

Going Deeper into First-Person Activity Recognition

Minghuang Ma, Haoqi Fan and Kris M. Kitani
Carnegie Mellon University
Pittsburgh, PA 15213, USA

minghuam@andrew.cmu.edu haoqif@andrew.cmu.edu kkitani@cs.cmu.edu

Abstract

We bring together ideas from recent work on feature design for egocentric action recognition under one framework by exploring the use of deep convolutional neural networks (CNN). Recent work has shown that features such as hand appearance, object attributes, local hand motion and camera ego-motion are important for characterizing first-person actions. To integrate these ideas under one framework, we propose a twin stream network architecture, where one stream analyzes appearance information and the other stream analyzes motion information. Our appearance stream encodes prior knowledge of the egocentric paradigm by explicitly training the network to segment hands and localize objects. By visualizing certain neuron activation of our network, we show that our proposed architecture naturally learns features that capture object attributes and hand-object configurations. Our extensive experiments on benchmark egocentric action datasets show that our deep architecture enables recognition rates that significantly outperform state-of-the-art techniques – an average 6.6% increase in accuracy over all datasets. Furthermore, by learning to recognize objects, actions and activities jointly, the performance of individual recognition tasks also increase by 30% (actions) and 14% (objects). We also include the results of extensive ablative analysis to highlight the importance of network design decisions.

1. Introduction

Recently there has been a renewed interest in the use of first-person point-of-view cameras to better understand human activity. In order to accurately recognize first-person activities, recent work in first-person activity understanding has highlighted the importance of taking into consideration both *appearance* and *motion* information. Since the majority of actions are centered around hand-object interactions in the first-person sensing scenario, it is important to capture *appearance* corresponding to such features as hand regions, grasp shape, object type or object attributes. Capturing *motion* information such as local hand movements and global

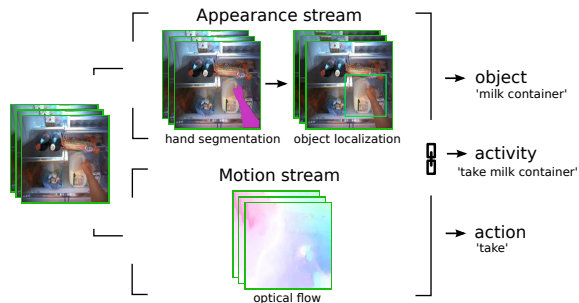


Figure 1: Approach overview. Our framework integrates both appearance and motion information. The appearance stream captures hand configurations and object attributes to recognize objects. The motion stream captures objects motion and head movement to recognize actions. The two streams are also learned jointly to recognize activities.

head motion, is another important visual cue as the temporal motion signature can be used to differentiate between complementary actions such as *take* and *put* or periodic actions such as the *cut with knife* action. It is also critical to reason about appearance and motion *jointly*. It has been shown in both third-person [11] and first-person activity analysis [20] that these two streams of activity information, appearance and motion, should be analyzed jointly to obtain best performance.

Based on these insights, we propose a deep learning architecture designed specifically for egocentric video, that integrates both action appearance and motion within a single model 1. More specifically, our proposed network has a two stream architecture composed of an appearance-based CNN that works on localized object of interest image frames and a motion-based CNN that uses stacked optical flow fields as input. Using the terminology of [5], we use late fusion with a fully-connected top layer to formulate a multi-task prediction network over *actions*, *objects* and *activities*. The term *action* describes motions such as *put*, *scoop* or *spread*. The term *object* refers to item such as bread, spoon or cup. The term *activity* is used to represent an action-object pair such as *take milk container*.

The appearance-based stream is customized for ego-

centric video analysis by explicitly training a hand segmentation network to enable an attention-based mechanism to focus on certain regions of the image near the hand. The appearance-based stream is also trained with object images cropped based on hand location to identify objects of manipulation. In this way, the appearance-based stream is enabled to encode such features such as hand-object configurations and object attributes.

The motion-based stream is a generalized CNN that takes as input a stack of optical-flow motion fields. This stream is trained to differentiate between action labels such as *put*, *take*, *close*, *scoop* and *spread*. Instead of compensating for camera ego-motion as a pre-processing step, we allow the network to automatically discover which motion patterns (camera, object or hand motion) are most useful for discriminating between action types. Results show that the network automatically learns to differentiate between different motion types.

We train the appearance stream and motion stream jointly as a multi-task learning problem. Our experiments show that by learning the parameters of our proposed network jointly, we are able to outperform state-of-the-art techniques by over 6.6% on the task of egocentric activity recognition without the use of gaze information, and in addition improve the accuracy of each sub-task (30% for action recognition and 14% object recognition).

