

Actor and Observer: Joint Modeling of First and Third-Person Videos

Gunnar A. Sigurdsson^{1*} Abhinav Gupta¹ Cordelia Schmid² Ali Farhadi³ Karteek Alahari²
¹Carnegie Mellon University ²Inria[†] ³Allen Institute for Artificial Intelligence
github.com/gsig/actor-observer

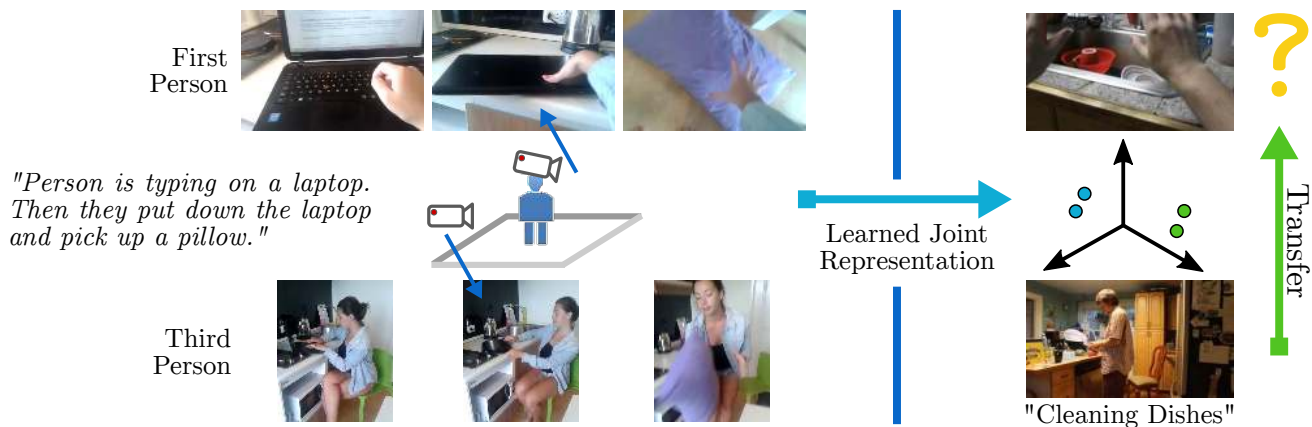


Figure 1: We explore how to reason jointly about first and third-person for understanding human actions. We collect paired data of first and third-person actions sharing the same script. Our model learns a representation from the relationship between these two modalities. We demonstrate multiple applications of this research direction, for example, transferring knowledge from the observer’s to the actor’s perspective.

Abstract

Several theories in cognitive neuroscience suggest that when people interact with the world, or simulate interactions, they do so from a first-person egocentric perspective, and seamlessly transfer knowledge between third-person (observer) and first-person (actor). Despite this, learning such models for human action recognition has not been achievable due to the lack of data. This paper takes a step in this direction, with the introduction of Charades-Ego, a large-scale dataset of paired first-person and third-person videos, involving 112 people, with 4000 paired videos. This enables learning the link between the two, actor and observer perspectives. Thereby, we address one of the biggest bottlenecks facing egocentric vision research, providing a link from first-person to the abundant third-person data on the web. We use this data to learn a joint representation of first and third-person videos, with only weak supervision, and show its effectiveness for transferring knowledge from the third-person to the first-person domain.

*Work was done while Gunnar was at Inria.

[†]Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France.

1. Introduction

What is an action? How do we represent and recognize actions? Most of the current research has focused on a data-driven approach using abundantly available third-person (observer’s perspective) videos. But can we really learn how to represent an action without understanding goals and intentions? Can we learn goals and intentions without simulating actions in our own mind? A popular theory in cognitive psychology, the Theory of Mind [30], suggests that humans have the ability to put themselves in each others’ shoes, and this is a fundamental attribute of human intelligence. In cognitive neuroscience, the presence of activations in mirror neurons and motor regions even for passive observations suggests the same [33].

When people interact with the world (or simulate these interactions), they do so from a first-person egocentric perspective [16]. Therefore, making strides towards human-like activity understanding might require creating a link between the two worlds of data: first-person and third-person. In recent years, the field of egocentric action understanding [14, 20, 22, 27, 32, 34] has bloomed due to a variety of practical applications, such as augmented/virtual reality. While first-person and third-person data represent the two

sides of the same coin, these two worlds are hardly connected. Apart from philosophical reasons, there are practical reasons for establishing this connection. If we can create a link, then we can use billions of easily available third-person videos to improve egocentric video understanding. Yet, there is no connection: why is that?

The reason for the lack of link is the lack of data! In order to establish the link between the first and third-person worlds, we need aligned first and third-person videos. In addition to this, we need a rich and diverse set of actors and actions in these aligned videos to generalize. As it turns out, aligned data is much harder to get. In fact, in the egocentric world, getting diverse actors and, thus, a diverse action dataset is itself a challenge that has not yet been solved. Most datasets are lab-collected and lack diversity as they contain only a few subjects [8, 10, 27].

