Searching Multi-Rate and Multi-Modal Temporal Enhanced Networks for Gesture Recognition

Zitong Yu¹, Graduate Student Member, IEEE, Benjia Zhou², Jun Wan³, Senior Member, IEEE,
Pichao Wang⁴, Member, IEEE, Haoyu Chen⁵, Xin Liu⁶, Member, IEEE,
Stan Z. Li, Fellow, IEEE, and Guoying Zhao⁷, Senior Member, IEEE

Abstract—Gesture recognition has attracted considerable attention owing to its great potential in applications. Although the great progress has been made recently in multi-modal learning methods, existing methods still lack effective integration to fully explore synergies among spatio-temporal modalities effectively for gesture recognition. The problems are partially due to the fact that the existing manually designed network architectures have low efficiency in the joint learning of multi-modalities. In this paper, we propose the first neural architecture search (NAS)-based method for RGB-D gesture recognition. The proposed method includes two key components: 1) enhanced temporal representation via the proposed 3D Central Difference Convolution (3D-CDC) family, which is able to capture rich temporal context via aggregating temporal difference information; and 2) optimized backbones for multi-sampling-rate branches and lateral connections among varied modalities. The resultant multi-modal multi-rate network provides a new perspective to understand the relationship between RGB and depth modalities and their temporal dynamics. Comprehensive experiments are performed on three benchmark datasets (IsoGD, NvGesture, and EgoGesture), demonstrating the state-of-the-art performance in both single- and multi-modality settings. The code is available at <https://github.com/ZitongYu/3DCDC-NAS>.

Index Terms—3D-CDC, NAS, RGB-D gesture recognition.

Manuscript received August 21, 2020; revised March 4, 2021; accepted May 25, 2021. Date of publication June 14, 2021; date of current version June 18, 2021. This work was supported in part by the Academy of Finland for project MiGA under Grant 316765, in part by the ICT 2023 Project under Grant 328115, in part by the Infotech Oulu, in part by the Chinese National Natural Science Foundation under Project 61961160704 and Project 61876179, in part by the Science and Technology Development Fund of Macau under Grant 0010/2019/AFJ and Grant 0025/2019/AKP, and in part by the External Cooperation Key Project of Chinese Academy Sciences under Grant 173211KYSB20200002. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Lucio Marcenoro. (Zitong Yu and Benjia Zhou contributed equally to this work.) (Corresponding authors: Jun Wan; Guoying Zhao.)

Zitong Yu, Haoyu Chen, and Guoying Zhao are with the Center for Machine Vision and Signal Analysis, University of Oulu, 90014 Oulu, Finland (e-mail: guoying.zhao@oulu.fi).

Benjia Zhou is with the Faculty of Information Technology, Macau University of Science and Technology, Macau 999078, China.

Jun Wan is with the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, University of Chinese Academy of Sciences, Beijing 100190, China (e-mail: jun.wan@ia.ac.cn).

Pichao Wang is with the Alibaba Group, Bellevue, WA 98004 USA.

Xin Liu is with the School of Electrical and Information Engineering, Tianjin University, Tianjin 300072, China, and also with the Computer Vision and Pattern Recognition Laboratory, School of Engineering Science, Lappeenranta-Lahti University of Technology (LUT), 53850 Lappeenranta, Finland.

Stan Z. Li is with the School of Engineering, Westlake University, Hangzhou 310012, China.

Digital Object Identifier 10.1109/TIP.2021.3087348

I. INTRODUCTION

AS ONE of the multi-modal video understanding problems, RGB-D based video gesture recognition [1]–[4] has been applied to many real-world applications, such as virtual reality [5] and human-computer interaction [6]. 3D convolutional neural network (3DCNN) [7]–[10] and long short-term memory (LSTM) [11] have been adopted to learn the spatio-temporal features for gesture recognition. However, the learned spatio-temporal representation is still easily contaminated by irrelevant factors (e.g., illumination and background). A feasible solution is to add an extra enhanced temporal feature learning module [12]–[14], which is computationally costly and tricky for the off-the-shelf 3DCNNs. **How to learn efficient spatio-temporal features in basic convolution operator for enriching temporal context is worth exploring for gesture recognition.**

As gestures have various temporal ranges, modeling such visual tempos would benefit for gesture recognition. Previous methods [15]–[17] attempt to construct the frame pyramid for such purpose, with each branch of the frame pyramid sampling the input frames at a different rate. However, the architecture (i.e., network structure) of each branch and relations (i.e., lateral connections) among the multi-rate branches are usually shared and hand-designed, which is sub-optimal for message propagation. Hence, **how to discover better-suited architectures and lateral connections among multi-rate branches is crucial.**

For RGB-D based gesture recognition, complementary feature learning from different data modalities is beneficial. For example, the depth data is easy to distinguish foregrounds (i.e., face, hands, and arms) from backgrounds while RGB data provides detailed texture/color appearances. However, most existing methods [7], [8], [18]–[20] conduct the multi-modal fusion via coarse strategies (e.g. score fusion or last layer concatenation), which may leverage the multi-modal information insufficiently. Thus, **to design more reasonable multi-modal fusion strategy is not a trivial work.**

Recently, both works about expert-designed [15], [21], [22] and automatically searched [23], [24] networks have been explored in video-based action recognition. Compared with action recognition, the main characteristic [25] of RGB-D gesture recognition derives from: 1) fine granularity within the body gestures. Actions tend to focus on the body information

with large motions (such as hugging and sports) while gestures are produced as part of deliberate actions and signs, involving the motion of the upper body, especially the arms, hands, and fingers. Thus, the resultant successive and temporally correlative movements are hard to be distinguished with only global or local spatio-temporal context; and 2) rich synergies between RGB-D modalities. The depth modality provides the foreground-based attention on upper body movement but with coarse and noisy local clues, while the RGB modality describes detailed appearance (e.g., local shape changes of the hands and fingers) but is easily influenced by the interference such as background and clothes. To efficiently explore RGB-D synergies is vital for reliable gesture recognition application.

Motivated by the above observations, we propose a novel spatio-temporal convolution family called 3D Central Difference Convolution (3D-CDC), to exploit the rich local motion and enhance the spatio-temporal representation. Moreover, over the 3D-CDC-based enhanced search space, Neural Architecture Search (NAS) is adopted to automatically discover the optimized multi-rate and multi-modal networks for RGB-D gesture recognition. Our contributions include:

- A novel spatio-temporal convolution family, 3D-CDC, is proposed, intending to capture rich temporal context via aggregating temporal difference information. Without introducing extra parameters, 3D-CDC can replace the vanilla 3D convolution, and plug and play in existing 3DCNNs for various modalities with enhanced temporal modeling capacity.
- We propose a two-stage NAS method to automatically discover well-suited backbones and lateral connections for the multi-rate and multi-modal networks, which effectively explores RGB-D-temporal synergies and represents global dynamics.
- To our best knowledge, this is the first approach that searches multi-rate and multi-modal architectures for RGB-D gesture recognition. Our searched architecture provides a new perspective to understand the relationship among multi-rate branches as well as modalities.
- Our proposed method achieves state-of-the-art (SOTA) performance on three benchmark datasets on both single- and multi-modality testing.

In the rest of the paper, Sec. II provides the related work. Sec. III formulates the 3D-CDC family, and introduce the two-stage multi-rate and multi-modal NAS algorithm. Sec. IV provides rigorous ablation studies and evaluates the performance of the proposed models on three benchmark datasets. Sec. V shows the visualization results and discusses transferability to the action recognition task. Finally, a conclusion with future directions is given in Sec. VI.

II. RELATED WORK

In this section, we first introduce some recent progress in multi-modal gesture recognition. Then, previous video-based NAS methods will be reviewed.

A. Multi-Modal Gesture Recognition

For video-based gesture recognition, it is challenging to track the motion of hands and arms owing to the

large degree of freedom. Many hand-crafted feature based [26]–[29] and deep learning-based methods [11], [30]–[33] are proposed to tackle this issue. As for the learning-based gesture recognition, on one hand, 3DCNNs including C3D [34], Res3D [35], I3D [21] and SlowFast [15] are utilized for gesture feature extractor. On the other hand, LSTM variants such as AttenConvLSTM [36], [36] and PreRNN [37] are introduced for temporal memory propagation. Based on the CNNs, several extended modules [12]–[14], [38] and convolution operators [39], [40] are developed for enhancing the spatio-temporal representation. However, most of them need extra structures and learnable parameters to modulate the original spatio-temporal features. In this paper, we propose 3D-CDC for modeling a rich temporal context, which is vital for describing fine-grained hands/arms motion. Among these methods, Lee *et al.*'s [13] motion feature network (MFNet) and Yu *et al.*'s [40], [41] central difference convolution (CDC) are the most similar to our work. Instead of the fixed motion filter used in MFNet and only the spatial context in CDC, our 3D-CDC learns the temporal gradient (motion) filters automatically in a unified 3D convolution operator.

Due to the development of the RGB-D cameras like Intel RealSense, RGB and depth modalities [7], [8], [42] are favorite to complementarily fuse and improve the performance. Besides the RGB-D data, some other modalities such as optical flow [7], [43], depth flow [42], skeleton [44], [45], dynamic images [46] and saliency video [9], [47]–[49], [49] are also employed for gesture recognition. However, the flow and saliency modalities need extra computational cost and might lose some vital information after pre-processing.

