

Improving the Performance of Unimodal Dynamic Hand-Gesture Recognition with Multimodal Training

Mahdi Abavisani
Rutgers University

mahdi.abavisani@rutgers.edu

Hamid Reza Vaezi Joze
Microsoft

hava@microsoft.com

Vishal M. Patel
Johns Hopkins University

vpatel136@jhu.edu

Abstract

We present an efficient approach for leveraging the knowledge from multiple modalities in training unimodal 3D convolutional neural networks (3D-CNNs) for the task of dynamic hand gesture recognition. Instead of explicitly combining multimodal information, which is commonplace in many state-of-the-art methods, we propose a different framework in which we embed the knowledge of multiple modalities in individual networks so that each unimodal network can achieve an improved performance. In particular, we dedicate separate networks per available modality and enforce them to collaborate and learn to develop networks with common semantics and better representations. We introduce a “spatiotemporal semantic alignment” loss (SSA) to align the content of the features from different networks. In addition, we regularize this loss with our proposed “focal regularization parameter” to avoid negative knowledge transfer. Experimental results show that our framework improves the test time recognition accuracy of unimodal networks, and provides the state-of-the-art performance on various dynamic hand gesture recognition datasets.

1. Introduction

Recent advances in computer vision and pattern recognition have made hand gesture recognition an accessible and important interaction tool for different types of applications including human-computer interaction [36], sign language recognition [7], and gaming and virtual reality control [22]. In particular, recent developments in deep 3-D convolutional neural networks (3D-CNNs) with video sequences have significantly improved the performance of dynamic hand gesture recognition [27, 24, 26].

Most state-of-the-art hand gesture recognition methods exploit multiple sensors such as visible RGB cameras, depth camera or compute an extra modality like optical flow to improve their performances [21, 50, 47, 23]. Multi-

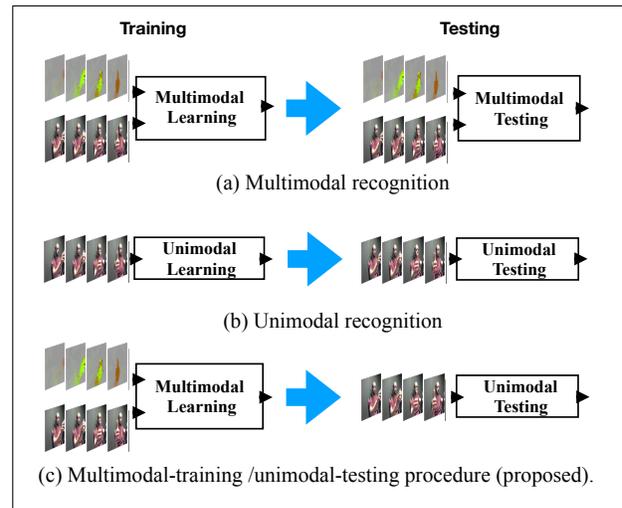


Figure 1. Training and testing schemes of different types of recognition systems. (a) The system is trained and tested with multiple modalities. (b) The system is trained and tested with a single modality. (c) The system leverages the benefits of multimodal training but can be ran as a unimodal system during testing.

modal recognition systems offer significant improvements to the accuracy of hand gesture recognition [25]. A multimodal recognition system is trained with multiple streams of data and classifies the multimodal observations during testing [35, 4] (Figure 1 (a)). On the other hand, a unimodal recognition system is trained and tested using only a single modality data (Figure 1 (b)). This paper introduces a third type of framework which leverages the knowledge from multimodal data during training and improves the performance of a unimodal system during testing. Figure 1 (c) gives an overview of the proposed framework.

The proposed approach uses separate 3D-CNNs per each stream of modality for primarily training them to recognize the dynamic hand gestures based on their input modality streams. The streams of modalities that are available in dynamic hand gesture recognition systems are often spatially and temporally aligned. For instance, the RGB and

depth maps captured with motion sensing devices and the optical flow calculated from the RGB streams are usually aligned. Hence, we encourage the individual modality networks to derive a common understanding for the spatiotemporal contents of different modalities. We do this by sharing their knowledge throughout the learning process by minimizing the introduced *spatiotemporal semantic alignment* (SSA) loss.

We further improve the learning process by regularizing the SSA loss with an adaptive regularization parameter. We call this regularization parameter, the *focal regularization parameter*. This parameter prevents the transfer of negative knowledge. In other words, it makes sure that the knowledge is transferred from more accurate modality networks to less accurate networks and not the other way. Once the networks are trained, during inference, each network has learned to recognize the hand gestures from its dedicated modality, but also has gained the knowledge transferred from the other modalities that assists in providing the better performance.

In summary, this paper makes the following contributions. First, we propose a new framework for single modality networks in dynamic hand gesture recognition task to learn from multiple modalities. This framework results in a *Multimodal Training / Unimodal Testing (MTUT)* scheme. Second, we introduce the SSA loss to share the knowledge of single modality networks. Third, we develop the *focal regularization parameter* for avoiding negative transfer. In our experiments, we show that learning with our method improves the test time performance of unimodal networks.

2. Related Work

Dynamic Hand Gesture Recognition: Dynamic hand-gesture recognition methods can be categorized on the basis of the video analysis approaches they use. Many hand-gesture methods have been developed based on extracting handcrafted features [48, 37, 42, 29]. These methods often derive properties such as appearance, motion cues or body-skeleton to perform gesture classification. Recent advances in action recognition methods and the introduction of various large video datasets have made it possible to efficiently classify unprocessed streams of visual data with spatiotemporal deep neural network architectures [9, 44, 38].