Perhaps more importantly, the trained network also helps to better understand and to reconfirm the value of key features needed to discriminate between various egocentric activities. We include visualizations of neuron activations and show that the network has learned intuitive features such as hand-object configurations, object attributes and hand motion signatures isolated from global motion.

Contributions: (1) we formulate a deep learning architecture customized for ego-centric vision; (2) we obtain state-of-the-art performance propelling the field towards higher performance; (3) we provide ablative analysis of design choices to help understand how each component contributes to performance; and (4) visualization and analysis of the resulting network to understand what is being learned at the intermediate layers of the network. The related work is summarized as follows.

Human Activity Recognition: Traditionally, in video-based human activity understanding research [1, 24], many approaches make use of local visual features like HOG [17], HOF [17] and MBH [34] to encode appearance information. These features are typically extracted from spatio-temporal keypoints [16] but can also be extracted over dense trajectories [33, 35], which can improve recognition performance. Most recently, it has been shown that the visual feature representation can be learned automatically using a deep convolutional neural network for image understanding tasks [15]. In the realm of action recognition, Si-

mony and Zisserman [30] proposed a two-stream network to capture spatial appearance on still images and temporal motion between frames. Ji *et al.* [12] used 3D convolutions to extract both spatial and temporal features using a one stream network. Wang *et al.* [36] further develops trajectory-pooled deep-convolutional descriptor (TDD) to incorporate both specially designed features and deep-learned features to achieve state-of-the-art results.

First-Person Video Analysis: In a similar fashion to third-person activity analysis, the first-person vision community has also explored various types of visual features for representing human activity. Kitani *et al.* [14] used optical flow-based global motion descriptors to discover ego-action in sports videos. Spriggs *et al.* [32] performed activity segmentation GIST descriptors. Pirsiavash *et al.* [27] developed a composition of HOG features to model object and hand appearance during an activity. Bambach *et al.* [2] used hand regions to understand activity. Fathi *et al.* proposed mid-level motion features and gaze for recognizing ego-centric activities in [6, 7]. To encode first-person videos using those features, the most prevalent representations are BoW and improved Fisher Vector [25]. In [20], Li *et al.* performed a systemic evaluation of features and provided a list of best practices of combining different cues to achieve state-of-the-art results for activity recognition. Similar to third-person vision activity recognition research, there has also been a number of attempts to use CNN for understanding activities in first-person videos. Ryoo *et al.* [29] develops a new pooled feature representation and shows superior performance using CNN as a appearance feature extractor. Poleg *et al.* [28] proposes to use temporal convolutions over optical flow motion fields to index first-person videos. However, a framework to integrate the success of ego-centric features and the power of CNNs is still missing due to challenges of feature diversity and limited training data. In this paper, we aim to design such a framework to address the problem of ego-centric activity recognition.

2. Egocentric Activity Deep Network

We describe our proposed deep network architecture for recognizing activity labels from short video clips taken by an egocentric camera. As we have argued above, the manifestation of an activity can be decomposed into observed appearance (hand and objects) and observed motion (local hand movement and user ego-motion). Based on this decomposition, we develop two base networks: (1) ObjectNet takes a single image as input to determine the *appearance* features of the activity and is trained using *object* labels; (2) ActionNet takes a sequence of optical flow fields to determine the *motion* features of the activity and is trained using *action* labels. Taking the output of both of these networks, we use a late fusion step to concatenate the output of the two networks and uses the joint representation to pre-

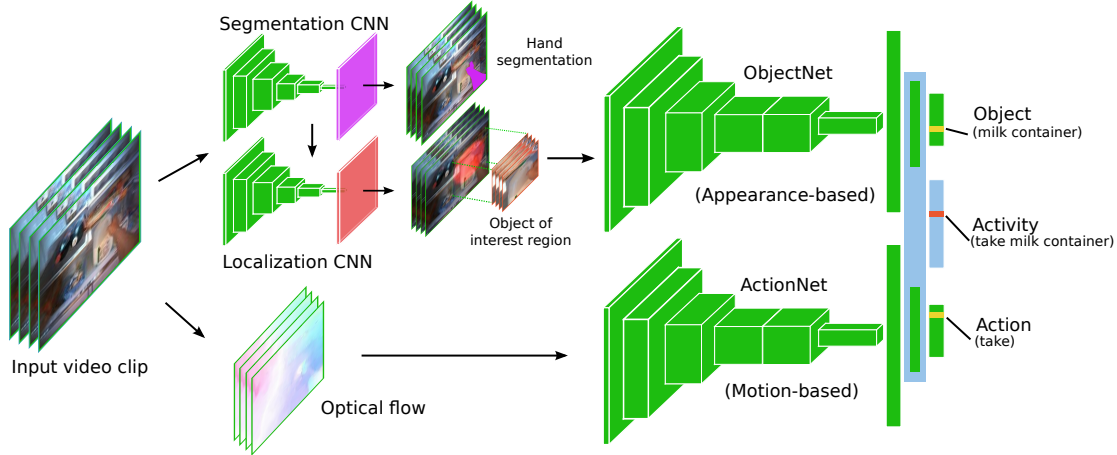


Figure 2: Framework architecture for action, object and activity recognition. Hand segmentation network is first trained to capture hand appearance. It is then fine-tuned to a localization network to localize object of interest. Object CNN and motion CNN are then trained separately to recognize objects and actions. Finally, the two networks are fine-tuned jointly with a triplet loss function to recognize objects, actions and activities. This proposed network beats all baseline models.




