Using Cross-Model EgoSupervision to Learn Cooperative Basketball Intention

Gedas Bertasius University of Pennsylvania gberta@seas.upenn.edu

Abstract

We present a first-person method for cooperative basketball intention prediction: we predict with whom the camera wearer will cooperate in the near future from unlabeled first-person images. This is a challenging task that requires inferring the camera wearer's visual attention, and decoding the social cues of other players. Our key observation is that a first-person view provides strong cues to infer the camera wearer's momentary visual attention, and his/her intentions. We exploit this observation by proposing a new cross-model EgoSupervision learning scheme that allows us to predict with whom the camera wearer will cooperate in the near future, without using manually labeled intention labels. Our cross-model EgoSupervision operates by transforming the outputs of a pretrained pose-estimation network, into pseudo ground truth labels, which are then used as a supervisory signal to train a new network for a cooperative intention task. We evaluate our method, and show that it achieves similar or even better accuracy than the fully supervised methods do.

1. Introduction

Consider a dynamic scene such as Figure 1, where you, as the camera wearer, are playing basketball. You need to make a decision with whom you will cooperate to maximize the overall benefit for your team. Looking ahead at your teammates, you make a conscious decision and then 2-3 seconds afterwards you perform a cooperative action such as passing the ball.

In a team sport such as basketball, an effective cooperation among teammates is essential. Thus, in this paper, we aim to investigate whether we can use a single first-person image to infer with whom the camera wearer will cooperate 2-3 seconds from now? This is a challenging task because predicting camera wearer's cooperative intention requires 1) inferring his/her momentary visual attention, 2) decoding dominant social signals expressed by other players who want to cooperate, and 3) knowing who your teammates are when the players are not wearing any team-specific uniJianbo Shi University of Pennsylvania jshi@seas.upenn.edu



Figure 1: With whom will I cooperate after 2-3 seconds? Given an **unlabeled** set of first-person basketball images, we predict with whom the camera wearer will cooperate 2 seconds from now. We refer to this problem as a cooperative basketball intention prediction.

forms.

To make this problem even more challenging we ask a question: "Can we infer cooperative basketball intention without manually labeled first-person data?". Building an unsupervised learning framework is important because manually collecting basketball intention labels is a costly and a time consuming process. In the context of a cooperative basketball intention task, an annotator needs to have highly specific basketball domain knowledge. Such a requirement limits the scalability of the annotation process because such annotators are difficult to find and costly to employ.

However, we conjecture that we can learn cooperative basketball intention in an unsupervised fashion by exploiting the signal provided by the first-person camera. What people see reflects how they are going to act. A firstperson camera placed on a basketball player's head allows



Figure 2: The illustration of our cross-model EgoSupervision training scheme. As our base model we use a multi-person pose estimation network from [6], which predicts 1) pose estimates of all people in a given first-person image and 2) the bounding boxes around each person. Next, we feed these outputs to an EgoTransformer, which transforms them such that the transformed output would approximately capture the camera wearer's attention and intentions. Then, we use such transformed output as a supervisory signal to train the network for our cooperative basketball intention task.

us to indirectly tap into that person's mind and reason about his/her internal state based on what the camera wearer sees. To do so we propose a novel cross-model EgoSupervision learning scheme, which allows us to learn the camera wearer's intention without the manually labeled intention data. Our cross-model EgoSupervision scheme works as follows. First we transform the output of a pretrained poseestimation network such that it would approximately reflect the camera wearer's internal state such as his/her visual attention and intentions. Then, we use such transformed output as a supervisory signal to train another network for our cooperative basketball intention task. We show that such a learning scheme allows us to train our model without manually annotated intention labels, and achieve similar or even better results as the fully supervised methods do.