Various 3D-CNN-based hand gesture recognition methods have been introduced in the literature. A 3D-CNN-based method was introduced in [24] that integrates normalized depth and image gradient values to recognize dynamic hand gestures. In [25], a 3D-CNN was proposed that fuses streams of data from multiple sensors including short-range radar, color and depth sensors for recognition. A real-time method is proposed in [26] that simultaneously detects and classifies gestures in videos. Camgoz et al. [6] suggested a user-independent system based on the

spatiotemporal encoding of 3D-CNNs. Miao et al. proposed ResC3D [23], a 3D-CNN architecture that combines multimodal data and exploits an attention model. Furthermore, some CNN-based models also use recurrent architectures to capture the temporal information [50, 8, 11, 52].

The main focus of this paper is to improve the performance of hand gesture recognition methods that are built upon 3D-CNNs. As will be described later, we assume that our networks have 4-D feature maps that contain positional, temporal and channel dimensions.

Transfer Learning: In transfer learning, first, an agent is independently trained on a source task, then another agent uses the knowledge of the source agent by repurposing the learned features or transferring them to improve its learning on a target task [32, 43]. This technique has been shown to be successful in many different types of applications [5, 30, 19, 17, 49, 34]. While our method is closely related to transfer learning, our learning agents (i.e. modality networks) are trained simultaneously, and the transfer occurs both ways among the networks. Thus, it is better categorized as a multi-task learning framework [10, 31], where each network has three tasks of providing the knowledge to the other networks, receiving the knowledge from them, and finally classifying based on their dedicated input streams.

Multimodal Fusion: In multimodal fusion, the model explicitly receives the data from multiple modalities and learns to fuse them [28, 3, 33]. The fusion can be achieved at feature level (i.e. early fusion), decision level (i.e. late fusion) or intermediately [35, 2]. Once the model is trained, during testing, it receives the data from multiple modalities for classification [35, 28]. While our method is related to multimodal fusion, it is not a fusion method. We do not explicitly fuse the representations from different modalities. Instead, we improve the representation learning of our individual modality networks by leveraging the knowledge from different modalities. During inference, we do not necessarily need multiple modalities but rather each individual modality network works independently to classify data.

3. Proposed Method

In our proposed model, per each modality, one 3D-CNN is trained. Assuming that the stream of data is available in M modalities, we have M classifier networks with similar architectures that classify based on their corresponding input. During training, while each network is primarily trained with the data from its corresponding modality, we aim to improve the learning process by transferring the knowledge among the networks of different modalities. The transferred knowledge works as an extra supervision in addition to the class labels.

We share the knowledge of networks by aligning the semantics of the deep representations they provide for the inputs. We do this by selecting an in-depth layer in the net-

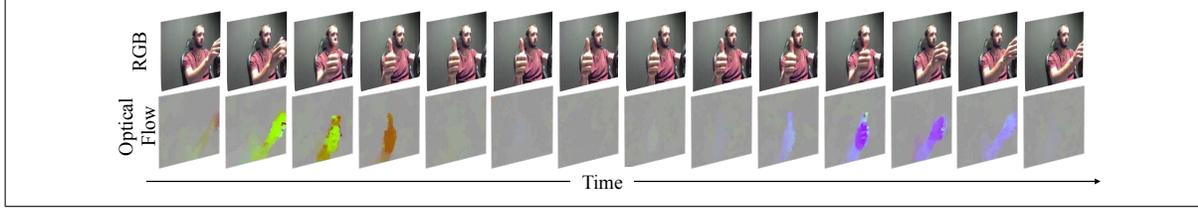


Figure 2. An example of the RGB and optical flow streams from the NVGesture Dataset [26]. As can be seen, while for the stationary frames RGB provides better representation, optical flow provides better representation for the dynamic frames.

works and enforcing them to share a common correlation across the in-depth layers of all the modality networks. This is done by minimizing the distance between their correlation matrices in the training stage. In addition, we regularize this loss by an adaptive parameter which ensures that the loss is serving as a one-way gate that only transfers the knowledge from more accurate modality networks to those with less accuracy, and not the other way.

3.1. Spatiotemporal Semantic Alignment

In an ideal case, all the M classifier modality networks of our model should have the same understanding for an input video. Even though they are coming in different modalities, their inputs are representing the same phenomena. In addition, since we assume that different modalities of the input videos are aligned over the time and spatial positions, in an ideal case the networks are expected to have the same understanding and share semantics for spatial positions and frames of the input videos across the different modalities. However, in practice, some spatiotemporal features may be better captured in one modality as compared to some other modalities. For instance, in the stream of visible RGB and optical flow frames shown in Figure 2, it can be observed that for static frames the RGB modality provides better information, while for dynamic frames optical flow has less noisy information. This results in different semantic understanding across the individual modality networks.

Thus, it is desirable to design a collaborative framework that encourages the networks to learn a common understanding across different modalities for the same input scene. This way, if in a training iteration one of the networks cannot learn a proper representation for a certain region or time in its feature maps, it can use the knowledge from the other networks to improve its representations. An iterative occurrence of this event during the training process leads the networks to develop better representations in a collaborative manner.

Let $\mathbf{F}_m, \mathbf{F}_n \in \mathbb{R}^{W \times H \times T \times C}$ be the in-depth feature maps of two networks corresponding to the m th modality and n th modality, where W, H, T and C denote width, heights, the number of frames and channels of the feature maps, respectively. An in-depth feature map should contain

high-level content representations (semantics) [16]. The element $\mathbf{f}_{i,j,t}^m \in \mathbb{R}^C$ in \mathbf{F}_m represents the content for a certain block of time and spatial position. It is reasonable to expect the network m to develop correlated elements in \mathbf{F}_m for spatiotemporal blocks with similar contents and semantics in the input. Thus, in an ideal case, the correlated elements in \mathbf{F}_m should have correlated counterpart elements in \mathbf{F}_n .

