networks.

TABLE 14

Receptive field search improves SOTA attention/convolution models for object detection and instance segmentation tasks on the COCO val dataset. Following the official implementation of PVT [145], [146] and ConvNeXt [147], PVTv2-B0 and ConvNeXt-T adopt the Mask RCNN detector and Cascade Mask RCNN detector, respectively.

Object det.	P	mAP	mAP ₅₀	mAP ₇₅	mAP _s	mAP _m	mAP _l
PVTv2-B0		38.2	60.5	40.7	22.9	40.9	49.6
RF-PVT		38.8	60.9	41.8	23.6	41.2	50.8
RF-PVT	✓	39.1	60.8	42.7	23.3	41.8	51.4
ConvNeXt-T		50.4	69.1	54.8	33.9	54.5	65.1
RF-ConvNeXt		50.6	69.2	54.8	34.1	54.0	65.5
RF-ConvNeXt	✓	50.9	69.5	55.5	34.3	54.6	65.8
Instance seg.		mAP	mAP ₅₀	mAP ₇₅	mAP _s	mAP _m	mAP _l
PVTv2-B0		36.2	57.8	38.6	18.0	38.4	51.9
RF-PVT		36.8	58.4	39.5	18.7	39.0	52.7
RF-PVT	✓	37.1	58.5	40.0	17.8	39.3	53.7
ConvNeXt-T		43.7	66.5	47.3	24.2	47.1	62.1
RF-ConvNeXt		44.0	66.8	47.5	24.8	47.0	62.1
RF-ConvNeXt	✓	44.3	67.3	47.8	24.7	47.4	62.6

TABLE 15

Performance of receptive field search on semantic segmentation using PVTv2-B0 backbone with the Semantic FPN method [146].

	P	Pascal VOC [125] mIoU	mAcc	ADE20K [126] mIoU	mAcc
PVTv2-B0		73.7	85.0	37.5	48.3
RF-PVT		74.4	86.0	38.0	48.6
RF-PVT	✓	74.4	85.9	37.8	48.7

Receptive Field Search Improves SOTA Models. We show recent SOTA models, *e.g.*, PVT [145], [146] and ConvNeXt [147], still benefit from our receptive field searching method. PVTv2 [145], [146] is a self-attention based pyramid vision transformer, which uses both global self-attentions and depth-wise convolutions. We apply the receptive field search to convolutions of PVTv2 for object detection, instance segmentation, and semantic segmentation tasks, achieving stable improvements over the strong PVTv2-B0 baseline. As shown in Tab. 14, the single-branch RF-PVTv2 has the gain of 0.6% box mAP and 0.6% mask mAP for object detection and instance segmentation tasks. The parallel RF-PVTv2 further improves 0.3% box mAP and 0.3% mask mAP for these two tasks. Tab. 15 shows the searched single-branch structure improves the baselines with 0.7% and 0.5% on Pascal VOC and ADE20K datasets. ConvNeXt [147] is a SOTA convolutional model that outperforms many SOTA attention-based models. Even though ConvNeXt manually tunes the kernel size of convolutions to support a larger range of receptive fields, the receptive field search still further improves the performance on object detection and instance segmentation. As shown in Tab. 14,

TABLE 16

Receptive field search improves hand-crafted multiple receptive fields models, *i.e.*, Res2Net [16] and HRNet [148], [149], for object detection and instance segmentation tasks on the COCO val dataset. The Cascade Mask RCNN method is used as the detector.

Object det.	P	mAP	mAP ₅₀	mAP ₇₅	mAP _s	mAP _m	mAP _l
Res2Net-101		46.3	64.4	50.5	27.2	50.3	60.5
RF-Res2Net		46.9	65.8	51.2	28.4	50.7	62.1
RF-Res2Net	✓	47.9	66.6	52.2	29.7	51.9	62.8
HRNetV2p-W18		41.6	58.7	45.4	23.5	44.7	54.9
RF-HRNet		42.9	60.8	46.7	25.9	46.2	54.8
RF-HRNet	✓	43.7	61.9	47.7	26.5	47.3	56.7
Instance seg.		mAP	mAP ₅₀	mAP ₇₅	mAP _s	mAP _m	mAP _l
Res2Net-101		40.0	61.7	43.3	22.2	43.8	54.1
RF-Res2Net		40.7	63.2	43.9	20.4	44.0	59.0
RF-Res2Net	✓	41.5	64.0	44.9	21.3	44.6	59.5
HRNetV2p-W18		36.4	56.3	39.3	19.1	39.1	49.5
RF-HRNet		37.6	58.3	40.4	19.0	40.2	53.9
RF-HRNet	✓	38.1	59.3	41.0	19.4	40.7	55.3

the parallel version of searched structure has considerable gains of 0.5% box mAP and 0.6% mask mAP over the strong ConvNeXt-T model. The performance gains over these two strong SOTA models prove the effectiveness of our receptive field searching method.

