

Learning Spatiotemporal Features for Infrared Action Recognition with 3D Convolutional Neural Networks

Zhuolin Jiang, Viktor Rozgic, Sancar Adali
Raytheon BBN Technologies, Cambridge, MA, 02138
{zjiang, vrozgic, sadali}@bbn.com

Abstract

Infrared (IR) imaging has the potential to enable more robust action recognition systems compared to visible spectrum cameras due to lower sensitivity to lighting conditions and appearance variability. While the action recognition task on videos collected from visible spectrum imaging has received much attention, action recognition in IR videos is significantly less explored. Our objective is to exploit imaging data in this modality for the action recognition task.

In this work, we propose a novel two-stream 3D convolutional neural network (CNN) architecture by introducing the discriminative code layer and the corresponding discriminative code loss function. The proposed network processes IR image and the IR-based optical flow field sequences. We pretrain the 3D CNN model on the visible spectrum Sports-1M action dataset and finetune it on the Infrared Action Recognition (InfAR) dataset. To our best knowledge, this is the first application of the 3D CNN to action recognition in the IR domain. We conduct an elaborate analysis of different fusion schemes (weighted average, single and double-layer neural nets) applied to different 3D CNN outputs. Experimental results demonstrate that our approach can achieve state-of-the-art average precision (AP) performances on the InfAR dataset: (1) the proposed two-stream 3D CNN achieves the best reported 77.5% AP, and (2) our 3D CNN model applied to the optical flow fields achieves the best reported single stream 75.42% AP.

1. Introduction

Deep convolutional neural networks (CNN) have shown remarkable success for various computer vision tasks in static images, such as object detection [6], recognition [10] and segmentation [16]. Encouraged by this success, researchers have proposed some CNN-based algorithms for action recognition in visible spectrum videos [28, 11, 3, 22, 27, 18, 4, 8]. One promising approach is using a two-stream CNN architecture developed by [22], which consists of a spatial stream network for learning salient appearance features from video frames, and a temporal stream network

for learning motion patterns. The prediction is computed by averaging the outputs of two networks. This architecture showed improved performance over traditional action recognition approaches such as improved dense trajectories features [24]. However, as pointed out by [23], 2D convolutions in a temporal network applied on the multi-frame stacking of optical flow fields (treating them as different channels) generate 2D representations; and the temporal network loses temporal information important for action recognition after the first convolution layer. To address this, [23] introduced a 3D CNN which takes multiple RGB frames as inputs and performs 3D convolution and pooling, preserving temporal information. The 3D CNN models can process appearance and motion information simultaneously, hence it is able to learn spatiotemporal features for action recognition.

For action recognition in infrared videos, there is limited work that uses deep CNNs to combine spatial and temporal cues for spatiotemporal feature learning [5, 29]. [5] applied two-stream CNNs for infrared action recognition. However, the CNN stream that processes the infrared image sequence achieved worse performance than several hand-crafted low-level features such as spatio-temporal interest point [12] and dense SIFT [17]. There are two potential reasons for this: 1) the infrared InfAR dataset is not large enough to learn spatiotemporal features leading to severe overfitting; and 2) 2D CNN loses temporal information contained in an input video volume as pointed in [23], so it can not properly model the temporal action patterns.

In this paper, we propose a two-stream 3D CNN architecture to learn spatio-temporal features for infrared action recognition. The two-stream 3D CNN contains two separate recognition networks (IR and optical-flow nets) combined using late fusion. In order to reduce the chance of overfitting and learn discriminative spatio-temporal features, we incorporate a discriminative code loss introduced in [8], and combine it with softmax classification loss to form the objective function used for network training. For faster convergence during training, we pretrained 3D CNN model parameters on the large-scale Sports-1M action dataset [9]

with videos from the visible light spectrum and finetuned them on the infrared dataset. The results are surprisingly good.

Our main contributions are the following:

- We develop a two-stream 3D CNN architecture to learn spatiotemporal features from infrared videos. This two-stream model learns representations that capture spatial and temporal information simultaneously.
- We add a discriminative code layer (DCL) on top of the last fully-connected layer and combine the discriminative code loss with softmax classification loss to train the 3D CNN. This discriminative code layer generates class-specific representations for infrared videos.
- We achieve state-of-the-art performances on the InfAR dataset. We find that a single 3D CNN with DCL layer trained using the optical flow field images can achieve an excellent infrared action recognition performance.

1.1. Related work

Many popular video feature extraction and classification approaches have been developed for action recognition in the visible spectrum, including low-level features (e.g., spatio-temporal interest point (STIP) [12], scale-invariance feature transform (SIFT) [17], optical flow fields [14], improved dense trajectory feature (IDT) [24]), and high-level semantic concepts (e.g., human action attributes [15, 20] and action parts [26]). In recent years, approaches based on convolutional neural networks have been proposed for action recognition [11, 3, 22, 27, 18, 4, 25]. The two-stream CNN architecture [22] achieved impressive recognition performances. [27] explored deeper two-stream network architectures and refined technical details yielding further performance improvements. [4] investigated an effective fusion strategy over space and time for two stream networks. However, all state-of-the-art approaches focus on the action recognition in the visible light spectrum.

