

Fully-Coupled Two-Stream Spatiotemporal Networks for Extremely Low Resolution Action Recognition

Mingze Xu^{1*} Aidean Sharghi^{2*} Xin Chen^{3†} David J. Crandall¹

¹Indiana University, Bloomington, IN

²University of Central Florida, Orlando, FL

³Midea Corporate Research Center, San Jose, CA

{mx6, djcran}@indiana.edu, chen1.xin@midea.com

Abstract

A major emerging challenge is how to protect people’s privacy as cameras and computer vision are increasingly integrated into our daily lives, including in smart devices inside homes. A potential solution is to capture and record just the minimum amount of information needed to perform a task of interest. In this paper, we propose a fully-coupled two-stream spatiotemporal architecture for reliable human action recognition on extremely low resolution (e.g., 12×16 pixel) videos. We provide an efficient method to extract spatial and temporal features and to aggregate them into a robust feature representation for an entire action video sequence. We also consider how to incorporate high resolution videos during training in order to build better low resolution action recognition models. We evaluate on two publicly-available datasets, showing significant improvements over the state-of-the-art.

1. Introduction

Cameras are seemingly everywhere, from the traffic cameras in cities and highways to the surveillance systems in businesses and public places. Increasingly we allow cameras even into the most private spaces in our lives: gaming consoles like Microsoft Kinect [1] watch our living rooms, “smart home” devices like Amazon Echo Look [2] and Nest Cam [3] monitor our homes, “smart toys” like the Fisher-Price Smart Toy Bear [4] entertain our children, “smart appliances” like the Samsung Family Hub [5] watch and respond to our everyday actions, and cameras in mobile devices like smartphones and tablets see us even in our bedrooms. While these cameras have the promise of making our lives safer and simpler, including making possible more

*Part of this work was done when Mingze Xu and Aidean Sharghi were interns at Midea Corporate Research Center, San Jose, CA.

†Corresponding author.




Figure 1. Sample frames of extremely low resolution (12×16 pixel) videos streams from the HMDB51 dataset. Original high resolution frames are shown in red.

natural, context-aware interactions with technology, they also record highly sensitive information about people and their private environments.

To make matters worse, processing for many of today’s devices is often performed by remote servers “in the cloud.” This means that even if a user trusts that a device is using recorded video solely for legitimate purposes (which is already a leap of faith, given cases of private data being used for marketing and other unscrupulous purposes [32]), and that the device has been adequately protected against hacking (another leap of faith [7, 17, 39]), a user must also trust a

remote cloud provider and the interlying network that their data must traverse. Even a secure and trusted cloud provider may be forced to share data with a government by a court order or intelligence agency request [30].

One way of addressing the privacy challenge is to transmit just the minimum amount of information needed for a computer vision task to be accurately performed. In the security community, solutions based on selective encryption of image content [20] and firmware-enforced access control [25] have been proposed to keep video data safer from hacks, but service providers must still be trusted with high-fidelity video content. Another strategy is to remove private details in imagery before they leave the device. While techniques like selective image blurring [8], obscuring [34], and various other transformations [19, 27, 33] have been studied [18], these all require accurately defining and detecting private content, which is in itself a highly non-trivial problem (and thus these techniques are often not effective in practice [31]).

Perhaps the most effective approach is to simply avoid collecting high-fidelity imagery to begin with. For example, low resolution imagery may prevent specific details of a scene from being identified – e.g. the appearance of particular people, or the identity of particular objects – while still preserving enough information for a task like scene type recognition [40]. Particularly important in many home applications of cameras is action and activity recognition, to help give smart devices high-level contextual information about what is going on in the environment and how to react and interact accordingly. Several recent papers have shown that very low resolution videos (around 16×12 pixels) preserve enough information for fine-grained action recognition [10, 12, 35, 36]. This is perhaps surprising, since even a human observer may have difficulty identifying actions from such little information (Figure 1). This raises the question of how much better action recognition on low resolution frames can progress?

Existing work on low resolution action recognition tends to focus on modeling the spatial (appearance) information in each individual frame. For example, a common approach is to use high resolution training videos to learn a transformation between high and low resolution frames, to help recover lost visual information [36]. This has been implemented by either semi-coupled networks sharing convolutional filters between high and low resolution inputs [10] or Multi-Siamese networks learning inherent properties of low resolution [35]. However, much (if not most) useful information about action recognition in low resolution video is in the *motion* information, not in any single frame.