In this paper, we address one of the biggest bottlenecks facing egocentric vision research. We introduce a large-scale and diverse egocentric dataset, Charades-Ego, collected using the Hollywood in Homes [37] methodology. We demonstrate an overview of the data collection and the learning process in Figure 1, and present examples from the dataset in Figure 2. Our new dataset has 112 actors performing 157 different types of actions. More importantly, we have the same actors perform the same sequence of actions from both first and third-person perspective. Thus, our dataset has semantically similar first and third-person videos. These “aligned” videos allow us to take the first steps in jointly modeling actions from first and third-person’s perspective. Specifically, our model, ActorObserverNet, aligns the two domains by learning a joint embedding in a weakly-supervised setting. We show a practical application of joint modeling: transferring knowledge from the third-person domain to the first-person domain for the task of zero-shot egocentric action recognition.

1.1. Related work

Action recognition from third-person perspective has been extensively studied in computer vision. The most common thread is to use hand-crafted features [17–19] or learn features for recognition using large-scale datasets [4, 38]. We refer the reader to [29, 43] for a detailed survey of these approaches, and in the following we focus on the work most relevant to our approach. Our work is inspired by methods that attempt to go beyond modeling appearances [14, 42]. Our core hypothesis is that modeling goals and intentions requires looking beyond the third-person perspective.

Egocentric understanding of activities. Given recent availability of head-mounted cameras of various types, there has been a significant amount of work in understanding first-person egocentric data [9, 20, 22, 23, 27, 34]. This unique insight into people’s behaviour gives rise to interesting applications such as predicting where people will

look [22], and how they will interact with the environment [31]. Furthermore, it has recently been shown that egocentric training data provides strong features for tasks such as object detection [14].

Datasets for egocentric understanding. Egocentric video understanding has unique challenges as datasets [8, 10, 20, 27] are smaller by an order of magnitude than their third-person equivalents [7, 37]. This is due to numerous difficulties in collecting such data, e.g., availability, complexity, and privacy. Recent datasets have targeted this issue by using micro-videos from the internet, which include both third and first-person videos [25]. While they contain both first and third-person videos, there are no paired videos that can be used to learn the connection between these two domains. In contrast, our dataset contains corresponding first and third-person data, enabling a joint study.

Unsupervised and self-supervised representation learning. In this work, we use the multi-modal nature of the data to learn a robust representation across those modalities. It allows us to learn a representation from the data alone, without any explicit supervision. This draws inspiration from recent work on using other cues for representation learning, such as visual invariance for self-supervised learning of features [1, 14, 21, 24, 26, 39, 41, 42]. For example, this visual invariance can be obtained by tracking how objects change in videos [42] or from consecutive video frames [24]. Typically, this kind of invariance is harnessed via deep metric learning with Siamese (triplet) architectures [5, 11–13, 40, 45].

Data for joint modeling of first and third person. To learn to seamlessly transfer between the first and third-person perspectives we require paired data of these two domains. Some recent work has explored data collected from multiple viewpoints for a fine-grained understanding human actions [15]. Due to the difficulty of acquiring such data, this is generally done in a small-scale lab setting [8, 15], with reconstruction using structure-from-motion techniques [15], or matching camera and head motion of the exact same event [28, 44]. Most related to our work is that of Fan et al. [8] which collects 7 pairs of videos in a lab setting, and learns to match camera wearers between third and first-person. In contrast, we look at thousands of diverse videos collected by people in their homes.

2. Charades-Ego

In order to link first-person and third-person data, we need to build a dataset that has videos shot in first and third-person views. We also need the videos to be semantically aligned, i.e., the same set of actions should appear in each video pair. Collection in a controlled lab setting is difficult to scale, and very few pairs of videos of this type are available on the web. In fact, collection of diverse egocentric data is a big issue due to privacy concerns. So how do we scale such a collection?

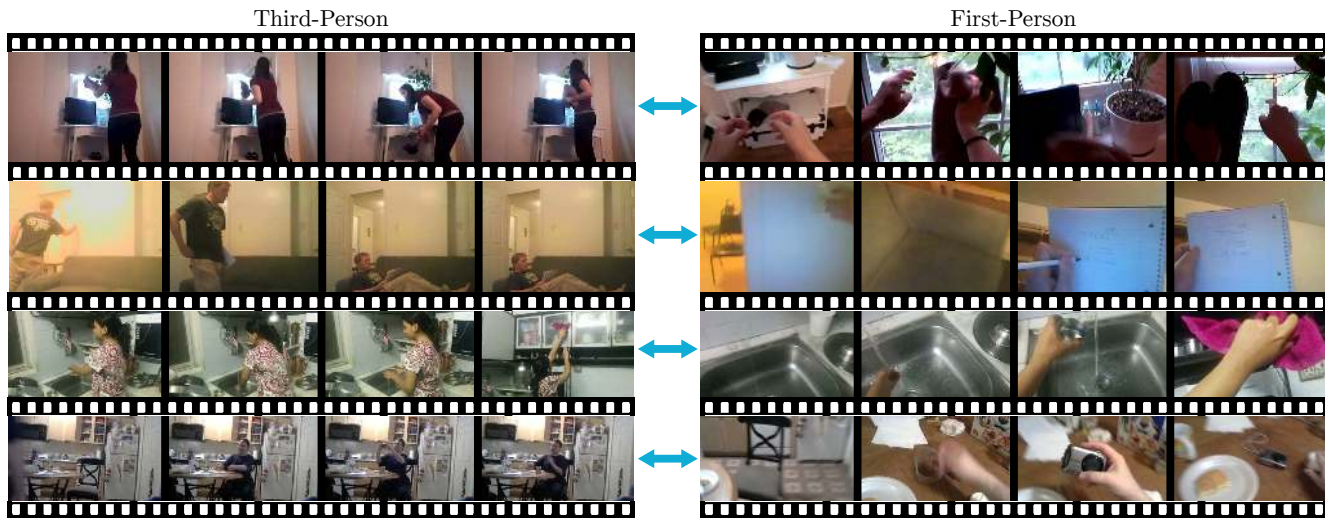


Figure 2: Examples from Charades-Ego, showing third-person (left) and the corresponding first-person (right) video frames.