In terms of the multi-modal fusion strategy, decision-level fusion [9], [50], [51] and feature-level fusion [7], [8], [20], [42] methods have developed for integrating mutual knowledge from varied modalities. Despite achieving SOTA performance, the existing multi-modal fusion strategies for gesture recognition are designed manually and coarsely, which might be sub-optimal for message propagation between modalities. In this paper, we prefer to discover well-suitable multi-modal fusion strategies automatically via NAS.

B. Video-Based Neural Architecture Search

Our work is motivated by recent researches on NAS [52]–[60], while we focus on searching for multi-rate and multi-modal networks specially for RGB-D gesture recognition. The existing NAS methods could be summarized in three categories: 1) Reinforcement learning-based [55], [56], 2) Evolution-based [61], [62], and 3) Gradient-based [53], [57], [63]. From the perspective of NAS based video classification applications, single-modal based [64]–[66] and multi-modal based [23], [24] methods have been developed for the action recognition task. Unlike AssembleNet [24] which searches on RGB and optical flow modalities with single frame rate inputs, our work focuses on searching the synergies among multi-rate and multi-modal branches. Some works also consider searching CNN [67] or graph convolutional networks [68] on skeleton modality for video action

recognition, while our work focuses on multi-modal NAS for gesture recognition on RGB-D modalities.

In terms of the search space design, cell-based NASNet [56] space is favorite due to its flexibility and rich capacity. Besides, some novel operators, such as extended convolution [40], [69] and attention [66], [70], are introduced into search space, which is proved to be beneficial for searching more powerful architecture. However, there are still no operators specially designed for temporal enhancement.

To our best knowledge, no NAS based method has ever been proposed for RGB-D gesture recognition. To fill in the blank, we search multi-rate and multi-modal networks over the temporal enhanced search space for RGB-D gesture recognition.

III. METHODOLOGY

In this section, we first introduce the 3D-CDC family in Sec. III-A. Over the 3D-CDC based space, we then propose the two-stage multi-rate and multi-modal NAS in Sec. III-B.

A. Temporal Enhancement via 3D-CDC

In classical 3DCNNs, 3D convolution is the most fundamental operator for spatio-temporal feature representation. In this subsection, for simplicity, the 3D convolutions are described in 3D (without channel) while an extension to 4D is straightforward. There are two main steps in the vanilla 3D convolution: 1) *sampling* local receptive field cube $\mathcal{C}$ over the input feature map x ; 2) *aggregation* of sampled values via weighted summation with learnable weights w . Hence, the output feature map *Vanilla* can be formulated as

$$Vanilla(p_0) = \sum_{p_n \in \mathcal{C}} w(p_n) \cdot x(p_0 + p_n), \quad (1)$$

where p_0 denotes current location on both input and output feature maps while p_n enumerates the locations in $\mathcal{C}$. In this subsection, 3D convolution with kernel size $3 \times 3 \times 3$ and dilation 1 is used for demonstration, and the other configurations are analogous. The local receptive field cube for the 3D convolution is $\mathcal{C} = \{(-1, -1, -1), (-1, -1, 0), \dots, (0, 1, 1), (1, 1, 1)\}$.

1) *Spatio-Temporal Central Difference Convolution (3D-CDC-ST)*: Inspired by the CDC [40] which introduces spatial gradient cues into representation learning, we integrate spatio-temporal gradient information into a unified 3D convolution operator. It is worth noting that such spatial gradient and temporal difference designs are widely used in the dense optical flow [71] calculation. Instead of the optical flow only performed in RGB sequence level, networks with stacked spatio-temporal CDC would be regularized to learn more local motion context in both RGB sequence and deep feature level, which is able to model fine-grained temporal dynamics for gesture recognition.

Similarly, spatio-temporal CDC also consists of two steps, i.e., *sampling* and *aggregation*. The sampling step is similar to vanilla 3D convolution but the aggregation step is different: as illustrated in Fig. 1(a), spatio-temporal CDC prefers to aggregate the center-oriented spatio-temporal gradient of sampled

values. Eq. (1) becomes

$$CDC(p_0) = \sum_{p_n \in \mathcal{C}} w(p_n) \cdot (x(p_0 + p_n) - x(p_0)). \quad (2)$$

When $p_n = (0, 0, 0)$, the gradient value always equals to zero with respect to the central location p_0 itself. For the gesture recognition task, both spatio-temporal intensity-level semantic information and gradient-level difference message are crucial and complementary. The former one is good at global modeling and robust to sensor-based noise while the latter one focuses more on local appearance and motion details and might be influenced by noise. As a result, combining vanilla 3D convolution with 3D-CDC might be a feasible manner to provide more robust and discriminative modeling capacity. Therefore we generalize spatio-temporal CDC as

$$\begin{aligned} CDC_{ST}(p_0) &= \theta \cdot CDC(p_0) + (1 - \theta) \cdot Vanilla(p_0) \\ &= \underbrace{\sum_{p_n \in \mathcal{C}} w(p_n) \cdot x(p_0 + p_n)}_{\text{vanilla 3D convolution}} + \theta \cdot \underbrace{(-x(p_0) \cdot \sum_{p_n \in \mathcal{C}} w(p_n))}_{\text{spatio-temporal CD term}}, \quad (3) \end{aligned}$$

where hyperparameter $\theta \in [0, 1]$ tradeoffs the contribution between intensity-level and gradient-level information. Please note that $w(p_n)$ is shared between vanilla 3D convolution and spatio-temporal central difference (CD) term, thus no extra parameters are added.

2) *Temporal Central Difference Convolution (3D-CDC-T)*: Unlike the aforementioned spatio-temporal CDC considering both spatial and temporal gradient cues, we propose a version with only temporal central differences. As shown in Fig. 1(b), the sampled local receptive field cube $\mathcal{C}$ is separated into two kinds of regions: 1) the region in the current time step $\mathcal{R}'$, and 2) the regions in the adjacent time steps $\mathcal{R}''$. In the setting of a temporal CDC, the central difference term is only calculated from $\mathcal{R}''$. Thus the generalized temporal CDC can be formulated via modifying Eq. (3) as

$$\begin{aligned} CDC_T(p_0) &= \underbrace{\sum_{p_n \in \mathcal{C}} w(p_n) \cdot x(p_0 + p_n)}_{\text{vanilla 3D convolution}} + \theta \cdot \underbrace{(-x(p_0) \cdot \sum_{p_n \in \mathcal{R}''} w(p_n))}_{\text{temporal CD term}}. \quad (4) \end{aligned}$$

3) *Temporal Robust Central Difference Convolution (3D-CDC-TR)*: In consideration of the sensor noise especially in the depth modality, we also propose a version with the temporal robust central difference. Similarly, the temporal robust CDC only calculates the difference term from the regions in the adjacent time steps $\mathcal{R}''$. As shown in Fig. 1(c), the robust temporal center is represented by averaging the spatial centers of all time steps (i.e., p_0^{t-1} , p_0 and p_0^{t+1}) within $\mathcal{C}$. The robust temporal center-oriented gradient might be less sensitive to the pixel jitters from the adjacent time steps. The generalized temporal robust CDC can also be

Fig. 1. Our proposed 3D-CDC family with kernel size 3, which could be used as novel operators for NAS. (a) 3D-CDC-ST considers the central difference information in the whole local spatio-temporal regions. (b) 3D-CDC-T only calculates the central difference clues from the local spatio-temporal regions of the adjacent frames. (c) 3D-CDC-TR is similar to 3D-CDC-T but adopts the temporal central mean pooling before calculating the central difference term, which is more robust to temporal noise. The symbols $\ominus$ and $\oplus$ denote element-wise subtraction and mean operations, respectively. (d) Feature response of various convolutions in RGB modality. Compared with vanilla 3D convolution, the 3D-CDC family enhances the temporal context obviously.

Fig. 2. Architecture search space in the first stage. Single-modal (RGB or depth) multi-rate frames are adopted as inputs. Here we utilize 3 branches with different frame rates (e.g., uniformly sampling the original videos into 8, 16, and 32 frames, respectively), and it also can be extended to more branches according to the actual situation. In this search stage, inspired by [15], the architectures of all lateral connections are fixed with temporal convolutions. And the outputs of the lateral connections are concatenated with the features from the target branch. The channel numbers are doubled after each MaxPool layer. The architecture of the cells from multi-rate branches to be searched can be shared or unshared (see Sec. IV-C for ablation study).

formulated via modifying Eq. (3) as

$$\begin{aligned}
 & CDC_{TR}(p_0) \\
 &= \underbrace{\sum_{p_n \in \mathcal{C}} w(p_n) \cdot x(p_0 + p_n)}_{\text{vanilla 3D convolution}} \\
 &\quad + \underbrace{\theta \cdot (-\text{Avg}[x(p_0^{t-1}), x(p_0), x(p_0^{t+1})] \cdot \sum_{p_n \in \mathcal{R}''} w(p_n))}_{\text{temporal robust CD term}}. \quad (5)
 \end{aligned}$$

We will henceforth refer to these three generalized versions (i.e., Eq. (3), (4) and (5)) as 3D-CDC-ST, 3D-CDC-T and 3D-CDC-TR, respectively. The ablation studies about the 3D-CDC family and hyperparameter θ are conducted in Sec. IV-C.