2. Related Work

First-Person Vision. In the past, most first-person methods have focused on first-person object detection [29, 10, 40, 15, 2], or activity recognition [44, 43, 38, 31, 35, 13]. Several methods have employed first-person videos to summarize videos [29, 34] while recently the work in [46] proposed to predict the camera wearer's engagement detection from first-person videos. The work in [14] used a group of people wearing first-person cameras to infer their social interactions such as monologues, dialogues, or discussions. The method in [37] predicted physical forces experienced by the camera wearer, while the work in [26] recognized the activities performed in various extreme sports. Several recent methods [36, 45] also predicted the camera wearer's movement trajectories. Finally, first-person cameras have also been used for various robotics applications [41, 18]

In comparison to these prior methods, we propose a novel cooperative basketball intention prediction task, that allows us to study cooperative behaviors of the basketball players. Furthermore, we note that these prior first-person methods (except [26]) rely on manually annotated labels for their respective tasks whether it would be an objectdetection, activity recognition, intention prediction or some other task. Instead, in this work, we demonstrate that we can solve a challenging cooperative basketball intention prediction task without using annotated first-person intention labels, which are time consuming and costly to obtain.

Knowledge Transfer across Models. With the introduction of supervised CNN models [27], there has been a lot of interest in adapting generic set of features [11] for different tasks at hand [22, 3, 16, 47, 39, 42]. Recently, generic image classification features were successfully used for the tasks such as edge detection [3, 47], object detection [16, 39, 42], and semantic segmentation [4, 32, 33, 7]. More related to our work, a recent line of research investigated how to transfer knowledge across different models by a combination of parameter updates [1, 12, 24], transformation learning [28, 17], network distillation [21] or crossmodel supervision [23, 19]. The most similar to our work are the methods in [23, 19] that use cross-model supervision to transfer knowledge from one model to another.

All of the above methods focus on the third-person data. In contrast, we show how to exploit a first-person view to solve a novel camera wearer's cooperative intention prediction task without using manually labeled first-person data.

3. Learning Cooperative Basketball Intention

The goal of our cooperative basketball intention task is to predict with whom the camera wearer will cooperate in the near future. Formally, we aim to learn a function $g(I_i)$ that takes a single first-person image I_i as an input and outputs a per-pixel likelihood map, where each pixel indicates the cooperation probability. Ideally, we would want such function to produce high probability values at pixels around the person with whom the camera wearer will cooperate, and low probability values around all the other pixels.

We implement $g(I_i)$ via a fully convolutional neural network based on the architecture of a multi-person pose estimation network in [6]. Let \hat{y} denote a per-pixel mask that is given to our network as a target label. We refer to \hat{y} as a *pseudo* ground truth because we obtain it automatically instead of relying on the manually annotated intention labels. Then, we learn our cooperative basketball intention model by optimizing the following cross-entropy loss objective:

$$\mathcal{L}^{(i)} = -\sum_{j=1}^{N} \hat{y}_{j}^{(i)} \log g_{j}(I_{i}) + (1 - \hat{y}_{j}^{(i)}) \log (1 - g_{j}(I_{i})),$$
(1)

where $\hat{y}_j^{(i)}$ is the pseudo ground truth value of image I_i at pixel j, $g_j(I_i)$ refers to our network's output at pixel j, and N denotes the number of pixels in an image. We now explain how we obtain the pseudo ground truth data \hat{y} .

3.1. EgoTransformer

To construct a pseudo ground truth supervisory signal \hat{y} , we transform the output of a pretrained multi-person pose estimation network [6], such that it would approximately capture the camera wearer's internal state such as his/her visual attention, and intentions. We do so using our proposed EgoTransformer scheme.

Let $f(I_i)$ denote a pretrained fully convolutional network from [6] that takes a first-person image as an input, and outputs the 1) pose part estimates of every person in an image, and 2) their bounding-box detections. We note that the pretrained network f was never trained on any first-person images. Then, formally, let $B \in \mathbb{R}^{n \times 5}$ denote the bounding box of people detected by f. Each of n detected bounding boxes is parameterized by 5 numbers (x, y, h, w, c) denoting the top-left bounding-box coordinates (x, y), the height h, and width w of the bounding box, and its confidence value c. Additionally, let $P \in \mathbb{R}^{n \times 18 \times 2}$ denote the predicted (x, y) locations of 18 pose parts (see [6]) for each of n detected people.