We introduce the Charades-Ego dataset in this paper. The dataset is collected by following the methodology outlined by the “Hollywood in Homes” approach [37], originally used to collect the Charades dataset [35, 37], where workers on Amazon Mechanical Turk (AMT) are incentivized to record and upload their own videos. This in theory allows for the creation of any desired data.

In particular, to get data that is both in first and third-person we use publicly available scripts from the Charades dataset [37], and ask users to record two videos: (1) one with them acting out the script from the third-person; and (2) another one with them acting out the same script in the same way, with a camera fixed to their forehead. We ensure that all the 157 activity classes from Charades occur sufficiently often in our data. The users are given the choice to hold the camera to their foreheads, and do the activities with one hand, or create their own head mount and use two hands. We encouraged the latter option by incentivizing the users with an additional bonus for doing so.* This strategy worked well, with 59.4% of the submitted videos containing activities featuring both hands, courtesy of a home-made head mount holding the camera.

Specifically, we collected 4000[†] pairs of third and first-person videos (8000 videos in total), with over 364 pairs involving more than one person in the video. The videos are 31.2 seconds long on average. This data contains videos that follow the same structure semantically, i.e., instead of being identical, each video pair depicts activities performed by the same actor(s) in the same environment, and with the same style. This forces a model to latch onto the semantics of the scene, and not only landmarks. We eval-

*We compensated AMT workers \$1.5 for each video pair, and \$0.5 in additional bonus.

[†]Since the scripts are from the Charades dataset, each video pair has another third-person video from a different actor. We use this video also in our work.

uated the alignment of videos by asking workers to identify moments that are shared across the two videos, similar to the algorithmic task in Section 4.3, and found the median alignment error to be 1.3s (2.1s average). This offers a compromise between a synchronized lab setting to record both views simultaneously, and scalability. In fact, our dataset is one of the largest first-person datasets available [8, 10, 20, 27], has significantly more diversity (112 actors in many rooms), and most importantly, is the only large-scale dataset to offer pairs of first and third-person views that we can learn from. Examples from the dataset are presented in Figure 2. Our data is publicly available at github.com/gsig/actor-observer.

3. Jointly Modeling First and Third-Person

As shown in Figure 1, our aim is to learn a shared representation, i.e., a common embedding for data, from the corresponding frames of the first and the third-person domains. In the example in the figure, we have a full view of a person working on a laptop in third-person. We want to learn a representation where the corresponding first-person view, with a close-up of the laptop screen and a hand typing, has a similar representation. We can use the correspondence between first and third-person as supervision to learn this representation that can be effective for multiple tasks. The challenges in achieving this are: the views are very visually different, and many frames are uninformative, such as walls, doors, empty frames, and blurry frames. We now describe a model that tackles these challenges by learning how to select training data for learning a joint representation.

3.1. Formulation

The problem of modeling the two domains is a multi-modal learning problem, in that, data in the first-person view is distinct from data in the third-person view. Following the taxonomy of Baltrusaitis et al. [3] we formulate

this as learning a *coordinated representation* such that corresponding samples in both the first and third-person modalities are close-by in the joint representation. The next question is how to find the alignment or corresponding frames between the two domains. We define ground-truth alignment as frames from first and third-person being within Δ -seconds of each other, and non-alignment as frames being further than Δ' -seconds, to allow for a margin of error.

If a third-person frame x and a first-person frame z map to representations $f(x)$ and $g(z)$ respectively, we want to encourage similarity between $f(x) \sim g(z)$ if their timestamps t_x and t_z satisfy $|t_x - t_z| < \Delta$. If the two frames do not correspond, then we maximize the distance between their learned representations $f(x)$ and $g(z)$. One possible way to now learn a joint representation is to sample all the corresponding pairs of (x, z) , along with a non-corresponding first-person frame z' and use a triplet loss. However, this is not ideal for three reasons: (1) It is inefficient to sample all triplets of frames; (2) Our ground truth (correspondence criteria) is weak as videos are not perfectly synchronized. (3) We need to introduce a mechanism which selects samples that are informative (e.g., hand touching the laptop in Figure 1) and conclusive. These informative samples can also be non-corresponding pairs (negative).

We define the problem of learning the joint representation formally with our loss function l_θ . The loss is defined over triplets from the two modalities (x, z, z') . The overall objective function is given by:

$$L = \mathbb{E}_{(x, z, z') \sim P_\theta} [l_\theta(x, z, z')], \quad (1)$$

where l_θ is a triplet loss on top of ConvNet outputs, and θ is set of all the model parameters. The loss is computed over a *selector* P_θ . We also learn P_θ , a parameterized discrete distribution over data, that represents how to sample more informative data triplets (x, z, z') . Intuitively, this helps us find what samples are likely too hard to learn from. To avoid the degenerate solution where P_θ emphasizes only one sample, we constrain P_θ by reducing the complexity of the function approximator, as discussed in Section 3.2.