4) *Relation Between 3D-CDC and 3D Vanilla Convolution:* Compared to 3D vanilla convolution, 1) 3D-CDC-ST regularizes the spatio-temporal representation with more detailed

spatial cues and temporal dynamics, which might be suitable for scene-aware video understanding tasks; 2) 3D-CDC-T enhances the spatio-temporal representation with only rich temporal context, which might be assembled for video temporal reasoning tasks; and 3) 3D-CDC-TR introduces robust but slighter temporal evolution cues for spatio-temporal representation, which might perform robustly even in noisy scenarios. In particular, 3D-CDC-ST, 3D-CDC-T, and 3D-CDC-TR degrade to vanilla 3D convolution when $\theta = 0$. As illustrated in Fig. 1(d), the 3D-CDC family provides more details about the trajectory of the left arm, and such a local motion context is crucial to gesture recognition. More visualizations are shown in Sec. V-C.

B. Multi-Rate and Multi-Modal NAS

In order to seek the best-suited backbones and lateral connections for multi-rate and multi-modal networks, we propose a two-stage NAS method to 1) search backbones for

Algorithm 1 Two-Stage Multi-Rate and Multi-Modal NAS**Stage1:** Fix lateral connections, and search backbonesFor multi-rate backbones, create a mixed operation $\tilde{o}_b^{(i,j)}$ parametrized by $\alpha_b^{(i,j)}$ for each edge (i, j)

```

1 : for each of modalities  $\mathcal{M}$  do
2 :   Fix the lateral connections among multi-rate branches
3 :   while not converged do
4 :     Update architecture  $\alpha_b$  by descending  $\nabla_{\alpha_b} \mathcal{L}_{val}(\Phi_b, \alpha_b)$ 
5 :     Update weights  $\Phi_b$  by descending  $\nabla_{\Phi_b} \mathcal{L}_{train}(\Phi_b, \alpha_b)$ 
6 :   end
7 :   Derive the final backbone of the current modality based on
   the learned  $\alpha_b$ 
8 : end

```

Stage2: Fix backbones, and search lateral connectionsFor lateral connections, create a mixed operation $\tilde{o}_c^{(i,j)}$ parametrized by $\alpha_c^{(i,j)}$ for each edge (i, j) 9 : Initialize and fix the multi-rate and multi-modal backbones searched in **Stage 1**

```

10 : while not converged do
11 :   Update architecture  $\alpha_c$  by descending  $\nabla_{\alpha_c} \mathcal{L}_{val}(\Phi_c, \alpha_c)$ 
12 :   Update weights  $\Phi_c$  by descending  $\nabla_{\Phi_c} \mathcal{L}_{train}(\Phi_c, \alpha_c)$ 
13 : end
14 : Derive the final lateral connections based on the learned  $\alpha_c$ 

```

multi-rate single-modal networks first, and then 2) search lateral connections for multi-modal networks based on the searched backbones. The iterative procedure is outlined in Algorithm 1. More technical details can be referred to two gradient-based NAS methods [53], [57].

1) *Stage 1: Searching Backbones for Multi-Rate Single-Modal Networks:* In SlowFast Networks [15], low- and high-rate branches are utilized for complementarily capturing dynamic visual tempos. However, the coarse hand-defined architecture with vanilla convolutions limits its representation capacity. Here we search optimal rate-aware backbones over the temporal enhanced search space.

As illustrated in Fig. 2, in the first stage, our goal is to search for cells to form multi-rate backbones for single-modal gesture recognition. As for the cell-level structure, similar to [53], each cell is represented as a directed acyclic graph (DAG) of K nodes $\{n_i\}_{i=0}^{K-1}$, where each node represents a network layer. We denote the operation space as $\mathcal{O}_b$, which consists of seven designed candidate operations: ‘Zero’, ‘Identity’, ‘Conv_1 × 3 × 3’, ‘CDC-T-06_3 × 1 × 1’, ‘CDC-T-06_3 × 3 × 3’, ‘CDC-TR-03_3 × 1 × 1’ and ‘CDC-TR-03_3 × 3 × 3’. To be specific, ‘CDC-T-06’ and ‘CDC-TR-03’ denote the 3D-CDC-T with $\theta = 0.6$ and 3D-CDC-TR with $\theta = 0.3$, respectively. These settings of θ are based on the ablation study results in Section IV-C. We also consider a vanilla operation space with vanilla 3D convolutions instead of 3D-CDC for comparison.

The multi-rate backbones for each of modalities $\mathcal{M}$ to be searched have the architecture parameters $\alpha_b^{(i,j)}$. Each edge (i, j) of DAG represents the information flow from node n_i to node n_j , which consists of the candidate operations weighted by the architecture parameter $\alpha_b^{(i,j)}$. Specially, each edge (i, j) can be formulated by a function $\tilde{o}_b^{(i,j)}$ where $\tilde{o}_b^{(i,j)}(n_i) = \sum_{o_b \in \mathcal{O}_b} \eta_{o_b}^{(i,j)} \cdot o_b(n_i)$. Softmax $\eta_{o_b}^{(i,j)} = \frac{\exp(\alpha_{o_b}^{(i,j)})}{\sum_{o'_b \in \mathcal{O}_b} \exp(\alpha_{o'_b}^{(i,j)})}$ is

utilized to relax architecture parameter $\alpha_b^{(i,j)}$ into operation weight $o_b \in \mathcal{O}_b$. The intermediate node can be denoted as $n_j = \sum_{i < j} \tilde{o}_b^{(i,j)}(n_i)$. The output node n_{K-1} is depth-wise concatenation of all the intermediate nodes excluding the input nodes.

In the searching stage, cross-entropy loss is utilized for the training loss $\mathcal{L}_{train}$ and validation loss $\mathcal{L}_{val}$. Network parameters Φ_b and architecture parameters α_b are learned via solving the bi-level optimization problem:

$$\begin{aligned} \min_{\alpha_b} \mathcal{L}_{val}(\Phi_b^*(\alpha_b), \alpha_b), \\ \text{s.t. } \Phi_b^*(\alpha) = \arg \min_{\Phi_b} \mathcal{L}_{train}(\Phi_b, \alpha_b). \end{aligned} \quad (6)$$

When the searching is converged, we derive the final architectures via: 1) setting $o_b^{(i,j)} = \arg \max_{o_b \in \mathcal{O}_b, o_b \neq \text{zero}} \eta_{o_b}^{(i,j)}$, and 2) for each intermediate node, choosing two incoming edges with the two largest values of $\max_{o_b \in \mathcal{O}_b, o_b \neq \text{zero}} \eta_{o_b}^{(i,j)}$.

2) *Stage 2: Searching Lateral Connections for Multi-Rate Multi-Modal Networks:* The lateral connections from most existing multi-rate [15] or multi-modal [7], [9], [20] spatio-temporal networks are designed manually, which might be sub-optimal for information exchange. Here we propose to search best-suited lateral connections among rate-aware and modality-aware branches, intending to effectively explore RGB-D-temporal synergies. In the second stage, our goal is to search for lateral connections among the multi-rate and multi-modal branches. As most of the definitions and the search procedure are similar to those in the first stage, here we only show the two main differences from the first stage.

On one hand, the composition of architecture search space is different. As shown in Fig. 3 (see low-level connections for example), the lateral connections search space can be represented as a bidirectional graph of $K' = 6$ nodes (branches) within the modalities $\mathcal{M}$. Specially, the lateral connections from the lower frame rate branches to the higher ones are not considered because we assume that the lower frame rate branches always have less information than the higher ones. Thus, there are total 18 edges (lateral connections) inside the bidirectional graph and each edge consists of the candidate operations weighted by its corresponding architecture parameter α_c . The final output of each node is the depth-wise concatenation of all outputs of the incoming edges.

On the other hand, the design of the operation space is different. The operation space for lateral connections is denoted as $\mathcal{O}_c$, which consists of two parts: 1) ‘Zero’, ‘Conv_5 × 1 × 1’, ‘CDC-T-06_5 × 1 × 1’, ‘CDC-TR-03_5 × 1 × 1’ with $\text{stride} = (4, 1, 1)$ for the edges from the high frame rate branches to the low frame rate branches; and 2) ‘Zero’, ‘Conv_3 × 1 × 1’, ‘CDC-T-06_3 × 1 × 1’, ‘CDC-TR-03_3 × 1 × 1’ with $\text{stride} = (2, 1, 1)$ for the others. Specially, $\text{stride} = 1$ is utilized for edges between the branches of different modalities with same frame rate. When the search is converged, for each edge (i, j) , only the operation in $\mathcal{O}_c$ with the largest $\alpha_c^{(i,j)}$ is adopted. With the two-stage multi-rate and multi-modal NAS in Algorithm 1, both the final backbones and lateral connections are derived.

Fig. 3. Architecture search space in the second stage. Multi-modal and multi-rate frames are adopted as inputs. Here we utilize 3 branches with different frame rates for two modalities (RGB and depth), respectively. In this search stage, all cells are initialized with the searched structures in the first stage and then fixed. The architecture of the low-, mid- and high-level lateral connections to be searched can be shared or unshared.

IV. EXPERIMENTS

In this part, we first give details for benchmark datasets and experimental setup. Then, we thoroughly evaluate the impacts of 3D-CDC family, multi-rate configuration, and two-stage NAS. Finally, we evaluate and compare our results with state-of-the-art methods on three benchmark datasets.