Compared with visible spectrum imaging, the infrared imaging has a nice property that it can work well even under poor light conditions, which is useful for nighttime surveillance [29, 5, 23]. Still, limited works address the action recognition in infrared spectrum. [29] trained a SVM classifier trained on visible light spectrum based bag-of-visual-words video representation and adapted it to the infrared domain. Recently, [5] applied a two-stream 2D CNN architecture to infrared videos. This architecture employs a motion-history-image (MHI) stream network and an optical-flow stream network to extract image-level features. When the MHI stream is replaced with the raw infrared image stream, the action recognition performance becomes very poor, showing that the two-stream 2D CNN highly relies on the MHI stream network.

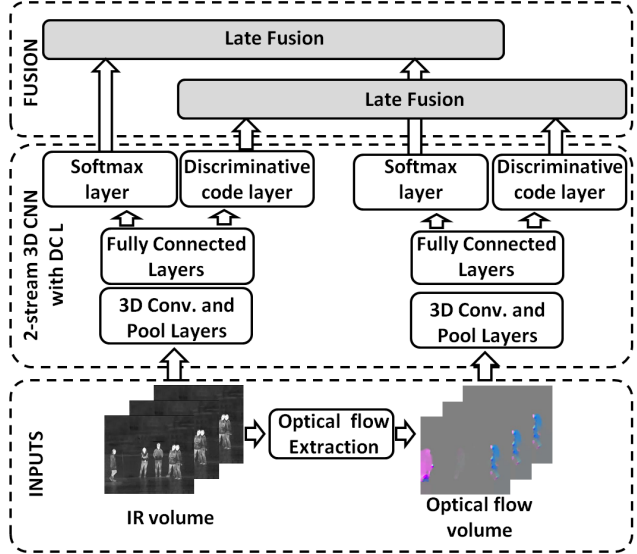


Figure 1. The IR action recognition pipeline: a two-stream 3D CNN trained with softmax classification loss and discriminative code loss, and their outputs aggregated using late fusion.

Different to [5], we introduce a two-stream 3D CNNs, which can model the video appearance and motion information simultaneously from an infrared video. To our best knowledge, 3D CNN architectures have not been explored for action recognition in infrared videos. In addition, we integrated the discriminative code loss from [8] to the objective function of network training. This makes the learned representations more discriminative and good for classification tasks, specifically action recognition. These two sources of novelties enable an action recognition system that advances state-of-the-art performance for action recognition in IR videos, as evidenced in our experiments.

2. Learning Spatiotemporal Features with 3D Convolutional Neural Networks

Our pipeline for action recognition in IR videos is presented in Figure 1. The pipeline inputs are IR video clips obtained by splitting IR frame sequences into non-overlapping segments of consecutive frames. Motivated by the success of the two-stream CNN architectures operating on visible spectrum videos and derived optical flow fields, we convert each consecutive pair of IR frames into optical flow field. We resize the IR and the optical flow images to the same height and width before feeding the IR and the optical flow clips to their corresponding networks. In the InfAR dataset experiments, we used a frame size of 128×171 , creating the input IR and optical flow clips with the same dimension $3 \times t \times 128 \times 171$, where 3 is the number of channels of an IR or flow image¹ and t is the temporal clip

¹We compute the optical flow image according to [2]. We stack the x, y components of a flow field, and compute its magnitude as the third channel

length in number of frames.

We process the corresponding IR and optical flow clips using a novel two-stream 3D CNN architecture. Both streams are processed using the same CNN model which extends the existing 3D-CNN architecture [23] by introduction of the additional *discriminative code layer*. The *discriminative code loss* associated with the discriminative code layer is combined with the softmax classification loss to train the 3D CNN.

We fuse the probabilistic outputs from the softmax (or the discriminative code) layers of the proposed two-stream network by using the weighted average, the single-layer neural network (NN) fusion or the two-layer NN fusion.

2.1. 3D Convolutional Neural Network with Discriminative Code Layer

Compared to 2D CNN architecture, 3D CNN architecture is better suited for the action recognition task, because it models spatial and temporal information jointly using 3D convolution and 3D pooling operations. While 2D convolution architectures process multiple input frames as different input channels and transform them into 2D representations; 3D convolution transforms input volume into 3D representation preserving temporal information. To our best knowledge, 3D CNNs have not yet been applied to the task of action recognition in IR videos.