The joint model from optimizing the loss and the selector can be used to generate the other view, given either first or third-person view. We illustrate this in Figure 3, where we find the closest first-person frames in the training set, given a third-person query frame. We see that the model is able to connect the two views from the two individual frames, and hallucinate what the person is seeing.

Our setup is related to previous formulations in self-supervised and unsupervised learning, where the pairs (x, z) are often chosen with domain-specific heuristics, e.g., temporal [14, 42] and spatial [6] proximity. Triplet loss is a common choice for the loss l_θ for these tasks [6, 13, 14, 42]. We will now address how we model our loss function with a ConvNet, and optimize it with stochastic gradient descent.

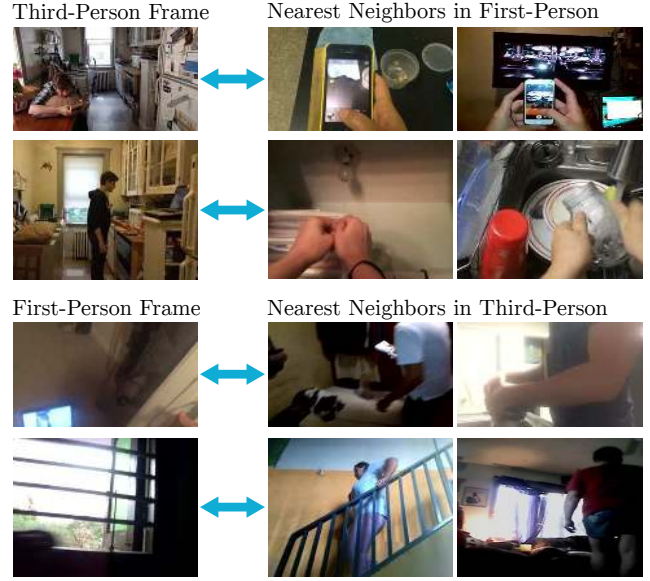


Figure 3: Using our joint first and third-person model we can hallucinate how a scene might look through the eyes of the actor in the scene. The top two rows show nearest neighbours (on the right) from first-person videos. The bottom two rows show the observer’s perspective, given a first-person video frame.

3.2. Optimizing the objective

Optimizing the objective involves learning parameters of both the triplet loss l_θ , as well as the selector P_θ . This correlated training can diverge. We address this by using importance sampling to rewrite the objective L (1) to an equivalent form. We move the distribution of interest P_θ to the objective and sample from a different fixed distribution Q as follows:

$$L = \mathbb{E}_{(x, z, z') \sim Q} \left[\frac{p_\theta(x, z, z')}{q(x, z, z')} l_\theta(x, z, z') \right]. \quad (2)$$

We choose Q to be a uniform distribution over the domain of possible triplets: $\{(x, z, z') \mid |t_x - t_z| < \Delta, |t_x - t'_z| > \Delta'\}$. We uniformly sample frames from first and third-person videos, but re-weight the loss based on the informativeness of the triplet. Here, $p_\theta(x, z, z')$ is the value of the selector for the triplet choice (x, z, z') .

Instead of modeling the informativeness of the *whole* triplet, we make a simplifying assumption. We assume the selector P_θ factorizes as $p_\theta(x, z, z') = p_\theta(x)p_\theta(z)p_\theta(z')$. Further, we constrain P_θ such the probability of selecting any given frame in that video sums to one for a given video. This has similarities with the concept of “bags” in multiple instance learning [2], where we only know whether a given set (bag) of examples contains positive examples, but not if all the examples in the set are positive. Similarly, here we learn a distribution that determines how to select the useful

examples from a set, where our sets are videos. We use a ConvNet architecture to realize our objective.

3.3. Architecture of ActorObserverNet

The ConvNet implementation of our model is presented in Figure 4. It consists of three streams: one for third-person, and two for first-person with some shared parameters. The streams are combined with a L2-based distance metric [13] that enforces small distance between corresponding samples, and large distance between non-corresponding ones:

$$l_{\theta}(x, z, z') = \frac{e^{\|x-z\|_2}}{e^{\|x-z\|_2} + e^{\|x-z'\|_2}}. \quad (3)$$

The computation of the selector value, $p_{\theta}(x, z, z')$, for a triplet (x, z, z') is also done by the three streams. The selector values are the result of a 4096×1 fully-connected layer, followed by a scaled tanh nonlinearity[‡] for each stream. We then define a novel non-linearity, VideoSoftmax, to compute the per-video normalized distribution over frames in different batches, which are then multiplied together to form $p_{\theta}(x)p_{\theta}(z)p_{\theta}(z')$. Once we have the different components of the loss in (2) we add a loss layer (“Final loss” in the figure). This layer combines the triplet loss l_{θ} with the selector output p_{θ} and implements the loss in (2). All the layers are implemented to be compatible with SGD [36]. More details are provided in the supplementary material.

VideoSoftmax layer. The distribution P_{θ} is modeled with a novel layer which computes a probability distribution across multiple samples corresponding to the same video, even if they occur in different batches. The selector value for a frame x is given by:

$$p_{\theta}(x) = \frac{e^{f_{\theta}(x)}}{\sum_{x' \in \mathcal{V}} e^{f_{\theta}(x')}} \quad (4)$$

where $f_{\theta}(x)$ is the input to the layer and denominator is the sum of $e^{f_{\theta}(x')}$ computed over all frames x' in the same video \mathcal{V} . This intuitively works like a softmax function, but across frames in the same video.