A. Datasets

We evaluate our method on three widely used RGB-D gesture datasets: Chalearn IsoGD [2], [25], NVGesture [1] and EgoGesture [4] datasets. The **Chalearn IsoGD dataset** [2] contains 47,933 RGB-D gesture videos divided into 249 kinds of gestures and is performed by 21 individuals. The dataset contains 35878 training, 5784 validation, and 6271 testing samples. We follow the SOTA methods [36], [42], [51] to evaluate performance on the validation set. The **NVGesture dataset** [1] focuses on touchless driver controlling. It contains 1532 dynamic gestures fallen into 25 classes. It includes 1050 samples for training and 482 for testing. The videos are recorded with three modalities (RGB, depth, and infrared). For fair evaluations with SOTA methods, infrared modality is not used in our experiments. The **EgoGesture dataset** [4] is a large multi-modal egocentric hand gesture dataset. It contains 24,161 hand gesture clips of 83 classes of gestures, performed by 50 subjects. Videos in this dataset are captured with an Intel RealSense SR300 device in RGB-D modalities across multiple indoor and outdoor scenes.

B. Implementation Details

Our proposed method is implemented with Pytorch. Cell nodes $K = 4$ and $K' = 6$ are used as the default setting. The optical flow is extracted by pyflow [72] - a python wrapper for dense optical flow [71]. In the search phase, partial channel connection and edge normalization [57] are adopted. The initial channel numbers for low, mid, and high

Fig. 4. Impact of 3D-CDC for RGB, depth and optical flow.

frame rate branches are 24, 16, and 8, respectively, which double after searching. There are 8 cells for each branch in the search stage, which increases to 12 cells after searching. SGD optimizer with learning rate $lr=1e-2$ and weight decay $wd=5e-5$ is utilized when training the network weights. The architecture parameters are trained with Adam optimizer with $lr=6e-4$ and $wd=1e-3$. The lr decays with factor 0.5 in the 20^{th} epoch. We search 30 epochs on the training set of IsoGD dataset [2] with batch size 20 while architecture parameters are not updated in the first 10 epochs. Especially, $\mathcal{L}_{train}$ is calculated on the first half of the training set while $\mathcal{L}_{val}$ on the latter part. The whole two-stage NAS costs 12 days on four P100s. In the training phase, models are trained with SGD optimizer with initial $lr=1e-2$ and $wd=5e-5$. The lr decays with factor 0.1 when the validation accuracy has not improved within 3 epochs. Random horizontal flip and spatial crops are utilized for data augmentation. We train models with batch size 12 for maximum 80 epochs on four RTX-2080Ti GPUs.

C. Ablation Study

All ablation studies are trained from scratch and evaluated on the validation set of the IsoGD dataset.

Fig. 5. The ablation study of the multi-rate networks. ‘8 + 16 + 32 frames’ means that there are three branches with different frame rate (i.e., temporally downsampling to 8, 16 and 32 frames, respectively) as inputs. ‘CDC-T_0.6_RGB’ and ‘CDC-TR_0.3_Depth’ denotes using 3D-CDC-T with $\theta = 0.6$ for RGB inputs and 3D-CDC-TR with $\theta = 0.3$ for depth inputs, respectively.

1) *Impact of 3D-CDC for Modalities:* In these experiments, we use C3D [34] as the backbone and sequence size $3 \times 16 \times 112 \times 112$ as the inputs. According to Eq. (3), (4) and (5), θ controls the contribution of the temporal difference cues. As illustrated in Fig. 4, 3D-CDC-T improves the accuracy of RGB modality dramatically. Compared with vanilla 3D convolution (i.e., $\theta = 0$), 3D-CDC-T gains 7.3% when $\theta = 0.6$, which indicates the advantages of temporal difference context. One highlight is that, assembled with 3D-CDC-T, the RGB modality is able to obtain comparable accuracy (44.06% vs. 47.02%) with optical flow modality, indicating an excellent dynamic modeling capacity of 3D-CDC-T. In terms of the depth and optical flow modalities, the best performance (42.86% and 52%) could be achieved when using 3D-CDC-TR with $\theta = 0.3$ and $\theta = 0.4$, respectively. Compared with 3D-CDC-T, 3D-CDC-TR is more robust for depth and optical flow modalities because it alleviates sensor noises and pre-processing artifacts between frames in these two modalities. By the observation, the 3D-CDC-ST performs relatively poorly. The reason might be that the gesture recognition task prefers more temporal reasoning context than spatial gradient cues and appearance details. According to their enhanced temporal representation abilities, 3D-CDC-T with $\theta = 0.6$ and 3D-CDC-TR with $\theta = 0.3$ are considered in our NAS operation space.

2) *Impact of Multi-Rate Branches:* As gestures have various temporal scales, modeling such visual tempos of different gestures facilitates their recognition. Here we conduct the ablation study with C3D [34] network to explore how the branches with different frame rates influence the gesture recognition task. As illustrated in Fig. 5, for the single rate network, the higher the frame rate it has, the better performance it will be. This is because a higher frame rate usually has less sampling temporal information loss, which has richer fine-grained temporal cues for gesture recognition. Furthermore, we could find that the

TABLE I
COMPARISON AMONG VARIOUS CONFIGURATIONS OF THE TWO-STAGE NAS FOR VARIED MODALITIES. THE UPPER PART IS ABOUT THE FIRST STAGE NAS1 WHILE BOTTOM PART IS ABOUT THE SECOND STAGE NAS2. THE EVALUATION METRIC IS ACCURACY (%)

Configuration	RGB	Depth	RGB-D
w/o CDC, w/o NAS	44.07	43.10	-
w/o CDC, w/ NAS1	44.86	43.57	-
w/ CDC, w/o NAS	47.26	44.28	-
w/ CDC, w/ NAS1	48.38	45.08	-
w/ CDC, w/ NAS2_Fixed	-	-	42.62
w/ CDC, w/ NAS2_Shared	-	-	46.43
w/ CDC, w/ NAS2_Unshared	-	-	50.17

performance could be further improved when cooperated with the multi-rate branches (e.g., ‘16 + 32 frames’ and ‘8 + 16 + 32 frames’). In terms of the impact of multi-rate branches for different modalities, it is obvious that multi-rate branches contribute more to RGB than depth modality. When assembling with 3D-CDC-T for RGB or 3D-CDC-TR for depth, the trends of multi-rate networks are analogous as the vanilla cases but achieving holistic performance gains (due to the excellent representation capacity of 3D-CDC).

We also reimplement SlowFast [15] Networks (trained from scratch) with ‘8 + 32 frames’ multi-rate setting on the IsoGD dataset. However, it only achieves respective 22.28% and 40.69% accuracy on RGB and depth inputs, which indicates the importance of suitable architecture design for multi-rate networks in the gesture recognition task.

3) *Effectiveness of the Two-Stage NAS:* Based on the best multi-rate setting (‘8 + 16 + 32 frames’), we study the two-stage NAS for both single and multiple modalities. The first stage NAS (NAS1) intends to find well-suited multi-rate single-modal networks. As illustrated in Tab. I, when searching over the vanilla search space w/o CDC, the architectures found by NAS1 improve 0.79% and 0.47% accuracy (compared with multi-rate C3D) for RGB and depth inputs, respectively. Moreover, the gains consistently occur when searching over 3D-CDC based search space for both RGB (+1.12%) and depth (+0.8%) modalities.

Based on the searched multi-rate networks for RGB and depth modalities, ‘NAS2_Fixed’ utilizes late fusion directly without searching the lateral connections between two modalities. Out of expectation, it performs even worse than the single-modal NAS1 searched models. It means that simply late fusion will encounter the problem of insufficient information exchange in feature levels. With the second stage NAS (NAS2), ‘NAS2_Unshared’ achieves more than 50% accuracy, which indicates the advantages of NAS that mines the efficient integration of multi-rate and multi-modal branches. Furthermore, compared with ‘NAS2_Shared’ searching the shared lateral connections for low-mid-high levels, ‘NAS2_Unshared’ performs better (+3.74%), which implies the importance of the specific design for interactions of each level.

TABLE II
RESULTS ON THE VALIDATION SET OF ISOGD [2]

Method	Modality	Accuracy (%)
Hu <i>et al.</i> [73]	RGB	44.88
Res3D [8]	RGB	45.07
Zhang <i>et al.</i> [51]	RGB	51.31
Zhang <i>et al.</i> [11]	RGB	55.98
Zhu <i>et al.</i> [36]	RGB	57.42
NAS1 (Ours)	RGB	58.88
Narayana <i>et al.</i> [42]	Depth	27.98
Res3D [8]	Depth	48.44
Zhang <i>et al.</i> [51]	Depth	49.81
Zhu <i>et al.</i> [36]	Depth	54.18
NAS1 (Ours)	Depth	55.68
Zhu <i>et al.</i> [74]	RGB-D	51.02
Hu <i>et al.</i> [73]	RGB-D	54.14
Zhang <i>et al.</i> [51]	RGB-D	55.29
Zhu <i>et al.</i> [36]	RGB-D	61.05
HybridNet w/o hand [75]	RGB-D	61.14
3D-CDCN [76]	RGB-D	62.04
HybridNet w/ hand [75]	RGB-D	64.61
NAS2 (Ours)	RGB-D	60.04
NAS1+NAS2 (Ours)	RGB-D	65.54
Li <i>et al.</i> [43]	RGB-D-Flow	54.50
Zhang <i>et al.</i> [51]	RGB-D-Flow	58.65
Miao <i>et al.</i> [7]	RGB-D-Flow	64.40
AssembleNet [24]	RGB-Flow	53.30
AssembleNet++ [77]	RGB-Flow	54.70
AssembleNet [24]	Flow-Depth	52.00
AssembleNet++ [77]	Flow-Depth	53.60
NAS2 (Ours)	RGB-Flow	61.22
NAS2 (Ours)	Flow-Depth	62.47
NAS2_all (Ours)	RGB-D-Flow	66.23
FOANet w/o hand[42]	RGB-D-Flow-DFlow	61.40
FOANet w/ hand [42]	RGB-D-Flow-DFlow	80.96

D. Comparison With State-of-the-Art Methods

After studying the components in Sec. IV-C, we evaluate our models on three benchmark datasets. Note that in this subsection, our models are firstly pre-trained on the Jester [78] gesture dataset, which is similar to [36], [51], [79].

