

Am I a Baller? Basketball Performance Assessment from First-Person Videos

Gedas Bertasius¹, Hyun Soo Park², Stella X. Yu³, Jianbo Shi¹

¹University of Pennsylvania, ²University of Minnesota, ³UC Berkeley ICSI
 {gberta, jshi}@seas.upenn.edu hspark@umn.edu stella.yu@berkeley.edu

Abstract

This paper presents a method to assess a basketball player’s performance from his/her first-person video. A key challenge lies in the fact that the evaluation metric is highly subjective and specific to a particular evaluator. We leverage the first-person camera to address this challenge. The spatiotemporal visual semantics provided by a first-person view allows us to reason about the camera wearer’s actions while he/she is participating in an unscripted basketball game. Our method takes a player’s first-person video and provides a performance measure that is specific to an evaluator’s preference.

To achieve this goal, we first use a convolutional LSTM network to detect atomic basketball events from first-person videos. Our network’s ability to zoom-in to the salient regions addresses the issue of a severe camera wearer’s head movement in first-person videos. The detected atomic events are then passed through the Gaussian mixtures to construct a highly non-linear visual spatiotemporal basketball assessment feature. Finally, we use this feature to learn a basketball assessment model from pairs of labeled first-person basketball videos, for which a basketball expert indicates, which of the two players is better.

We demonstrate that despite not knowing the basketball evaluator’s criterion, our model learns to accurately assess the players in real-world games. Furthermore, our model can also discover basketball events that contribute positively and negatively to a player’s performance.

1. Introduction

A gifted offensive college basketball player, Kris Jenkins (Villanova), made a three point buzzer beater against UNC (2015-2016 season), and recorded one of the greatest endings in NCAA championship history. He was arguably one of the best players in the entire NCAA tournament. A question is “what makes him stand out from his peer players?”. His stats, e.g., average points and rebounds per game, can be a measure to evaluate his excellence. However, these measures do not capture every basketball aspect that a coach

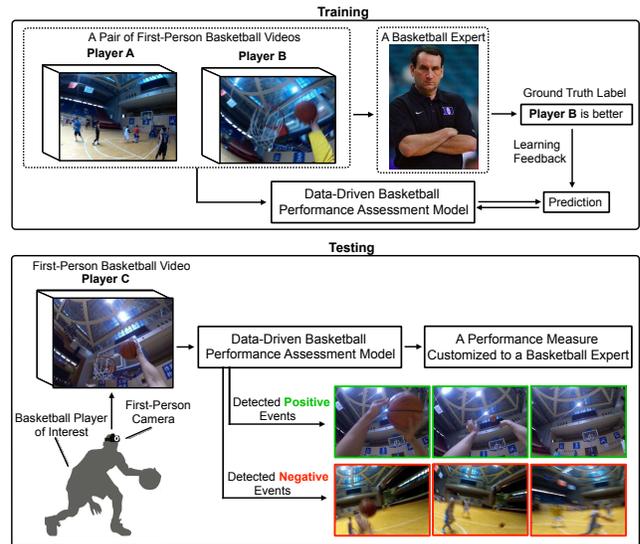


Figure 1: Our goal is to assess a basketball player’s performance based on an evaluator’s criterion from an unscripted first-person basketball video of a player. During training, we learn such a model from the pairs of weakly labeled first-person basketball videos. During testing, our model predicts a performance measure customized to a particular evaluator from a first-person basketball video. Additionally, our model can also discover basketball events that contribute positively and negatively to a player’s performance.

may want to use for assessing his potential impact in the future team, which is difficult to measure quantitatively. NBA coaches and scouts are eager to catch every nuance of a basketball player’s abilities by watching a large number of his basketball videos.

Now consider a college recruitment process where there is a massive number of high school players. In such conditions, the searching task for the best players becomes much more challenging, more expensive and also more labor intense. More importantly, the recruiters need to measure and evaluate a sequence of atomic decision makings, e.g., when does a player shoot, whether he makes a shot, how good is

his passing ability, etc. There exists neither universal measure nor golden standard to do this, i.e., most scouts and coaches have their own subjective evaluation criterion.

In this paper, we address a problem of computational basketball player assessment customized to a coach's or scout's evaluation criterion. Our conjecture is that a first-person video captures a player's basketball actions and his/her basketball decision making in a form of the camera motion and visual semantics of the scene. A key challenge of first-person videos is that it immediately violates primary assumptions made for third-person recognition systems: first-person videos are highly unstable and jittery and visual semantics does not appear as iconic as in third-person [9].

Our first-person approach innovates the traditional assessment methods, e.g., watching hours of third-person videos taken by non professional videographers and assessing the players in them. In contrast, a first-person video records what the player sees, which directly tells us what is happening to the player himself, e.g., the body pose of a point guard who is about to pass at HD resolution while a third-person video produces a limited visual access to such subtle signals. Furthermore, the 3D camera egomotion of the first person video reflects the decision making of how the player responds to the team configuration, e.g., can I drive towards the basket and successfully finish a layup? Finally, a first-person camera eliminates the tracking and player association tasks of the third-person video analysis, which prevents applications of computational approaches for amateur games¹.