Since triplet loss l_{θ} is weighted by the output of the selector, the gradient updates with respect to the triplet loss are simply a weighted version of the original gradient. The gradient for optimizing the loss in (2) with respect to the selector in (4) is (with slight abuse of notation for simplicity):

$$\frac{\partial L}{\partial f} \propto p_{\theta}(x, z, z')(l_{\theta}(x, z, z') - L), \quad (5)$$

where the gradient is with respect to the input of the VideoSoftmax layer f , so we can account for the other samples in the denominator of (4). Q is defined as a constant

[‡]The choice of Tanh nonlinearity makes the network more stable than unbounded alternatives like ReLU.

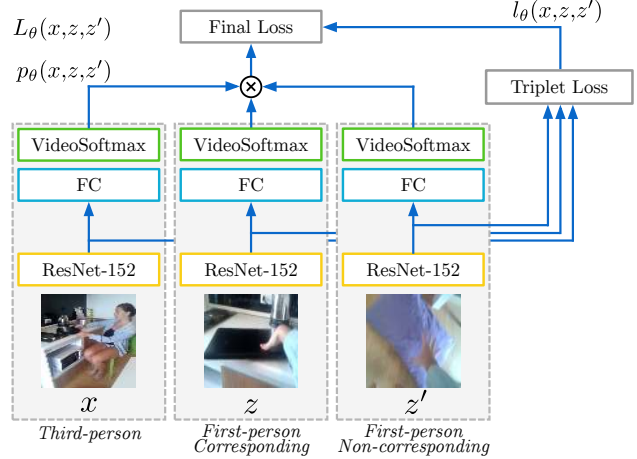


Figure 4: Illustration of our ActorObserverNet. The model has separate streams for first and third-person. Given a triplet of frames from these two modalities, the model computes their fc7 features, which are used to compare and learn their similarity. The FC and the VideoSoftmax layers also compute the likelihood of this sample with respect to the selector P_{θ} .

over the domain, and can be ignored in the derivation. The intuition is that this decreases the weight of the samples that are above the loss L (1), and increases it otherwise. Our method is related to mining easy examples. The selector learns to predict the relative weight of each triplet, i.e., instead of using the loss directly to select triplets (as in mining hard examples). The gradient is then scaled by the magnitude of the weight. The average loss L is computed across all the frames; see supplementary material for more details.

4. Experiments

We demonstrate the effectiveness of our joint modeling of first and third-person data through several applications, and also analyze what the model is learning. In Section 4.2 we evaluate the ability of the joint model to discriminate correct first and third-person pairs from the incorrect ones. We investigate how well the model localizes a given first-person moment in a third-person video, from the same as well as users, by temporally aligning a one-second moment between the two videos (Section 4.3). Finally, in Section 4.4 we present results for transferring third-person knowledge into the first-person modality, by evaluating zero-shot first-person action recognition. We split the 8000 videos into 80% train/validation, and 20% test for our experiments.


4.1. Implementation details

Our model uses a ResNet-152 video frame classification architecture, pretrained on the Charades dataset [37], and shares parameters between both the first and third-person streams. This is inspired by the two-stream model [38], which is a common baseline architecture even in ego-centric

Good First-Person



Bad First-Person



Good Third-Person



Bad Third-Person



Figure 5: A selection of frames, from third and first-person videos, the model assigns the highest and the lowest weights, i.e., $p_\theta(x)$ and $p_\theta(z)$ from (2) respectively. This provides intuition into what the model is confident to learn from.

videos [8, 23]. The scale of random crops for data augmentation in training was set to 80–100% for first-person frames, compared to the default 8–100% for third-person frames. We set the parameter Δ for the maximum distance to determine a positive pair as one second (average alignment error in the dataset), and the parameter Δ' for the negative pair as 10 seconds. More details about the triplet network are available in the supplementary material.

We sample the training data triplets, in the form of a positive pair with first and third-person frames, which correspond to each other, and a negative pair with the same third-person frame and an unrelated first-person frame from the same video. This sampling is done randomly following the uniform distribution Q in (2). The scales of tanh are constrained to be positive. For the experiments in Sections 4.3 and 4.4, the parameters of the fully connected layers for the two first-person streams are shared. Our code is implemented in the PyTorch machine learning framework and is available at github.com/gsig/actor-observer.

4.2. Mapping third-person to first-person

The first problem we analyze is learning to model first and third-person data jointly, which is our underlying core problem. We evaluate the joint model by finding a corresponding first-person frame, given a third-person frame, under two settings: (1) using the whole test set (‘All test data’); and (2) when the model assigns weights to each sample (‘Choose $X\%$ of test data’). In the second case, the triplets with the top 5%, 10%, or 50% highest weights are evaluated. Each triplet contains a given third-person frame, and a positive and negative first-person frames. This allows the model to choose examples from the test set to evaluate.