Figure 3: Pipeline for localization network training. Hand segmentation network is first trained using images and binary hand masks. Localization network is then fine-tuned from hand segmentation network using images and object location heatmaps synthesized from object locations.

dict three outputs, namely, *action*, *object* and *activity*. More formally, given a short video sequence of N image frames $\mathbf{I} = \{I_1, \dots, I_N\}$, our network predicts three output labels: $\{y_{\text{object}}, y_{\text{action}}, y_{\text{activity}}\}$. The architecture of the entire network is illustrated in Figure 2.

2.1. ObjectNet: Recognizing Objects from Appearance

As shown in [10, 27], recognizing objects in videos is an important aspect of ego-centric activity understanding. We aim to predict the object label y_{object} in this section. To do so, we are particularly interested in the object being interacted with or manipulated – the *object of interest*. However, detecting all objects accurately in the scene is difficult. It also provides limited information about the interested object. Our proposed model will first localize and then recognize the object of interest.

Although we can assume that the object of interest is often located at the center of the subject’s reachable region, it is not always present at the center of the camera image due to head motion. Instead, we observe that the object of interest most frequently appears in the vicinity of hands. A

similar observation was also made in [19, 9]. Besides hand location, hand pose and shape is also important to estimate the manipulation points as shown in [19]. We therefore seek to segment the hands out of the image and use hand appearance to predict the location of the object of interest. We first train a pixel-to-pixel hand segmentation network using raw images and binary hand masks. This network will output a hand probability map. To predict object of interest location using this hand representation, a naive approach is to build a regression model on top. For instance, we can train another CNN or a regressor using features from the hand segmentation network. However, our experiments with this approach achieve low performance due to limited training data. The prediction tends to favor the image center as it is where the object of interest occurs most frequently. Our final pipeline is illustrated in Figure 3. After training the hand segmentation network, we fine-tune a localization network to predict a pixel-level object occurrence probability map. Inspired by previous work in pose estimation [26], we synthesize a heatmap by placing a 2D Gaussian distribution at the location of the object of interest. We use this heatmap as ground-truth and use per-pixel Euclidean loss to train the network. To transfer the hand representation learned from the segmentation network, we initialize the localization network with the weights from the segmentation network and then fine-tune the localization network with the new loss layer. The details of the segmentation and localization CNNs are listed as follows.

(1) Hand segmentation network: For training data, we can either use annotated ground-truth hand masks or output of pixel-level hand detectors like [18]. For the network architecture, we use a low resolution FCN32-s as in [21] as it is a relatively smaller model and converges faster. The loss function for the segmentation network is the sum of




Figure 4: Training data examples for localization CNN. (a) Raw video images with annotated object locations (in red). (b) Ground-truth hand masks which can be annotated manually or generated using hand detectors such as [18]. (c) Synthesized location heat-maps by placing a Gaussian bump at the object location.

per-pixel two-class softmax losses.

(2) Object localization network: For training data, we first manually annotate object of interest locations in the training images of the hand segmentation network. We then synthesize the location heatmaps using a Gaussian distribution as discussed above. Examples of training data are shown in Figure 4. We use the same FCN32-s network architecture and replace the softmax layer with a per-pixel Euclidean loss layer.

To this extent, we have trained an object localization network that outputs a per-pixel occurrence probability map of the object of interest. To generate the final object region proposals, we first run the localization network on input image sequence $\mathbf{I} = \{I_1, I_2, \dots, I_N\}$ and generate probability maps of object locations $\mathbf{H} = \{H_1, H_2, \dots, H_N\}$. We then threshold each probability map and use the centroid of the largest blob as the predicted center of the object. We then crop the object out of the raw image at the predicted center using a fixed-size bounding box. We fix the crop size and ignore the scale difference by observing that the object of interest is always within the reachable distance of the subject. In this way, we generate a sequence of cropped object region images $\mathbf{O} = \{O_1, O_2, \dots, O_N\}$ as the input of the object recognition CNN. The localization result is stable on a per-frame basis, hence there is no temporal smoothness adopted.