Our system takes a first-person video of basketball players and outputs a basketball assessment metric that is specific to an evaluator's preference. The evaluator provides the comparative weak labels of the performance of the players, e.g., the player A is better than B based on his own subjective criteria.

Our method first uses a convolutional LSTM to detect atomic basketball events from a first-person video. Our network's ability to localize the most informative regions in a first-person image, is essential for first-person videos where the camera undergoes severe head movement, which causes videos to be blurry. These atomic events are then passed through the Gaussian mixtures to produce a highly non-linear visual spatiotemporal basketball assessment feature. Finally, our basketball assessment model is learned from the pairs of labeled first-person basketball videos by minimizing a hinge loss function. We learn such a basketball skill assessment model from our new 10.3 hour long first-person basketball dataset that captures 48 distinct college level basketball players in an unscripted basketball game.

Impact Ample money and effort have been invested in

¹Usage of commercial tracking systems using multiple calibrated cameras is limited due to a high cost [1]

recruiting, assessing, and drafting basketball players every year. However, limited progress has been made on developing computational models that can be used to automatically assess an athlete's performance in a particular sport [39, 23]. As wearable technology advances, cameras can be non-invasively worn by players, which delivers a vivid sense of their dynamics, e.g., Spanish Liga ACB has demonstrated a possibility of a jersey camera that allows you to put yourself in the court [2]. This trend will open up a new opportunity to share experiences and evaluate performance across players in different continents without bias and discrimination. Our work takes a first step towards enabling a computational analysis for such first-person data.

Contribution To the best of our knowledge, this is the first paper that addresses practical behavioral assessment tasks using first-person vision specific to an evaluator's preference. The core technical contributions of the paper include 1) a basketball assessment model that assesses the players based on an evaluator's assessment criterion, which we learn from the pairs of weakly labeled first-person basketball videos; 2) a predictive network that learns the visual semantics of important actions and localizes salient regions of first-person images to handle unstable first-person videos and 3) a new 10.3 hour long first-person basketball video dataset capturing 48 players in an unscripted basketball game.

2. Related Work

Talent wins games, but teamwork and intelligence wins championships. — Michael Jordan

Accurate diagnosis and evaluation of athletes is a key factor to build a synergic teamwork. However, it is highly subjective and task dependent, and the psychological and financial cost of such process is enormous. A large body of sport analytics and kinesiology has studied a computational approaches to provide a quantitative measure of the performance [39, 23, 38, 19, 25].

Kinematic abstraction (position, orientation, velocity, and trajectory) of the players offers a global centric representation of team behaviors, which allows a detailed analysis of the game such as the probability of shoot success, rebound, and future movement prediction [38, 19, 35]. Not only an individual performance, but also team performance can be measured through the kinematic abstraction [25, 23].

These kinematic data are often obtained by multiple third-person videos [1, 19, 23] where the players and ball are detected using recognition algorithms combined with multiple view geometry [14]. Tracking and data association is a key issue where the role of the players provides a strong cue to disambiguate appearance based tracking [22]. Events such as ball movement, can be also recognized using a spatiotemporal analysis [24]. As players behave strategi-

cally and collectively, their group movement can be predicted [17] and the ball can be localized without detection. Various computational models have been used for such tasks, e.g., Dynamic Bayesian Network [36], hierarchical LSTM [15], attention based LSTM [31] learned from a large collection of third-person videos.

Unlike third-person videos, first-person cameras closely capture what the players see. Such property is beneficial to understand activities highly correlated with visual attention, e.g., object manipulation and social communications. Important objects to the camera wearer are detected and segmented [20, 8, 32, 12, 5], which can be used to compress life-log videos [20, 21]. As visual attention is also related with the intent of the camera wearer, her/his future movement can be predicted [27]. Beyond individual behaviors, joint attention is a primary indicator of social interactions, which can be directly computed from first-person videos [11, 29], and further used for human-robot interactions [33, 13].

In sports, the complex interactions with a scene in first-person videos can be learned through spatiotemporal visual patterns. For instance, the scene can tell us about the activity [18] and the egomotion can tell us about the physical dynamics of activity [28]. Joint attention still exists in team sports which can be described by the team formation [30] and future behaviors [35].

Unlike previous work that mainly focuses on recognizing and tracking objects, activities, and joint attention, we take one step further: performance assessment based on the evaluator’s preference. We introduce a computational model that exhibits strong predictive power when applied on the real world first-person basketball video data.