The correlations between all the elements of \mathbf{F}_m is expressed by its correlation matrix defined as follows

$$\text{corr}(\mathbf{F}_m) = \hat{\mathbf{F}}_m \hat{\mathbf{F}}_m^T \in \mathbb{R}^{d \times d}, \quad (1)$$

where $\hat{\mathbf{F}}_m \in \mathbb{R}^{d \times C}$ contains the normalized elements of \mathbf{F}_m in its rows, and $d = WHT$ is the number of elements in \mathbf{F}_m . The element $\mathbf{f}_{i,j,t}^m$ is normalized as $\hat{\mathbf{f}}_{i,j,t}^m = \tilde{\mathbf{f}}_{i,j,t}^m / \|\tilde{\mathbf{f}}_{i,j,t}^m\|$ where $\|\tilde{\mathbf{f}}_{i,j,t}^m\|$ is the magnitude of $\tilde{\mathbf{f}}_{i,j,t}^m$, and $\tilde{\mathbf{f}}_{i,j,t}^m$ is calculated by $\hat{\mathbf{f}}_{i,j,t}^m = \frac{\mathbf{f}_{i,j,t}^m - \mu_{i,j,t}}{\sigma_{i,j,t}}$, where $\mu_{i,j,t}$ and $\sigma_{i,j,t}$ are respectively the sample mean and variance of the element. We encourage the networks of the m th and the n th modalities to share a common correlation matrix for the feature maps of \mathbf{F}_m and \mathbf{F}_n so that they can have similar understanding for the input video while being free to have different styles. We do this by minimizing their *spatiotemporal semantic alignment* loss defined as

$$\ell_{SSA}^{m,n} = \rho^{m,n} \|\text{corr}(\mathbf{F}_m) - \text{corr}(\mathbf{F}_n)\|_F^2, \quad (2)$$

where $\rho^{m,n}$ is an adaptive regularization parameter defined in Section 3.2.

The *spatiotemporal semantic alignment* loss is closely related to the covariance matrix alignment of the source and target feature maps in domain adaptation methods [41, 40]. In addition, in some style transfer methods, the Gram matrices of feature maps are aligned [15, 16]. Aligning the Gram matrices, as opposed to our approach, discards the positional information and aligns the styles. In contrast, our method aligns the positional and temporal information and discards the style.

3.2. Avoiding Negative Transfer

As discussed earlier, some modalities may provide weak features as compared to the others. In addition, even the strong modalities may sometimes have corrupted or hard

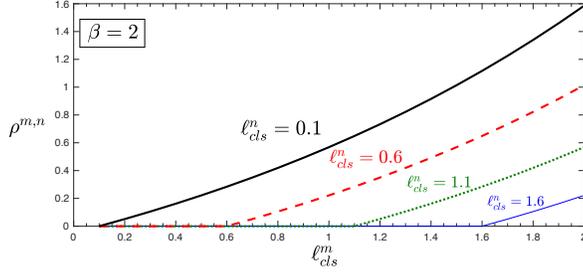


Figure 3. The value of focal regularization parameter ($\rho^{m,n}$) when $\beta = 2$ for different values of classification losses, ℓ_{cls}^m and ℓ_{cls}^n . Proportional to the classification performances of networks m and n , this parameter scales the SSA loss to focus on transferring positive knowledge.

examples in their training set. In these cases, aligning the spatiotemporal semantics of the representations from the other networks to the semantics of a weak network may lead to a decrease in the performance. In such a case, a negative transfer has occurred. It is desirable to develop a method that produces positive knowledge transfer between the networks while avoiding negative transfer. Such a method in our framework should enforce the networks to only mimic the semantics of more accurate networks in learning the representations for their hard examples. To address this issue, we regularize our SSA loss with an adaptive regularization parameter termed as *focal regularization parameter*. This parameter is denoted as $\rho^{m,n}$ in equation (2).

In order to measure the performance of the network modalities, we can use their classification loss values. Assume ℓ_{cls}^m and ℓ_{cls}^n are the classification losses of the networks m and n that respectively correspond to the m th and n th modalities. In addition, let $\Delta\ell = \ell_{cls}^m - \ell_{cls}^n$ be their difference. A positive $\Delta\ell$ indicates that network n works better than network m . Hence, in training the network m , for large positive values of $\Delta\ell$, we want large values for $\rho^{m,n}$ to enforce the network to mimic the representations of the network n . As $\Delta\ell \rightarrow 0^+$, network n becomes less an assist. Hence, we aim to have smaller $\rho^{m,n}$ s to focus more on the classification task. Finally, negative $\Delta\ell$ indicates that the network n does not have better representations than the network m , and therefore $\rho^{m,n}$ should be zero to avoid the negative transfer. To address these properties, we define the *focal regularization parameter* as follows

$$\rho^{m,n} = \mathbf{S}(e^{\beta\Delta\ell} - 1) = \begin{cases} e^{\beta\Delta\ell} - 1 & \Delta\ell > 0 \\ 0 & \Delta\ell \leq 0 \end{cases} \quad (3)$$

where β is a positive focusing parameter, and $\mathbf{S}(\cdot)$ is the thresholding function at zero.