Our two-stream 3D CNN architecture is based on the 3D CNN model with the *discriminative code layer* presented in Figure 2. The 3D CNN follows the architecture proposed in [23]. The network has eight 3D convolution layers, combined using five 3D max-pooling layers. The last pooling layer is followed by two fully connected layers (fc6 and fc7). The details on sizes of convolutional kernels, numbers of filters in different convolutional layers, sizes of the max-pooling and fully connected layers are provided in Section 3.1.

The proposed 3D CNN architecture has two output layers, a *softmax layer* that generates m -dimensional one-hot encoding of m activity categories and a *discriminative code layer* that generates discriminative codes for inputs signals. Assuming the discriminative code layer with N neurons, the training goal is to make it generate N dimensional p -hot encoding² of activity categories. The p -hot encoding represents a sample from k^{th} action category ($k = 1, \dots, m$) as a binary vector with coordinates $[p(k-1) + 1 : pk]$ equal to one and rest of coordinates equal to zero. Intuitively, the group of neurons activates only when a sample from the corresponding category is presented. In order to achieve this,

of the flow image. The flow values of x, y components are centered around 128 and scaled such that flow values fall between 0 and 255.

²here we assume that N can be exactly divided by m , i.e., $N = pm$. If $N = pm + k$ ($k < m$) we allocate the remaining k neurons to classes with high intra-class appearance variations.

we introduce the *discriminative code loss* associated with the discriminative code layer activations. This loss encourages groups of output neurons to activate simultaneously encoding the category label.

Let's assume that the 3D CNN architecture has $n+1$ layers, n levels including all convolution layers, pooling layers and fully connected layers, and the layer $n+1$ including the softmax layer and discriminative code outputs. The output of the i^{th} layer is denoted as $\mathbf{x}^{(i)}$, where $\mathbf{x}^{(0)}$ represents input. Therefore, the network architecture can be concisely expressed as:

$$\mathbf{x}^{(i)} = F(\mathbf{W}^{(i)}\mathbf{x}^{(i-1)}), \quad i = 1, 2, \dots, n \quad (1)$$

$$\mathbf{x}_d^{(n+1)} = \mathbf{A}\mathbf{x}^{(n)} \quad (2)$$

$$\mathbf{x}_c^{(n+1)} = \text{softmax}(\mathbf{W}\mathbf{x}^{(n)}) \quad (3)$$

where $\mathbf{W}^{(i)}$ represents the network parameters³ of the i^{th} layer for convolution and fully-connected layers, $\mathbf{W}^{(i)}\mathbf{x}^{(i-1)}$ is a linear operation (e.g. convolution in a convolution layer, or a linear transformation in fully-connected layer), $F(\cdot)$ is a non-linear activation function (e.g. ReLU). \mathbf{A} contains parameters of a linear transformation implemented by the *discriminative code layer* and \mathbf{W} contains the *softmax layer* parameters. $\mathbf{x}_d^{(n+1)}$ is the *predicted code* while $\mathbf{x}_c^{(n+1)}$ is the predicted class score vector.

The overall loss function used in network training is a linear combination of the softmax classification loss (multinomial logistic loss) L_c , and *discriminative code loss* L_d with a cost-balancing hyper parameter α .

$$L = L_c + \alpha L_d \quad (4)$$

$$L_c = L_c(\mathbf{x}_c^{(n+1)}, y) \quad (5)$$

$$L_d = L_d(\mathbf{x}_d^{(n+1)}, y) \quad (6)$$

The L_d cost component can be defined as:

$$L_d = \|\mathbf{q}^{(n)} - \mathbf{A}\mathbf{x}^{(n)}\|_2^2, \quad (7)$$

where the binary vector $\mathbf{q}^{(n)} = [q_1^{(n)}, \dots, q_j^{(n)}, \dots, q_N^{(n)}]^T \in \{0, 1\}^N$ denotes the p -hot label encoding (or target discriminative code), which indicates the ideal activations of neurons (j denotes the index of neuron). Each neuron is associated with a certain class label and, ideally, only activates to samples from that class. Therefore, when a sample is from the k^{th} action category, $q_j^{(n)} = 1$ if and only if the j^{th} neuron is assigned to class k , and neurons associated to other classes should not be activated so that the corresponding entries in $\mathbf{q}^{(n)}$ are zero. Note that \mathbf{A} is the only parameter³ to be learned in this cost component.

³For simplicity, we ignore the bias term for convolution layers, fully connected layers, softmax layer and discriminative code layer.

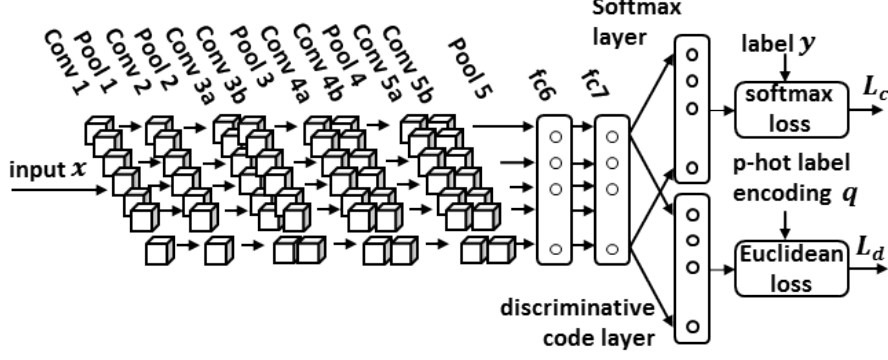


Figure 2. 3D CNN is trained using the cost function that combines the softmax classification loss and the discriminative code loss.