In this paper, we propose a fully-coupled two-stream spatiotemporal network architecture to better take advantage of both local and global temporal features for action recognition in low resolution video. Our architecture incor-

porates motion information at three levels: (1) a two-stream network incorporates stacked optical flow images to capture subtle spatial changes in low resolution videos; (2) a 3D Convolution (C3D) network computes temporal features within local video intervals; (3) a Recurrent Neural Network (RNN) uses the extracted C3D features from videos and optical flow fields to model more robust longer-range features. Our experiments on two challenging datasets (HMDB51 and DogCentric) show that our model significantly outperforms the previous state-of-the-art.

2. Related Work

Most state-of-the-art techniques for action recognition in video use deep learning methods. At a very high level, there are two important types of evidence about action: appearance (spatial) features within individual frames, and motion (temporal) features that cue on distinctive movement patterns. Karpathy *et al.* [26] were among the first to study deep learning-based action recognition, proposing a multiresolution CNN architecture that operates on individual frames without explicitly modeling temporal information. Simonyan and Zisserman [38] used a two-stream CNN framework to incorporate both feature types, with one stream taking RGB image frames as input and the other taking pre-computed stacked optical flows. The additional stream significantly improved action recognition accuracy, indicating the importance of motion features. Tran *et al.* [41] avoided the need for precomputing optical flow features through their 3D convolution (C3D) framework, which allows deep networks to learn temporal features in an end-to-end manner.

Diba *et al.* [13] combined these two ideas into two-stream C3D networks, and also proposed a more robust fusion method for better temporal information encoding. Zhu *et al.* [43] avoid pre-computing optical flow, instead learning the motion features in an end-to-end framework with a hidden two-stream network. That approach is about ten times faster than having to pre-compute optical flow, but accuracy is somewhat weaker. Most of these papers capture motion information over relatively short temporal intervals, although several recent papers generate action proposals for longer videos with a combination of C3D and Recurrent Neural Networks (RNNs) [9, 14, 16].

Several papers have focused on recognition in extremely low resolution (LR) imagery. This problem is considered more difficult, of course, since there is simply less visual information available at low resolutions [44]. Wang *et al.* [42] address low resolution object recognition by taking advantage of high resolution training images to learn a transformation between the two resolutions. Dai *et al.* [12] adapted this idea to action recognition in the video domain, specifically focusing on extracting and learning better low resolution features from limited information. Ryoo *et al.* [36]

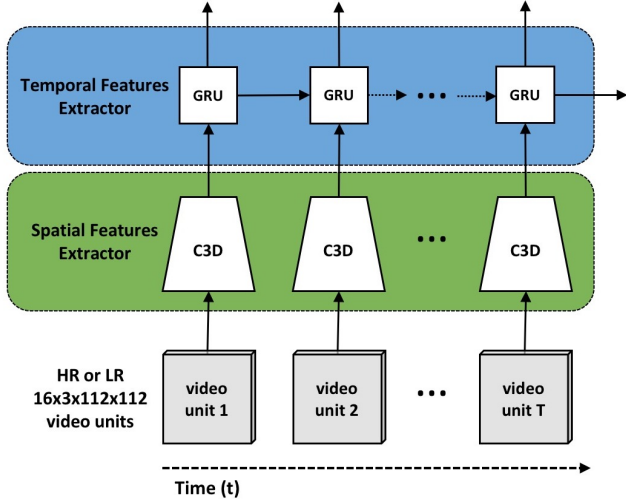


Figure 2. Visualization of our spatiotemporal features extractor, which uses a C3D network to capture spatial and temporal features for video units and an RNN to encode motion information across the entire video stream.

defined low resolution to be around 16×16 to 12×16 pixels and decomposed high resolution training videos into multiple LR training videos by learning different resolution transformations. Chen *et al.* [10] used two-stream semi-coupled networks to design an end-to-end training network on both visual and motion information. By observing that two LR images taken from exactly the same scene can contain totally different pixel values, recent follow-up work by Ryoo *et al.* [35] achieved state-of-the-art LR action recognition performance by learn an embedding representation with Multi-Siamese networks.

We build on these previous methods that focus mostly on modeling and encoding spatial features from low resolution video frames, and propose an action recognition approach that incorporates stronger motion information. In particular, our models captures motion information within different temporal neighborhoods, including both sequential dependencies between consecutive frames and more global temporal features. We find that this, combined with a fully-coupled network that learns from high resolution training videos, yields stronger models that significantly outperform state-of-the-art methods.