Figure 3 visualizes values of $\rho^{m,n}$ for various ℓ_{cls}^m s and $\ell_{cls}^n \in [0, 2]$, when $\beta = 2$. As can be seen, the parameter is dynamically scaled, where the scaling factor decays to

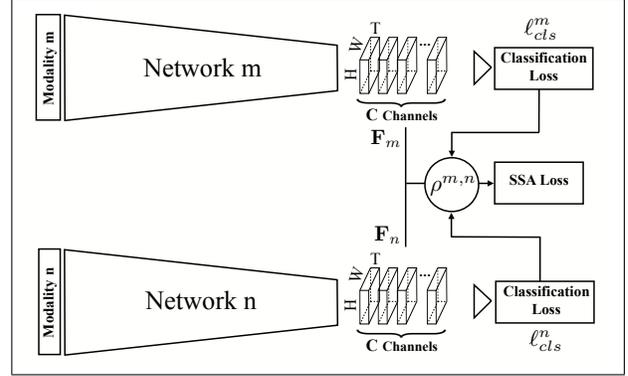


Figure 4. Training network m with the knowledge of network n . Training network m , is primarily done with respect to its classifier loss (ℓ_{cls}^m), but comparing with ℓ_{cls}^n , $\rho^{m,n}$ determines if involving the SSA loss is necessary, and if yes, it regularizes this loss with respect to the difference between the performances of two networks. Note that in the test time, both networks perform independently.

zero as confidence in the classification performance of the current network modality increases (measured using $\Delta\ell$). This scaling factor can automatically down-weight the contribution of the shared knowledge if the performance of the modality network n is degraded (measured by ℓ_{cls}^n).

The focal regularization parameter $\rho^{m,n}$ is used as the regularization factor when aligning the correlation matrix of \mathbf{F}^m in m th modality network to the correlation matrix of \mathbf{F}^n in n th modality network.

3.3. Full Objective of the Modality Networks

Combining the aforementioned objectives, our full objective for training the network corresponding to the m th modality in an M -modality task is as follows

$$\ell^m = \ell_{cls}^m + \lambda \sum_{n=1}^M \ell_{SSA}^{m,n} \quad (4)$$

where λ is a positive regularization parameter. Note that for $n = m$, $\rho^{m,n} = 0$ and thus $\ell_{SSA}^{m,n} = 0$.

Figure 4 shows an overview of how the representations for the n th modality affects on learning the representation in the m th modality. Since $\rho^{m,n}$ is differentiable, the training can be done in an end-to-end manner.

Our model encourages the networks to improve their representation learning in the training stage. During testing, each network performs separately. Thus, once the networks are trained, one can use an individual modality network to acquire efficient recognition. However, it is worth mentioning that with our framework, applying a decision level modality fusion in the test stage is also possible. In fact, our experiments show that the proposed method not only improves the performance of unimodal networks, but it can also improve the fusion performance.

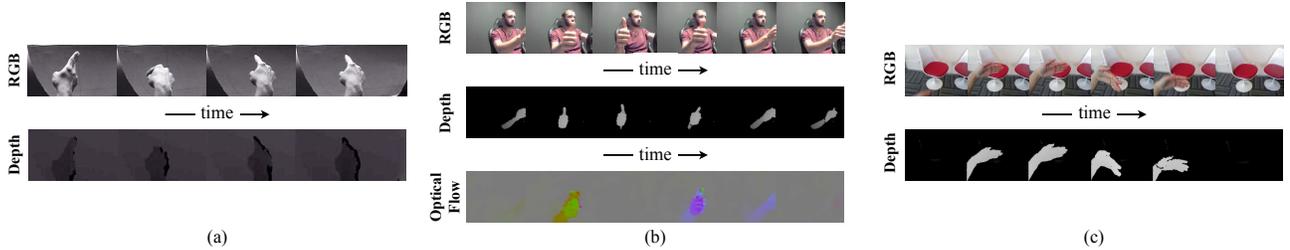


Figure 5. Sample sequences from different modalities of used datasets. (a) VIVA hand gesture dataset [29]. (b) NVGesture dataset [26]. (c) EgoGesture [8, 51]. As can be seen, the modalities in VIVA and EgoGesture datasets are well-aligned, while the depth map is not quite aligned with RGB and Optical flow maps in NVGesture.

4. Experimental Results

In this section, we evaluate our method against state-of-the-art dynamic hand gesture methods. We conduct our experiments on three publicly available multimodal dynamic hand gesture datasets. The following datasets are used in our experiments.

- *VIVA hand gestures dataset* [29] is a multimodal dynamic hand gesture dataset specifically designed with difficult settings of cluttered background, volatile illumination, and frequent occlusion for studying natural human activities in real-world driving settings. This dataset was captured using a Microsoft Kinect device, and contains 885 visible RGB and depth video sequences (RGB-D) of 19 hand gesture classes, collected from 8 subjects.
- *EgoGesture dataset* [8, 51] is a large multimodal hand gesture dataset collected for the task of egocentric gesture recognition. This dataset contains 24,161 hand gesture clips of 83 classes of gestures, performed by 50 subjects. Videos in this dataset include both static and dynamic gestures captured with an Intel RealSense SR300 device in RGB-D modalities across multiple indoor and outdoor scenes.
- *NVGestures dataset* [26] has been captured with multiple sensors and from multiple viewpoints for studying human-computer interfaces. It contains 1532 dynamic hand gestures recorded from 20 subjects inside a car simulator with artificial lighting conditions. This dataset includes 25 classes of hand gestures. The gestures were recorded with SoftKinetic DS325 device as the RGB-D sensor and DUO-3D for the infrared streams. In addition, the optical flow and infrared disparity map modalities can be calculated from the RGB and infrared streams, respectively. We use RGB, depth and optical flow modalities in our experiments. Note that IR streams in this dataset do not share the same view with RGB, depth and optical flow modalities. The optical flow is calculated using the method presented in [14].

Figure 5 (a), (b), and (c) show sample frames from the different modalities of these datasets that are used in our

experiments. Note that the RGB and depth modalities are well-aligned in the VIVA and EgoGesture datasets, but are not completely aligned in the NVGestures dataset.