Then our goal is to come up with a transformation function $T(B^{(i)}, P^{(i)})$ that takes these two outputs and transforms them into a per-pixel pseudo ground truth mask $\hat{y}^{(i)}$ for our cooperative basketball intention prediction task.



First-Person RGB Pseudo GT Ground Truth

Figure 3: Qualitative comparison of the pseudo ground truth labels obtained via an EgoTransformer versus the actual ground truth. Note that while the pseudo ground truth is not always correct (see the third row), in most cases, it successfully assigns high values around the player with whom the camera wearer will cooperate (see the first two rows).

We do so by exploiting three different characteristics encoded in a first-person view: 1) egocentric location prior, 2) egocentric size prior, and 3) egocentric pose prior. All of these characteristics can be used to reason about the camera wearer's internal state.

For instance, the location where another person is detected in a first-person image can be used to assess how likely the camera wearer is looking at that person [31, 2]. The size of another person in a first-person image can be used to infer how far the camera wearer is from that person, and hence, how likely will the camera wearer interact with that person (the nearer the more likely). Finally, most person-to-person interactions involve people looking at each other, which imposes a certain pose prior. We can then use such a pose prior to predict whether two people will cooperate with each other in the near future based on whether another person is looking at the camera wearer at present.

We express our pseudo ground truth data \hat{y} using these three characteristics using what we refer to as an EgoTransformer scheme:

$$\hat{y} = \left[\sum_{j=1}^{n} V(B_j, \phi_{size}(B_j)) \cdot V(B_j, \phi_{pose}(B_j))\right] \cdot \phi_{loc}(B)$$
(2)

where *n* denotes the number of detected bounding boxes in a given image, B_j depicts a j^{th} bounding box, *V* is a function that takes two inputs: 1) a bounding box B_j , and 2) a scalar value *v*, and outputs a $H \times W$ dimensional mask by assigning every pixel inside this bounding box B_j to *v*, and



Figure 4: The qualitative cooperative basketball intention prediction results. Despite not using any manually annotated first-person labels during training, in most cases, our cross-model EgoSupervision method correctly predicts with whom the camera wearer will cooperate (the first two rows). In the third row, we also illustrate two cases where our method fails to produce correct predictions.

zeros to all the pixels outside B_j . Here, H and W depict the height and the width of the original input image. Finally, $\phi_{size}(B_j) \in \mathbb{R}^{1 \times 1}$ and $\phi_{pose}(B_j) \in \mathbb{R}^{1 \times 1}$ are scalars that capture the size and pose priors associated with a bounding box B_j , while $\phi_{loc} \in \mathbb{R}^{H \times W}$ is a first-person location prior of the same dimensions as the original input image.

Intuitively, the formulation above operates by first assigning a specific value to each of the detected bounding boxes. This yields a $H \times W$ dimensional prediction map where every pixel that does not belong to any bounding boxes is assigned a zero value. Then, this prediction map is multiplied with the location prior $\phi_{loc} \in \mathbb{R}^{H \times W}$ (using elementwise multiplication). Finally, all the values are normalized to be in range [0, 1], which produces our final pseudo ground truth labels. We now explain each of the components in more detail.

Egocentric Location Prior. The location of the camera wearer's visual attention is essential for inferring his/her cooperative intentions. We know that a first-person camera is aligned with the person's head direction, and thus, it captures exactly what the camera wearer sees. As a result, the way the camera wearer positions himself with respect to other players affects the location where these players will be mapped in a first-person image.

Instead of assuming any specific location a-priori (e.g. a center prior), as is done in [31, 29], we find the egocentric location prior directly from the data. As before, Let $B \in \mathbb{R}^{n \times 5}$ denote the bounding boxes detected by a pretrained network. Then we can compute $\phi_{loc} \in \mathbb{R}^{H \times W}$ as follows:

$$\phi_{loc}(B) = \sum_{j=1}^{n} V(B_j^{(i)}, c_j^{(i)}) \cdot \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{n} V(B_j^{(i)}, c_j^{(i)}))$$

where $c_j^{(i)}$ is the predicted confidence of the j^{th} bounding box in the i^{th} image. Intuitively, the first term $\sum_{j=1}^{n} V(B_j, c_j^{(i)})$ depicts a $H \times W$ dimensional mask that is obtained by assigning confidence values to all pixels in their respective bounding boxes in the current image, and zero values to the pixels outside the bounding boxes. The second term $\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{n} V(B_j, c_j^{(i)})$ also depicts a $H \times W$ dimensional mask that is obtained using this same procedure but across the entire training training dataset rather than a single image. In other words, the second term captures the locations in a first-person image where the bounding box predictions are usually most dense.