3. Technical Approach

We now present our approach for action recognition on extremely low resolution videos. Specifically, we first introduce the basic architecture of our spatiotemporal feature extractor, which uses a combination of 3D Convolutional (C3D) Networks and Recurrent Neural Networks (RNNs). Then, we discuss how to learn transferable features from

high to low resolution videos, which is based on the assumption that high resolution *training* videos are available. Finally, we explore four fusion methods to efficiently combine visual and motion information during recognition.

3.1. Spatiotemporal Features Extractor

We assume we are given an input video sequence with L frames that has captured a person or people performing a single action (e.g., kicking a ball, shaking hands, chewing food, etc.) and our goal is to recognize this action. Unlike prior work that has viewed each frame as a separate processing unit, we discretize the video into $T = L/\delta$ non-overlapping *video units* V_1, \dots, V_T , each containing δ consecutive frames. Figure 2 illustrates the general architecture of the extractor network.

Spatial Features Extractor. Motivated by Buch *et al.*'s work [9] on high resolution video, we propose a feature extractor to characterize appearance information in low resolution video units. In particular, we use the C3D network, which has proven to be well-suited for modeling sequential inputs such as videos [41]. Since C3D uses 3D convolution and pooling operations that operate over both spatial and temporal dimensions, it is able to capture motion information within each input video unit. We successively feed each video unit V_i to C3D, and extract its feature representation x_t from the last fully-connected layer.

Temporal Features Extractor. While the C3D network is able to encode local temporal features within each video unit, it cannot model across the multiple units of a video sequence. We thus introduce a Recurrent Neural Network (RNN) to capture global sequence dependencies of the input video and cue on motion information (e.g., trajectories). Although Long Short-Term Memory (LSTM) [23] networks are the most widely used RNNs in video applications, we found that Gated Recurrent Units (GRUs) [11] performed better in our application. The basic mechanism behind GRUs is similar to LSTMs, except that they do not use memory units to control the information flow, and they have fewer parameters which makes them slightly easier and faster to use both in training and testing.

More formally, the GRU is trained to take the extracted C3D features x_1, x_2, \dots, x_T and output a sequence of hidden states h_1, h_2, \dots, h_T , with learnable hidden layer parameters W , U , and b . Then, the GRU cell iterates the following operations for $t \in [1, T]$:

$$\begin{aligned}
 z_t &= \text{sigmoid}(W_z x_t + U_z h_{t-1} + b_z), \\
 r_t &= \text{sigmoid}(W_r x_t + U_r h_{t-1} + b_r), \\
 n_t &= \tanh(W_h x_t + U_h (r_t \circ h_{t-1}) + b_h), \text{ and} \\
 h_t &= z_t \circ h_{t-1} + (1 - z_t) \circ n_t,
 \end{aligned} \tag{1}$$