Receptive Field Search Improves Multi-scale Models. We show the advantage of receptive field searching over popular hand-crafted multi-scale models with multiple receptive fields, *e.g.*, HRNet [148], [149] and Res2Net [16]. HRNet parallel processes features of multiple resolutions to form the multi-scale representation. Res2Net constructs hierarchical residual-like connections within a block to enable fine-grained multiple receptive fields. Despite their good hand-crafted multi-scale ability, we show in Tab. 16 that searched receptive fields constantly improve their performance on object detection and instance segmentation tasks. For HRNet, the single/multiple-branch RF-HRNet improves 1.3%/2.1% box mAP on object detection and 1.2%/1.7% mask mAP on instance segmentation. The single/multiple-branch RF-Res2Net also improves the Res2Net with 0.6%/1.6% box mAP on object detection and 0.7%/1.5% mask mAP on instance segmentation. Therefore, our receptive field searching method can further improve the hand-crafted multi-scale models with better receptive field combinations.

Comparison with Attention Mechanisms. As discussed in the related work, attention mechanisms theoretically can form arbitrary receptive fields [39]–[41]. However, we are not aware of the actual receptive field representation ability of attention mechanisms. Therefore, we propose to compare our receptive field search method with the non-local module [40] on the instance segmentation task. Following the implementation in [40], we insert the non-local blocks to each residual block in stage 4 of the ResNet50 backbone. As shown in Tab. 10, on the COCO testing set, the non-local based Mask-RCNN improves the Mask-RCNN baseline with 1.1 mask mAP and 1.3 box mAP. The Mask-RCNN with searched parallel receptive fields outperforms non-local based Mask-RCNN by 1.1 mask mAP 1.2 box mAP, showing that the searched receptive fields provide better representation than non-

TABLE 17

Performance of local search on semantic segmentation with Auto-deeplab [23] and Cityscapes datasets [150]. Local-P means the local-searched structure with parallel receptive fields as described in Sec. 3.3.

	P	mIoU	mAcc
Auto-deeplab [23]		76.0	83.7
RF-Auto-deeplab		76.3	84.2
RF-Auto-deeplab	✓	76.7	84.3

local modules. Therefore, though the non-local module improves the model performance, it cannot provide effective receptive fields as strong as our searched receptive fields. Then, we apply the local search to the non-local based network, and the hyper-parameters for searching are kept consistent with the standard ResNet. We observe a further performance gain with both searched single-branch receptive fields and parallel-branch receptive fields. The single-branch version has the gain of 0.4 mask mAP and 0.6 box mAP, and the parallel-branch version has the gain of 1.3 mask mAP and 1.4 box mAP. The receptive field searching scheme can further improve the performance of the non-local based model, indicating that the non-local module may not be able to cover all effective receptive fields even with its dense connections among pixels.

Searching Receptive Fields over the Searched Network. Auto-deeplab [23] searches the feature resolution for different stages of the semantic segmentation network. To verify if it is possible to adjust the receptive fields of Auto-deeplab further, we conduct the local search on top of the Auto-deeplab. We follow the implementation of Auto-deeplab [23] to search on the Cityscapes [150] dataset. As shown in Tab. 17, both single/multiple-branch RF-Auto-deeplab brings performance gain over Auto-deeplab in mIoU. Therefore, our receptive field searching scheme can further benefit the semantic segmentation model with searched feature resolutions.

6 CONCLUSION

We propose a global-to-local search scheme to search for effective receptive field combinations at a coarse-to-fine scheme. The global search discovers effective receptive field combinations with better performance than hand designings but completely different patterns. The expectation-guided iterative local search scheme enables searching fine-grained receptive field combinations in the dense search space. RF-Next models, enhanced with receptive field search scheme, can be plugged into multiple tasks, *e.g.*, action segmentation, sequence modeling [134], [151], segmentation [4], [5], [152]–[154], object detection [155], [156] methods to boost the performance further.

Acknowledgement

This research was supported by the Major Project for New Generation of AI under Grant No. 2018AAA0100400, NSFC (61922046), and S&T innovation project from Chinese Ministry of Education.