With the cropped image sequence of objects of interest $\{\mathbf{O}_i\}$, we then train the object CNN using the model of CNN-M-2048 [3] to recognize the objects. We choose this network architecture due to its high performance on ImageNet image classification. Since a better architecture of base network is out of the scope of this work, we use this network as our base network in this paper unless otherwise mentioned. We adapt it to different tasks (e.g. action recognition) with minimum modifications in this paper. For object recognition, we train the network using $\{(O_i, y_{\text{object}})\}$ pairs as training data and softmax as the loss function. At

testing time, we run the network on the cropped object image O_i to predict object class scores. We then calculate the mean score of all frames in a sequence for each activity class and select the activity label with largest mean score as the final predicted label of object.

Up until now, we have trained a localization network to localize the object of interest by explicitly incorporating hand appearance. Using the cropped images of the localized object of interest, we have trained an object recognition network to predict the object label y_{object} . We will show later that this recognition pipeline also captures important appearance cues such as object attributes and hand appearance. We now move forward to the motion stream of our framework.

2.2. ActionNet: Recognizing Actions from Motion

In this section, our target is to predict the action label y_{action} from motion. Unlike straightforward appearance cues like hands and objects discussed in previous section, motion features in ego-centric videos are more complex because the head motion might cancel the object and hand motion. Although Wang *et al.* [35] shows that compensation of camera motion improves accuracy in traditional action recognition tasks, for ego-centric videos, background motion is often a good estimation of head motion and thus an important cue for recognizing actions [20]. Instead of decoupling foreground (object and hand) motion and background (camera) motion and calculating features separately, we aim to use CNN to capture different local motion features and temporal features together implicitly.

In order to train a CNN network with motion input, we follow [30] to use optical flow images to represent motion information. In particular, given a video sequence of N frames $\mathbf{I} = \{I_1, I_2, \dots, I_N\}$ and corresponding action label y_{action} , we first calculate optical flow of each two consecutive frames and encode the horizontal and vertical flow separately in $\mathbf{U} = \{U_1, U_2, \dots, U_{N-1}\}$ and $\mathbf{V} = \{V_1, V_2, \dots, V_{N-1}\}$. To incorporate temporal information, we use a fixed length of L frames and stack corresponding optical flow images together as input samples of the network noted as $\mathbf{X} = \{X_1, \dots, X_{N-L+1}\}$ where $X_i = \{U_i, V_i, \dots, U_{i+L-1}, V_{i+L-1}\}$.

With motion represented in optical flow images, we train the motion CNN using $\{(X_i, y_{\text{action}})\}$ pairs as training data and softmax as the loss function. At testing time, we run the network on input motion data X_i to predict the scores for each action class. We then average the scores for all frames in the sequence and pick the action class with maximum average score as the predicted label of the action. With the learned representation of objects and actions, we now move to the next section for activity recognition.

2.3. Fusion Layer: Recognizing Activity

In this section, we seek to recognize the activity label y_{activity} given the representations learned from the two network streams in previous sections. A natural approach is to use the two networks as feature extractors and training a classifier using activity labels. However, this approach ignores the co-relation between actions, objects and activities. For instance, if we are confident that the action is *stir* from repeated circular motion, it is highly probable that the object is *tea* or *coffee*. In the other way, if we know the object is *tea* or *coffee*, the probability that the action is *cut* or *fold* should be very low. Based on this intuition, we fuse the action and object networks together as one network by concatenating the second last fully connected layers from the two networks and add another fully connected layer on top. We then add another loss layer for activity on top. The final fused network therefore has three weighted losses: action, object and activity loss. Then weighted sum of three losses is calculated as the overall loss. We can set the weights empirically by the relative importance of three tasks and train one network to learn activity, action and object simultaneously. The loss function for the final network can be formulated as $L_{\text{network}} = w_1 \cdot L_{\text{action}} + w_2 \cdot L_{\text{object}} + w_3 \cdot L_{\text{activity}}$.

To train the fused network, we transfer the weights of the trained motion CNN and object CNN and fine-tune it to recognize the activity. Specifically, given a video sequence of N frames $\mathbf{I} = \{I_1, I_2, \dots, I_N\}$, we follow section 2.1 to localize the objects of interest and get a sequence of object images $\mathbf{O} = \{O_1, O_2, \dots, O_N\}$. We follow section 2.2 to calculate optical flow image pairs $\{\mathbf{U}, \mathbf{V}\}$ and stack them using a fixed length of L frames into $\mathbf{X} = \{X_1, \dots, X_{N-L+1}\}$ where $X_i = (U_i, V_i)$. At training time, for each optical flow blob X_i , we randomly pick a object image O_j where $i \leq j < i + L$ and form the training data pair $(X_i, O_j, y_{\text{action}}, y_{\text{object}}, y_{\text{activity}})$. This is also a way to augment the training data to avoid over-fitting. At testing time, we pick the center object image frame such that $j = (2i + L)/2$ as the annotated boundary of an activity sequence is loose. We run the network on all data pairs to predict the scores for activity. We then average the scores and use the activity class with maximum average score as the predicted activity label.