where \circ is the Hadamard product, and z_t , r_t , n_t , and h_t

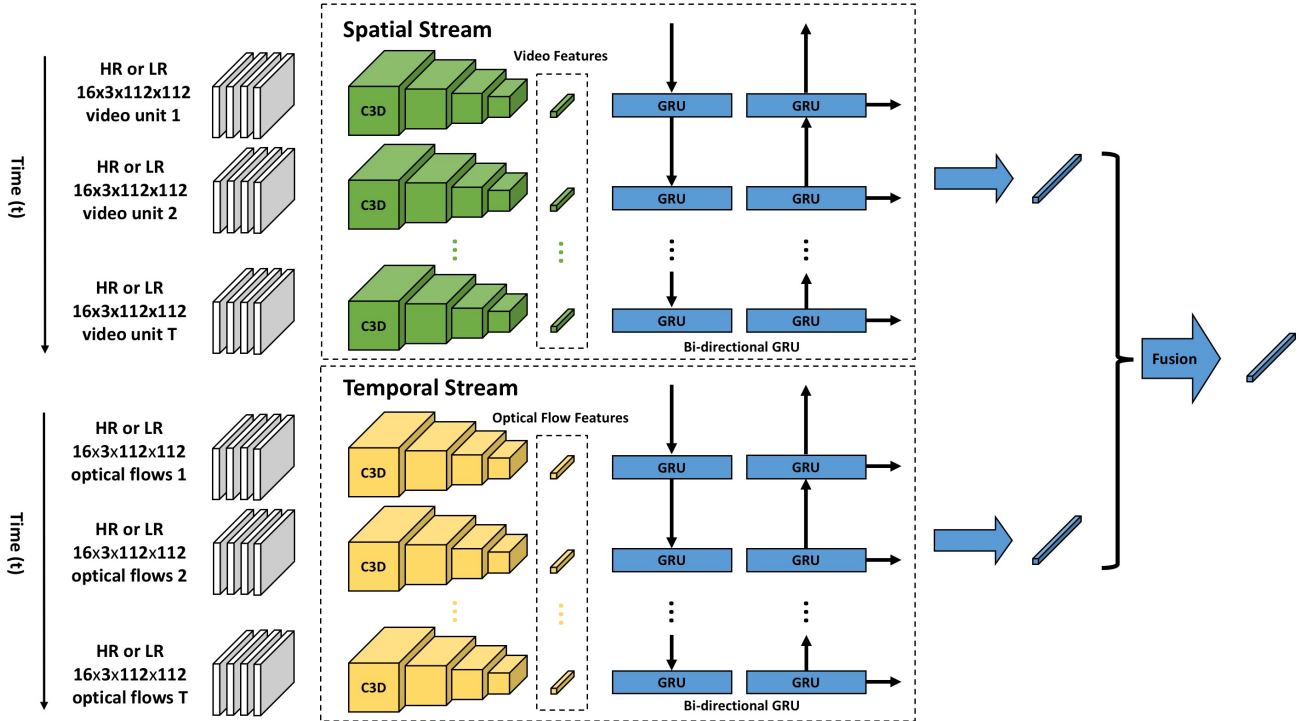


Figure 3. Visualization of our fully-coupled two-stream spatiotemporal networks. We feed RGB frames into the spatial stream (green) and corresponding stacked optical flow fields to the temporal stream (yellow). The GRU networks (blue) compute spatiotemporal information for the entire video using the extracted C3D features as inputs. In training, both C3D and GRU are fully-coupled with convolution filters shared between high and low resolution training videos, whereas only low resolution videos are used in testing.

are the reset, input, new gates, and hidden state for time t , respectively.

The hidden state h_t at each time t is a feature vector representing encoded spatiotemporal information corresponding to the first t video units, so h_T incorporates features of the entire video. We use h_T as input to a fully-connected layer with a softmax classifier to output the confidence score of each action class. We call the spatiotemporal feature extractor the *spatial stream* of our two-stream networks.

3.2. Fully-coupled Networks

Low resolution recognition approaches in both image and video domains have achieved better performance by learning transferable features from high to low resolutions. This process can be done either using unsupervised pre-training on super resolution sub-networks [42] or with partially-coupled networks [10, 42] which are more flexible for knowledge transformation. Inspired by these results, we propose a fully-coupled network architecture where all parameters of both the C3D and GRU networks are shared between high and low resolutions in the (single) training stage. The key idea is that by viewing high and low resolution video frames as two different domains, the fully-coupled

network architecture is able to extract features across them. Since high resolution video contains much more visual information, training on both resolutions helps improve learning spatial features; using high resolution in training can be thought of as data augmentation, since different techniques for sub-sampling produce different low resolution exemplars from the same original high resolution image.

3.3. Two-stream Networks

In this subsection, we extend our single-stream networks to two-stream networks by adding a similar architecture but with optical flow fields as the input. Since motion features between consecutive low resolution video frames are often quite small, our model benefits from optical flow images to learn pixel-level correspondences of temporal features. We call this module the *temporal stream*. Figure 3 shows the main architecture of our two-stream networks.

In particular, we compute optical flow fields for both high and low resolution videos, following Chen *et al.* [10] and using the public MATLAB toolbox of Liu [29]. The optical flows are encoded as 3-channel HSL images, where hue and saturation are converted from optical flow vectors (x- and y-components) into polar coordinates and lightness

is set to one. Before computing optical flow, we downsample the high resolution frames and upsample the low resolution frames to a common size of 112×112 pixels.

We also explore four widely used fusion methods, which enable the model to leverage joint visual and motion information more effectively. Following the notational convention of Feichtenhofer *et al.* [15], a fusion function $f : \mathbf{x}_t^a, \mathbf{x}_t^b \rightarrow \mathbf{y}_t$ fuses two feature maps $\mathbf{x}_t^a \in \mathbb{R}^{H \times W \times D}$ and $\mathbf{x}_t^b \in \mathbb{R}^{H' \times W' \times D'}$, at state t , to output a map $\mathbf{x}_t^b \in \mathbb{R}^{H'' \times W'' \times D''}$, where W , H , and D denote the width, height and number of channels, respectively. Different papers apply the fusion function to feature maps at different points in a deep network, while our model only applies it to the final hidden state h_T . Since the hidden state of a GRU is a D -dimensional vector, for simplicity, we drop the width, height and t subscripts and assume that $D = D' = D''$.