We conjecture, that $\phi_{loc}(I_i)$ can then be used to approximate the camera wearer's visual attention location, which is essential for inferring the camera wearer's cooperative intentions.

Egocentric Size Prior. Spatial 3D cues provides important information to infer the camera wearer's intentions [36, 45]. For instance, the camera wearer is more likely to cooperate with a player who is near him/her. We propose to capture this intuition, by exploiting an egocentric size prior. We know that the size of a bounding box in a first-person image is inversely related to the distance between the camera wearer and the person in the bounding box. Thus, let h_j be the height of the bounding box B_j . Then we express the egocentric size prior $\phi_{size}(B_j) \in \mathbb{R}^{1\times 1}$ for a given bounding box as:

$$\phi_{size}(B_j) = \exp\left(-\frac{\sigma}{h_j}\right)$$

where σ denotes a hyperparameter controlling how much to penalize small bounding boxes. Such a formulation al-



Figure 5: Several qualitative examples from the top 4 performing subjects in our conducted human study. Each subject specified their prediction by clicking on the person, with whom he/she thought the camera wearer was going to cooperate. We then placed a fixed size Gaussian around the location of the click. Note that based on these results, we can conclude that some instances of this task are quite difficult even for humans, i.e. in these examples, there is no general consensus among the subjects' responses.

lows us to capture the intuition that the camera wearer is more likely to cooperate with players who are physically closer to him/her.

Egocentric Pose Prior. In basketball, people tend to look at each other to express their intentions before actually performing cooperative actions. Detecting whether a particular person is facing the camera wearer can be easily done by examining the x coordinates of the paired body parts such as eyes, arms, legs, etc of a person detected in a first-person image. For instance, if a particular person is facing the camera wearer then, we will observe that for most of his/her paired parts visible in a first-person image the following will be true: $x(right_part) < x(left_part)$. In other words, the right parts of that person's body will have smaller x coordinate value in a first-person image, than the left parts. We use this intuition to encode the egocentric pose prior $\phi_{pose}(B_j) \in \mathbb{R}^{1 \times 1}$ for a given bounding box B_j as follows:

$$\phi_{pose}(B_j) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} 1\{x(right_part) < x(left_part)\}$$

where \mathcal{P} is the set of all paired parts, and $1\{x(right_part) < x(left_part)\}$ is an indicator function that returns 1 if the x coordinate of the right part in a first-person image is smaller than the x coordinate of the left part. The computed value $\phi_{pose}(B_j)$ can then be viewed as a confidence that a person in the bounding box B_j is facing the camera wearer, which is an important cue for inferring the camera wearer's cooperative intentions.

Pseudo Ground Truth. We then combine all the above discussed components into a unified framework using the Equation 2. Such a formulation allows us to automatically construct pseudo ground truth labels from the outputs of a pretrained multi-person pose estimation network. We illustrate several examples of our obtained pseudo ground truth

labels in Figure 3. Notice that while our computed pseudo ground truth is not always correct, in many cases it correctly captures the player with whom the camera wearer will cooperate in the near future. In our experimental section, we will demonstrate that despite the imperfections of our pseudo ground truth labels, we can use them to obtain a model that is almost as good as the model trained in a fully supervised fashion using manually annotated cooperation labels.

3.2. Cross-Model EgoSupervision

After obtaining the pseudo ground truth data \hat{y} , we train our cooperative basketball intention FCN using the crossmodel EgoSupervision scheme as shown in Figure 2. We employ a multi-person pose estimation network from [6] as our base model, which is used to predict the 1) pose estimates of all people in a given image and 2) their bounding boxes. The parameters inside the base network are fixed throughout the entire training procedure. At each iteration, the outputs from the base network are fed to the EgoTransformer, which transforms them into the pseudo ground truth cooperate intention labels. These pseudo ground truth labels are then used as a supervisory signal to train our cooperative basketball intention FCN using a sigmoid cross entropy per-pixel loss as illustrated in Equation 1.