2.2. Network Training

The network parameters $(\mathbf{W}^{(i)})_{i=1,2,\dots,n}, A, W$ are trained via back-propagation using the mini-batch stochastic gradient descent method. Compared to the parameter update equations for a multi-layer CNN [13] without the discriminative cost loss, the gradient term, *i.e.* $\frac{\partial L}{\partial \mathbf{x}^{(n)}}$ changes and two gradient terms $\frac{\partial L}{\partial \mathbf{A}}$ and $\frac{\partial L}{\partial \mathbf{x}_d^{(n+1)}}$ are introduced, since $\mathbf{x}^{(n)}$, \mathbf{A} and $\mathbf{x}_d^{(n+1)}$ are related to the discriminative code loss L_d .

From Equations (5) and (6), we can calculate $\frac{\partial L}{\partial \mathbf{x}_d^{(n+1)}}$ and $\frac{\partial L}{\partial \mathbf{x}_c^{(n+1)}}$. Then we can obtain $\frac{\partial L}{\partial \mathbf{A}}$, $\frac{\partial L}{\partial \mathbf{W}}$ and $\frac{\partial L}{\partial \mathbf{x}^{(n)}}$ by applying the chain rule:

$$\frac{\partial L}{\partial \mathbf{x}_d^{(n+1)}} = \alpha \frac{\partial L_d}{\partial \mathbf{x}_d^{(n+1)}}, \quad \frac{\partial L}{\partial \mathbf{x}_c^{(n+1)}} = \frac{\partial L_c}{\partial \mathbf{x}_c^{(n+1)}} \quad (8)$$

$$\frac{\partial L}{\partial \mathbf{A}} = 2\alpha(\mathbf{A}\mathbf{x}^{(n)} - \mathbf{q}^{(n)})\mathbf{x}^{(n)\top}, \quad \frac{\partial L}{\partial \mathbf{W}} = \frac{\partial L_c}{\partial \mathbf{W}} \quad (9)$$

$$\frac{\partial L}{\partial \mathbf{x}^{(n)}} = \frac{\partial L}{\partial \mathbf{x}_c^{(n+1)}} \frac{\partial \mathbf{x}_c^{(n+1)}}{\partial \mathbf{x}^{(n)}} + 2\alpha(\mathbf{A}\mathbf{x}^{(n)} - \mathbf{q}^{(n)})^\top \mathbf{A} \quad (10)$$

Once the partial derivative of L with respect to $\mathbf{x}^{(n)}$ is known, the partial derivative of L with respect to $\mathbf{W}^{(i)}$ and $\mathbf{x}^{(i-1)}$ can be computed using the backward recurrence:

$$\begin{aligned} \frac{\partial L}{\partial \mathbf{W}^{(i)}} &= \frac{\partial L}{\partial \mathbf{x}^{(i)}} \frac{\partial \mathbf{x}^{(i)}}{\partial \mathbf{W}^{(i)}}, \\ \frac{\partial L}{\partial \mathbf{x}^{(i-1)}} &= \frac{\partial L}{\partial \mathbf{x}^{(i)}} \frac{\partial \mathbf{x}^{(i)}}{\partial \mathbf{x}^{(i-1)}}, \quad \forall i \in \{1, \dots, n\} \end{aligned} \quad (11)$$

where $\frac{\partial \mathbf{x}^{(i)}}{\partial \mathbf{W}^{(i)}}$ and $\frac{\partial \mathbf{x}^{(i)}}{\partial \mathbf{x}^{(i-1)}}$ can be computed from Equation (1).

2.3. Fusion

As a part of the two-stream pipeline, we perform fusion of the probabilistic outputs from the IR and the optic flow

nets. While the softmax layer directly provides a probabilistic output, we propose a method to convert the discriminative code layer outputs to a multinomial distribution over action classes. Given the predicted code of a test sample, we find its k nearest neighbors from each class in the training set and calculate the average distances from the sample to its k neighbor training samples from each class. Then, we convert a set of average distances to sample-to-class similarity weights using a Gaussian kernel. Finally, we obtain a probability vector over action categories by ℓ_1 normalization of the similarity weights.