3. Experiments

We briefly introduce the datasets in Section 3.1 and describe the details for training networks in Section 3.2. We then present experimental results for the three tasks of object recognition (Section 3.3), action recognition (Section 3.4) and activity recognition (Section 3.5).

3.1. Dataset

We run experiments on three public datasets: GTEA, GTEA gaze (Gaze) and GTEA gaze+ (Gaze+) as these datasets were collected using a head-mount camera and

most of the activities involve hand-object interactions. The annotation label for each activity contains a verb (action) and a set of nouns (object). We perform all our experiments using leave-one-subject-out cross-validation. We also report results on fixed-splits following previous work.

GTEA: This dataset [9] contains 7 types of activities performed by 4 different subjects. There are 71 activity categories and 525 instances in the original labels. We report comparative results on two subsets used in previous work [8, 20, 9]: 71 classes and 61 classes. **Gaze:** This dataset [6] contains 17 sequences performed by 14 different subjects. There are 40 activity categories and 331 instances in the original labels. We report results on two subsets used in previous works [20, 6]: 40 classes and 25 classes. **Gaze+:** This dataset [6] contains 7 types of activities performed by 10 different subjects. We report results on a 44 classes subset with 1958 action instances following [20].

3.2. Network Training

For network architecture, we use FCN32-s [21] for hand segmentation and object localization. We use CNN-M-2048 [3] for action and object recognition. Due to the limited sizes of the three public datasets, we adopt the fine-tuning [23] approach to initialize our networks. Specifically, we use available pre-trained models from three large-scale datasets: UCF101[31], Pascal-Context[22] and ImageNet[4] for motion, hand segmentation and object CNN respectively.

Data augmentation. To further address the problem of limited data, we apply data augmentation [15] to improve generalization of CNN networks. **Crop:** All of our network inputs are resized to $K \times C \times 256 \times 256$, where K is batch size, C is input channels. We randomly crop them to $K \times C \times 224 \times 224$ at training time. **Flip:** We randomly mirror input images horizontally. For optical flow frames (U_i, V_i) , we mirror them to $(255 - U_i, V_i)$. **Replication:** We also replicate training data by repeating minority classes to match with majority classes at training time.

Training details. We use a modified version of Caffe [13] and Nvidia Titan X GPU to train our networks. We use stochastic gradient descent with momentum as our optimization method. We use a fixed learning rate of $\gamma = 1e - 8$ for fine-tuning hand segmentation and object localization CNNs, $\gamma = 5e - 4$ for motion CNN and $\gamma = 1e - 4$ for object CNN. For joint training, we lower the learning rate of two sub-networks by a factor of 10. We use batch sizes of 16, 128, 180 for localization, object and motion CNNs respectively.

3.3. ObjectNet Performance

We evaluate the localization network and object recognition network of the ObjectNet stream.

Localizing object of interest. As described in Section 2.1 and illustrated in Figure 3, we first train a hand segmen-




Figure 5: Object localization using Hand Information. Visualization of object location probability map (red) and object bounding box (green). (a: GTEA, b: Gaze+)

tation to learn the bottom layers of the object localization network. The intuition behind this training procedure is to purposefully bias the object localization network to use the hands as evidence to infer an object bounding box. We first train a hand segmentation network using the model of [21] to capture hand appearance information. We then swap out the top layer for segmentation with a top layer optimized for object localization (*i.e.*, fine-tune the network to repurpose it for object localization). The network has an input size of $K \times 3 \times 256 \times 256$ where K is the batch size. After five convolutional ($conv1 - conv5$) layers of 2×2 pooling operations, the input image dimension is down-sampled to $1/32$ of the original size. The final deconvolution layer up-samples it back to the original size of $K \times 2 \times 256 \times 256$. We use raw images and hand masks provided with GTEA and Gaze as training data for the hand segmentation network. Since Gaze+ is not annotated with hand masks, we use [18] to detect hands and use the result to train the network. Once the segmentation network is trained, we use manually annotated training images of object locations to re-purpose the network for localization. Instead of using raw object locations (exact center position of the object), we place a Gaussian bump at the center position to create a heat-map representation as described in Section 2.1.

Figure 5 shows qualitative results of the localization network. The localization network successfully predicts the key object of interest out of other irrelevant objects in the scene. Notice that the result is strongly tied to the hand as the network is *pre-train* for hand segmentation. The results also show that the model can deal seamlessly with different hand configurations like one-hand or two-hand scenarios.