REFERENCES

- [1] S. Gao, Q. Han, Z.-Y. Li, P. Peng, L. Wang, and M.-M. Cheng, "Global2local: Efficient structure search for video action segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. **1, 2**
- [2] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Adv. Neural Inform. Process. Syst.*, 2015. **1, 3, 10, 11**
- [3] M.-M. Cheng, S. Gao, A. Borji, Y.-Q. Tan, Z. Lin, and M. Wang, "A highly efficient model to study the semantics of salient object detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. **1**
- [4] K. He, G. Gkioxari, P. Dollar, and R. Girshick, "Mask r-cnn," *Int. Conf. Comput. Vis.*, Oct 2017. **1, 3, 10, 11, 12, 15**
- [5] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, 2018. **1, 2, 12, 15**
- [6] S. Gao, Z.-Y. Li, M.-H. Yang, M.-M. Cheng, J. Han, and P. Torr, "Large-scale unsupervised semantic segmentation," *arXiv preprint arXiv:2106.03149*, 2021. **1**
- [7] S.-J. Li, Y. AbuFarha, Y. Liu, M.-M. Cheng, and J. Gall, "Ms-tcn++: Multi-stage temporal convolutional network for action segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, pp. 1–1, 2020. **1, 2, 3, 7, 8**
- [8] M.-H. Chen, B. Li, Y. Bao, G. AlRegib, and Z. Kira, "Action segmentation with joint self-supervised temporal domain adaptation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 9454–9463. **1, 7, 8**
- [9] W. Ping, K. Peng, K. Zhao, and Z. Song, "Waveflow: A compact flow-based model for raw audio," in *International Conference on Machine Learning (ICML)*, ser. Proceedings of Machine Learning Research, vol. 119. PMLR, 2020, pp. 7706–7716. **1, 3**
- [10] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural speech synthesis with transformer network," in *AAAI*, 2019, pp. 6706–6713. **1, 3**
- [11] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks for action segmentation and detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 156–165. **1, 3, 7, 8**
- [12] Y. A. Farha and J. Gall, "Ms-tcn: Multi-stage temporal convolutional network for action segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 3575–3584. **1, 2, 3, 4, 7, 8, 9**
- [13] Z. Wang, Z. Gao, L. Wang, Z. Li, and G. Wu, "Boundary-aware cascade networks for temporal action segmentation," in *Eur. Conf. Comput. Vis.*, 2020. **1, 3, 7, 8**
- [14] M. Fayyaz and J. Gall, "Sct: Set constrained temporal transformer for set supervised action segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 501–510. **1**
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 770–778. **1, 8, 11, 12**
- [16] S. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. Torr, "Res2net: A new multi-scale backbone architecture," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021. **1, 3, 15**
- [17] S. Gao, Q. Han, D. Li, P. Peng, M.-M. Cheng, and P. Peng, "Representative batch normalization with feature calibration," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. **1**
- [18] J. Li and S. Todorovic, "Set-constrained viterbi for set-supervised action segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 10 820–10 829. **1**
- [19] Y. Huang, Y. Sugano, and Y. Sato, "Improving action segmentation via graph-based temporal reasoning," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 14 024–14 034. **1, 8**
- [20] H. Liu, K. Simonyan, and Y. Yang, "Darts: Differentiable architecture search," in *Int. Conf. Learn. Represent.*, 2019. **1, 4, 5, 6, 8, 9**
- [21] H. Cai, L. Zhu, and S. Han, "ProxylessNAS: Direct neural architecture search on target task and hardware," in *Int. Conf. Learn. Represent.*, 2019. **1, 4, 5**
- [22] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan *et al.*, "Searching for mobilenetv3," in *Int. Conf. Comput. Vis.*, 2019, pp. 1314–1324. **1**
- [23] C. Liu, L.-C. Chen, F. Schroff, H. Adam, W. Hua, A. L. Yuille, and L. Fei-Fei, "Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 82–92. **1, 4, 15**
- [24] L. Xie and A. Yuille, "Genetic cnn," in *Int. Conf. Comput. Vis.*, 2017, pp. 1379–1388. **1, 3, 4**
- [25] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," in *AAAI*, vol. 33, 2019, pp. 4780–4789. **1, 3, 4**
- [26] H. Liu, K. Simonyan, O. Vinyals, C. Fernando, and K. Kavukcuoglu, "Hierarchical representations for efficient architecture search," in *Int. Conf. Learn. Represent.*, 2018. **1, 3, 4**
- [27] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *ICASSP*. IEEE, 2019, pp. 3617–3621. **2, 3**
- [28] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016. **2, 3, 13**
- [29] W. Luo, Y. Li, R. Urtasun, and R. Zemel, "Understanding the effective receptive field in deep convolutional neural networks," in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2016, pp. 4905–4913. **2**
- [30] M. Loog and F. Lauze, "Supervised scale-regularized linear convolutionary filters," in *BMVC*, 2017. **2**
- [31] R. Wang, M. Gong, and D. Tao, "Receptive field size versus model depth for single image super-resolution," *IEEE Transactions on Image Processing*, vol. 29, pp. 1669–1682, 2019. **2**
- [32] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," *arXiv preprint arXiv:1706.05587*, 2017. **2, 12, 13**
- [33] Y. Chen, H. Fan, B. Xu, Z. Yan, Y. Kalantidis, M. Rohrbach, S. Yan, and J. Feng, "Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 3435–3444. **2**
- [34] N. Tomen, S.-L. Pintea, and J. Van Gemert, "Deep continuous networks," in *International Conference on Machine Learning (ICML)*. PMLR, 2021, pp. 10 324–10 335. **2**
- [35] G. Seif and D. Androutsos, "Large receptive field networks for high-scale image super-resolution," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 763–772. **2**
- [36] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9. **2**
- [37] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826. **2**
- [38] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Thirty-first AAAI conference on artificial intelligence*, 2017. **2**
- [39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008. **2, 15**
- [40] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7794–7803. **2, 11, 15**
- [41] Y. Chen, M. Rohrbach, Z. Yan, Y. Shuicheng, J. Feng, and Y. Kalantidis, "Graph-based global reasoning networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 433–442. **2, 15**
- [42] M. Soleymani, J. Lichtenauer, T. Pun, and M. Pantic, "A multimodal database for affect recognition and implicit tagging," *IEEE Trans. Affect. Comput.*, vol. 3, no. 1, pp. 42–55, 2011. **2**
- [43] R. T. Collins, A. J. Lipton, and T. Kanade, "Introduction to the special section on video surveillance," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 745–746, 2000. **2**
- [44] R. T. Collins, A. J. Lipton, T. Kanade, H. Fujiyoshi, D. Duggins, Y. Tsin, D. Tolliver, N. Enomoto, O. Hasegawa, P. Burt *et al.*, "A system for video surveillance and monitoring," *VSAM final report*, vol. 2000, pp. 1–68, 2000. **2**
- [45] V. Saligrama and Z. Chen, "Video anomaly detection based on local statistical aggregates," in *IEEE Conf. Comput. Vis. Pattern Recog.* IEEE, 2012, pp. 2112–2119. **2**
- [46] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Adv. Neural Inform. Process. Syst.*, 2014, pp. 568–576. **2**
- [47] C. Feichtenhofer, A. Pinz, and R. P. Wildes, "Spatiotemporal multiplier networks for video action recognition," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 4768–4777. **2**
- [48] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 6299–6308. **2, 7**