We fuse the probability outputs of the IR and the optical flow nets using a simple weighted average approach. In addition, we apply two neural network based methods to fuse the predicted codes from IR and optical flow nets: (1) we concatenate predicted codes from from the IR and flow nets and use the obtained vector as an input to a single softmax layer neural network, which outputs probability estimates for action classes (**single-layer NN fusion**); and (2) we use the concatenated predicted codes as inputs to the two-layer neural network consisting of one 1×1 convolution layer and the softmax output layer (**two-layer NN fusion**).

3. Experiments

We evaluate our approach on the recently released InfAR video dataset [5], which is collected using infrared cameras. This dataset contains videos of 12 different action classes⁴ with 50 videos in each class. Figure 3 shows video examples from the dataset. First, using this dataset, we evaluate the performance of three widely used low-level descriptor features for the action class prediction task : dense SIFT (D-SIFT) [1], opponent SIFT (O-SIFT) [21], and motion features - improved dense trajectories features (IDT) [24]. Then, we evaluate the prediction performance of semantic feature vector produced by 2,784 concept detectors, each of

⁴The action classes include: ‘0-fight’, ‘1-handclapping’, ‘2-handshake’, ‘3-hug’, ‘4-jog’, ‘5-jump’, ‘6-punch’, ‘7-push’, ‘8-skip’, ‘9-walk’, ‘10-wave1’ and ‘11-wave2’.




Figure 3. Video samples for 12 action classes from the InfAR action dataset [5].

which outputs a concept score, given the low-level feature vector (e.g. D-SIFT). Finally, we evaluate the two-stream 3D CNNs with the discriminative code layer using different output fusion strategies.

3.1. Experimental Settings

To extract low-level image-based features such as dense SIFT and opponent SIFT, we uniformly sample 50 frames per video. We use Fisher vector (FV) encoding [19] to obtain video-level representations from these local low-level descriptors. The video-level features are computed by spatio-temporal pooling of the frame-based FV features. In addition to these low-level video representations, we extract features whose dimensions correspond to measures of evidence of high-level concepts in videos. For extracting semantic concept features, we trained 2,784 concept detectors utilizing the VideoStory Dataset [7]. The detectors are trained to predict high-level concept features from different FV-encoded features (D-SIFT, O-SIFT and IDT). Finally, we trained multiple linear multi-class SVM classifiers to predict actions from different low-level features and concept features.

The architectures of IR net and Flow net are identical and follow [23]. Assume $C(k, n, s)$ is a convolutional layer with kernel size $k \times k \times k$, n filters and stride $s \times s \times s$, $P(k1, k2, s1, s2)$ is a max-pooling layer with kernel temporal size $k1$, kernel spatial size $k2 \times k2$, temporal stride $s1$ and spatial stride $s2$. $FC(n)$ is a fully connected layer with n filters. $SM(n)$ is a softmax layer with n filters. $DC(n)$ is a discriminative code layer with n filters. The main architecture follows: $C(3, 64, 1) - P(1, 2, 1, 2) - C(3, 128, 1) - P(2, 2, 2, 2) - C(3, 256, 1) - C(3, 256, 1) - P(2, 2, 2, 2) - C(3, 512, 1) - C(3, 512, 1) - P(2, 2, 2, 2) - C(3, 512, 1) - C(3, 512, 1) - P(2, 2, 2, 2) - FC(4096) - FC(4096) - SM(12)(DC(4096))$.

The temporal length t of each IR and flow clip is 16 frames. The parameter α in Eq 4 is set to be 0.02 in our experiments. Both IR stream and Flow stream nets are pre-trained on the large-scale Sports-1M dataset [9] and fine-tuned on the InfAR dataset. The learning rate, training batch size, weight decay coefficient and maximum itera-

tions are set as 0.0001, 30, 0.0005 and 10,000, respectively. To extract video-level CNN features, we split a video into 16 frame long clips without overlapping between two consecutive clips. To get the video-level representations, we simply averaged representations extracted from each clip belonging to the video. To classify actions, we use two ways: (1) employ the softmax output layer to produce the confidence scores for all action classes for each video clip and use their average to predict the video-level class label (**softmax**); (2) employ the k -NN classification based on the video-level representation, which is computed as the average of predicted codes of all video clips belonging to the video (k -NN).

We follow the standard setting in [5], we randomly select 30 samples from each category as training, and the rest for testing. We repeat the experiments five times and report their performance average as the final performance in this paper. For the evaluation metrics, we used average precision (AP) as in [5], which is the average of recognition precisions of all actions.

3.2. Comparisons with Other Approaches

3.2.1 Low-level and high-level semantic features

We evaluate two static appearance features (D-SIFT and O-SIFT), one motion feature (IDT) and the corresponding high-level semantic concept features, all extracted from infrared videos. We perform early SVM fusion by concatenating concept feature vectors obtained using concept detectors on different low-level features. Finally, we perform late fusion by averaging the posterior scores of SVMs trained on different features.