3. Basketball Performance Assessment Model

We define a measure of performance assessment using a first-person video:

$$S(\mathcal{V}) = \frac{\sum_{t=1}^T p_t^{(1)} \mathbf{w}^\top \phi(\mathbf{V}_t, \mathbf{x})}{\sum_{t=1}^T p_t^{(1)}} \quad (1)$$

where \mathcal{V} is a first-person video of T number of frames, ϕ is a visual spatiotemporal basketball assessment feature, and \mathbf{w} is a weight vector of performance regressor. $\mathbf{V}_t \subset \mathcal{V}$ is a segmented video starting at the t^{th} frame with a fixed length, T_s . $p_t^{(1)} \in [0, 1]$ is a relevance of \mathbf{V}_t to evaluate a given player’s performance. $\mathbf{x} \in \mathbb{R}^2$ is the 2D coordinate of the basketball player, i.e., the projection of 3D camera pose computed by structure from motion [14] onto the canonical basketball court. In Figure 2, we provide a detailed illustration of our basketball assessment prediction framework.

3.1. Visual Spatiotemporal Assessment Feature

Our first goal is to use a first-person basketball video to build a powerful feature representation that could be used for an effective player’s performance assessment. We identify three key challenges related to building such a representation from first-person basketball videos: 1) our system needs to handle severe camera wearer’s head motion, 2) we need to have an interpretable basketball representation in terms of its atomic events, and 3) our feature representation has to be highly discriminative for a player’s performance prediction task.

To address these problems, we propose to represent the visual feature of the segmented video, \mathbf{V}_t , as follows, where each function below addresses one of the listed challenges:

$$\phi(\mathbf{V}_t, \mathbf{x}) = f_{\text{gm}}(f_{\text{event}}(f_{\text{crop}}(\mathbf{V}_t), \mathbf{x})), \quad (2)$$

where f_{crop} is a function that handles a severe camera wearer’s head motion by producing a cropped video by zooming in on the important regions, f_{event} is a function that computes the probability of atomic basketball events, and f_{gm} is a Gaussian mixture function that computes a highly non-linear visual feature of the video.

Zooming-In. A key property of f_{crop} is the ability to zoom-in to relevant pixels which allows to learn an effective visual representation for the basketball performance assessment. Using this regional cropping, we minimize the effect of jittery and unstable nature of first person videos that causes larger variation of visual data. In our experimental section, we demonstrate that using f_{crop} in our model substantially improves the prediction performance. Thus, initially we process a first-person video to produce a cropped video:

$$\bar{\mathbf{V}}_t = f_{\text{crop}}(\mathbf{V}_t; \mathbf{w}_{\text{crop}}),$$

where f_{crop} is parametrized by \mathbf{w}_{crop} , $\bar{\mathbf{V}}_t$ is the cropped video with fixed size $C_w \times C_w \times 3 \times T_s$, and C_w is the width and height of the cropping window.

We predict the center of the cropping window by learning \mathbf{w}_{crop} using a fully convolutional network [7]. To do this, we train the network to predict the location of a ball, which is typically where most players are looking at. Afterwards, for each frame in a video, we compute a weighted average of XY location coordinates weighted by the detected ball probabilities and then crop a fixed size patch around such a weighted average location. We illustrate some of the qualitative zoom-in examples in Figure 6.

Atomic Basketball Event Detection. To build an interpretable representation in terms of atomic basketball events, we predict basketball events of 1) somebody shooting a ball, 2) the camera wearer possessing the ball, and 3) a made shot respectively. Note that the cropped video focuses on

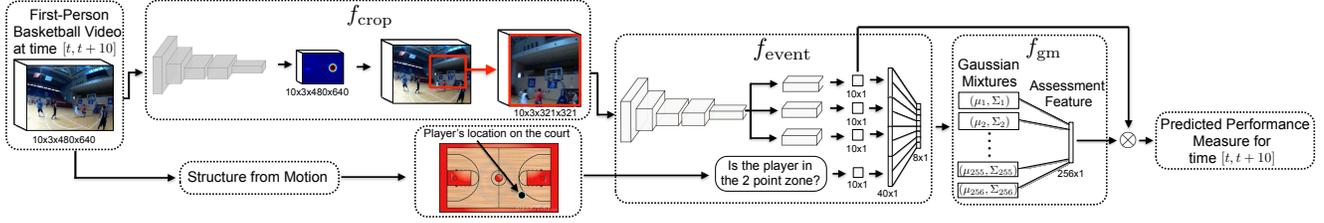


Figure 2: A detailed illustration of our basketball assessment prediction scheme. Given a video segment from time interval $[t, t + 10]$, we first feed it through a function f_{crop} , which zooms-in to the relevant parts of a video. We then apply f_{event} to predict 4 atomic basketball events from a zoomed-in video and a player’s (x, y) location on the court. We then feed these predictions through a Gaussian mixture function f_{gm} , which produces a highly non-linear visual spatiotemporal assessment feature. Finally, we use this feature to compute a player’s assessment measure by multiplying it with linear weights learned from the data, and with a predicted relevance indicator for a given video segment.