For all the datasets, we compare our method against two state-of-the-art action recognition networks, I3D [9] and C3D [44], as well as state-of-the-art dynamic hand gesture recognition methods that were reported on the used datasets. In the tables, we report the results of our method as “*Multimodal Training Unimodal Testing*” (MTUT).

Implementation Details: In the design of our method, we adopt the architecture of I3D network as the backbone network of our modality networks, and employ its suggested implementation details [9]. This network is an inflated version of Inception-V1 [18], which contains several 3D convolutional layers followed with 3D max-pooling layers and inflated Inception-V1 submodules. The detailed architecture can be found in [9]. We select the output of the last inflated Inception submodule, “*Mixed_5c*”, as the in-depth feature map in our modality networks for applying the SSA loss (2). In all the experiments λ is set to 50×10^{-3} , and $\beta = 2$. The threshold function in the focal regularization parameter is implemented by a ReLU layer. For all the experiments with our method and I3D benchmarks, unless otherwise stated, we start with the publicly available ImageNet [12] + Kinetics [20] pre-trained networks.

We set the momentum to 0.9, and optimize the objective function with the standard SGD optimizer. We start with the base learning rate of 10^{-2} with a $10\times$ reduction when the loss is saturated. We use a batch size of 6 containing 64-frames snippets in the training stage. The models were implemented in Tensor-Flow 1.9 [1]. For our method, we start with a stage of pretraining with only applying the classification losses on the modality networks for 60 epochs, and then continue training with the SSA loss for another 15 epochs.

We employ the following spacial and temporal data augmentations during the training stage. For special augmentation, videos are resized to have the smaller video size of 256 pixels, and then randomly cropped with a 224×224 patch. In addition, the resulting video is randomly but consistently flipped horizontally. For temporal augmentation,

Method	Testing modality	
	RGB	Depth
HOG+HOG2 [29]	52.3	58.6
CNN:LRN [24]	57.0	65.0
C3D [44]	71.26	68.32
I3D [9]	78.25	74.46
MTUT (ours)	81.33	81.31

Table 1. 8-fold cross-subject average accuracies of different hand gesture methods on the VIVA hand gesture dataset [29]. The top performer is denoted by boldface.

64 consecutive frames are picked randomly from the videos. Shorter videos are randomly padded with zero frames on both sides to obtain 64 frames. During testing, we use 224×224 center crops, apply the models convolutionally over the full video, and average predictions.

Note that we follow the above mentioned implementation details identically for the experiments with both the I3D method [9], and our method. The only difference between the I3D method and our MTUT is in their learning objective. In our case, it consists of the introduced constraints as well.

4.1. VIVA Hand Gestures Dataset

In this set of experiments, we compare our method on the VIVA dataset against a hand-crafted approach (HOG+HOG2 [29]), a recurrent CNN-based method (CNN:LRN [24]), a C3D [44] model which were pre-trained on Sports-1M dataset [19] as well as the I3D method that currently holds the best results in action recognition [9]. All the results are reported by averaging the classification accuracies over 8-fold cross-subject cross-validation.

Table 1 shows the performance of the dynamic hand gesture methods tested on the visible and depth modalities of the VIVA dataset. As can be seen from this table, the I3D network performs significantly better than *HOG+HOG2* and *CNN:LRN*. This is in part due to the knowledge that I3D contains from its pretraining on ImageNet and Kinematic datasets. Nonetheless, we observe that our method that shares the same architecture and settings with the I3D networks and only differs in the learning procedure has significantly improved the I3D method by a 3.08% boost in the performance of RGB’s network and 6.85% improvement on the performance of the depth’s network. This experiment shows that our method is able to integrate the complementary information between two different modalities to learn efficient representations that can improve their individual performances.

4.2. EgoGesture Dataset

We assess the performance of our method along with various hand gesture recognition methods published on the

Method	Testing modality	
	RGB	Depth
VGG16 [39]	62.5	62.3
VGG16 + LSTM [13]	74.7	77.7
C3D [44]	86.4	88.1
C3D+LSTM+RSTTM [8]	89.3	90.6
I3D [9]	90.33	89.47
MTUT (ours)	92.48	91.96

Table 2. Accuracies of different hand gesture methods on the EgoGesture dataset [8]. The top performer is denoted by boldface.

large-scale hand gesture dataset, EgoGesture [8]. Table 2 compares unimodal test accuracies of different hand gesture methods. VGG16 [39] is a frame-based recognition method, and VGG16+LSTM [13] combines this method with a recurrent architecture to leverage the temporal information as well. As can be seen, the 3D-CNN-based methods, C3D, C3D+LSTM+RSTMM [8], and I3D, outperform the VGG16-based methods. However, among the 3D-CNN-based methods, our method outperforms the top performers in the RGB domain by 2.15% and in the Depth domain by 1.09%.

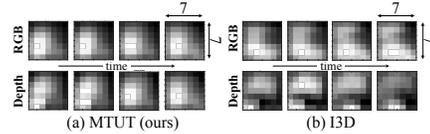


Figure 6. Visualization of the feature maps corresponding to the layer “Mixed_5c” in different networks for a sample input from EgoGesture dataset. These figures show the sequence of average feature maps (over 1024 channels) in (a) the RGB and depth networks trained with the I3D method. (b) the RGB and depth networks trained with our method. Intensity displays the magnitude.

In Figure 6, we visualize a set of feature maps from the RGB and depth networks trained with the I3D and our method. We feed a given input from the EgoGesture dataset to different networks and calculate the average of feature maps over the channels in the layer “Mixed_5c”. We display the resulting sequence in four 7×7 blocks. Here the temporal dimension is four and the spatial content is 7×7 . Layer “Mixed_5c” is the layer in the I3D architecture in which we apply the SSA loss to. We observe that the networks trained with our model have learned to detect similar structures for the given input (Figure 6 (a)). On the other hand, the networks trained with the I3D model are not bounded to develop similar structures. Thus, even though the input of the two modalities represent the same content, the feature maps may detect different structures (Figure 6 (b)).