We consider four different fusion techniques:

1. **Sum Fusion**, $\mathbf{y}^{\text{sum}} = f^{\text{sum}}(\mathbf{x}^a, \mathbf{x}^b)$, computes the sum of two feature vectors over each channel d ,

$$y_d^{\text{sum}} = x_d^a + x_d^b, \quad (2)$$

where $1 < d < D$ and $\mathbf{x}^a, \mathbf{x}^b, \mathbf{y} \in \mathbb{R}^D$. Sum fusion is based on the assumption that dimensions of the two output vectors of a two-stream network correspond to one another.

2. **Max fusion**, $\mathbf{y}^{\text{max}} = f^{\text{max}}(\mathbf{x}^a, \mathbf{x}^b)$, computes the maximum of two feature vectors over each channel d ,

$$y_d^{\text{max}} = \max\{x_d^a, x_d^b\}. \quad (3)$$

In contrast to sum fusion, max fusion only keeps the feature with the higher response, but again assumes that corresponding dimensions are comparable.

3. **Concatenation Fusion**, $\mathbf{y}^{\text{cat}} = f^{\text{cat}}(\mathbf{x}^a, \mathbf{x}^b)$, stacks two feature vectors along the channel dimension,

$$y_{2d}^{\text{cat}} = x_d^a \quad \text{and} \quad y_{2d-1}^{\text{cat}} = x_d^b, \quad (4)$$

where $\mathbf{y} \in \mathbb{R}^{2D}$. Instead of defining the correspondence between two feature vectors by hand, concatenation fusion leaves this to be learned by the subsequent convolution layers.

4. Finally, **Convolution Fusion**, $\mathbf{y}^{\text{conv}} = f^{\text{conv}}(\mathbf{x}^a, \mathbf{x}^b)$, is similar to concatenation fusion, but is convolved with a learnable bank of filters $\mathbf{f} \in \mathbb{R}^{2D \times D^o}$ and biases $b \in \mathbb{R}^{D^o}$ are appended after the concatenation fusion,

$$y^{\text{conv}} = y^{\text{cat}} * \mathbf{f} + b, \quad (5)$$

where $*$ denotes convolution and D^o denotes the number of output channels. Filter \mathbf{f} provides a flexible way

for the fusion method to measure the weighted combination of \mathbf{x}^a and \mathbf{x}^b , and is able to project the channel dimension from $2D$ to D^o . In order to permit a fair comparison with other methods, we set $D^o = D$ and $D^o = 2D$ in our experiments.

4. Experiments

We now evaluate and report results of each of the above methods on our problem of action recognition in very low resolution video sequences.

4.1. Datasets

We evaluated our proposed techniques on low resolution versions of the HMDB51 [28] and DogCentric [24] benchmarks, which are among the most widely used datasets in action recognition at extremely low resolutions. The HMDB51 (Human Motion Database) dataset consists of 7,000 video clips from a variety of sources ranging from movies to YouTube videos, and is annotated with 51 action classes such as eating, smiling, clapping, bike riding, shaking hands, etc. Each clip is approximately 2-3 seconds long and is recorded at 30 frames per second. We follow prior work and report accuracies averaged over three splits of training and testing data. The DogCentric dataset is collected from wearable cameras mounted on dogs, and consists of about 200 videos categorized into 10 different actions (which include activities performed by the dog itself (e.g., drinking, walking, etc.) and interactions between the dog and people. In contrast to HMDB51 which captures actions from a third-person perspective, these videos capture actions of the camera wearer from a first-person viewpoint.

Low Resolution Videos. The above two datasets were recorded at a resolution of 240×320 pixels. Since our C3D network expects frames at a resolution of 112×112 pixels, we subsample the high resolution images (used only during training) to this size. To generate low resolution training and testing data, we resize the original videos to 12×16 using average downsampling, then upscale them back to 112×112 resolution using bi-cubic interpolation. We do not introduce any extra evidence into the interpolation operation to ensure the 112×112 videos have no more information than the 12×16 videos. Figure 4 shows several corresponding low and high resolution frames as examples.

4.2. Implementation Details

Our learning process consists of two stages: (1) training the C3D network and extracting features for video units; and (2) training our fully-coupled two-stream networks with the concatenated C3D features of each video as inputs.