3.3. Implementation Details

For all of our experiments, we used a Caffe deep learning library [25]. As our base FCN model we used a multiperson pose estimation network from [6]. Inspired by the success of this method, we also used the same architecture for our cooperative basketball intention FCN. During training, we optimized the network for 5000 iterations with a learning rate of 10^{-7} , the momentum equal to 0.9, the weight decay of 0.0005, and the batch size of 15. The weights inside the base FCN network were fixed throughout the entire training procedure. To compute the egocentric

Human Subjects	Accuracy
Subject-4	0.802
Subject-2	0.895
Subject-3	0.901
Subject-5	0.904
Subject-1	0.927

Table 1: Quantitative human study results on our cooperative basketball intention task. We ask 5 subjects to predict a player in the first-person image, with whom they think the camera wearer will cooperate after 2 seconds. We then compute the accuracy as the fraction of correct responses. The results indicate that most subjects achieve the accuracy of about 90%. We conjecture that Subject-4 may be less familiar with the basketball game thus, the lower accuracy.

size prior mask we used $\sigma = 10$.

4. Cooperative Basketball Intention Dataset

We build upon the dataset from [5], which captures firstperson basketball videos of 48 distinct college-level players in an unscripted basketball game. The work in [5] studies a basketball performance assessment problem, and provides 401 training and 343 testing examples of basketball cooperations among players from 10.3 hours of videos.

To obtain ground truth labels corresponding to the specific players, with whom the camera wearer cooperated, we look at the video segments corresponding to all such cooperation. We then identify the player with whom the camera wearer cooperated, go back to the frame about 2 seconds before the cooperation happens, and label that player with a bounding box. The ground truth data is then generated by placing a Gaussian inside the bounding box, according to the height and width of the bounding box.

Once again we note that these labels are only used for the evaluation purposes, and also to train other baseline models. In comparison, our method learns to detect the players with whom the camera wearer will cooperate, without relying on manually annotated intention labels.

5. Experimental Results

In this section, we present quantitative and qualitative results for our cooperative basketball intention prediction task. To compute the accuracy of each method, we select the player in the image with the maximum predicted probability as the the final prediction and then compute the fraction of all the correct predictions across the entire testing dataset.

5.1. Human Study

First, to see how well humans can predict cooperative basketball intention from first-person images, we conduct a human study consisting of 5 human subjects. Each subject

Method	Accuracy
DCL [30]	0.222
MPP-pretrained [6]	0.586
DeepLab [‡] [9]	0.644
Pseudo GT	0.665
ResNet-50 [‡] [20]	0.675
PSPNet [‡] [48]	0.695
ResNet-101 [‡] [20]	0.706
DeepLab-v2 [‡] [8]	0.757
MPP-finetuned [‡] [6]	0.778
CMES	0.775

Table 2: The quantitative cooperative basketball intention results evaluated as the fraction of correct predictions. We compare our Cross-Model EgoSupervision (CMES) scheme with a variety of supervised methods (marked by ‡). These results indicate that even without using manually annotated intention labels, our method outperforms most supervised methods, and produces almost identical performance as our main baseline "MPP-finetuned".

is shown 343 testing images one at a time, and asked to click on the player in an image, with whom he/she thinks the camera wearer will cooperate 2 seconds from now. Then the accuracy of each subject is evaluated as the fraction of correct responses.

We present these results in Table 1, and demonstrate that this task is not trivial even for humans: most of the subjects achieve about 90% accuracy on our task, which is solid but not perfect. We also point out that we did not collect information on how familiar each subject was with basketball. However, based on the results, we conjecture that Subject-4 who achieved almost 10% lower accuracy than the other subjects was probably not very familiar with basketball, which contributed to his lower performance. In Figure 5, we also visualize the qualitative examples that human subjects found the most difficult, i.e. in these instances, the predictions among the subjects differed substantially.