1) *Results on IsoGD*: As shown in Table II, although the existing methods [8], [11], [51] adopt 3DCNNs to learn from single RGB or depth modality, it is still challenging to represent the discriminative and robust spatio-temporal features with vanilla 3D convolutions and coarsely designed architecture. With the enhanced temporal representation capacity via 3D-CDC and multi-rate collaboration, our proposed multi-rate single-modal NAS method ‘NAS1’ obtains the best accuracy on every single modality. This exactly demonstrates the superiority of the searched architecture. In terms of the RGB-D gesture recognition, our searched architecture with two-stage NAS (NAS2) obtains more than 1% and 4% accuracy gains compared with the ‘NAS1’ with RGB and Depth modality, respectively. It demonstrates the effectiveness of RGB-D-temporal synergies at earlier stages. Similar to [36] ensembling the results from varied modalities, our ‘NAS1+NAS2’ boosts the accuracy to 65.54%.

To evaluate the modality generalization of the architecture searched from RGB-D, we retrain the same model ‘NAS2’ with RGB-Flow and Flow-Depth modalities separately and

TABLE III
RESULTS ON THE NVGESTURE [1] DATASET

Method	Modality	Accuracy (%)
HOG+HOG ² [3]	RGB	24.50
Simonyan <i>et al.</i> [32]	RGB	54.60
Wang <i>et al.</i> [80]	RGB	59.10
C3D [34]	RGB	69.30
R3DCNN [1]	RGB	74.10
GPM [81]	RGB	75.90
PreRNN [37]	RGB	76.50
Transformer [82]	RGB	76.50
ResNeXt-101 [79]	RGB	78.63
MTUT [10]	RGB	81.33
NAS1 (Ours)	RGB	83.61
HOG+HOG ² [3]	Depth	36.30
SNV [83]	Depth	70.70
C3D [34]	Depth	78.80
R3DCNN [1]	Depth	80.30
Transformer [82]	Depth	83.00
ResNeXt-101 [79]	Depth	83.82
PreRNN [37]	Depth	84.40
MTUT [10]	Depth	84.85
GPM [81]	Depth	85.50
NAS1 (Ours)	Depth	86.10
HOG+HOG ² [3]	RGB-D	36.90
I3D [21]	RGB-D	83.82
Transformer [82]	RGB-D	84.60
PreRNN [37]	RGB-D	85.00
MTUT [10]	RGB-D	85.48
GPM [81]	RGB-D	86.10
NAS2 (Ours)	RGB-D	86.93
NAS1+NAS2 (Ours)	RGB-D	88.38

also achieve comparable performance (61.22% and 62.47%, respectively). To our best knowledge, it is the first to explore the modality generalization issues for multi-modal NAS. We also compare with other multi-modal NAS-based architectures AssembleNet [24] and AssembleNet++ [77], which are initially designed for RGB-Flow based action recognition. Due to the task gap between action and gesture recognition, AssembleNet and AssembleNet++ perform poorly (53.3% and 54.7%, respectively) in RGB-Flow gesture recognition on IsoGD dataset. In terms of modality generalization, AssembleNet and AssembleNet++ achieve 52.0% and 53.6% accuracy on Flow-Depth modalities, indicating the moderate transferability of the architectures. In contrast, our proposed multimodal NAS2 (searched on RGB-D) could generalize well on both RGB-Flow (61.22%) and Flow-Depth (62.47%), which further demonstrates the excellent transferability of our method.

Finally, the best accuracy (66.23%) could be achieved when ensembling the scores from all three ‘NAS2’ models among RGB-D-Flow modalities. Although the FOANet [42] reports the best performance (80.96%) on IsoGD, the high accuracy is achieved by fusing 12 channels (i.e., global/left/right channels for four modalities) with manual hand detection. Note that without hand detection preprocessing, our ‘NAS2_all’ outperforms FOANet (66.23% vs. 61.4%) by a large margin using only RGB-D-Flow modalities. This exactly demonstrates the superiority of our searched architectures.

Fig. 6. The confusion matrices obtained by comparing the ground-truth labels and the predicted labels from the NAS1 and NAS2 networks trained on the NVGesture dataset. Best seen on the computer, in color and zoomed in.

TABLE IV
RESULTS ON THE EgoGesture [84] DATASET

Method	Modality	Accuracy (%)
VGG16+LSTM [32]	RGB	74.7
C3D+LSTM+RSTTM [34]	RGB	89.3
SeST [85]	RGB	93.20
ResNeXt-101 [79]	RGB	93.75
NAS1 (Ours)	RGB	93.31
VGG16+LSTM [83]	Depth	77.7
C3D+LSTM+RSTTM [1]	Depth	90.6
SeST [85]	Depth	93.35
ResNeXt-101 [79]	Depth	94.03
NAS1 (Ours)	Depth	94.13
VGG16+LSTM [21]	RGB-D	81.4
C3D+LSTM+RSTTM [10]	RGB-D	92.2
I3D [81]	RGB-D	92.78
MTUT ^F [81]	RGB-D	93.87
NAS2 (Ours)	RGB-D	94.38
NAS1+NAS2 (Ours)	RGB-D	95.52

2) *Results on NVGesture*: Table III compares the performance of our method with SOTA methods on the NVGesture dataset. It can be seen that our approach performs the best for both single-modal and multi-modal testing, which indicates 1) our searched architecture is able to represent discriminative spatio-temporal features for single/multi-modal gesture recognition; and 2) the architecture searched on the source dataset (IsoGD) via the proposed two-stage NAS transfers well on the target dataset (NVGesture), demonstrating the excellent generalization ability of the proposed NAS method.

Fig. 6 evaluates the coherence between the predicted labels from the searched ‘NAS1’ and ‘NAS2’ architectures, and the ground truths on the NVGesture dataset. The coherence is calculated by their confusion matrices. We observe that with RGB-D-temporal synergies, ‘NAS2’ has less confusion between the input classes and provides generally a more diagonalized confusion matrix. This improvement is better observed in the first three and last six classes.

TABLE V
RESULTS ON THE RGB-D ACTION RECOGNITION DATASET THU-READ [84] WITH THE CROSS-SUBJECT PROTOCOL. THE ARCHITECTURES OF ‘NAS1’ AND ‘NAS2’ ARE SEARCHED ON ISOGD AND THEN RETRAINED/EVALUATED ON THU-READ

Method	Modality	Accuracy(%)
HOG [86]	RGB	39.93
HOF [87]	RGB	46.27
Appearance Stream [88]	RGB	41.90
C3D [34]	RGB	66.25
SlowFast [15]	RGB	69.58
NAS1 (Ours)	RGB	71.25
HOG [86]	Depth	45.83
HOF [87]	Depth	43.96
Depth Stream [88]	Depth	34.06
C3D [34]	Depth	63.75
SlowFast [15]	Depth	68.75
NAS1 (Ours)	Depth	69.58
MDNN[84]	RGB-D-Flow	62.92
C3D [34]	RGB-D	75.83
SlowFast [15]	RGB-D	76.25
NAS2 (Ours)	RGB-D	73.85
NAS1+NAS2 (Ours)	RGB-D	78.38

Fig. 7. Impacts of the 3D-CDC family on two benchmark action recognition datasets (a) UCF101 [89], and (b) HMDB51 [90]. We use 3D-ResNet18 [35] as the backbone and sequence size $3 \times 16 \times 112 \times 112$ as the inputs. All experiments are trained from scratch for fair comparison.

3) *Results on EgoGesture*: We also evaluate the robustness of our searched architectures on egocentric (first-person view) gesture recognition. Compared with the third-person view gesture recognition, the main challenges include the complex scene background and motion blurriness caused by the subject walking. It can be seen in Table IV that our approach achieves the best performance (‘NAS2’ with 94.38% and ‘NAS1+NAS2’ with 95.52%) on the EgoGesture dataset using RGB-D modalities, which indicates the strong generalization ability of the proposed 3D-CDC and two-stage NAS

Fig. 8. The searched architecture from (a) the first stage NAS, and (b) the second stage NAS. The three rows in (a) represent the searched cell structure in the low, mid, and high frame branches, respectively.

method. Note that ResNeXt-101 [79] needs an extra detector to capture the key segments for preprocessing. As our proposed multi-rate and multi-modal network recognizes the gesture on original video clips directly, the performance might be further boosted with the gesture detector.

V. DISCUSSION AND VISUALIZATION

In this section, we first discuss the transferability of the proposed two-stage NAS and 3D-CDC on action recognition task, which is interesting and necessary because it exists task-oriented biases (gesture recognition is less relied on the

Fig. 9. Features visualization from C3D assembled with varied convolutions on the IsoGD dataset. With (a) RGB and (b) Depth modality inputs, the four rows represent the neural activation with 3D vanilla convolution, 3D-CDC-ST, 3D-CDC-T, and 3D-CDC-TR, respectively. Best view when zoom in.

scene but more related to the fine-grained temporal cues when compared with the action recognition). Then, we give the visualization of the searched architecture and feature response.