From Table 1 we see that the original problem (‘All test data’) is extremely challenging, even for state-of-the-art representations. The baseline results are obtained with models using fc7 features from either ResNet-152 trained on ImageNet or a two-stream network (RGB stream using

	Random	ImageNet ResNet-152	Charades Two-Stream	ActorObserverNet
Same person				
All test data	50.0	53.6	55.5	51.7
Choose 50% of test data	50.0	55.7	60.2	73.9
Choose 10% of test data	50.0	57.9	68.8	97.2
Choose 5% of test data	50.0	56.5	71.9	96.8
Different persons				
All test data	50.0	50.6	51.7	50.4
Choose 50% of test data	50.0	50.4	51.6	76.3
Choose 10% of test data	50.0	49.6	50.8	98.8
Choose 5% of test data	50.0	45.6	51.4	98.3

Table 1: Given a third-person frame, we determine whether a first-person frame corresponds to it. Results are shown as correspondence classification accuracy (in %). Higher is better. See Section 4.2 for details.

ResNet-152 from [37]) trained on Charades to compute the loss. The baselines use the difference in distance between positive and negative pairs as the weight used to pick what samples to evaluate on in the second setting.

The results of the two-Stream network (‘Charades Two-Stream’ in the table) and our ActorObserverNet using all test data (‘All test data’) are similar, but still only slightly better than random chance. This is expected, since many of the frames correspond to occluded human actions, people looking at walls, blurry frames, etc., as seen in Figure 5. On the other hand, our full model, which learns to weight the frames (‘Choose $X\%$ of test data’ in the table), outperforms all the other methods significantly. Note that our model assigns a weight for each image frame independently, and in essence, learns if it is a good candidate for mapping. We observe similar behavior when we do the mapping with third and first-person videos containing the same action performed by different people (‘Different persons’ in the table).

Figure 5 shows a qualitative analysis to understand what the model is learning. Here, we illustrate the good and the bad frames chosen by the model, according to the learned weights, both in the third and first-person cases. We ob-



Figure 6: Conv5 activations of ActorObserverNet. The colors range from blue to red, denoting low to high activations. We observe the network attending to hands, objects, and the field of view.




Figure 7: By backpropagating the similarity loss to the image layer, we can visualize what regions the model is learning from. The colors range from blue to red, denoting low to high importance.

serve that the model learns to ignore frames without objects and people, and blurry, feature-less frames, such as the ones seen in the bottom row in the figure. Furthermore, our model prefers first-person frames that include hands, and third-person frames with the person performing an action, such as answering a phone or drinking; see frames in the top row in the figure.

Quantitatively, we found that 68% of high-ranked and only 15% of low-ranked frames contained hands. This is further highlighted in Figures 6 and 7 where we visualize conv5 activations, and gradients at the image layer, respectively. We observe the network attending to hands, objects, and the field of view. Figure 8 illustrates the selection over a video sequence. Here, we include the selector value of $p_\theta(z)$ for each frame in a first-person video. The images highlight points in the graph with particularly useful/useless frames. In general, we see that the weights vary across the video, but the high points correspond to useful moments in the first-person video (top row of images), for example, with a clear view of hands manipulating objects.

4.3. Alignment and localization

In the second experiment we align a given first-person moment in time, i.e., a set of frames in a one-second time interval, with a third-person video, and evaluate this tempo-

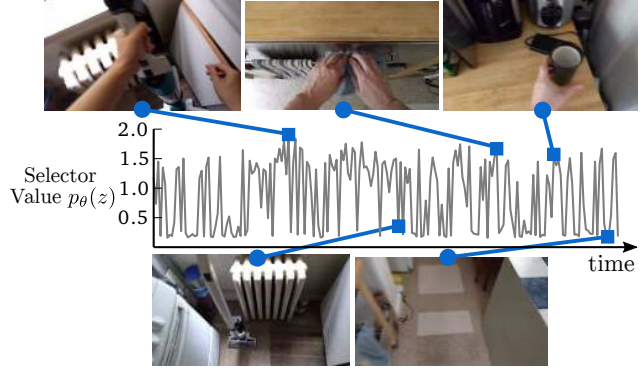


Figure 8: Our model learns to assign weights to all the frames in both third and first-person videos. Here we show the selector value $p_\theta(z)$ (the importance of each frame) for a sample first-person video, and highlight frames assigned with high and low values. See Section 4.2 for details.

	Random Chance	Human	ImageNet ResNet-152	Charades Two-Stream	ActorObserverNet
Same person	11.0	1.3	8.3	6.5	5.2
Different persons	11.0	1.3	8.7	7.0	6.1

Table 2: Alignment error in seconds for our method ‘ActorObserverNet’ and baselines. Lower is better. See Section 4.3 for details.

ral localization. In other words, our task is to find any one-second moment that is shared between those first and third-person perspectives, thus capturing their *semantic similarity*. This allows for evaluation despite uninformative frames and approximate alignment. For evaluation, we assume that the ground truth alignment can be approximated by temporally scaling the first-person video to have the same length as the third-person video.

If m denotes all the possible one-second moments in a first-person and n in a third-person video, there are $m \times n$ ways to pick a pair of potentially aligned moments. Our goal is to pick the pair that has the best alignment from this set. The moments are shuffled so there is no temporal context. We evaluate this chosen pair by measuring how close these moments are temporally, in seconds, as shown in Table 2. To this end, we use our learned model, and find one-second intervals in both videos that have the lowest sum of distances between the frames within this moment. We use L2 distance between fc7 features in these experiments.