For the C3D networks, we followed Tran *et al.* [41] and used their publicly available pre-trained model and fine-tuned on the HMDB51 and DogCentric datasets in both




Figure 4. Examples of video frames from the HMDB51 (left) and DogCentric (right) datasets. *First row*: high resolution images resized to 112×112 pixels; *Second row*: low resolution (12×16 pixel) images upsampled to 112×112 with bi-cubic interpolation; *Third and fourth rows*: optical flow fields for high resolution and low resolution images, respectively, calculated from images rescaled to 112×112 .

high and low resolution. Since we also extracted C3D features for optical flow inputs, we fine-tuned another C3D model with high and low resolution optical flows. The network architecture had 8 convolution layers, 5 max-pooling layers, and 2 fully-connected layers. The length of each video unit was set to 16 frames. We used the output of the first fully-connected layer *fc6* (which had 4096 dimensions) and stacked them together to form the video descriptor.

We implemented the fully-coupled two-stream networks in PyTorch [6] and used the C3D features and optical flow as inputs. As discussed before, we used mixed (high and low resolution) data in the training stage, but no high resolution information is used in testing. We used the root-mean-square propagation (RMSprop) [22] update rule to learn the network parameters with fixed learning rate 10^{-3} and weight decay 0.0005. The whole training process stopped at 50 epochs, with the batch size set to 256. All our experiments were conducted on a system with a Nvidia Titan X Pascal graphics card.

4.3. Evaluation

Table 2 shows results of the evaluation. Our full model featuring pre-trained C3D networks, the bi-directional GRU network, and the fully-coupled two-stream architecture with sum fusion achieves 44.96% accuracy on the low resolution HMDB51 dataset and 73.19% on the low resolution DogCentric dataset. Of course, these results are significantly worse than the best results on the high resolution versions of these datasets (e.g around 80.7% for HMDB51 [?]). (We also tested our best model on action recognition on high resolution videos, and easily achieved over 68% accuracy without any explicitly tuning on network architecture and hyperparameters.) However, as shown in Tables 2 and 3, our results do beat all state-of-the-art approaches on low resolution video, including Pooled Time Series (PoT) (which uses a combination of HOG, HOF, CNN features) [37], Inverse Super Resolution (ISR) [36], Semi-coupled Two-stream Fusion ConvNets [10], and Multi-Siamese Embedding CNNs [35]. Our best result outperforms these methods by 7.2% (3.3% without pre-training) on the HMDB51 dataset and 3.7% on the DogCentric dataset.

Network Architecture				
Type	C3D	RNN	Fusion	Accuracy
Temporal Network	pre-trained	w/o GRU	Sum Fusion	21.24%
	pre-trained	uni-directional GRU	Sum Fusion	24.11%
Spatial Network	pre-trained	w/o GRU	Sum Fusion	34.90%
	pre-trained	uni-directional GRU	Sum Fusion	39.15%
Two-stream Network	pre-trained	w/o GRU	Sum Fusion	38.56%
	pre-trained	uni-directional GRU	Sum Fusion	43.38%
	w/o pre-trained	bi-directional GRU	Sum Fusion	41.04%
	pre-trained	bi-directional GRU	Sum Fusion	44.96%
	pre-trained	bi-directional GRU	Max Fusion	42.02%
	pre-trained	bi-directional GRU	Concatenate Fusion	43.13%
	pre-trained	bi-directional GRU	Convolution Fusion	43.46%

Table 1. Evaluation results of each component of our network architecture on the HMDB51 dataset.




Figure 5. Confusion matrix on the HMDB51 dataset using our best model. The x-axis denotes the predicted labels and the y-axis represents the ground truth labels for 51 action classes.

To see how well our model performs on different categories, we also present a confusion matrix in Figure 5.

4.4. Discussion

To evaluate the contribution of each component of our model, we also implemented multiple simpler baselines.

C3D Pre-training. To more fairly compare our model with that of Chen *et al.* [10], we experimented with training the C3D networks from scratch on HMDB51, using the same architecture and hyperparameters as the pre-trained network. As shown in Tables 1 and 2, although the result

is about 3.9% worse without pre-training on the Sport-1M dataset, our model still achieves state-of-the-arts results.

GRU Networks. To measure the contribution of the GRU networks to our overall approach, we tried replacing them with two fully-connected layers. To reduce the interference from the two-stream networks, we also implemented two one-stream networks each with the spatial and temporal features as inputs, and compared the results between the ones with and without the GRU, respectively. The results are presented in Table 1. It is clear that the models with GRU outperform by about 4.5%. We also tested uni- and bi-directional GRU architectures, and found that bi-directional GRUs perform slightly better, as shown in Table 1.