Table 1 summarizes the recognition performances with these approaches. The high-level concept features achieved similar or better performance compared with the corresponding low-level features. In addition, the early fusion of all concept features provided similar results as the late fusion approach that combined the prediction scores from six SVM classifiers.

Method	AP (%)
D-SIFT [1]	46.7
D-SIFT based concepts	46.7
O-SIFT [21]	47.5
O-SIFT based concepts	47.1
IDT [24]	43.3
IDT based concepts	44.6
Early fusion of all concepts	47.5
Late fusion of all features	47.9

Table 1. Recognition performance comparisons in terms of average precisions (%) using low level features and their corresponding semantic concept features.

Method	AP (%)
IR net without DCL	48.75
IR net (softmax)	52.91
IR net (k -NN)	54.58
Flow net without DCL	69.58
Flow net (softmax)	72.91
Flow net (k -NN)	75.42
Two-stream-CNN-1 [5]	32.08
Two-stream-CNN-2 [5]	76.66

Table 2. Recognition results of 3D-CNNs trained with or without discriminative code loss, and using different classification methods. The results of ‘Two-stream CNN-1’ and ‘Two-stream CNN-2’ are copied from the original paper [5].

3.2.2 Discriminative features from single 3D CNN

Our two stream 3D-CNN is based on the C3D architecture in [23]. It consists of an IR net taking 16-frame clips as inputs and a flow net taking 16-frame sequences of the optical flow fields. We trained a 3D CNN in two ways: (1) using softmax classification loss only; (2) using both softmax classification loss and the discriminative code loss (DCL).

We first trained IR and Flow nets using softmax classification only. We called these two networks as ‘IR net without DCL’ and ‘Flow net without DCL’, respectively. Then we maintain the softmax layer in both networks but add the discriminative code layer to the fc_7 layer to train two stream networks, which are referred to as ‘IR net’ and ‘Flow net’. For IR and flow nets, we can employ softmax or k -NN classification methods. The performances of different training and classification methods are presented in Table 2.

As shown in Table 1 and Table 2, the ‘IR net without DCL’ can still obtain marginally better results than the late fusion of all low-level features and their concept features. The ‘Flow net without DCL’ can obtain around 20% improvement on average precision compared to its partner stream, which demonstrate that the motion information is very important for action recognition. The ‘ k -NN’ classification method achieved better performance than the ‘softmax’ method, due to the increase of inter-class distances in

the generated discriminative code space.

We compare our results with two-stream CNN results reported in [5]. ‘Two-stream CNN-1’ uses a two-stream 2D convolutional architecture [22] consisting of an IR stream and a flow stream as ours. ‘Two-stream CNN-2’ is similar to ‘Two-stream CNN-1’, but the IR stream is replaced with a stream taking optical flow motion history images as inputs. As pointed out in [23], 3D CNN can model appearance and motion simultaneously, hence our IR net outperforms the two-stream-CNN-1, which is trained with IR images and optical flow fields. The two-stream-CNN-2 achieves good results due to the use of motion history images. However, the result of our flow net, which is based on optical flow field only, is still comparable to two-stream-CNN-2.

3.2.3 Fusion of two 3D CNNs

We investigate different fusion strategies for the outputs from IR and flow nets:

Late fusion1. We fuse the softmax probability outputs of IR and Flow nets using simple weighted average rule, where the weight is 2 for flow net and 1 for IR net.

Late fusion2. Given a predicted code vector from the discriminative code layer, we compute the average distances to each class i based on its k neighbor training samples from each class, and convert them to a similarity vector using a Gaussian kernel function : $\exp(-\gamma d_i^2)$, where d_i is the average distance to class i and γ is a normalization factor. Finally, we compute the probability vector by ℓ_1 normalization on the similarity vector. k is set to 5 in our experiment. The parameter γ is set to be 0.05. We combine the probability vectors from different streams using a simple weighted average rule.

Single-layer NN fusion. We first concatenate discriminative codes from IR and Flow nets. Then, we construct a single softmax layer ‘shallow’ neural network with concatenated discriminative codes as inputs and action classes as outputs.

Two-layer NN fusion. We train a two-layer neural network consisting of one 1×1 convolution layer and a softmax output layer, with concatenated discriminative codes from IR and Flow nets as inputs. The convolution layer can play a role in selecting good features from the input concatenated features. We used 50 filters for the convolutional layer.

Table 3 presents the results of these fusion strategies. Simple weighted rule for ‘Late fusion1’ and ‘Late fusion2’ can lead to performance improvement over single stream CNN. The more complex single-layer NN and two-layer NN fusion methods do not outperform the simple weighted average fusion rule, likely due to a limited size of the InfAR dataset, since there is insufficient training samples to learn the parameters of convolution and softmax layers in these

Method	AP (%)
Late fusion 1	74
Late fusion 2	77.5
Single-layer NN fusion	71.25
Two-layer NN fusion	70.42

Table 3. Recognition performances of fusion with 3D CNN features from IR and Flow nets. Please refer to the section ‘Fusion of two 3D CNNs’ for the details of these fusion methods.