We present our alignment results in Table 2, and compare with other methods. These results are reported as median alignment error in seconds. The performance of fc7 features from the ImageNet ResNet-152 network is close to that of a random metric (11.0s). ‘Two-Stream’, which refers to the performance of RGB features from the two-stream network trained on the Charades dataset, performs better. Our ‘ActorObserverNet’ outperforms all these methods.

We visualize the temporal alignment between a pair of

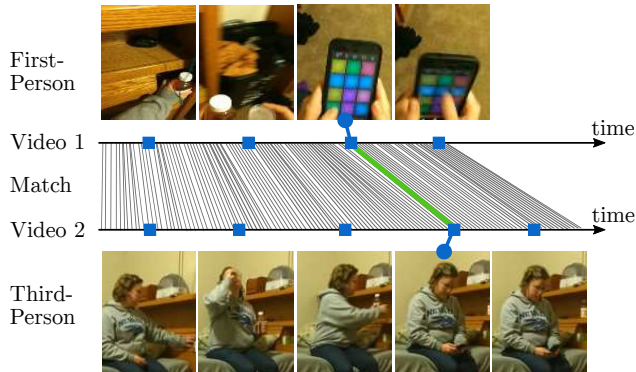


Figure 9: Our model matches corresponding moments between two videos. We find the moment in the third-person video (bottom row) that best matches (shown in green) our one second first-person moment (top row), along with other possible matches (gray). (*Best viewed in pdf.*)

	Random	Charades VGG-16	Charades ResNet-152	ActorObserverNet
Accuracy	8.9	17.8	22.7	25.9

Table 3: Egocentric action recognition in the zero-shot learning setup. We show the video-level mAP on our Charades-Ego dataset. Higher is better. See Section 4.4 for details.

videos in Figure 9. We highlight in green the best moment in the video chosen by the model: the person looking at their cell phone in the third-person view, and a close-up of the cell phone in the first-person view.

4.4. Zero-shot first-person action recognition

Since our ActorObserverNet model learns to map between third and first-person videos, we use it to transfer knowledge acquired from a dataset of third-person videos, annotated with action labels, to the first-person perspective. In essence, we evaluate first-person action recognition in a zero-shot setting. We annotated first-person videos in the test set with the 157 categories from Charades [37] to evaluate this setup. Following the evaluation setup from Charades, we use the video-level multi-class mean average precision (mAP) measure.

In order to transfer knowledge from the third-person to the first-person perspective, we add a classification loss to the third-person model after the fc7 layer. To train this framework, we use third-person training examples from the Charades dataset, in addition to the training set from our Charades-Ego dataset. Note that the third-person videos from Charades are annotated with action labels, while our data only has unlabelled first/third person pairs. Thus, we use the mapping loss in (2) when updating the network parameters due to first/third person pair, and the RGB component of the two-stream classification loss for an update due

to a Charades third-person example.

Our model now learns to not only map both first and third-person frames to a shared representation, but also a third-person activity classifier on top of that shared representation. At test time, we make a prediction for each frame in a first-person test video, and then combine predictions over all the video frames with mean pooling. We present the results in Table 3.

Baseline results. The performance of random chance is 8.9% on the Charades-Ego dataset. We also compare to the RGB two-stream model trained on Charades (third-person videos), using both VGG-16 and ResNet-152 architectures, which achieve 18.6% and 22.8% mAP respectively, on the Charades test set. Both are publicly available [37], and show a 8.9% and 13.8% improvement respectively, over random chance on our first-person videos.

Our results. Our ActorObserverNet further improves over the state-of-the-art two-stream network by 3.2%. This shows that our model can transfer knowledge effectively from the third-person to the first-person domain.

To further analyze whether the gain in performance is due to a better network, or third to first-person transfer, we evaluated our network on the Charades test set. It achieves 23.5% on third-person videos, which is only 0.7% higher than the original model, which suggests that the performance gain is mainly due to the new understanding of how third-person relates to first-person view.

5. Summary

We proposed a framework towards linking the first and third-person worlds, through our novel Charades-Ego dataset, containing pairs of first and third-person videos. This type of data is a first big step in bringing the fields of third-person and first-person activity recognition together. Our model learns how to jointly represent those two domains by learning a robust triplet loss. Semantic equivalence in data allows it to relate the two perspectives from different people. Our results on mapping third-person to first-person, alignment of videos from the two domains, and zero-shot first-person action recognition clearly demonstrate the benefits of linking the two perspectives.

Acknowledgments. This work was supported by Intel via the Intel Science and Technology Center for Visual Cloud Systems, Sloan Fellowship to AG, the Inria associate team GAYA, the ERC advanced grant ALLEGRO, gifts from Amazon and Intel, and the Indo-French project EVEREST (no. 5302-1) funded by CEFIPRA. The authors would like to thank Achal Dave, Vicky Kalogeiton, Kris Kitani, Nick Rhinehart, Jardin du Thé for their invaluable suggestions and advice, and the Amazon Mechanical Turk workers for their time.