- [49] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *Int. Conf. Comput. Vis.*, 2019, pp. 6202–6211. **2**
- [50] A. Fathi and J. M. Rehg, "Modeling actions through state changes," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2013, pp. 2579–2586. **2**
- [51] A. Fathi, A. Farhadi, and J. M. Rehg, "Understanding egocentric activities," in *Int. Conf. Comput. Vis.* IEEE, 2011, pp. 407–414. **2**
- [52] A. Fathi, X. Ren, and J. M. Rehg, "Learning to recognize objects in egocentric activities," in *IEEE Conf. Comput. Vis. Pattern Recog.* IEEE, 2011, pp. 3281–3288. **2, 7**
- [53] M. Rohrbach, S. Amin, M. Andriluka, and B. Schiele, "A database for fine grained activity detection of cooking activities," in *IEEE Conf. Comput. Vis. Pattern Recog.* IEEE, 2012, pp. 1194–1201. **2**
- [54] S. Karaman, L. Seidenari, and A. Del Bimbo, "Fast saliency based pooling of fisher encoded dense trajectories," in *Eur. Conf. Comput. Vis.*, vol. 1, 2014, p. 5. **2**
- [55] S. Bhattacharya, M. M. Kalayeh, R. Sukthankar, and M. Shah, "Recognition of complex events: Exploiting temporal dynamics between underlying concepts," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 2235–2242. **2**
- [56] N. N. Vo and A. F. Bobick, "From stochastic grammar to bayes network: Probabilistic parsing of complex activity," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 2641–2648. **2**
- [57] K. Tang, L. Fei-Fei, and D. Koller, "Learning latent temporal structure for complex event detection," in *IEEE Conf. Comput. Vis. Pattern Recog.* IEEE, 2012, pp. 1250–1257. **2**
- [58] H. Kuehne, J. Gall, and T. Serre, "An end-to-end generative framework for video segmentation and recognition," in *Winter Conference on Applications of Computer Vision.* IEEE, 2016, pp. 1–8. **2, 8**
- [59] H. Kuehne, A. Richard, and J. Gall, "Weakly supervised learning of actions from transcripts," *Comput. Vis. and Image Understanding*, vol. 163, pp. 78–89, 2017. **2, 3**
- [60] A. Richard, H. Kuehne, and J. Gall, "Weakly supervised action learning with rnn based fine-to-coarse modeling," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 754–763. **2, 3, 8**
- [61] H. Kuehne, A. Richard, and J. Gall, "A hybrid rnn-hmm approach for weakly supervised temporal action segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2018. **2**
- [62] Y. Cheng, Q. Fan, S. Pankanti, and A. Choudhary, "Temporal sequence modeling for video event detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 2227–2234. **2**
- [63] A. Richard and J. Gall, "Temporal action detection using a statistical language model," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 3131–3140. **3**
- [64] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao, "A multi-stream bi-directional recurrent neural network for fine-grained action detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 1961–1970. **3, 8**
- [65] S. Singh, C. Arora, and C. Jawahar, "First person action recognition using deep learned descriptors," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 2620–2628. **3**
- [66] L. Ding and C. Xu, "Tricornet: A hybrid temporal convolutional and recurrent network for video action segmentation," *arXiv preprint arXiv:1705.07818*, 2017. **3, 8**
- [67] H. Gammulle, T. Fernando, S. Denman, S. Sridharan, and C. Fookes, "Coupled generative adversarial network for continuous fine-grained action segmentation," in *Winter Conference on Applications of Computer Vision.* IEEE, 2019, pp. 200–209. **3**
- [68] P. Lei and S. Todorovic, "Temporal deformable residual networks for action segmentation in videos," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 6742–6751. **3, 8**
- [69] L. Ding and C. Xu, "Weakly-supervised action segmentation with iterative soft boundary assignment," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 6508–6516. **3, 8**
- [70] F. Sener and A. Yao, "Unsupervised learning and segmentation of complex activities from video," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 8368–8376. **3**
- [71] T. Dutoit, *An introduction to text-to-speech synthesis.* Springer Science & Business Media, 1997, vol. 3. **3**
- [72] P. Taylor, *Text-to-speech synthesis.* Cambridge university press, 2009. **3**
- [73] S. Kim, S. Lee, J. Song, J. Kim, and S. Yoon, "Flowavenet : A generative flow for raw audio," in *International Conference on Machine Learning (ICML)*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 2019, pp. 3370–3378. **3**