A. Task Transferability

1) *Generalization to RGB-D Action Recognition:* In order to validate the generalization ability of our 3D-CDC based two-stage NAS, we transfer the searched architecture ('NAS1' and 'NAS2') to another multi-modal video understanding task, i.e., RGB-D action recognition. Here one of the largest RGB-D egocentric dataset THU-READ [84] is used for experiments. It consists of 40 different actions and 1920 videos. We adopt the released leave-one-split-out protocol. For fair comparison, C3D [34], SlowFast [15], 'NAS1', and 'NAS2' are pre-trained on Jester gesture dataset. Table V compares the performance

of our method with SOTA methods on THU-READ. It can be seen that our approach outperforms the mainstream 3DCNNs (e.g., C3D [34] and SlowFast [15]) with a convincing margin, indicating that the architecture searched on the source task (gesture recognition) could be generalized well on the target video understanding task (e.g., action recognition).

2) *Impact of 3D-CDC for RGB Action Recognition:* Here we explore the effectiveness of 3D-CDC for scene-based action recognition tasks. Fig. 7 illustrates the results of two classical scene-related action datasets (UCF101 [89] and HMDB51 [90]). It is obvious that compared with the 3D vanilla convolution ($\theta = 0$), 3D-CDC-ST improves the performance dramatically in both datasets (+3% for UCF101 when $\theta = 0.6$ and +2.1% for HMDB51 when $\theta = 0.4, 0.6, 0.8$). The reason might be twofold. On one hand,

TABLE VI

COMPUTATIONAL COMPLEXITY OF THE SEARCHED ARCHITECTURES ON A SINGLE GeForce RTX 2080 Ti GPU. THE INPUT TENSOR SIZES FOR NAS1 AND NAS2 ARE $3 \times 32 \times 112 \times 112$ AND $6 \times 32 \times 112 \times 112$, RESPECTIVELY

Method	#Params	#FLOPS	Runtime
NAS1 (RGB)	127.1M	120.2G	172ms
NAS1 (Depth)	93.9M	105.0G	146ms
NAS2 (RGB-D)	251.4M	232.1G	364ms

an enhanced spatio-temporal context is helpful to represent scene-aware appearance and motions. On the other hand, the spatio-temporal difference term can be regarded as a regularization term to alleviate overfitting. Another highlight is that, without extra parameters, 3D-ResNet18 assembled with 3D-CDC-ST outperforms that with Spatio-Temporal Channel Correlation (STC) Block [38] by +2.1% on UCF101 split 1. In contrast, 3D-CDC-T performs the worst because of its weak spatial context representation capacity and vulnerability to the scene noises, which are vital in these two scene-aware datasets.

B. Visualization of the Searched Architecture

Fig. 8 shows the searched cells and lateral connections with the proposed two-stage NAS. It can be seen from Fig. 8(a) that there are more ‘CDC-T-06’ based operators in all three RGB branches while more ‘CDC-TR-03’ based operators in the depth branches. This consists with the observations in our ablation study of ‘Impact of 3D-CDC for Modalities’ in Section 4.1. In Fig. 8(b), it is interesting to find that the lower-level lateral connections are sparser (i.e., with more ‘Zero’ operators) and the high-level lateral connections have more learnable operators (i.e., convolution operators). This might inspire the video understanding community to design more reasonable multi-modal networks in the future.

C. Feature Visualization

The neural activation (before Pool3 in C3D) are visualized in Fig. 9. It can be seen from Fig. 9(a) that the proposed 3D-CDC-ST, 3D-CDC-T, and 3D-CDC-TR enhance the spatio-temporal representation and enforce the model to focus more on the trajectories of arms and hands. As for the depth modality shown in Fig. 9(b), all the convolutions are able to make the accurate attention on the movements from arms and hands due to the benefits from the foregrounds provided by the depth inputs. Despite more robust representation achieved by the 3D-CDC family, the interferences from the sensor-based noise and undesirable movements (e.g., head and hair) still occur. Thus, it is necessary to explore RGB-D-temporal synergies to overcome such limitation.

D. Model Size and Computational Cost

Table VI displays the model size and computational cost & runtime of the searched architectures. Compared with the RGB modality (the first row of Table VI), the searched architecture on Depth modality (the second row of Table VI) is more

lightweight and efficient, and has only about 85% model size as well as computational cost & runtime. It is reasonable as the information hidden in depth modality is relatively sparser and less than the RGB modality with full color channels. Both NAS1 and NAS2 models could be performed less than half a second on a single RTX 2080 Ti GPU with the input sizes $3 \times 32 \times 112 \times 112$ and $6 \times 32 \times 112 \times 112$, respectively. Despite with state-of-the-art performance, the searched architectures are large and inefficient, which limits the real-world applications in gesture recognition. Thus, our future work is to search more efficient mobile-level architectures with more compact 3D-CDC search space and size/latency regularization.

VI. CONCLUSION

We present a novel 3D convolution family called 3D-CDC for enhancing the spatio-temporal representation for video understanding tasks. Over the 3D-CDC search space, we propose a two-stage NAS method to discover well-suited multi-rate and multi-modal networks with RGB-D-temporal synergies. Extensive experiments show the effectiveness of our method. Future directions include: 1) exploring 3D-CDC family on other video understanding tasks (e.g., temporal localization); 2) searching temporal synergies with more modalities (e.g., audio and skeleton).

ACKNOWLEDGMENT

The authors wish to acknowledge the CSC-IT Center for Science, Finland, for computational resources.

REFERENCES

- [1] P. Molchanov, X. Yang, S. Gupta, K. Kim, S. Tyree, and J. Kautz, “Online detection and classification of dynamic hand gestures with recurrent 3D convolutional neural networks,” in *Proc. CVPR*, Jun. 2016, pp. 4207–4215.
- [2] J. Wan, S. Z. Li, Y. Zhao, S. Zhou, I. Guyon, and S. Escalera, “ChaLearn looking at people RGB-D isolated and continuous datasets for gesture recognition,” in *Proc. CVPR Workshops*, Jun. 2016, pp. 56–64.
- [3] E. Ohn-Bar and M. M. Trivedi, “Hand gesture recognition in real time for automotive interfaces: A multimodal vision-based approach and evaluations,” *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 6, pp. 2368–2377, Dec. 2014.
- [4] Y. Zhang, C. Cao, J. Cheng, and H. Lu, “EgoGesture: A new dataset and benchmark for egocentric hand gesture recognition,” *IEEE Trans. Multimedia*, vol. 20, no. 5, pp. 1038–1050, May 2018.
- [5] J. Weissmann and R. Salomon, “Gesture recognition for virtual reality applications using data gloves and neural networks,” in *Proc. IJCNN*, vol. 3, 1999, pp. 2043–2046.
- [6] S. S. Rautaray and A. Agrawal, “Vision based hand gesture recognition for human computer interaction: A survey,” *Artif. Intell. Rev.*, vol. 43, no. 1, pp. 1–54, Jan. 2015.
- [7] Q. Miao *et al.*, “Multimodal gesture recognition based on the ResC3D network,” in *Proc. ICCV Workshops*, Oct. 2017, pp. 3047–3055.
- [8] Y. Li *et al.*, “Large-scale gesture recognition with a fusion of RGB-D data based on the C3D model,” in *Proc. ICPR*, Dec. 2016, pp. 25–30.
- [9] H. Wang, P. Wang, Z. Song, and W. Li, “Large-scale multimodal gesture segmentation and recognition based on convolutional neural networks,” in *Proc. ICCV*, Oct. 2017, pp. 3138–3146.
- [10] M. Abavisani, H. R. V. Joze, and V. M. Patel, “Improving the performance of unimodal dynamic hand-gesture recognition with multimodal training,” in *Proc. CVPR*, Jun. 2019, pp. 1165–1174.
- [11] L. Zhang, G. Zhu, L. Mei, P. Shen, S. A. A. Shah, and M. Bennamoun, “Attention in convolutional LSTM for gesture recognition,” in *Proc. NeurIPS*, 2018, pp. 1953–1962.
- [12] S. Sun, Z. Kuang, L. Sheng, W. Ouyang, and W. Zhang, “Optical flow guided feature: A fast and robust motion representation for video action recognition,” in *Proc. CVPR*, Jun. 2018, pp. 1390–1399.