the ball and its visual context, which allows to learn the visual semantics of each atomic event more effectively. To do this we use a multi-path convolutional LSTM network, where each pathway predicts its respective atomic basketball event. We note that such a multi-path architecture is beneficial as it allows each pathway to focus on learning a single atomic basketball concept. In contrast, we observed that training a similar network with a single pathway failed to produce accurate predictions for all three atomic events. Given a cropped video, our multi-path network is jointly trained to minimize the following cross-entropy loss:

$$\mathcal{L}_{\text{event}} = - \sum_{t=1}^{T_s} \sum_{b=1}^3 y_t^{(b)} \log p_t^{(b)} + (1 - y_t^{(b)}) \log (1 - p_t^{(b)}),$$

where $p_t^{(b)}$ depicts a network’s prediction for an atomic basketball event b at a time step t ; $y_t^{(b)} \in \{0, 1\}$ is a binary atomic basketball event ground truth value for frame t and basketball event b .

We also note that because many important basketball events occur when somebody shoots the ball [3, 4], the detected probability $p_t^{(1)}$ is also later used in Equation (1), as a relevance indicator for each video segment, \mathbf{V}_t .

As our fourth atomic basketball event $p_t^{(4)}$, we use a binary value indicating whether a player is in the 2 point or 3 point zone, which is obtained from a player’s (x, y) location coordinates on the court.

We then split each of the 4 basketball event predictions in half across the temporal dimension, and perform temporal max pooling for each of the 8 blocks. All the pooled values are then concatenated into a single vector \mathbf{b}_t :

$$\mathbf{b}_t = f_{\text{event}}(\bar{\mathbf{V}}_t, \mathbf{x}; \mathbf{w}_{\text{event}})$$

Gaussian Mixtures. To build a representation that is discriminative, and yet generalizable, we construct a highly non-linear feature that works well with a linear classifier. To achieve these goals we employ Gaussian mixtures, that

transform the atomic basketball event feature, into a complex basketball assessment feature, which we will show to be very effective in our assessment model. Formally, given a vector \mathbf{b}_t over T_s , we compute the visual spatiotemporal assessment features for a given video segment as:

$$\phi_t = f_{\text{gm}}(\mathbf{b}_t; \{\mu_n, \Sigma_n\}_{n=1}^N)$$

where f_{gm} is parametrized by Gaussian mixtures, $\{\mu_n, \Sigma_n\}_{n=1}^N$, and N is the number of mixtures. Each mixture j is defined by a function $z(y_{t_1}^{(1)}, y_{t_2}^{(1)}, \dots, y_{t_1}^{(4)}, y_{t_2}^{(4)}) = j$. Here $y_{t_1}^{(i)}, y_{t_2}^{(i)} \in \{0, 1\}$ refer to the binary ground truth values associated with an atomic basketball event $i \in \{1, 2, 3, 4\}$; the index t_1 indicates the first half of an input video segment, whereas t_2 indicates the second half. Every possible combination of these values define one of the $2^8 = 256$ Gaussian mixtures. We learn the parameters of each Gaussian mixture using maximum likelihood from the training data with diagonal covariances.

3.2. Basketball Assessment Prediction

We learn a linear weight \mathbf{w} in Equation (1) based on the comparative assessment of players provided by a former professional basketball player in Section 4. We minimize the following hinge loss:

$$\mathcal{L}_{\mathbf{w}} = \sum_{i=1}^D \max \left(0, \left(\frac{1}{2} - Y_i \right) (S(\mathcal{V}_1^i) - S(\mathcal{V}_2^i)) \right), \quad (3)$$

where $Y_i = 1$ if a basketball expert declared Player 1 to be better than Player 2; otherwise $Y_i = 0$. $S(\mathcal{V}_1^i), S(\mathcal{V}_2^i)$ depict our predicted performance measure for Players 1, and 2 respectively, \mathcal{V}_1^i and \mathcal{V}_2^i are the first-person basketball videos of Player 1 and Player 2 respectively, and D is the number of data points. Then based on Equation 1, we can compute the subgradients of this loss function with respect to w and find w by minimizing it via a standard gradient descent. In Figure 3, we provide an illustration of such a learning framework.