5.2. Quantitative Results

In Table 2, we present quantitative cooperative basketball intention results of our method and several other baselines. As our baselines, we use a collection of methods that were successfully used for other computer vision tasks such as image classification, semantic segmentation or saliency detection. These include a 1) Deep Contrast Saliency (DCL) method [30], 2-3) several variations of highly successful DeepLab semantic segmentation systems [9, 8] adapted to our task, 4-5) image classification ResNets [20] adapted to our task, 6) one of the top performing semantic segmentation systems PSPNet [48], 7-8) a pretrained and finetuned multi-person pose estimation (MPP) network [6], and 9) a pseudo ground truth obtained from

	Accuracy	
Method	pseudo GT	Trained Model
no ϕ_{loc}	0.481	0.560
no ϕ_{pose}	0.557	0.694
no ϕ_{size}	0.571	0.731
Ours-Full	0.665	0.775

Table 3: The quantitative ablation studies documenting the importance of each component in our EgoTransformer scheme. We separately remove each of ϕ_{loc} , ϕ_{size} , ϕ_{pose} and investigate how the accuracy changes. The second column in the table denotes the accuracy of a pseudo ground truth, while the third column depicts the accuracy of our trained model. Based on these results, we can conclude that each component of our EgoTransformer is essential for an accurate cooperative basketball intention prediction.

our EgoTransformer.

Note that our Cross-Model EgoSupervision (CMES) method is based on an MPP network architecture [6], and thus, as our main baseline we use the "MPP-finetuned" method, which uses the manually labeled bounding box intention labels to infer with whom the camera wearer will interact. In contrast to this baseline, our CMES method is only trained on the automatically generated pseudo ground truth labels. We note that the supervised methods employing manually labeled data are marked with [‡]. We now discuss several interesting observations based on these results.

Comparison with the Supervised Methods. Based on the results, we observe that despite not using manually annotated bounding box intention labels, our method outperforms a number of supervised baselines and achieves almost equivalent results to our main baseline "MPP-finetuned", which was trained using manually annotated cooperative intention labels. Thus, these results indicate the effectiveness of our cross-model EgoSupervision scheme.

Comparison with the Pseudo Ground Truth. One interesting and a bit surprising observation from Table 2, is that our cross-model EgoSupervision model achieves substantially better accuracy than the pseudo ground truth, which was used to optimize our model. We conjecture that this happens due to the following reasons. The pseudo ground truth labels are constructed using three different signals: 1) an egocentric location prior, 2) an egocentric size prior, and 3) an egocentric pose prior. Note, that our constructed pseudo ground truth does not incorporate any visual appearance information, i.e. it does not consider how the players look like. In contrast, during training, our network, learns what are the visual appearance cues indicative of the players with high pseudo ground truth values. Arguably, such visual cues provide a stronger signal for a cooperative intention recognition, which then leads to a substantially better performance than the pseudo ground truth labels.



First-Person RGB FCN Activations Ground Truth

Figure 6: The visualization of the activation values inside the second to last layer in our trained network. Note that the network produces high activation values around the faces of the players in the camera wearer's field of view. This makes intuitive sense, as facial expressions provide the most informative cues for a cooperative basketball intention task.

5.3. Qualitative Results

In Figure 4, we present our qualitative results, where we show that in most cases, our model successfully learns to predict with whom the camera wearer will cooperate. Furthermore, to gain a better understanding of what the network learned, in Figure 6, we visualize the activations inside the second to last FCN's layer. Note that our network has high activation values around the faces of people with whom the camera wearer intends to cooperate. This makes intuitive sense, as face is probably the most useful cue to recognize the camera wearer's intention to cooperate.

5.4. Ablation Experiments

In Table 3, we present the results analyzing the behavior of our EgoTransformer scheme. Earlier we discussed that to implement our EgoTransformer scheme we exploit three characteristics: 1) egocentric location prior ϕ_{loc} , 2) egocentric size prior ϕ_{size} , and 3) egocentric pose prior ϕ_{pose} . We want to investigate how much each of these priors affect 1) the quality of our generated pseudo ground truth data, and 2) the quality of our model trained using such pseudo ground truth. To do this, we run experiments with three baselines where for each baseline we remove one of ϕ_{loc} , ϕ_{size} , or ϕ_{pose} components. We denote these three baselines as "no ϕ_{loc} ", "no ϕ_{size} " and "no ϕ_{pose} " respectively. Finally, we include the results of our model using the full EgoTransformer scheme.