Recognizing object of interest. The localized object images are used to train the object CNN. Table 1 compares the performance of our proposed methods with [5]. Our proposed method dramatically outperforms [5] by 14%. As seen in Table 1 the boost in performance can be attributed to improved localization through the use of hand segmentation-based pre-training.

We visualize the activations of the object recognition network and present two important findings: **(1) Hands are important for object recognition:** Although the localization network is targeted for object of interest, the cropped image also contains a large portion of hands. We visualize the activations of the $conv5$ layer and find that the 50^{th}

Object Recognition	GTEA(71)	Gaze(40)	Gaze+(44)
Fathi <i>et al.</i> [9]	61.36	N/A	N/A
Object CNN	67.74	38.05	61.87
Joint training (Ours)	76.15	55.55	74.34

Table 1: Average object recognition accuracy. Proposed method performs 14% better than the baseline. Joint training of motion and object networks improves accuracy across all datasets.




Figure 6: (a) Top 5 training images with strongest activations from the 50^{th} neuron unit in the $conv5$ layer. (b) 5 test images (top row) and 13×13 activations (bottom row) of the same unit. The visualization shows that this unit responds strongly to hand regions. The object network is capturing hand appearance.

neuron unit responds particularly strongly to training images with large hand regions as shown in Figure 6. We further test the network with test images shown in Figure 6. We observe that the strongest activations overlap with hand regions. We therefore conclude that the object recognition network is learning appearance features from hand regions to help recognize objects. When there is no hand in the scene, the localization network will predict no interacting object. Since some of the iterating objects as tea bags and utensils are small, it is extremely challenging to locate them using an traditional object detector. The hands, their shape and their motion can act as a type of proxy small objects. **(2) Object attributes are important for object recognition:** Figure 7 shows examples of a particular neuron unit responding to particular object attributes like color, texture and shape. In Figure 7b, we observe that this specific neuron is activated when it observes round objects.

3.4. ActionNet performance

We first evaluate the ActionNet to recognize actions. In our experiments, we crop and resize input images to 256×256 and calculate optical flow using OpenCV GPU implementation of [37]. We clip and normalize the flow values from $[-20, 20]$ to $[0, 255]$. We found empirically that $L = 10$ optical flow frames generates good performance.

Table 2 compares our proposed method with the baseline in [5]. While our motion network significantly improves the average recognition accuracy, we are also interested in understanding what the network is learning. Our visual-




Figure 7: Neuron activation in the *conv5* layer for test images. Neuron responding to: (a) transparent bottle, (b) edges of container, (c) cups, (d) white round shapes.

Method & dataset	GTEA(71)	Gaze(40)	Gaze+(44)
Fathi <i>et al.</i> [5]	47.70	N/A	N/A
Motion CNN	75.85	33.65	62.62
Joint training	78.33	36.27	65.05

Table 2: Average action recognition accuracy. Proposed method performs 30% better than the baseline. Joint training of motion and object networks improves accuracy across all datasets.

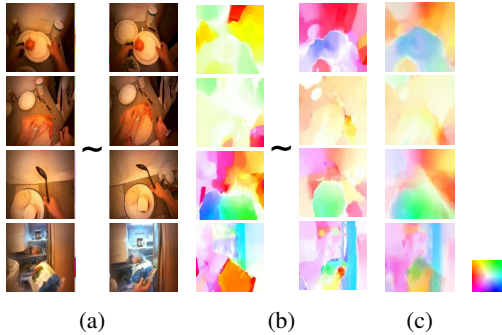


Figure 8: Top 4 training sequences with strongest activations for the 346th neuron unit in *conv5* layer. (a) Start/end image frames, (b) Start/end optical flow images, (c) Average optical flow for each sequence. From top to bottom, ground-truth activity labels are *put_cupPlateBowl*, *put_knife*, *put_cupPlateBowl* and *put_lettuce_container*.

ization shows two important discoveries: (1) our motion network automatically identifies foreground (objects and hands) motion out of complex background (camera) motion (2) our motion network automatically encodes temporal motion patterns.

(1) Camera motion compensation is important for action recognition: As summarized in [20], motion compensation is important for ego-centric action understanding as it provides more reliable foreground motion features. Through visualization, we discover that the network is automatically learning to identify foreground objects and hands. Figure 8 shows top 4 training sequences that activate a particular neuron unit most strongly in the *conv5* layer. All these sequences have the same action verb *put* despite the diversity in camera egomotion. This shows that the network automatically learns to ignore background camera motion for




Figure 9: 4 test sequences and activations of the 346th neuron unit in *conv5* layer. (a) Start/end image frames, (b) Start/end optical flow images, (c) Average optical flow, (d) 13×13 activation maps of the neuron unit using the optical flow sequence, (e) Overlay of activation map on the end image frame, (f) 13×13 activation maps of the neuron unit using reversed optical flow sequence. From top to bottom, ground-truth activity labels are *put_milk_container*, *put_milk_container*, *put_cupPlateBowl*, *put_tomato_cupPlateBowl*.

this neuron. We further test the network with a few test sequences of *put* actions. The results (in Figure 9) agree with our observation in the following aspects: (1) Activation of the same unit is very strong on all these test *put* actions compared with other actions; (2) The strongest activation location coincides roughly with foreground objects and hands location in Figure 9e.