- [13] M. Lee, S. Lee, S. Son, G. Park, and N. Kwak, "Motion feature network: Fixed motion filter for action recognition," in *Proc. ECCV*, 2018, pp. 387–403.
- [14] A. Piergiovanni and M. S. Ryoo, "Representation flow for action recognition," in *Proc. CVPR*, Jun. 2019, pp. 9945–9953.
- [15] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "SlowFast networks for video recognition," in *Proc. ICCV*, Oct. 2019, pp. 6202–6211.
- [16] D. Zhang, X. Dai, and Y. F. Wang, "Dynamic temporal pyramid network: A closer look at multi-scale modeling for activity detection," in *Proc. Asian Conf. Comput. Vis.* Cham, Switzerland: Springer, Dec. 2018, pp. 712–728.
- [17] Y. Wang, M. Long, J. Wang, and P. S. Yu, "Spatiotemporal pyramid network for video action recognition," in *Proc. CVPR*, Jul. 2017, pp. 1529–1538.
- [18] N. Neverova, C. Wolf, G. Taylor, and F. Nebout, "ModDrop: Adaptive multi-modal gesture recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 8, pp. 1692–1706, Aug. 2016.
- [19] V. Pitsikalis, A. Katsamanis, S. Theodorakis, and P. Maragos, "Multi-modal gesture recognition via multiple hypotheses rescoring," in *Gesture Recognition*. Cham, Switzerland: Springer, 2017, pp. 467–496.
- [20] A. Roitberg, T. Pollert, M. Haurilet, M. Martin, and R. Stiefelhagen, "Analysis of deep fusion strategies for multi-modal gesture recognition," in *Proc. CVPR Workshops*, Jun. 2019, pp. 1–9.
- [21] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," in *Proc. CVPR*, Jul. 2017, pp. 6299–6308.
- [22] C. Li *et al.*, "Deep manifold structure transfer for action recognition," *IEEE Trans. Image Process.*, vol. 28, no. 9, pp. 4646–4658, Sep. 2019.
- [23] J.-M. Pérez-Rúa, V. Vielzeuf, S. Pateux, M. Baccouche, and F. Jurie, "MFAS: Multimodal fusion architecture search," in *Proc. CVPR*, Jun. 2019, pp. 6966–6975.
- [24] M. S. Ryoo, A. Piergiovanni, M. Tan, and A. Angelova, "AssembleNet: Searching for multi-stream neural connectivity in video architectures," in *Proc. ICLR*, 2020, pp. 1–15.
- [25] J. Wan *et al.*, "ChalLearn looking at people: IsoGD and ConGD large-scale RGB-D gesture recognition," 2019, *arXiv:1907.12193*. [Online]. Available: <http://arxiv.org/abs/1907.12193>
- [26] J. Wan, Q. Ruan, W. Li, G. An, and R. Zhao, "3D SMOsIFT: Three-dimensional sparse motion scale invariant transform for activity recognition from RGB-D videos," *J. Electron. Imag.*, vol. 23, no. 2, Apr. 2014, Art. no. 023017.
- [27] J. Wan, Q. Ruan, W. Li, and S. Deng, "One-shot learning gesture recognition from RGB-D data using bag of features," *J. Mach. Learn. Res.*, vol. 14, pp. 2549–2582, Sep. 2013.
- [28] M. R. Malgiredy, I. Inwogu, and V. Govindaraju, "A temporal Bayesian model for classifying, detecting and localizing activities in video sequences," in *Proc. CVPR Workshops*, Jun. 2012, pp. 43–48.
- [29] J. Wan, G. Guo, and S. Z. Li, "Explore efficient local features from RGB-D data for one-shot learning gesture recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 8, pp. 1626–1639, Aug. 2016.
- [30] X. Ji, J. Cheng, D. Tao, X. Wu, and W. Feng, "The spatial Laplacian and temporal energy pyramid representation for human action recognition using depth sequences," *Knowl.-Based Syst.*, vol. 122, pp. 64–74, Apr. 2017.
- [31] Z. Liu, X. Chai, Z. Liu, and X. Chen, "Continuous gesture recognition with hand-oriented spatiotemporal feature," in *Proc. ICCV Workshops*, Oct. 2017, pp. 3056–3064.
- [32] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. NeurIPS*, 2014, pp. 568–576.
- [33] P. Wang, W. Li, J. Wan, P. Ogunbona, and X. Liu, "Cooperative training of deep aggregation networks for RGB-D action recognition," in *Proc. AAAI*, 2018, pp. 1–8.
- [34] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3D convolutional networks," in *Proc. ICCV*, Dec. 2015, pp. 4489–4497.
- [35] K. Hara, H. Kataoka, and Y. Satoh, "Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet?" in *Proc. CVPR*, Jun. 2018, pp. 6546–6555.
- [36] G. Zhu *et al.*, "Redundancy and attention in convolutional LSTM for gesture recognition," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 31, no. 4, pp. 1323–1335, Apr. 2020.
- [37] X. Yang, P. Molchanov, and J. Kautz, "Making convolutional networks recurrent for visual sequence learning," in *Proc. CVPR*, Jun. 2018, pp. 6469–6478.
- [38] A. Diba *et al.*, "Spatio-temporal channel correlation networks for action classification," in *Proc. ECCV*, 2018, pp. 284–299.
- [39] S. Kumawat and S. Raman, "LP-3DCNN: Unveiling local phase in 3D convolutional neural networks," in *Proc. CVPR*, Jun. 2019, pp. 4903–4912.
- [40] Z. Yu *et al.*, "Searching central difference convolutional networks for face anti-spoofing," in *Proc. CVPR*, Jun. 2020, pp. 5295–5305.
- [41] Z. Yu *et al.*, "Multi-modal face anti-spoofing based on central difference networks," in *Proc. CVPR Workshops*, Jun. 2020, pp. 650–651.
- [42] P. Narayana, J. R. Beveridge, and B. A. Draper, "Gesture recognition: Focus on the hands," in *Proc. CVPR*, Jun. 2018, pp. 5235–5244.
- [43] Y. Li *et al.*, "Large-scale gesture recognition with a fusion of RGB-D data based on optical flow and the C3D model," *Pattern Recognit. Lett.*, vol. 119, pp. 187–194, Mar. 2019.
- [44] X. Liu, H. Shi, X. Hong, H. Chen, D. Tao, and G. Zhao, "3D skeletal gesture recognition via hidden states exploration," *IEEE Trans. Image Process.*, vol. 29, pp. 4583–4597, 2020.
- [45] X. Liu and G. Zhao, "3D skeletal gesture recognition via discriminative coding on time-warping invariant Riemannian trajectories," *IEEE Trans. Multimedia*, early access, Jun. 22, 2020, doi: [10.1109/TMM.2020.3003783](https://doi.org/10.1109/TMM.2020.3003783).
- [46] S. Mukherjee, L. Anvitha, and T. M. Lahari, "Human activity recognition in RGB-D videos by dynamic images," *Multimedia Tools Appl.*, vol. 79, nos. 27–28, pp. 19787–19801, Jul. 2020.
- [47] B. Zhang, Y. Yang, C. Chen, L. Yang, J. Han, and L. Shao, "Action recognition using 3D histograms of texture and a multi-class boosting classifier," *IEEE Trans. Image Process.*, vol. 26, no. 10, pp. 4648–4660, Oct. 2017.
- [48] Y. Li *et al.*, "Large-scale gesture recognition with a fusion of RGB-D data based on saliency theory and C3D model," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 28, no. 10, pp. 2956–2964, Oct. 2018.
- [49] J. Duan, J. Wan, S. Zhou, X. Guo, and S. Z. Li, "A unified framework for multi-modal isolated gesture recognition," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 14, no. 1, pp. 1–16, Apr. 2018.
- [50] G. Zhu, L. Zhang, L. Mei, J. Shao, J. Song, and P. Shen, "Large-scale isolated gesture recognition using pyramidal 3D convolutional networks," in *Proc. ICPR*, Dec. 2016, pp. 19–24.
- [51] L. Zhang, G. Zhu, P. Shen, J. Song, S. A. Shah, and M. Bennamoun, "Learning spatiotemporal features using 3DCNN and convolutional LSTM for gesture recognition," in *Proc. ICCV Workshops*, Oct. 2017, pp. 3120–3128.
- [52] X. Dong and Y. Yang, "Searching for a robust neural architecture in four GPU hours," in *Proc. CVPR*, 2019, pp. 1761–1770.
- [53] H. Liu, K. Simonyan, and Y. Yang, "Darts: Differentiable architecture search," in *Proc. ICLR*, 2019, pp. 1–13.
- [54] H. Pham, M. Y. Guan, B. Zoph, Q. Le, and J. Dean, "Efficient neural architecture search via parameters sharing," in *Proc. ICML*, 2018, pp. 4095–4104.
- [55] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," in *Proc. ICLR*, 2017, pp. 1–16.
- [56] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in *Proc. CVPR*, Jun. 2018, pp. 8697–8710.
- [57] Y. Xu *et al.*, "PC-DARTS: Partial channel connections for memory-efficient architecture search," in *Proc. ICLR*, 2019, pp. 1–13.
- [58] Z. Yu *et al.*, "Auto-Fas: Searching lightweight networks for face anti-spoofing," in *Proc. ICASSP*, May 2020, pp. 996–1000.
- [59] Z. Yu, X. Li, X. Niu, J. Shi, and G. Zhao, "AutoHR: A strong end-to-end baseline for remote heart rate measurement with neural searching," *IEEE Signal Process. Lett.*, vol. 27, pp. 1245–1249, 2020.
- [60] Z. Yu, J. Wan, Y. Qin, X. Li, S. Z. Li, and G. Zhao, "NAS-FAS: Static-dynamic central difference network search for face anti-spoofing," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Nov. 20, 2020, doi: [10.1109/TPAMI.2020.3036338](https://doi.org/10.1109/TPAMI.2020.3036338).
- [61] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 4780–4789.
- [62] E. Real *et al.*, "Large-scale evolution of image classifiers," in *Proc. ICML*, 2017, pp. 2902–2911.
- [63] H. Cai, L. Zhu, and S. Han, "ProxylessNAS: Direct neural architecture search on target task and hardware," in *Proc. ICLR*, 2019, pp. 1–13.
- [64] W. Peng, X. Hong, and G. Zhao, "Video action recognition via neural architecture searching," in *Proc. ICIP*, Sep. 2019, pp. 11–15.
- [65] Z. Qiu, T. Yao, Y. Zhang, Y. Zhang, and T. Mei, "Scheduled differentiable architecture search for visual recognition," 2019, *arXiv:1909.10236*. [Online]. Available: <http://arxiv.org/abs/1909.10236>