Based on the results in Table 3, we first observe that each of these components have a significant impact to the quality

of pseudo ground truth that we obtain. Specifically, using our full model yields 9.4% better pseudo ground truth results than the second best baseline. Additionally, note that the network trained to the pseudo ground truth of our full model achieves 4.4% higher accuracy than the second best baseline. These results indicate that each component in our EgoTransformer scheme is crucial for learning a high quality cooperative intention model.

6. Conclusions

In this work, we present a new task of predicting cooperative basketball intention from a single first-person image. We demonstrate that a first-person image provides strong cues to infer the camera wearer's intentions based on what he/she sees. We use this observation to design a new cross-model EgoSupervision learning scheme that allows us to predict with whom the camera wearer will cooperate, without using manually labeled intention labels. We demonstrate that despite not using such labels, our method achieves similar or even better results than fully supervised methods.

We believe that our proposed cross-model EgoSupervision scheme could be applied on various other first-person vision tasks without the need to manually collect labels for each of such tasks. In the long run, a learning scheme such as ours could effectively replace the supervised methods, which require costly and time consuming annotation process.

References

- Y. Aytar and A. Zisserman. Tabula rasa: Model transfer for object category detection. In *IEEE International Conference on Computer Vision*, 2011. 2
- [2] Gedas Bertasius, Hyun Soo Park, Stella X. Yu, and Jianbo Shi. First-person action-object detection with egonet. In *Proceedings of Robotics: Science and Systems*, July 2017. 2, 3
- [3] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Deepedge: A multi-scale bifurcated deep network for top-down contour detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 2
- [4] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Semantic segmentation with boundary neural fields. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [5] Gedas Bertasius, Stella X. Yu, Hyun Soo Park, and Jianbo Shi. Am I a baller? basketball performance assessment from first-person videos. *CoRR*, abs/1611.05365, 2016. 6
- [6] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. *CoRR*, abs/1611.08050, 2016. 2, 3, 5, 6, 7
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015. 2
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. arXiv:1606.00915, 2016. 6

- [9] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L. Yuille. Attention to scale: Scale-aware semantic image segmentation. *CVPR*, 2016. 6
- [10] Dima Damen, Teesid Leelasawassuk, Osian Haines, Andrew Calway, and Walterio Mayol-Cuevas. You-do, i-learn: Discovering task relevant objects and their modes of interaction from multi-user egocentric video. In *Proceedings of the British Machine Vision Conference*. BMVA Press, 2014. 2
- [11] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International Conference in Machine Learning (ICML)*, 2014. 2
- [12] Lixin Duan, Dong Xu, and Ivor W. Tsang. Learning with augmented features for heterogeneous domain adaptation. In *Proceedings of the International Conference on Machine Learning*, pages 711–718, Edinburgh, Scotland, June 2012. Omnipress. 2
- [13] Alireza Fathi, Ali Farhadi, and James M. Rehg. Understanding egocentric activities. In *ICCV*. 2
- [14] Alireza Fathi, Jessica K. Hodgins, and James M. Rehg. Social interactions: A first-person perspective. 2
- [15] Alireza Fathi, Xiaofeng Ren, and James M. Rehg. Learning to recognize objects in egocentric activities. In *CVPR*, pages 3281–3288. IEEE Computer Society, 2011. 2
- [16] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2
- [17] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012, pages 2066–2073, 2012. 2
- [18] Ilaria Gori, J. K. Aggarwal, and Michael S. Ryoo. Building unified human descriptors for multi-type activity recognition. *CoRR*, abs/1507.02558, 2015. 2
- [19] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. arXiv preprint arXiv:1512.03385, 2015. 6
- [21] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531, 2015. 2
- [22] Judy Hoffman, Sergio Guadarrama, Eric S Tzeng, Ronghang Hu, Jeff Donahue, Ross Girshick, Trevor Darrell, and Kate Saenko. Lsda: Large scale detection through adaptation. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, Advances in Neural Information Processing Systems 27, pages 3536–3544. Curran Associates, Inc., 2014. 2
- [23] Judy Hoffman, Saurabh Gupta, and Trevor Darrell. Learning with side information through modality hallucination. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [24] Judy Hoffman, Erik Rodner, Jeff Donahue, Kate Saenko, and Trevor Darrell. Efficient learning of domain-invariant image representations. In *International Conference on Learning Representations*, 2013. 2
- [25] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. arXiv:1408.5093, 2014. 5