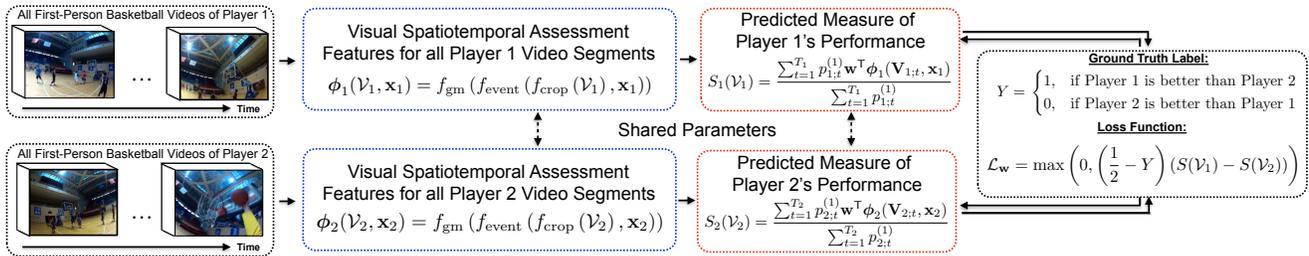


Figure 3: An illustration of our training procedure to learn the linear weights w that are used to assess a given basketball player’s performance. As an input we take a pair of labeled first-person basketball videos with a label provided by a basketball expert indicating, which of the two players is better. Then, we compute visual spatiotemporal basketball assessment features for all input video segments, and use them to learn weights w by minimizing our formulated hinge loss function.

Why Linear Classifier? We only have 250 labeled pairs for learning the weights, which is a small amount of training data. Thus, making a classifier more complex typically results in a severe overfitting. Through our experiments, we discovered that linear weights work the best.

3.3. Implementation Details

For all of our experiments involving CNNs, we used a Caffe library [16]. Both networks were based on DeepLab’s [7] architecture and were trained for 4000 iterations with a learning rate of 10^{-8} , 0.9 momentum, the weight decay of $5 \cdot 10^{-5}$, and 30 samples per batch. The LSTM layers inside the atomic basketball event network spanned 10 consecutive frames in the video input. Each pathway in the atomic basketball event network was composed of two 1024 dimensional convolution layers with kernel size 1×1 and a 1024 dimensional LSTM layer. The networks were trained using standard data augmentation. To learn the weights w we used a learning rate of 0.001 and ran gradient descent optimization for 100 iterations.

4. First-Person Basketball Dataset

We present a first person basketball dataset composed of 10.3 hours of videos with 48 college players. Each video is about 13 minutes long captured by GoPro Hero 3 Black Edition mounted with a head strip. It is recorded at 1280×960 with 100 fps. We record 48 videos during the two days, with a different group of people playing each day. We use 24 videos from the first day for training and 24 videos from the second day for testing. We extract the video frames at 5 fps to get 98,452 frames for training, and 87,393 frames for testing.

We ask a former professional basketball player (played in an European national team) to label which player performs better given a pair of first-person videos. Total 500 pairs are used: 250 for training and 250 for testing. Note that there were no players overlapping between the training and testing splits.

We also label three simple basketball events: 1) somebody shooting a ball, 2) the camera wearer possessing the

	Atomic Events			
	$p^{(1)}$	$p^{(2)}$	$p^{(3)}$	mean
Tran et al. [37]	0.312	0.428	0.193	0.311
Singh et al [34]	0.469	0.649	0.185	0.434
Bertasius et al [6]	0.548	0.723	0.289	0.520
Ma et al [26]	0.622	0.718	0.364	0.568
Ours: no LSTM & no zoom-in	0.711	0.705	0.192	0.536
Ours: no zoom-in	0.693	0.710	0.248	0.550
Ours: single path	0.678	0.754	0.308	0.580
Ours: no LSTM	0.718	0.746	0.397	0.620
Ours	0.724	0.756	0.395	0.625

Table 1: The quantitative results for atomic basketball event detection on our first-person basketball dataset according to max F-score (MF) metric. These results show that our method 1) outperforms prior first-person methods and 2) that each component plays a critical role in our system.

ball, and 3) a made shot. These are the key atomic events that drive a basketball game. In total, we obtain 3,734, 4,502, and 2,175 annotations for each of these three events respectively.

Furthermore, to train a ball detector we label the location of a ball at 5,073 images by clicking once on the location. We then place a fixed sized Gaussian around those locations and use it as a ground truth label.

5. Experimental Results

5.1. Quantitative Results

Atomic Basketball Event Detection. In Table 1, we first illustrate our results for atomic basketball event detection task. The results are evaluated according to the maximum F-score (MF) metric by thresholding the predicted atomic event probabilities at small intervals and then computing a precision and recall curve. First, we compare our model’s predictions with several recent first-person activity recognition baselines [34, 6, 26] and also with the successful video activity recognition baseline C3D [37]. We show that our model outperforms all of these baselines for each atomic event.