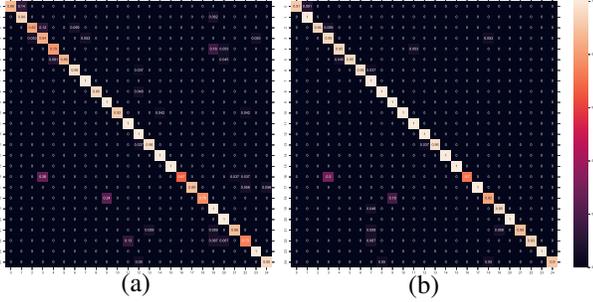


Figure 7. The confusion matrices obtained by comparing the grand-truth labels and the predicted labels from the RGB network trained on the NVGesture dataset by (a) I3D [9] model, and (b) our model. Best seen on the computer, in color and zoomed in.

4.3. NVGesture Dataset

In order to test our method on tasks with more than two modalities, in this section, we report the classification results on the RGB, depth and optical flow modalities of the NVGesture dataset [26]. The RGB and optical flow modalities are well-aligned in this dataset, however, the depth map includes a larger field of view (see Figure 5 (b)).

Table 3 tabulates the results of our method in comparison with the recent state-of-the-art methods: HOG+HOG2, improved dense trajectories (iDT) [46], R3DCNN [26], two-stream CNNs [38], and C3D as well as human labeling accuracy. The iDT [46] method is often recognized as the best performing hand-crafted method [45]. However, we observe that similar to the previous experiments the 3D-CNN-based methods outperform the other hand gesture recognition methods, and among them, our method provides the top performance in all the modalities. This table confirms that our method can improve the unimodal test performance by leveraging the knowledge from multiple modalities in the training stage. This is despite the fact that the depth map in this dataset is not completely aligned with the RGB and optical flow maps.

Figure 7 evaluates the coherence between the predicted labels and ground-truths in our method and compares it with I3D for the RGB modality of the NVGesture dataset. This coherence is calculated by their confusion matrices. We observe that our method has less confusion between the input classes and provides generally a more diagonalized confusion matrix. This improvement is better observed in the first six classes.

4.4. Effect of Unimodal Improvements on Multimodal Fusion

As previously discussed, our method is designed for embedding knowledge from multiple modalities in unimodal networks for improving their unimodal test performance. In this section, we examine if the enhanced unimodal networks trained by our approach can also improve the accuracy of a

Method	Testing modality		
	RGB	Depth	Opt. Flow
HOG+HOG2 [29]	24.5	36.3	-
Two Stream CNNs [38]	54.6	-	68.0
C3D [44]	69.3	78.8	-
iDT [46]	59.1	-	76.8
R3DCNN [26]	74.1	80.3	77.8
I3D [9]	78.42	82.28	83.19
MTUT (ours)	81.33	84.85	83.40
Human labeling accuracy:		88.4	

Table 3. Accuracies of different unimodal hand gesture methods on the NVGesture dataset [26]. The top performer is denoted by boldface.

decision level fusion that is calculated from the average of unimodal predictions. The decision level fusion of different modality streams is currently the most common fusion technique in the top performer dynamic action recognition methods [9, 44, 38].

In Table 4 and Table 5 we compare the multimodal-fusion versions of our method ($MTUT^F$) to state-of-the-art multimodal hand gesture recognition systems tested on the VIVA hand gesture and EgoGesture datasets, respectively. As can be seen, our method shows the top multimodal fusion performance on both datasets. These tables show that if multiple modalities are available at the test time, the improved performance of unimodal networks gained by training with our method can also result in an improved multimodal fusion performance in the test time.

Similarly, in Table 6 we report the multimodal fusion results on the NVGesture dataset. Note that since this dataset includes three modalities, based on the modalities we include in the training stage, we report multiple versions of our method. We report the version of our method that includes all three modalities in the training stage as $MTUT_{all}^F$, and the versions that only involve (RGB+Depth) and (RGB+Optical-Flow) in their training as $MTUT_{RGB-D}^F$ and $MTUT_{RGB-OF}^F$, respectively. While all versions of our method outperform the other multimodal fusion methods in Table 6, the performances of $MTUT_{RGB-D}^F$ and $MTUT_{all}^F$ in the fusion of RGB+Depth is worth highlighting. $MTUT_{all}^F$ in this experiment has also been trained on the absent modality, the optical flow, while $MTUT_{RGB-D}^F$ has been only trained on the RGB and Depth modalities. We observe that $MTUT_{all}^F$ has successfully integrated the knowledge of the absent modality and provided a better performance at the test time.

4.5. Analysis of the Network

To understand the effects of some of our model choices, we explore the performance of some variations of our model on the VIVA dataset. In particular, we compare our method with and without the *focal regularization parameter* and the

Method	Fused modalities	Accuracy
HOG+HOG2 [29]	RGB+Depth	64.5
CNN:LRN [24]	RGB+Depth	74.4
CNN:LRN:HRN [24]	RGB+Depth	77.5
C3D [44]	RGB+Depth	77.4
I3D [9]	RGB+Depth	83.10
MTUT ^F (ours)	RGB+Depth	86.08

Table 4. Accuracies of different multimodal fusion-based hand gesture methods on the VIVA dataset [29]. The top performer is denoted by boldface.