References

- [1] P. Agrawal, J. Carreira, and J. Malik. Learning to see by moving. In *ICCV*, 2015. 2
- [2] S. Andrews, I. Tsochantaridis, and T. Hofmann. Support vector machines for multiple-instance learning. In *NIPS*, 2003. 4
- [3] T. Baltrušaitis, C. Ahuja, and L.-P. Morency. Multimodal machine learning: A survey and taxonomy. *PAMI*, 2018. 3
- [4] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 2
- [5] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005. 2
- [6] C. Doersch, A. Gupta, and A. A. Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 4
- [7] B. G. Fabian Caba Heilbron, Victor Escorcia and J. C. Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 2
- [8] C. Fan, J. Lee, M. Xu, K. K. Singh, Y. J. Lee, D. J. Crandall, and M. S. Ryoo. Identifying first-person camera wearers in third-person videos. In *CVPR*, 2017. 2, 3, 6
- [9] A. Fathi, A. Farhadi, and J. M. Rehg. Understanding egocentric activities. In *ICCV*, 2011. 2
- [10] A. Fathi, X. Ren, and J. M. Rehg. Learning to recognize objects in egocentric activities. In *CVPR*, 2011. 2, 3
- [11] Y. Gong, Y. Jia, T. K. Leung, A. Toshev, and S. Ioffe. Deep convolutional ranking for multilabel image annotation. *ICLR*, 2014. 2
- [12] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 2
- [13] E. Hoffer, I. Hubara, and N. Ailon. Spatial contrasting for deep unsupervised learning. *NIPS Workshop*, 2016. 2, 4, 5
- [14] D. Jayaraman and K. Grauman. Learning image representations tied to ego-motion. In *ICCV*, 2015. 1, 2, 4
- [15] H. Joo, H. Liu, L. Tan, L. Gui, B. Nabbe, I. Matthews, T. Kanade, S. Nobuhara, and Y. Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *ICCV*, 2015. 2
- [16] T. Kanade and M. Hebert. First-person vision. *Proc. IEEE*, 2012. 1
- [17] A. Klaser, M. Marszalek, and C. Schmid. A spatio-temporal descriptor based on 3d-gradients. In *BMVC*, 2008. 2
- [18] I. Laptev. On space-time interest points. *IJCV*, 2005. 2
- [19] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *CVPR*, 2008. 2
- [20] Y. J. Lee, J. Ghosh, and K. Grauman. Discovering important people and objects for egocentric video summarization. In *CVPR*, 2012. 1, 2, 3
- [21] Y. Li, M. Paluri, J. M. Rehg, and P. Dollár. Unsupervised learning of edges. In *CVPR*, 2016. 2
- [22] Y. Li, Z. Ye, and J. M. Rehg. Delving into egocentric actions. In *CVPR*, 2015. 1, 2
- [23] M. Ma, H. Fan, and K. M. Kitani. Going deeper into first-person activity recognition. In *CVPR*, 2016. 2, 6
- [24] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. In *ICLR*, 2016. 2
- [25] P. X. Nguyen, G. Rogez, C. Fowlkes, and D. Ramanan. The open world of micro-videos. *arXiv*, 2016. 2
- [26] D. Pathak, R. Girshick, P. Dollár, T. Darrell, and B. Hariharan. Learning features by watching objects move. In *CVPR*, 2017. 2
- [27] H. Pirsiavash and D. Ramanan. Detecting activities of daily living in first-person camera views. In *CVPR*, 2012. 1, 2, 3
- [28] Y. Poleg, C. Arora, and S. Peleg. Head motion signatures from egocentric videos. In *ACCV*, 2014. 2
- [29] R. Poppe. A survey on vision-based human action recognition. *IVC*, 2010. 2
- [30] D. Premack and G. Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1978. 1
- [31] N. Rhinehart and K. M. Kitani. Learning action maps of large environments via first-person vision. In *CVPR*, 2016. 2
- [32] N. Rhinehart and K. M. Kitani. First-person activity forecasting with online inverse reinforcement learning. In *ICCV*, 2017. 1
- [33] G. Rizzolatti and L. Craighero. The mirror-neuron system. *Annu. Rev. Neurosci.*, 2004. 1
- [34] M. S. Ryoo and L. Matthies. First-person activity recognition: What are they doing to me? In *CVPR*, 2013. 1, 2
- [35] G. A. Sigurdsson, J. Choi, A. Farhadi, and A. Gupta. Charades challenge 2017. <http://vuchallenge.org/charades.html>, 2017. 3
- [36] G. A. Sigurdsson, S. Divvala, A. Farhadi, and A. Gupta. Asynchronous temporal fields for action recognition. In *CVPR*, 2017. 5
- [37] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016. 2, 3, 5, 6, 8
- [38] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014. 2, 5
- [39] N. Srivastava, E. Mansimov, and R. Salakhutdinov. Unsupervised learning of video representations using LSTMs. *ICML*, 2015. 2
- [40] J. Wang, Y. Song, T. Leung, C. Rosenberg, J. Wang, J. Philbin, B. Chen, and Y. Wu. Learning fine-grained image similarity with deep ranking. In *CVPR*, 2014. 2
- [41] W. Wang, R. Arora, K. Livescu, and J. Bilmes. On deep multi-view representation learning. In *ICML*, 2015. 2
- [42] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015. 2, 4
- [43] D. Weinland, R. Ronfard, and E. Boyer. A survey of vision-based methods for action representation, segmentation and recognition. *CVIU*, 2011. 2
- [44] R. Yonetani, K. M. Kitani, and Y. Sato. Ego-surfing first person videos. In *CVPR*, 2015. 2
- [45] S. Zagoruyko and N. Komodakis. Learning to compare image patches via convolutional neural networks. In *CVPR*, 2015. 2