Furthermore, to justify our model’s design choices, in

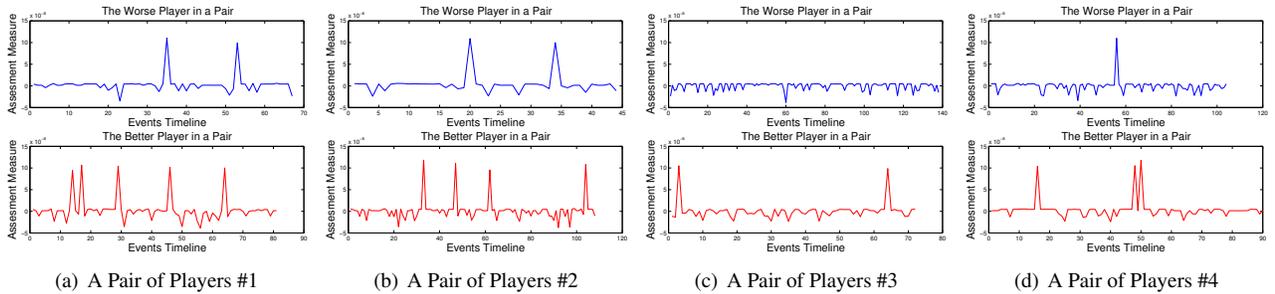


Figure 4: We randomly select 4 pairs of basketball players, and visualize how our assessment model evaluates each player over time. The red plot denotes the better player in a pair, whereas the blue plot depicts the worse player. The y -axis in the plot illustrates our predicted performance measure for an event occurring at a specific time in a player’s first-person video.

	Accuracy	
	Pred. Events	GT Events
LRCN [10] 2-pt made shot detector	fail	-
LRCN [10] 3-pt made shot detector	fail	-
Ours: no GMs	0.477	-
Ours: no $p^{(3)}$	0.496	-
Ours: no $p^{(2)}$	0.515	-
Ours: no $p^{(1)}$	0.536	-
Ours: single GM-top2	0.537	-
Ours: all weights w set to 1	0.583	-
Ours: single GM-top1	0.609	-
Ours: no $p^{(4)}$	0.649	-
Ours	0.765	0.793

Table 2: The quantitative results for our basketball assessment task. We evaluate our method on 250 labeled pairs of players, and predict, which of the two players in a pair is better. We then compute the accuracy as the fraction of correct predictions. We report the results of various baselines in two settings: 1) using our predicted atomic events, and 2) using ground truth atomic events. These results show that 1) our model achieves best results, 2) that each of our proposed components is important, and 3) that our system is pretty robust to atomic event recognition errors.

Table 1 we also include several experiments studying the effect of 1) a multi-path architecture, 2) LSTM layers, and 3) zooming-in scheme. Our experiments indicate that each of these components is crucial for achieving a solid atomic event recognition accuracy, i.e. the system achieves the best performance when all three of these components are included in the model.

Basketball Assessment Results. In Table 2, we present our results for assessing 24 basketball players from our testing dataset. To test our method’s accuracy we evaluate our method on 250 labeled pairs of players, where a label provided by a basketball expert indicates, which of the two players is better. For each player, our method produces an assessment measure indicating, which player is better (the higher the better). To obtain the accuracy, we compute the fraction of correct predictions across all 250 pairs.

We note that to the best of our knowledge, we are the first

to formally investigate a basketball performance assessment task from a first-person video. Thus, there are no well established prior baselines for this task. As a result, we include the following list of baselines for a comparison.

First, we include two basketball activity baselines: the detectors of 1) 2-point and 2) 3-point shots made by the camera wearer. We label all instances in our dataset where these activities occur and discover ≈ 100 of such instances. Note that such a small number of instances is not a flaw of our dataset, but instead an inherent characteristic of our task. Such basketball activities belong to a long-tail data distribution, i.e. they occur pretty rarely, and thus, it is difficult to train supervised classifiers for such activity recognition. We then train an LRCN [10] model as 1) a 2 point made shot detector, and 2) a 3 point made shot detector. We report that due to a small amount of training data, in all cases the network severely overfit the training data and did not learn any meaningful pattern.

Furthermore, to justify each of our proposed components in the model, in Table 2 we also include several ablation baselines. First, we study how 1) Gaussian Mixtures (GM) and 2) the process of learning the weights affect the performance assessment accuracy. We do it 1) with our predicted and 2) with the ground truth atomic events. We show that in both cases, each of our proposed components is beneficial. In addition, we also observe that our system is robust to atomic event recognition errors: the accuracy when using the ground truth atomic events is only 2.8% better compared to our original model.

We also present the performance assessment results when we remove one of the four atomic events from our system. We show that our method performs the best when all four atomic events are used, suggesting that each atomic event is useful. Finally, as two extra baselines we manually select two Gaussian mixtures with the largest weight magnitudes and use each of their predictions independently (denoted as single GM-top1,2 in Table 2). We show that our full model outperforms all the other baselines, thus, indicating that each of our proposed component in our model is crucial for an accurate player performance assessment.