Two-stream Networks. After evaluating the GRU networks, we now turn to the two-stream architecture, where we believe that the pixel-level motion information acquired from optical flow can improve the model’s ability in temporal feature encoding. As shown in Table 1, it is clear that the two-stream networks significantly outperform one-stream networks with both spatial (4.2% better) and temporal features (19.2% better), respectively. We also assessed the effect of our four different fusion methods: (1) sum fusion, (2) max fusion, (3) concatenate fusion, and (4) convolution fusion, as discussed in the previous section. The results summarized in Table 1 show that sum fusion achieves the best performance.

5. Conclusion

We presented a new Convolutional Neural Network framework for action recognition on extremely low resolution videos. We achieved state-of-the-art results on the HMDB51 and DogCentric datasets with a combination of four important components: (1) a fully-coupled network ar-

Approach	Accuracy
3-layer CNN [36]	20.81%
ResNet-31 [21]	22.37%
PoT (HOG + HOF + CNN) [37]	26.57%
ISR [36]	28.68%
Semi-coupled Two-stream ConvNets [10]	29.20%
Multi-Siamese Embedding CNN [35]	37.70%
Ours (w/o pre-trained C3D)	41.04%
Ours (w/ pre-trained C3D)	44.96%

Table 2. Performance of our model compared to the state-of-the-art results on the HMDB51 dataset.

Approach	Accuracy
PoT (HOG + HOF + CNN) [37]	64.60%
ISR [36]	67.36%
Multi-Siamese Embedding CNN [35]	69.43%
Ours (w/ pre-trained C3D)	73.19%

Table 3. Performance of our model compared to the state-of-the-art results on the DogCentric dataset.

chitecture to leverage high resolution images in training in order to learn a cross-domain transformation between low and high resolution feature spaces; (2) 3D convolutional components which extract compact and efficient spatiotemporal features for short video units; (3) a Recurrent Neural Network (RNN) which considers long-range temporal motion information; and (4) two network streams having both image frames and stacked dense optical flow fields as input, in order to take into account detailed motion features between adjacent video frames. We hope this paper inspires more work on extremely low resolution action recognition, and in methods to learn spatial-temporal features through sequential images more generally.

6. Acknowledgements

This work was supported by the Midea Corporate Research Center University Program, the National Science Foundation under grants CNS-1408730 and CAREER IIS-1253549, and the IU Office of the Vice Provost for Research, the College of Arts and Sciences, and the School of Informatics, Computing, and Engineering through the Emerging Areas of Research Project “Learning: Brains, Machines, and Children.”

References

- [1] <http://www.xbox.com/en-US/xbox-one/accessories/kinect>.
- [2] <https://www.amazon.com/Echo-Hands-Free-Camera-Style-Assistant/dp/B0186JAEWK>.
- [3] <https://nest.com/cameras/>.
- [4] <http://www.smarttoy.com>.
- [5] <http://www.samsung.com/us/explore/family-hub-refrigerator/overview/>.
- [6] <http://pytorch.org/>.
- [7] Bouvet. Investigation of privacy and security issues with smart toys. Technical report, Norwegian Consumer Council, 2016.
- [8] M. Boyle, C. Edwards, and S. Greenberg. The effects of filtered video on awareness and privacy. In *ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2000.
- [9] S. Buch, V. Escorcia, C. Shen, B. Ghanem, and J. C. Niebles. SST: Single-stream temporal action proposals. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [10] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. *arXiv:1705.07750*, 2017.
- [11] J. Chen, J. Wu, J. Konrad, and P. Ishwar. Semi-coupled two-stream fusion convnets for action recognition at extremely low resolutions. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2017.
- [12] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Gated feedback recurrent neural networks. In *International Conference on Machine Learning (ICML)*, 2015.
- [13] J. Dai, B. Saghaei, J. Wu, J. Konrad, and P. Ishwar. Towards privacy-preserving recognition of human activities. In *IEEE International Conference on Image Processing (ICIP)*, 2015.
- [14] A. Diba, V. Sharma, and L. Van Gool. Deep temporal linear encoding networks. *arXiv:1611.06678*, 2016.
- [15] V. Escorcia, F. C. Heilbron, J. C. Niebles, and B. Ghanem. Daps: Deep action proposals for action understanding. In *European Conference on Computer Vision (ECCV)*, 2016.
- [16] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] J. Gao, Z. Yang, C. Sun, K. Chen, and R. Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. *arXiv:1703.06189*, 2017.
- [18] G. Gottsegen. Former NSA employee: This hack gains access to your Mac’s webcam. *CNet*, 2016.
- [19] R. Hasan, E. Hassan, Y. Li, K. Caine, D. J. Crandall, R. Hoyle, and A. Kapadia. Viewer experience of obscuring scene elements in photos to enhance privacy. In *ACM CHI Conference on Human Factors in Computing Systems (CHI)*, 2018.
- [20] E. Hassan, R. Hasan, P. Shaffer, D. Crandall, and A. Kapadia. Cartooning for enhanced privacy in lifelogging and streaming video. In *CVPR Workshop on the Bright and Dark Sides of Computer Vision: Challenges and Opportunities for Privacy and Security*, 2017.
- [21] J. He, B. Liu, D. Kong, X. Bao, N. Wang, H. Jin, and G. Kesidis. Puppies: Transformation-supported personalized privacy preserving partial image sharing. In *IEEE International Conference on Dependable Systems and Networks (DSN)*, 2014.