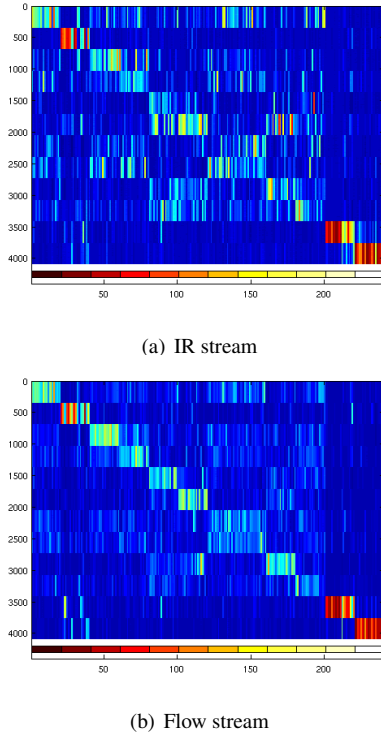


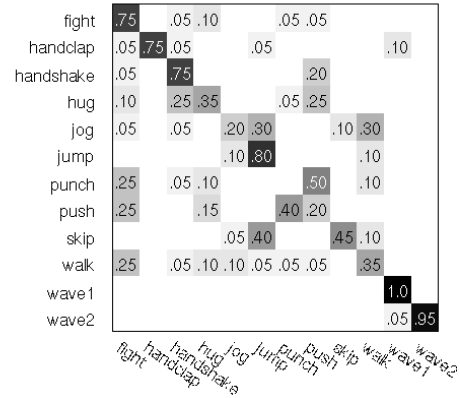
Figure 4. Visualization of learned discriminative codes of testing videos. Y axis indicates the dimensions of predicted codes while each position in X axis correspond to one test video. Blue color indicates ‘low value’ while red color indicates ‘high value’. Each color in the color bar located at the bottom of each subfigure represents one action class for a subset of testing videos.

two fusion networks.

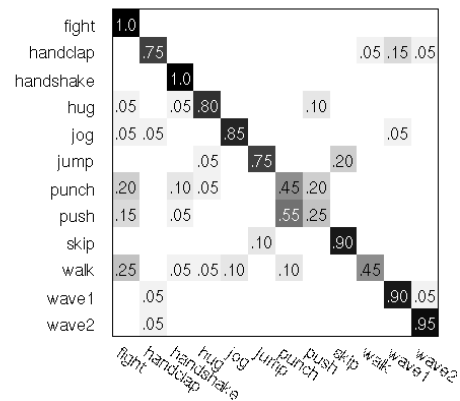
If we use a simple average rule to combine the confidence scores for all actions for each video from the method ‘Late fusion2’ with the scores generated from the method ‘Early fusion of all concepts’ in section 3.2.1, we can achieve 79.2% AP. This means that high-level concept features can provide complementary information to 3D CNN features for action recognition.

3.3. Discussion

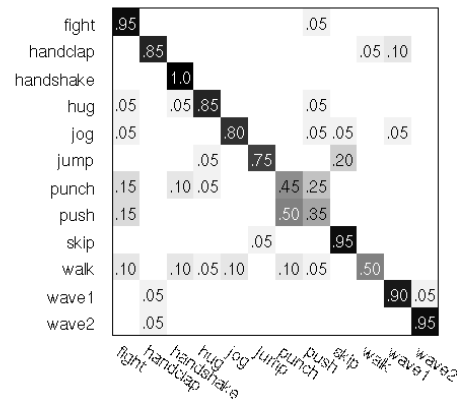
Figure 4 visualizes the learned discriminative codes of testing videos using both IR and Flow nets. The discriminative predicted code matrix of all testing videos should be block-diagonal. Y axis indicates the dimension of predicted code, and each position in X axis correspond to one test



(a) IR net



(b) Flow net




(c) Late fusion of IR and flow nets


Figure 5. Confusion matrices for infrared action recognition. The 3D CNNs, ‘IR net’ and ‘Flow net’, are trained with softmax loss and discriminative code loss. Weighted average rule is used to fuse IR net’s and Flow net’s prediction scores.

video. As can be seen from the figure, the videos from the same class have similar representations, while videos from different classes have different representations.

Figure 5 shows confusion matrices obtained using IR net, flow net and late fusion of two nets. The misclassifications are mainly related to ‘push’ and ‘punch’ categories



(a)



(b)

Figure 6. Effects of parameter selection of α and k on the average precision performance on the InfAR dataset. (a) Effects of parameter selection of cost-balancing hyper-parameter α . (b) Effects of parameter selection of k -NN neighborhood size k .


which are visually similar. ‘walk’ is misclassified as ‘fight’, which is possibly caused by presence of moving people in the background. The fusion of two nets helps correcting some typical misclassifications, such as the confusion between ‘handclap’ and ‘wave1’ classes.