Figure 5: A visualization of basketball activities that we discovered by manually inspecting Gaussian mixtures associated with the largest basketball assessment model weights w . Each row in the figure depicts a separate event, and the columns illustrate the time lapse of the event (from left to right). We discover that the two most positive Gaussian mixtures correspond to the events of a player making a 2 point and a 3 point shot respectively (the first two rows), while the mixture with the most negative weight captures an event when a player misses a 2 point shot (last row).

5.2. Qualitative Results

In addition, in Figure 4, we also include a more dynamic visualization of how our assessment model works over time. To do this, we randomly select 4 pairs of basketball players, and visualize how our model evaluates each player over time. The red plot in each pair denotes the better player, whereas the blue plot depicts the worse player. The y -axis in the plot illustrates our predicted performance measure for an event occurring at a specific time in a player’s first-person video.

Furthermore, in Figure 6 we also include examples of short sequences, illustrating 1) a player’s actions that contributed most positively to his/her performance assessment and also 2) actions that contributed most negatively. We select these action sequences by picking the first-person video sequences with a largest positive and negative values of the terms inside the summation of Equation 1 (which also correspond to positive and negative peaks from Figure 4). Such terms depict each video segment’s contribution to the overall basketball skill assessment measure.

We would like to note that it is quite difficult to include such results in an image format, because 1) images are static and thus, they cannot capture the full content of the videos; 2) images in the paper, appear at a very low-resolution compared to the original 480×640 videos, which makes it more difficult to understand what kind of events are depicted in these images. To address some of these issues, in our supplementary material, we include even more of such qualitative examples in a video format.

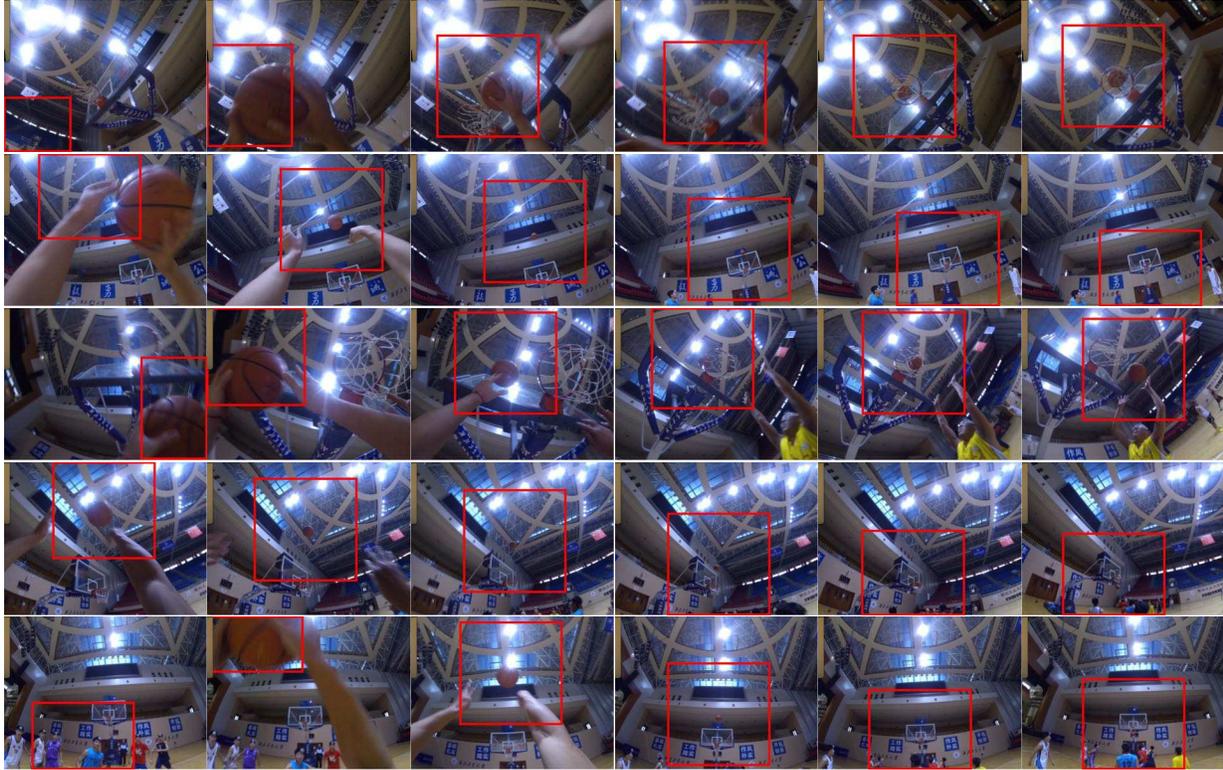
Understanding the Feature Representation. Earlier, we claimed that Gaussian mixtures produce a highly non-linear feature representation. We now want to get a better

insight into what it represents. To do so we analyze the learned weights w , and then manually inspect the Gaussian mixtures associated with the largest magnitude weights in w . Upon doing so we discover that the two mixtures with the most positive weights learn to capture basketball activities when camera wearer makes a 2 point shot, and a 3 point shot respectively. Conversely, the mixtures with the two most negative weights represent the activities of the camera missing a 2 point shot, and the camera wearer’s defender making a shot respectively. In Figure 5, we include several sequences corresponding to such discovered activities.