- [74] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, "A survey on neural speech synthesis," *arXiv preprint arXiv:2106.15561*, 2021. **3**
- [75] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *ICASSP.* IEEE, 2015, pp. 5206–5210. **3**
- [76] S. Ö. Arik, M. Chrzanowski, A. Coates, G. Diamos, A. Gibiansky, Y. Kang, X. Li, J. Miller, A. Ng, J. Raiman *et al.*, "Deep voice: Real-time neural text-to-speech," in *International Conference on Machine Learning (ICML).* PMLR, 2017, pp. 195–204. **3**
- [77] T. Yoshimura, K. Tokuda, T. Masuko, T. Kobayashi, and T. Kitamura, "Simultaneous modeling of spectrum, pitch and duration in hmm-based speech synthesis," in *Sixth European Conference on Speech Communication and Technology*, 1999. **3**
- [78] H. Ze, A. Senior, and M. Schuster, "Statistical parametric speech synthesis using deep neural networks," in *ICASSP.* IEEE, 2013, pp. 7962–7966. **3**
- [79] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio *et al.*, "Tacotron: Towards end-to-end speech synthesis," in *Interspeech*, 2017, pp. 4006–4010. **3, 13**
- [80] S. Qiao, L.-C. Chen, and A. Yuille, "Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution," *arXiv preprint arXiv:2006.02334*, 2020. **3**
- [81] P. Sun, R. Zhang, Y. Jiang, T. Kong, C. Xu, W. Zhan, M. Tomizuka, L. Li, Z. Yuan, C. Wang, and P. Luo, "Sparse r-cnn: End-to-end object detection with learnable proposals," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 14 454–14 463. **3**
- [82] Y. Xiong, H. Liu, S. Gupta, B. Akin, G. Bender, Y. Wang, P.-J. Kindermans, M. Tan, V. Singh, and B. Chen, "Mobilenets: Searching for object detection architectures for mobile accelerators," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 3825–3834. **3**
- [83] J. Wang, L. Song, Z. Li, H. Sun, J. Sun, and N. Zheng, "End-to-end object detection with fully convolutional network," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 15 849–15 858. **3**
- [84] T. Liang, Y. Wang, Z. Tang, G. Hu, and H. Ling, "Opanas: One-shot path aggregation network architecture search for object detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 10 195–10 203. **3**
- [85] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *Eur. Conf. Comput. Vis.* Springer, 2016, pp. 21–37. **3**
- [86] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 779–788. **3**
- [87] H. Law and J. Deng, "Cornersnet: Detecting objects as paired keypoints," in *Eur. Conf. Comput. Vis.*, 2018, pp. 734–750. **3**
- [88] Z. Cai and N. Vasconcelos, "Cascade r-cnn: High quality object detection and instance segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, p. 1–1, 2019. [Online]. Available: <http://dx.doi.org/10.1109/tpami.2019.2956516> **3**
- [89] K. Chen, J. Pang, J. Wang, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, "Hybrid task cascade for instance segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. **3**
- [90] Z. Zou, Z. Shi, Y. Guo, and J. Ye, "Object detection in 20 years: A survey," *arXiv preprint arXiv:1905.05055*, 2019. **3**
- [91] S. S. A. Zaidi, M. S. Ansari, A. Aslam, N. Kanwal, M. Asghar, and B. Lee, "A survey of modern deep learning based object detection models," *arXiv preprint arXiv:2104.11892*, 2021. **3**
- [92] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, and M. Pietikäinen, "Deep learning for generic object detection: A survey," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 261–318, 2020. **3**
- [93] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 580–587. **3**
- [94] M. Hu, Y. Li, L. Fang, and S. Wang, "A2-fpn: Attention aggregation based feature pyramid network for instance segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 15 343–15 352. **3**
- [95] X. Shen, J. Yang, C. Wei, B. Deng, J. Huang, X.-S. Hua, X. Cheng, and K. Liang, "Dct-mask: Discrete cosine transform mask representation for instance segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 8720–8729. **3**
- [96] H. Ding, S. Qiao, A. Yuille, and W. Shen, "Deeply shape-guided cascade for instance segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 8278–8288. **3**
- [97] T. Cheng, X. Wang, L. Huang, and W. Liu, "Boundary-preserving mask r-cnn," in *Eur. Conf. Comput. Vis.* Springer, 2020, pp. 660–676. **3**
- [98] Z. Hayder, X. He, and M. Salzmann, "Boundary-aware instance segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 5696–5704. **3**