- [26] Kris Makoto Kitani, Takahiro Okabe, Yoichi Sato, and Akihiro Sugimoto. Fast unsupervised ego-action learning for first-person sports videos. In CVPR, 2011. 2
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*. 2012. 2
- [28] B. Kulis, K. Saenko, and T. Darrell. What you saw is not what you get: Domain adaptation using asymmetric kernel transforms. In *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, CVPR '11, pages 1785–1792, 2011. 2
- [29] Yong Jae Lee and Kristen Grauman. Predicting important objects for egocentric video summarization. *IJCV*, 2015. 2, 4
- [30] G. Li and Y. Yu. Deep contrast learning for salient object detection. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), pages 478–487, June 2016. 6
- [31] Yin Li, Zhefan Ye, and James M. Rehg. Delving into egocentric actions. In CVPR. 2, 3, 4
- [32] Guosheng Lin, Chunhua Shen, Ian D. Reid, and Anton van den Hengel. Efficient piecewise training of deep structured models for semantic segmentation. *CoRR*, abs/1504.01013, 2015. 2
- [33] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *CoRR*, 2014. 2
- [34] Zheng Lu and Kristen Grauman. Story-driven summarization for egocentric video. In *CVPR*, 2013. 2
- [35] Kris Kitani Minghuang Ma. Going deeper into first-person activity recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [36] Hyun Soo Park, Jyh-Jing Hwang, Yedong Niu, and Jianbo Shi. Egocentric future localization. In CVPR, 2016. 2, 4
- [37] Hyun Soo Park, Jyh-Jing Hwang, and Jianbo Shi. Force from motion: Decoding physical sensation from a first person video. In *CVPR*, 2016. 2
- [38] Hamed Pirsiavash and Deva Ramanan. Detecting activities of daily living in first-person camera views. In CVPR, 2012. 2
- [39] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems (NIPS)*, 2015. 2
- [40] Xiaofeng Ren and Chunhui Gu. Figure-ground segmentation improves handled object recognition in egocentric video. In CVPR, 2010. 2
- [41] M. S. Ryoo, Thomas J. Fuchs, Lu Xia, J.K. Aggarwal, and Larry Matthies. Robot-centric activity prediction from first-person videos: What will they do to me? In *Proceedings of the Tenth Annual* ACM/IEEE International Conference on Human-Robot Interaction, HRI '15, pages 295–302, New York, NY, USA, 2015. ACM. 2
- [42] Pierre Sermanet, David Eigen, Xiang Zhang, Michael Mathieu, Rob Fergus, and Yann Lecun. Overfeat: Integrated recognition, localization and detection using convolutional networks. http://arxiv.org/abs/1312.6229. 2
- [43] Suriya Singh, Chetan Arora, and C. V. Jawahar. First person action recognition using deep learned descriptors. In *The IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), June 2016. 2
- [44] Bilge Soran, Ali Farhadi, and Linda Shapiro. Action Recognition in the Presence of One Egocentric and Multiple Static Cameras. 2015.
 2
- [45] Shan Su, Jung Pyo Hong, Jianbo Shi, and Hyun Soo Park. Predicting behaviors of basketball players from first person videos. In *CVPR*, 2017. 2, 4

- [46] Yu-Chuan Su and Kristen Grauman. Detecting engagement in egocentric video. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [47] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In ICCV, 2015. 2
- [48] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. *CoRR*, abs/1612.01105, 2016. 6