(2) Temporal motion patterns are important for action recognition: While instantaneous motion is an important cue for action recognition, it is also crucial to integrate temporal motion information as shown in [20, 29, 33, 35]. Figure 8 shows that the neuron unit is able to capture the movement of subjects during image sequences. We perform another experiment by reversing the order of the input optical flow images to observe how this neuron responds. Figure 9f shows the activation maps of the same neuron unit with respect to reversed optical flow frames. The weak responses on suggests that the temporal ordering has been encoded in this neuron. This is reasonable, as actions such as *put* and *take* can only be differentiated by preserving temporal ordering.

3.5. Activity recognition

We finally evaluate our framework for the task of activity recognition. In this experiment, we concatenate the two fully connected layers from the ActionNet and ObjectNet and add another fully connected layer on top. Then we fine-tune the two streams together with optical flow images, cropped object images and three weighted losses for three tasks. We compare our results with the state-of-the-art method from Li *et al.* [20] in Table 3. We also report results using the two-stream networks from Simonyan and Zisserman [30] without decomposing activity labels. The confusion matrices are shown in Figure 10. Our proposed method

	Methods	GTEA(61)*	GTEA(61)	GTEA(71)	Gaze(25)*	Gaze(40)*	Gaze+(44)
Li <i>et al.</i> [20]	O+M+E+H	61.10	59.10	59.20	53.20	35.70	60.50
	O+M+E+G	N/A	N/A	N/A	60.90	39.60	60.30
	O+E+H	66.80	64.00	62.10	51.10	35.10	57.40
S. & Z.[30]	temporal-cnn	34.30	30.92	30.33	38.76	22.01	44.45
	spatial-cnn	53.77	41.13	40.16	30.84	18.46	45.97
	temporal+spatial-svm	46.51	35.69	35.81	25.94	22.18	43.23
	temporal+spatial-joint	57.64	51.58	49.65	44.29	34.70	58.77
Ours	object-cnn	60.02	56.49	50.35	47.09	35.56	46.38
	motion+object-svm	53.01	50.45	47.07	28.42	16.00	34.75
	motion+object-joint	75.08	73.02	73.24	62.40	43.42	66.40

Table 3: Quantitative results for activity recognition. (a) Best results reported from Li *et al.* [20]. (b) Two-stream CNN [30] results with single streams, SVM-fusion and joint training. (c) Results from our proposed methods with localized object only, SVM-fusion and joint training. Our proposed joint training model significantly outperforms the two baseline approaches on all datasets. Note that even the network trained using only cropped object images (object-cnn) achieves very promising results thanks to our localization network. (*: fixed split, O: object, M: motion, E: egocentric, H: hand, G: gaze).




Figure 10: Confusion matrices of our proposed method for activity recognition. Improvement on the Gaze dataset is lower due to low video quality and inefficient data. (best view in color)

significantly improves the state-of-the-art performance on all datasets. We conclude that this is due to better representations of action and object from the base motion and appearance streams in our framework. We further analyze two main findings from our experiments.

(1) Joint training is effective: Instead of fixing ActionNet and ObjectNet, and only training stacked layers on top, we jointly train all three networks using three losses as discussed in Section 2.3. This avoids over-fitting the newly added top layers and leads to a joint representation of activities with actions and objects. In our experiments, we set $w_{\text{action}} = 0.2$, $w_{\text{object}} = 0.2$ and $w_{\text{activity}} = 1.0$. We set the activity loss weight higher for faster convergence of activity recognition. We also compare joint training with SVM fusion of two networks in Table 3. Joint training boosts the performance consistently by 27% over all datasets.

(2) Object localization is crucial: We seek to understand the importance of localizing objects by training a network using cropped object images and activity labels. We compare three networks for activity recognition with best results reported in [20]: (1) motion-cnn (temporal-cnn) using optical flow images and activity labels (2) spatial-cnn using raw images and activity labels (3) object-cnn using cropped object images and activity labels. The performance is lower than [20] on three networks as shown in Table 3 because we

are not using any motion or temporal information. However, the performance of object-cnn is surprisingly close, only 9.6% lower (25.5% lower with motion-cnn, 20.6% lower with spatial-cnn) on average. We conclude that localizing the key object of interest is crucial for egocentric activity understanding.