- [66] X. Wang *et al.*, “AttentionNAS: Spatiotemporal attention cell search for video classification,” 2020, *arXiv:2007.12034*. [Online]. Available: <http://arxiv.org/abs/2007.12034>
- [67] H. Zhang, Y. Hou, P. Wang, Z. Guo, and W. Li, “SAR-NAS: Skeleton-based action recognition via neural architecture searching,” *J. Vis. Commun. Image Represent.*, vol. 73, Nov. 2020, Art. no. 102942.
- [68] W. Peng, X. Hong, H. Chen, and G. Zhao, “Learning graph convolutional network for skeleton-based human action recognition by neural searching,” in *Proc. AAAI*, 2020, vol. 34, no. 3, pp. 2669–2676.
- [69] M. Tan and Q. V. Le, “MixConv: Mixed depthwise convolutional kernels,” 2019, *arXiv:1907.09595*. [Online]. Available: <http://arxiv.org/abs/1907.09595>
- [70] A. Howard *et al.*, “Searching for MobileNetV3,” in *Proc. ICCV*, Oct. 2019, pp. 1314–1324.
- [71] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert, “High accuracy optical flow estimation based on a theory for warping,” in *Proc. Eur. Conf. Comput. Vis.* Berlin, Germany: Springer, May 2004, pp. 25–36.
- [72] D. Pathak, R. Girshick, P. Dollár, T. Darrell, and B. Hariharan, “Learning features by watching objects move,” in *Proc. CVPR*, Jul. 2017, pp. 2701–2710.
- [73] T.-K. Hu, Y.-Y. Lin, and P.-C. Hsiu, “Learning adaptive hidden layers for mobile gesture recognition,” in *Proc. AAAI*, 2018, pp. 1–9.
- [74] G. Zhu, L. Zhang, P. Shen, and J. Song, “Multimodal gesture recognition using 3-D convolution and convolutional LSTM,” *IEEE Access*, vol. 5, pp. 4517–4524, 2017.
- [75] H. Wang, Z. Song, W. Li, and P. Wang, “A hybrid network for large-scale action recognition from RGB and depth modalities,” *Sensors*, vol. 20, no. 11, p. 3305, Jun. 2020.
- [76] A. Elboushaki, R. Hannane, K. Afdel, and L. Koutti, “MultiD-CNN: A multi-dimensional feature learning approach based on deep convolutional networks for gesture recognition in RGB-D image sequences,” *Expert Syst. Appl.*, vol. 139, Jan. 2020, Art. no. 112829.
- [77] M. S. Ryoo, A. J. Piergiovanni, J. Kangasputa, and A. Angelova, “AssembleNet++: Assembling modality representations via attention connections,” in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, Aug. 2020, pp. 654–671.
- [78] J. Materzynska, G. Berger, I. Bax, and R. Memisevic, “The jester dataset: A large-scale video dataset of human gestures,” in *Proc. ICCV Workshops*, Oct. 2019, pp. 1–9.
- [79] O. Köpüklü, A. Gunduz, N. Kose, and G. Rigoll, “Real-time hand gesture detection and classification using convolutional neural networks,” in *Proc. FG*, May 2019, pp. 1–8.
- [80] H. Wang, D. Oneata, J. Verbeek, and C. Schmid, “A robust and efficient video representation for action recognition,” *Int. J. Comput. Vis.*, vol. 119, no. 3, pp. 219–238, Sep. 2016.
- [81] V. Gupta, S. K. Dwivedi, R. Dabral, and A. Jain, “Progression modelling for online and early gesture detection,” in *Proc. 3DV*, Sep. 2019, pp. 289–297.
- [82] A. D’Eusano, A. Simoni, S. Pini, G. Borghi, R. Vezzani, and R. Cucchiara, “A transformer-based network for dynamic hand gesture recognition,” in *Proc. Int. Conf. 3D Vis.*, Nov. 2020, pp. 1–10.
- [83] X. Yang and Y. Tian, “Super normal vector for activity recognition using depth sequences,” in *Proc. CVPR*, Jun. 2014, pp. 804–811.
- [84] Y. Tang, Z. Wang, J. Lu, J. Feng, and J. Zhou, “Multi-stream deep neural networks for RGB-D egocentric action recognition,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 29, no. 10, pp. 3001–3015, Oct. 2019.
- [85] X. Tang, Z. Yan, J. Peng, B. Hao, H. Wang, and J. Li, “Selective spatiotemporal features learning for dynamic gesture recognition,” *Expert Syst. Appl.*, vol. 169, May 2021, Art. no. 114499.
- [86] H. Wang, M. M. Ullah, A. Klaser, I. Laptev, and C. Schmid, “Evaluation of local spatio-temporal features for action recognition,” in *Proc. Brit. Mach. Vis. Conf.*, 2009, pp. 1–12.
- [87] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld, “Learning realistic human actions from movies,” in *Proc. CVPR*, Jun. 2008, pp. 1–8.
- [88] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” 2014, *arXiv:1409.1556*. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [89] K. Soomro, A. Roshan Zamir, and M. Shah, “UCF101: A dataset of 101 human actions classes from videos in the wild,” 2012, *arXiv:1212.0402*. [Online]. Available: <http://arxiv.org/abs/1212.0402>
- [90] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre, “HMDB: A large video database for human motion recognition,” in *Proc. ICCV*, Nov. 2011, pp. 2556–2563.

Zitong Yu (Graduate Student Member, IEEE) received the M.S. degree from the University of Nantes, France, in 2016. He is currently pursuing the Ph.D. degree with Center for Machine Vision and Signal Analysis, University of Oulu, Finland. His research interests focus on remote physiological measurement, face anti-spoofing, and video understanding. He led the team and won the first place in the ChaLearn Multi-Modal Face Anti-Spoofing Attack Detection Challenge with CVPR 2020.

Benjia Zhou received the B.Eng. degree from the Minzu University of China, Beijing, China, in 2019. He is currently pursuing the M.S. degree with the Macau University of Science and Technology, Macau, China. His research interests focus on video recognition, sign language recognition, and neural architecture search.

Jun Wan (Senior Member, IEEE) received the B.S. degree from the China University of Geosciences, Beijing, China, in 2008, and the Ph.D. degree from the Institute of Information Science, Beijing Jiaotong University, Beijing, in 2015. Since January 2015, he has been a Faculty Member with the National Laboratory of Pattern Recognition (NLPR), Institute of Automation, Chinese Academy of Science (CASIA), China, where he currently serves as an Associate Professor. His main research interests include computer vision and machine learning. He is an Area Chair of ICME 2021. He served as the Co-Editor for special issues in IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE TRANSACTIONS ON BIOMETRICS, BEHAVIOR, AND IDENTITY SCIENCE, *Machine Vision and Applications*, and *Entropy*. He is an Associate Editor of the *IET Biometrics* from 2020 to 2022.

Pichao Wang (Member, IEEE) received the Ph.D. degree in computer science from the University of Wollongong, Australia. He is currently a Research Scientist with Alibaba DAMO Academy, USA. He has published 50+ peer-reviewed articles, including those in highly regarded journals and conferences, such as IEEE TRANSACTIONS ON MULTIMEDIA, IEEE TRANSACTIONS ON HUMAN-MACHINE SYSTEMS, CVPR, ICCV, AAAI, ACM MM, and so on. His current research interests include computer vision and machine learning. He has been an Associate Editor of *Computer Engineering* since 2019.

Haoyu Chen received the B.Sc. degree from the China University of Geosciences, Wuhan, China, in 2015, and the M.Sc. degree in computer sciences and engineering from the University of Oulu, Finland, in 2017, where he is currently pursuing the Ph.D. degree in computer science and engineering with Center for Machine Vision and Signal Analysis, under the supervision of Prof. G. Zhao. His research interests include action, gesture recognition, and emotional AI.

Xin Liu (Member, IEEE) received the Ph.D. degree in computer science from the University of Oulu, Finland, in 2019. He is currently an Associate Professor with the School of Electrical and Information Engineering, Tianjin University, China. His research interests include human behavior recognition and analysis, emotion understanding, and affective computing. He received the IEEE ICME Best Paper Award in 2017.

Guoying Zhao (Senior Member, IEEE) received the Ph.D. degree in computer science from the Chinese Academy of Sciences, Beijing, China, in 2005. She is currently a Professor with Center for Machine Vision and Signal Analysis, University of Oulu, Finland, where she has been a Senior Researcher, since 2005, and an Associate Professor, since 2014. She has authored or coauthored more than 240 articles in journals and conferences. Her articles have currently over 15 600 citations in Google scholar (H-index 57). She is a Fellow of the IAPR. She has served

as the area chairs for several conferences. She is an Associate Editor of *Pattern Recognition*, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, and *Image and Vision Computing* journal.

Stan Z. Li (Fellow, IEEE) received the B.Eng. degree from Hunan University, China, in 1982, the M.Eng. degree from the National University of Defense Technology, China, in 1985, and the Ph.D. degree from Surrey University, U.K., in 1991. He was a Researcher and the Director of the Center for Biometrics and Security Research, Institute of Automation, Chinese Academy of Sciences. He was a Researcher with Microsoft Research Asia and an Associate Professor with Nanyang Technological University, Singapore. He is a Chair Professor of

artificial intelligence with Westlake University, China. He has published over 500 articles with Google scholar index of over 47 000 and H-index of 127.