- [99] J. Liang, N. Homayounfar, W.-C. Ma, Y. Xiong, R. Hu, and R. Urtasun, "Polytransform: Deep polygon transformer for instance segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 9131–9140. **3**
- [100] Y. Yuan, J. Xie, X. Chen, and J. Wang, "Segfix: Model-agnostic boundary refinement for segmentation," in *Eur. Conf. Comput. Vis.* Springer, 2020, pp. 489–506. **3**
- [101] C. Tang, H. Chen, X. Li, J. Li, Z. Zhang, and X. Hu, "Look closer to segment better: Boundary patch refinement for instance segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021, pp. 13926–13935. **3**
- [102] M. Mitchell, *An introduction to genetic algorithms*. MIT press, 1998. **3, 4**
- [103] Y. Sun, B. Xue, M. Zhang, G. G. Yen, and J. Lv, "Automatically designing cnn architectures using the genetic algorithm for image classification," *IEEE Trans. Cybernetics*, 2020. **3, 4**
- [104] Z. Lu, I. Whalen, V. Boddeti, Y. Dhebar, K. Deb, E. Goodman, and W. Banzhaf, "Nsga-net: neural architecture search using multi-objective genetic algorithm," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2019, pp. 419–427. **3**
- [105] X. Chu, T. Zhou, B. Zhang, and J. Li, "Fair darts: Eliminating unfair advantages in differentiable architecture search," in *Eur. Conf. Comput. Vis.* Springer, 2020, pp. 465–480. **4, 9**
- [106] Y. Xu, L. Xie, X. Zhang, X. Chen, G.-J. Qi, Q. Tian, and H. Xiong, "Pc-darts: Partial channel connections for memory-efficient architecture search," in *Int. Conf. Learn. Represent.*, 2020. **4**
- [107] A. Wan, X. Dai, P. Zhang, Z. He, Y. Tian, S. Xie, B. Wu, M. Yu, T. Xu, K. Chen *et al.*, "Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 12965–12974. **4**
- [108] X. Chen, L. Xie, J. Wu, and Q. Tian, "Progressive differentiable architecture search: Bridging the depth gap between search and evaluation," in *Int. Conf. Comput. Vis.*, 2019, pp. 1294–1303. **4**
- [109] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *nature*, vol. 323, no. 6088, pp. 533–536, 1986. **4**
- [110] A. Zela, T. Elsken, T. Saikia, Y. Marrakchi, T. Brox, and F. Hutter, "Understanding and robustifying differentiable architecture search," in *Int. Conf. Learn. Represent.*, 2020. **4, 6**
- [111] X. Zhang, H. Xu, H. Mo, J. Tan, C. Yang, L. Wang, and W. Ren, "Dcnas: Densely connected neural architecture search for semantic image segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13956–13967. **4**
- [112] S. Stein and S. J. McKenna, "Combining embedded accelerometers with computer vision for recognizing food preparation activities," in *The ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2013, pp. 729–738. **7**
- [113] H. Kuehne, A. Arslan, and T. Serre, "The language of actions: Recovering the syntax and semantics of goal-directed human activities," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 780–787. **7**
- [114] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Adv. Neural Inform. Process. Syst.*, 2019, pp. 8026–8037. **7**
- [115] S.-M. Hu, D. Liang, G.-Y. Yang, G.-W. Yang, and W.-Y. Zhou, "Jittor: a novel deep learning framework with meta-operators and unified graph execution," *Science China Information Sciences*, vol. 63, no. 12, pp. 1–21, 2020. **7**
- [116] H. Li, Z. Lin, X. Shen, J. Brandt, and G. Hua, "A convolutional neural network cascade for face detection," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 5325–5334. **7**
- [117] C. Lea, A. Reiter, R. Vidal, and G. D. Hager, "Segmental spatiotemporal cnns for fine-grained action segmentation," in *Eur. Conf. Comput. Vis.* Springer, 2016, pp. 36–52. **8**
- [118] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Int. Conf. Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241. **8**
- [119] H. Liang, S. Zhang, J. Sun, X. He, W. Huang, K. Zhuang, and Z. Li, "Darts+: Improved differentiable architecture search with early stopping," *arXiv preprint arXiv:1909.06035*, 2019. **9**
- [120] X. Chu, T. Zhou, B. Zhang, and J. Li, "Fair DARTS: Eliminating Unfair Advantages in Differentiable Architecture Search," in *Eur. Conf. Comput. Vis.*, 2020. **9**
- [121] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Eur. Conf. Comput. Vis.*, 2014. **10, 11, 12**
- [122] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2009. **11, 12**
- [123] Z. Huang, X. Wang, L. Huang, C. Huang, Y. Wei, and W. Liu, "Ccnet: Criss-cross attention for semantic segmentation," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 603–612. **12**
- [124] M. Yang, K. Yu, C. Zhang, Z. Li, and K. Yang, "Denseaspp for semantic segmentation in street scenes," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 3684–3692. **12**
- [125] M. Everingham, L. Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *Int. J. Comput. Vis.*, vol. 88, pp. 303–338, 2009. **12, 13, 14**
- [126] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Scene parsing through ade20k dataset," in *IEEE Conf. Comput. Vis. Pattern Recog.*, July 2017. **12, 13, 14**
- [127] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *ICASSP*, 2019, pp. 3617–3621. **13, 14**
- [128] D. P. Kingma and P. Dhariwal, "Glow: Generative flow with invertible 1x1 convolutions," in *Adv. Neural Inform. Process. Syst.*, vol. 31, 2018. **13**
- [129] K. Ito and L. Johnson, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017. **13, 14**
- [130] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4779–4783. **13**
- [131] R. Kubichek, "Mel-cepstral distance measure for objective speech quality assessment," in *PACRIM*, vol. 1, 1993, pp. 125–128 vol.1. **13, 14**
- [132] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *ICASSP*, vol. 2, 2001, pp. 749–752 vol.2. **13, 14**
- [133] S. R. Quackenbush, T. P. Barnwell, and M. A. Clements, *Objective measures of speech quality*. Prentice-Hall, 1988. **13, 14**
- [134] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv:1803.01271*, 2018. **13, 14, 15**
- [135] Q. V. Le, N. Jaitly, and G. E. Hinton, "A simple way to initialize recurrent networks of rectified linear units," *arXiv preprint arXiv:1504.00941*, 2015. **13**
- [136] S. Zhang, Y. Wu, T. Che, Z. Lin, R. Memisevic, R. R. Salakhutdinov, and Y. Bengio, "Architectural complexity measures of recurrent neural networks," in *Adv. Neural Inform. Process. Syst.*, vol. 29, 2016, pp. 1822–1830. **13**
- [137] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998. **13**
- [138] S. Wisdom, T. Powers, J. Hershey, J. Le Roux, and L. Atlas, "Full-capacity unitary recurrent neural networks," in *Adv. Neural Inform. Process. Syst.*, vol. 29, 2016, pp. 4880–4888. **13**
- [139] T. Cooijmans, N. Ballas, C. Laurent, Ç. Gülçehre, and A. Courville, "Recurrent batch normalization," in *Int. Conf. Learn. Represent.*, 2016. **13**
- [140] D. Krueger, T. Maharaj, J. Kramár, M. Pezeshki, N. Ballas, N. R. Ke, A. Goyal, Y. Bengio, A. Courville, and C. Pal, "Zoneout: Regularizing rnns by randomly preserving hidden activations," in *Int. Conf. Learn. Represent.*, 2016. **13**
- [141] L. Jing, Y. Shen, T. Dubcek, J. Peurifoy, S. Skirlo, Y. LeCun, M. Tegmark, and M. Soljačić, "Tunable efficient unitary neural networks (eunn) and their application to rnns," in *International Conference on Machine Learning (ICML)*. PMLR, 2017, pp. 1733–1741. **13**
- [142] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," *arXiv preprint arXiv:1412.3555*, 2014. **14**
- [143] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, "How to construct deep recurrent neural networks," *arXiv preprint arXiv:1312.6026*, 2013. **14**
- [144] R. Jozefowicz, W. Zaremba, and I. Sutskever, "An empirical exploration of recurrent network architectures," in *International Conference on Machine Learning (ICML)*. PMLR, 2015, pp. 2342–2350. **14**
- [145] W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao, "Pyramid vision transformer: A versatile backbone for dense prediction without convolutions," in *Int. Conf. Comput. Vis.*, 2021, pp. 568–578. **14**