- [22] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] G. Hinton, N. Srivastava, and K. Swersky. Neural networks for machine learning, lecture 6a: Overview of mini-batch gradient descent.
- [24] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [25] Y. Iwashita, A. Takamine, R. Kurazume, and M. S. Ryoo. First-person animal activity recognition from egocentric videos. In *International Conference on Pattern Recognition (ICPR)*, 2014.
- [26] S. Jana, A. Narayanan, and V. Shmatikov. A Scanner Darkly: Protecting User Privacy from Perceptual Applications. In *IEEE Symposium on Security and Privacy (SP)*, 2013.
- [27] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [28] P. Korshunov and T. Ebrahimi. Using face morphing to protect privacy. In *IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, 2013.
- [29] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. Hmdb: a large video database for human motion recognition. In *IEEE International Conference on Computer Vision (ICCV)*, 2011.
- [30] C. Liu. *Beyond pixels: exploring new representations and applications for motion analysis*. PhD thesis, Massachusetts Institute of Technology, 2009.
- [31] C. Mele. Bid for access to Amazon Echo audio in murder case raises privacy concerns. *The New York Times*, Dec 28, 2016.
- [32] C. Neustaedter, S. Greenberg, and M. Boyle. Blur filtration fails to preserve privacy for home-based video conferencing. *ACM Transaction Computer Human Interactions*, 13(1):1–36, Mar. 2006.
- [33] N. Nguyen. If you have a smart TV, take a closer look at your privacy settings. *CNBC*, 2017.
- [34] J. R. Padilla-López, A. A. Chaaraoui, and F. Flórez-Revuelta. Visual privacy protection methods. *Expert Systems with Applications*, 42(9):4177–4195, June 2015.
- [35] N. Raval, A. Srivastava, K. Lebeck, L. Cox, and A. Machanavajjhala. MarkIt: Privacy Markers for Protecting Visual Secrets. In *ACM Joint International Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2014.
- [36] M. S. Ryoo, K. Kim, and H. J. Yang. Extreme low resolution activity recognition with multi-siamese embedding learning. *arXiv:1708.00999*, 2017.
- [37] M. S. Ryoo, B. Rothrock, C. Fleming, and H. J. Yang. Privacy-preserving human activity recognition from extreme low resolution. In *AAAI Conference on Artificial Intelligence*, 2017.
- [38] M. S. Ryoo, B. Rothrock, and L. Matthies. Pooled motion features for first-person videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [39] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, 2014.
- [40] R. Templeman, Z. Rahman, D. Crandall, and A. Kapadia. PlaceRaider: Virtual theft in physical spaces with smartphones. In *Network & Distributed System Security Symposium (NDSS)*, 2013.
- [41] A. Torralba, R. Fergus, and W. T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*, pages 1958–1970, 2008.
- [42] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. *arXiv:1412.0767*, 2014.
- [43] Z. Wang, S. Chang, Y. Yang, D. Liu, and T. S. Huang. Studying very low resolution recognition using deep networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [44] Y. Zhu, Z. Lan, S. Newsam, and A. G. Hauptmann. Hidden two-stream convolutional networks for action recognition. *arXiv:1704.00389*, 2017.
- [45] W. W. Zou and P. C. Yuen. Very low resolution face recognition problem. *IEEE Transactions on Image Processing*, 21(1):327–340, 2012.