Method	Fused modalities	Accuracy
VGG16 [39]	RGB+Depth	66.5
VGG16 + LSTM [13]	RGB+Depth	81.4
C3D [44]	RGB+Depth	89.7
C3D+LSTM+RSTTM [8]	RGB+Depth	92.2
I3D [9]	RGB+Depth	92.78
MTUT ^F (ours)	RGB+Depth	93.87

Table 5. Accuracies of different multimodal fusion hand gesture methods on the EgoGesture dataset [26]. The top performer is denoted by boldface.

Method	Fused modalities	Accuracy
HOG+HOG2 [29]	RGB+Depth	36.9
I3D [9]	RGB+Depth	83.82
MTUT ^F _{RGB-D} (ours)	RGB+Depth	85.48
MTUT ^F _{all} (ours)	RGB+Depth	86.10
Two Stream CNNs [38]	RGB+Opt. flow	65.6
iDT [46]	RGB+Opt. flow	73.4
I3D [9]	RGB+Opt. flow	84.43
MTUT ^F _{RGB-OF} (ours)	RGB+Opt. flow	85.48
MTUT ^F _{all} (ours)	RGB+Opt. flow	85.48
R3DCNN [26]	RGB+Depth+Opt. flow	83.8
I3D [9]	RGB+Depth+Opt. flow	85.68
MTUT ^F _{all} (ours)	RGB+Depth+Opt. flow	86.93
Human labeling accuracy:		88.4

Table 6. Accuracies of different multimodal fusion hand gesture methods on the NVGesture dataset [8]. The top performer is denoted by boldface.

SSA loss. Beside our I3D-based method, we analyze these variations on a different backbone network, C3D [44] as well. C3D is another recently proposed activity recognition architecture. We name this method MTUT_{C3D}. Besides, we use C3D+SSA and I3D+SSA to refer to versions of our method with C3D and I3D backbones that contain a variation of the SSA loss that does not have the *focal regularization parameter*. For MTUT_{C3D} and C3D+SSA, we apply the SSA loss on feature maps of the last maxpooling layer (“MaxPool3d_5”).

To provide a fair comparison setting, we train these networks from scratch on the VIVA dataset, and report their performances in Table 7. As can be seen, the top performer is our I3D-based network with both SSA and focal regular-

Method	Testing modality	
	RGB	Depth
C3D	53.05	55.65
C3D+SSA	53.73	54.52
MTUT _{C3D}	56.56	58.71
I3D	65.72	67.30
I3D+SSA	65.83	66.96
MTUT	68.43	71.26

Table 7. Comparison of variations of MTUT with C3D and I3D backbones trained from scratch.

ization parameter. Several interesting observations can be made from the results in Table 7. As the table reveals, the I3D-based methods generally perform better than the C3D-based methods. This coincides with the previous reports [9]. In addition, C3D+SSA and I3D+SSA methods in the case of RGB networks show improvements and in the case of depth modality have comparable results as compared to their base networks C3D and I3D, respectively. However, the top performers in both modalities are the full version of our method applied on these base networks. This clearly shows the importance of our focal regularization parameter in avoiding negative transfer when transferring the knowledge between the modalities. Note that C3D, I3D and MTUT are trained from scratch in this experiment, while in the Table 1 we reported their performance on the networks trained with pre-trained weights.

5. Conclusion

We presented a new framework to leverage the knowledge of multiple modalities when training unimodal networks that can independently work at the test time inference with improved accuracy. Our model trains separate 3D-CNNs per available modalities, and shares their knowledge by the introduced *spatiotemporal semantic alignment* loss. We also regularized this loss with a *focal regularization parameter* that ensures that only positive knowledge is transferred between the modality networks, and negative transfer is avoided. Our experiments confirmed that our method can provide remarkable improvements to the unimodal networks at the test time. We also showed that the enhanced unimodal networks that are trained with our method can contribute to an improved multimodal fusion performance at test time as well.

The incorporation of our method for multimodal learning in other applications is a topic of further research.

References

- [1] Martın Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In *Operating Sys-*

- tems Design and Implementation*, volume 16, pages 265–283, 2016.
- [2] Mahdi Abavisani and Vishal M Patel. Deep multimodal subspace clustering networks. *IEEE Journal of Selected Topics in Signal Processing*, 12(6):1601–1614, 2018.
 - [3] Mahdi Abavisani and Vishal M Patel. Multimodal sparse and low-rank subspace clustering. *Information Fusion*, 39:168–177, 2018.
 - [4] Behnam Babagholami-Mohamadabadi, Seyed Mahdi Roostaiyan, Ali Zarghami, and Mahdiah Soleymani Baghshah. Multi-modal distance metric learning: Abayesian non-parametric approach. In *European Conference on Computer Vision*, pages 63–77. Springer, 2014.
 - [5] Yoshua Bengio. Deep learning of representations for unsupervised and transfer learning. In *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, pages 17–36, 2012.
 - [6] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, and Richard Bowden. Using convolutional 3d neural networks for user-independent continuous gesture recognition. In *Pattern Recognition (ICPR), 2016 23rd International Conference on*, pages 49–54. IEEE, 2016.
 - [7] Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, and Richard Bowden. Subunets: End-to-end hand shape and continuous sign language recognition. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 3075–3084. IEEE, 2017.
 - [8] Congqi Cao, Yifan Zhang, Yi Wu, Hanqing Lu, and Jian Cheng. Egocentric gesture recognition using recurrent 3d convolutional neural networks with spatiotemporal transformer modules. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3763–3771, 2017.
 - [9] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 4724–4733. IEEE, 2017.
 - [10] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
 - [11] Runpeng Cui, Hu Liu, and Changshui Zhang. Recurrent convolutional neural networks for continuous sign language recognition by staged optimization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
 - [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. Ieee, 2009.
 - [13] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015.
 - [14] Gunnar Farneback. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*, pages 363–370. Springer, 2003.
 - [15] Leon Gatys, Alexander S Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 262–270, 2015.
 - [16] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.
 - [17] Jui-Ting Huang, Jinyu Li, Dong Yu, Li Deng, and Yifan Gong. Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 7304–7308. IEEE, 2013.
 - [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456, 2015.
 - [19] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014.
 - [20] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
 - [21] Chenyang Li, Xin Zhang, and Lianwen Jin. Lpsnet: A novel log path signature feature based hand gesture recognition framework. In *Computer Vision Workshop (ICCVW), 2017 IEEE International Conference on*, pages 631–639. IEEE, 2017.
 - [22] Zhihan Lv, Alaa Halawani, Shengzhong Feng, Shafiq Ur Rehman, and Haibo Li. Touch-less interactive augmented reality game on vision-based wearable device. *Personal and Ubiquitous Computing*, 19(3-4):551–567, 2015.
 - [23] Qiguang Miao, Yunan Li, Wanli Ouyang, Zhenxin Ma, Xin Xu, Weikang Shi, Xiaochun Cao, Zhipeng Liu, Xiujuan Chai, Zhuang Liu, et al. Multimodal gesture recognition based on the resc3d network. In *ICCV Workshops*, pages 3047–3055, 2017.
 - [24] Pavlo Molchanov, Shalini Gupta, Kihwan Kim, and Jan Kautz. Hand gesture recognition with 3d convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 1–7, 2015.
 - [25] Pavlo Molchanov, Shalini Gupta, Kihwan Kim, and Kari Pulli. Multi-sensor system for driver’s hand-gesture recognition. In *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*, volume 1, pages 1–8. IEEE, 2015.
 - [26] Pavlo Molchanov, Xiaodong Yang, Shalini Gupta, Kihwan Kim, Stephen Tyree, and Jan Kautz. Online detection and classification of dynamic hand gestures with recurrent 3d convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4207–4215, 2016.