4. Conclusion

We have developed a twin stream CNN network architecture to integrate features that characterize ego-centric activities. We demonstrated how our proposed network jointly learns to recognize actions, objects and activities. We evaluated our model on three public datasets and it significantly outperformed the state-of-the-art methods. We further analyzed what the networks were learning. Our visualizations showed that the networks learned important cues like hand appearance, object attribute, local hand motion and global ego-motion as designed. We believe this will help advance the field of ego-centric activity analysis.

Acknowledgement

This research was funded in part by a grant from the Pennsylvania Department of Health’s Commonwealth Universal Research Enhancement Program and CREST, JST.

References

- [1] J. K. Aggarwal and M. S. Ryoo. Human activity analysis: A review. *ACM Computing Surveys (CSUR)*, 43(3):16, 2011.
- [2] S. Bambach, S. Lee, D. J. Crandall, and C. Yu. Lending a hand: Detecting hands and recognizing activities in complex egocentric interactions. 2015.
- [3] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. *British Machine Vision Conference*, 2014.
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. 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.
- [5] A. Fathi, A. Farhadi, and J. M. Rehg. Understanding egocentric activities. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 407–414. IEEE, 2011.
- [6] A. Fathi, Y. Li, and J. M. Rehg. Learning to recognize daily actions using gaze. In *Computer Vision–ECCV 2012*, pages 314–327. Springer, 2012.
- [7] A. Fathi and G. Mori. Action recognition by learning mid-level motion features. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [8] A. Fathi and J. M. Rehg. Modeling actions through state changes. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 2579–2586. IEEE, 2013.
- [9] A. Fathi, X. Ren, and J. M. Rehg. Learning to recognize objects in egocentric activities. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference On*, pages 3281–3288. IEEE, 2011.
- [10] J. Ghosh, Y. J. Lee, and K. Grauman. Discovering important people and objects for egocentric video summarization. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1346–1353. IEEE, 2012.
- [11] A. Gupta, A. Kembhavi, and L. S. Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(10):1775–1789, 2009.
- [12] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(1):221–231, 2013.
- [13] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014.
- [14] K. M. Kitani, T. Okabe, Y. Sato, and A. Sugimoto. Fast unsupervised ego-action learning for first-person sports videos. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3241–3248. IEEE, 2011.
- [15] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [16] I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2-3):107–123, 2005.
- [17] I. Laptev, M. Marszałek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [18] C. Li and K. M. Kitani. Pixel-level hand detection in egocentric videos. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3570–3577. IEEE, 2013.
- [19] Y. Li, A. Fathi, and J. M. Rehg. Learning to predict gaze in egocentric video. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 3216–3223. IEEE, 2013.
- [20] Y. Li, Z. Ye, and J. M. Rehg. Delving into egocentric actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 287–295, 2015.
- [21] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [22] R. Mottaghi, X. Chen, X. Liu, N.-G. Cho, S.-W. Lee, S. Fidler, R. Urtasun, et al. The role of context for object detection and semantic segmentation in the wild. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 891–898. IEEE, 2014.
- [23] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1717–1724. IEEE, 2014.
- [24] X. Peng, L. Wang, X. Wang, and Y. Qiao. Bag of visual words and fusion methods for action recognition: Comprehensive study and good practice. *Computer Vision and Image Understanding*, 2014.
- [25] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *Computer Vision–ECCV 2010*, pages 143–156. Springer, 2010.
- [26] T. Pfister, J. Charles, and A. Zisserman. Flowing convnets for human pose estimation in videos. *Computer Vision (ICCV), International Conference on*, 2015.
- [27] H. Pirsiavash and D. Ramanan. Detecting activities of daily living in first-person camera views. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2847–2854. IEEE, 2012.
- [28] Y. Poleg, A. Ephrat, S. Peleg, and C. Arora. Compact cnn for indexing egocentric videos. *Workshop on Applications of Computer Vision WACV*, 2016.
- [29] M. S. Ryoo, B. Rothrock, and L. Matthies. Pooled motion features for first-person videos. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [30] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, pages 568–576, 2014.
- [31] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *CRCV-TR-12-01*, 2012.

- [32] E. H. Spriggs, F. De La Torre, and M. Hebert. Temporal segmentation and activity classification from first-person sensing. In *Computer Vision and Pattern Recognition Workshops, 2009. CVPR Workshops 2009. IEEE Computer Society Conference On*, pages 17–24. IEEE, 2009.
- [33] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3169–3176. IEEE, 2011.
- [34] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Dense trajectories and motion boundary descriptors for action recognition. *International journal of computer vision*, 103(1):60–79, 2013.
- [35] H. Wang and C. Schmid. Action recognition with improved trajectories. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 3551–3558. IEEE, 2013.
- [36] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [37] C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime tv-l 1 optical flow. In *Pattern Recognition*, pages 214–223. Springer, 2007.