- [146] —, “Pvtv2: Improved baselines with pyramid vision transformer,” *Computational Visual Media*, vol. 8, no. 3, pp. 1–10, 2022. [14](#)
- [147] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, “A convnet for the 2020s,” *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. [14](#)
- [148] K. Sun, B. Xiao, D. Liu, and J. Wang, “Deep high-resolution representation learning for human pose estimation,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. [15](#)
- [149] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang, W. Liu, and B. Xiao, “Deep high-resolution representation learning for visual recognition,” *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019. [15](#)
- [150] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016. [15](#)
- [151] Y. Chen, Y. Kang, Y. Chen, and Z. Wang, “Probabilistic forecasting with temporal convolutional neural network,” *Neurocomputing*, vol. 399, pp. 491–501, 2020. [15](#)
- [152] S. Gao, Y.-Q. Tan, M.-M. Cheng, C. Lu, Y. Chen, and S. Yan, “Highly efficient salient object detection with 100k parameters,” in *Eur. Conf. Comput. Vis.*, 2020. [15](#)
- [153] Y.-Q. Tan, S. Gao, X.-Y. Li, M.-M. Cheng, and B. Ren, “Vecroad: Point-based iterative graph exploration for road graphs extraction,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. [15](#)
- [154] Y. Gu, L. Wang, Z. Wang, Y. Liu, M.-M. Cheng, and S.-P. Lu, “Pyramid constrained self-attention network for fast video salient object detection,” in *AAAI*, 2020, pp. 10 869–10 876. [15](#)
- [155] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” *arXiv preprint arXiv:1506.01497*, 2015. [15](#)
- [156] Q. Han, K. Zhao, J. Xu, and M.-M. Cheng, “Deep hough transform for semantic line detection,” in *Eur. Conf. Comput. Vis.*, 2020. [15](#)