- [27] Natalia Neverova, Christian Wolf, Graham W Taylor, and Florian Nebout. Multi-scale deep learning for gesture detection and localization. In *Workshop at the European conference on computer vision*, pages 474–490. Springer, 2014.
- [28] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 689–696, 2011.
- [29] Eshed Ohn-Bar and Mohan Manubhai Trivedi. Hand gesture recognition in real time for automotive interfaces: A multimodal vision-based approach and evaluations. *IEEE transactions on intelligent transportation systems*, 15(6):2368–2377, 2014.
- [30] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1717–1724, 2014.
- [31] Poojan Oza and Vishal M Patel. Deep cnn-based multi-task learning for open-set recognition. *arXiv preprint arXiv:1903.03161*, 2019.
- [32] Sinno Jialin Pan, Qiang Yang, et al. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010.
- [33] Pramuditha Perera, Mahdi Abavisani, and Vishal M Patel. In2i: Unsupervised multi-image-to-image translation using generative adversarial networks. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 140–146. IEEE, 2018.
- [34] Pramuditha Perera and Vishal Patel. Deep transfer learning for multiple class novelty detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [35] Dhanesh Ramachandram and Graham W Taylor. Deep multimodal learning: A survey on recent advances and trends. *IEEE Signal Processing Magazine*, 34(6):96–108, 2017.
- [36] Siddharth S Rautaray and Anupam Agrawal. Vision based hand gesture recognition for human computer interaction: a survey. *Artificial Intelligence Review*, 43(1):1–54, 2015.
- [37] Xiaohui Shen, Gang Hua, Lance Williams, and Ying Wu. Dynamic hand gesture recognition: An exemplar-based approach from motion divergence fields. *Image and Vision Computing*, 30(3):227–235, 2012.
- [38] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014.
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [40] Baochen Sun, Jiashi Feng, and Kate Saenko. Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*, pages 153–171. Springer, 2017.
- [41] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*, pages 443–450. Springer, 2016.
- [42] Amir Tamrakar, Saad Ali, Qian Yu, Jingen Liu, Omar Javed, Ajay Divakaran, Hui Cheng, and Harpreet Sawhney. Evaluation of low-level features and their combinations for complex event detection in open source videos. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3681–3688. IEEE, 2012.
- [43] Lisa Torrey and Jude Shavlik. Transfer learning. In *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*, pages 242–264. IGI Global, 2010.
- [44] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [45] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018.
- [46] Heng Wang, Dan Oneata, Jakob Verbeek, and Cordelia Schmid. A robust and efficient video representation for action recognition. *International Journal of Computer Vision*, 119(3):219–238, 2016.
- [47] Huogen Wang, Pichao Wang, Zhanjie Song, and Wanqing Li. Large-scale multimodal gesture recognition using heterogeneous networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3129–3137, 2017.
- [48] Jiang Wang, Zicheng Liu, Jan Chorowski, Zhuoyuan Chen, and Ying Wu. Robust 3d action recognition with random occupancy patterns. In *Computer vision—ECCV 2012*, pages 872–885. Springer, 2012.
- [49] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, pages 3320–3328, 2014.
- [50] Liang Zhang, Guangming Zhu, Peiyi Shen, Juan Song, Syed Afaq Shah, and Mohammed Bennamoun. Learning spatiotemporal features using 3dcnn and convolutional lstm for gesture recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3120–3128, 2017.
- [51] Yifan Zhang, Congqi Cao, Jian Cheng, and Hanqing Lu. Egogesture: A new dataset and benchmark for egocentric hand gesture recognition. *IEEE Transactions on Multimedia*, 20(5):1038–1050, 2018.
- [52] Guangming Zhu, Liang Zhang, Peiyi Shen, and Juan Song. Multimodal gesture recognition using 3-d convolution and convolutional lstm. *IEEE Access*, 5:4517–4524, 2017.