For 5 action categories, we achieved precisions higher than 90% when using 30 positive video samples per category. Figure 7 shows some samples from three of these classes.


In Figure 6(a), we plot the performance curves for a range of parameter α in flow net. We observe that our approach is not sensitive to the selection of α . In addition, the simple k -NN classification scheme consistently outperforms the ‘softmax’ classification on the full α range, this is because the generated predicted codes using our approach are discriminative. In Figure 6(b), we show the performances using different k (recall k is the number of nearest neighbors for a k -NN classifier) for the IR and flow nets. Our approach is not sensitive to k due to the increase of inter-class distances in the discriminative code space.

4. Conclusion


We introduce a two-stream 3D convolutional network for action recognition in infrared videos. Each recognition



(a) Handshake, AP: 100%



(b) Fight, AP: 95%



(c) Wave2, AP: 95%

Figure 7. Video examples from classes with the high classification precision from the InfAR dataset. Each pair of images (IR and optical flow) is sampled from different testing videos and their corresponding optical flow fields.

stream (IR and flow nets), was trained with softmax classification loss and discriminative code loss making the extracted representations of infrared videos become more discriminative. Both nets are initialized by pretraining on visible spectrum videos, and finetuned on the infrared videos. Our experiments show that using even a single flow stream CNN can achieve state-of-the-art performance on the InfAR dataset. The goals of our future work are to extend the current approach to the cross-spectral feature learning and explore the domain adaptation techniques that can more effectively exploit the high resource spectrum for the action recognition in the low-resource spectrum.

Acknowledgement

This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center contract number D11PC20071. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government.

References

- [1] Y.-L. Boureau, F. Bach, Y. LeCun, and J. Ponce. Learning mid-level features for recognition. In *CVPR*, 2010. 4, 6
- [2] T. Brox, A. Bruhn, N. Papenberger, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision*, pages 25–36. Springer, 2004. 2
- [3] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 1, 2
- [4] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016. 1, 2
- [5] C. Gao, Y. Du, J. Liu, J. Lv, L. Yang, D. Meng, and A. Hauptmann. Infar dataset: Infrared action recognition at different times. *Neurocomputing*, 212:36–47, 2016. 1, 2, 4, 5, 6
- [6] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 1
- [7] A. Habibian, T. Mensink, and C. Snoek. Videostory: A new multimedia embedding for few-example recognition and translation of events. In *ACM MM*, 2014. 5
- [8] Z. Jiang, Y. Wang, L. Davis, W. Andrews, and V. Rozgic. Learning discriminative features via label consistent neural network. In *WACV*, 2017. 1, 2
- [9] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 1, 5
- [10] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 1
- [11] Z. Lan, M. Lin, X. L. A. G. Hauptmann, and B. Raj. Beyond gaussian pyramid: Multi-skip feature stacking for action recognition. In *CVPR*, 2015. 1, 2
- [12] I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64:107–123, 2005. 1, 2
- [13] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller. Efficient backpro. *Neural Networks: Tricks of the Trade*, 7700:9–48, 2012. 4
- [14] Z. Lin, Z. Jiang, and L. S. Davis. Recognizing actions by shape-motion prototype trees. In *ICCV*, 2009. 2
- [15] J. Liu, B. Kuipers, and S. Savarese. Recognizing human actions by attributes. In *CVPR*, 2011. 2
- [16] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1
- [17] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 1, 2
- [18] J. Y. Ng, M. J. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015. 1, 2
- [19] F. Perronnin, J. Sanchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *ECCV*, 2010. 5
- [20] Q. Qiu, Z. Jiang, and R. Chellappa. Sparse dictionary-based representation and recognition of action attributes. In *ICCV*, 2011. 2
- [21] K. Sande, T. Gevers, and C. Snoek. Evaluating color descriptors for object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1582–1596, 2010. 4, 6
- [22] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014. 1, 2, 6
- [23] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 1, 2, 3, 5, 6
- [24] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013. 1, 2, 4, 6
- [25] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. V. Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 2
- [26] B. Yao, X. Jiang, A. Khosla, A. L. Lin, L. Guibas, and L. Fei-Fei. Human action recognition by learning bases of action attributes and parts. In *ICCV*, 2011. 2
- [27] H. Ye, Z. Wu, R. Zhao, X. Wang, Y. Jiang, and X. Xue. Evaluating two-stream CNN for video classification. In *ICMR*, 2015. 1, 2
- [28] S. Zha, F. Luisier, W. Andrews, N. Srivastava, and R. Salakhutdinov. Exploiting image-trained CNN architectures for unconstrained video classification. In *BMVC*, 2015. 1
- [29] Y. Zhu and G. Guo. A study on visible to infrared action recognition. *IEEE Signal Processing Letters*, 20(9), 2013. 1, 2