Ming-Ming Cheng received his PhD degree from Tsinghua University in 2012, and then worked with Prof. Philip Torr in Oxford for 2 years. He is now a professor at Nankai University, leading the Media Computing Lab. His research interests includes computer vision and computer graphics. He received awards including ACM China Rising Star Award, IBM Global SUR Award, *etc.* He is a senior member of the IEEE and on the editorial boards of IEEE TPAMI and IEEE TIP.



Liang Wang received his Ph.D. degree from the Institute of Automation, Chinese Academy of Sciences (CASIA) in 2004. From 2004 to 2010, he has been working at Imperial College London, United Kingdom, Monash University, Australia, the University of Melbourne, Australia, and the University of Bath, United Kingdom, respectively. Currently, he is a full Professor at the National Lab of Pattern Recognition, CASIA. His major research interests include machine learning, pattern recognition and computer vision. He

is an associate editor of IEEE Transactions on SMC-B. He is currently an IAPR Fellow and Senior Member of IEEE.



Shanghua Gao is a Ph.D. candidate in Media Computing Lab at Nankai University. He is supervised via Prof. Ming-Ming Cheng. His research interests include computer vision and representation learning.



Zhong-Yu Li is a Ph.D. student from the college of computer science, Nankai university. He is supervised via Prof. Ming-Ming cheng. His research interests include deep learning, machine learning and computer vision.



Qi Han is a master student from the College of Computer Science, Nankai University, under the supervision of Prof. Ming-Ming Cheng. He received his bachelor degree from Xidian University in 2019. His research interests include deep learning and computer vision.