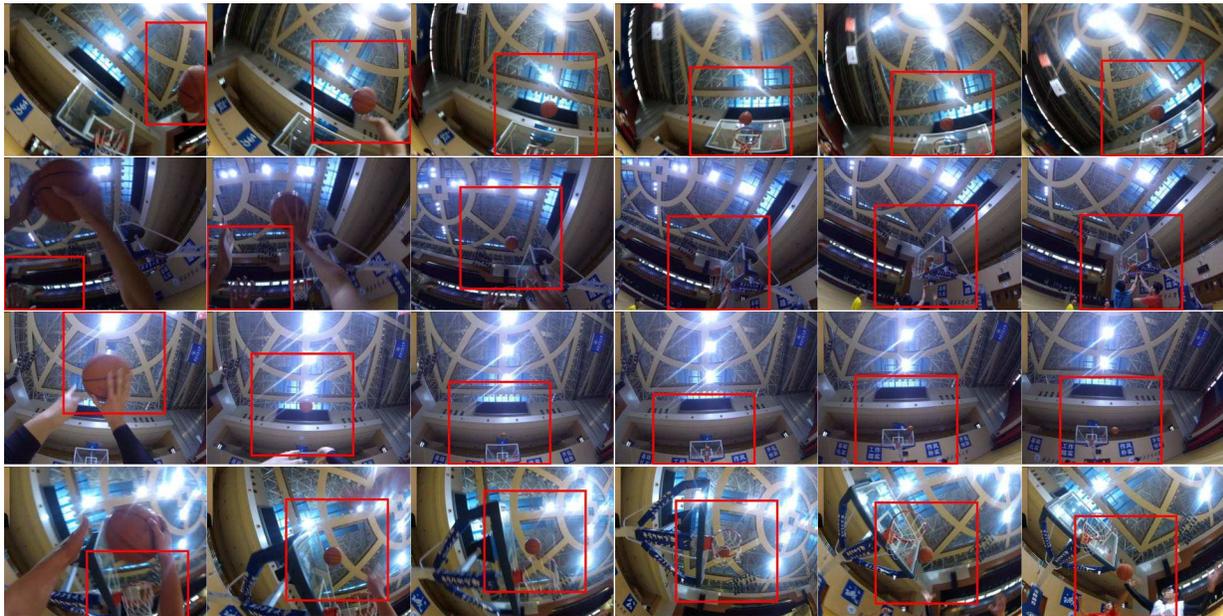
6. Conclusions

In this work, we introduced a basketball assessment model that evaluates a player’s performance from his/her first-person basketball video. We showed that we can learn powerful visual spatiotemporal assessment features from first-person videos, and then use them to learn our skill assessment model from the pairs of weakly labeled first-person basketball videos. We demonstrated that despite not knowing the labeler’s assessment criterion, our model learns to evaluate players with a solid accuracy. In addition, we can also use our model to discover the camera wearer’s activities that contribute positively or negatively to his/her performance assessment.

We also note that performance assessment is an important problem in many different areas not just basketball. These include musical instrument playing, job related activities, and even our daily moments such as cooking a meal. In our future work, we plan to investigate these new areas, and try to generalize our model to such activities too.



(a) The detected events that contributed most **positively** to a player's performance assessment score according to our model



(b) The detected events that contributed most **negatively** to a player's performance assessment score according to our model

Figure 6: A figure illustrating the events that contribute most positively (top figure) and most negatively (bottom figure) to a player's performance measure according to our model. The red box illustrates the location where our method zooms-in. Each row in the figure depicts a separate event, and the columns illustrate the time lapse of the event (from left to right). We note that among the detected positive events our method recognizes events such as assists, made layups, and made three pointers, whereas among the detected negative events, our method identifies events such as missed layups, and missed jumpshots. We present more of such video examples in the supplementary material.

References

- [1] <https://www.stats.com>. 2
- [2] <http://firstvision.tv/>. 2
- [3] <http://nbasavant.com>. 4
- [4] <https://www.sloansportsconference.com>. 4
- [5] Gedas Bertasius, Hyun Soo Park, Stella X. Yu, and Jianbo Shi. First person action-object detection with egonet. *CoRR*, abs/1603.04908, 2016. 3
- [6] 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. 5
- [7] Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L. Yuille. Attention to scale: Scale-aware semantic image segmentation. *CVPR*, 2016. 3, 5
- [8] 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. 3
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009. 2
- [10] Jeff Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 6
- [11] Alireza Fathi, Jessica K. Hodgins, and James M. Rehg. Social interactions: A first-person perspective. 3
- [12] Alireza Fathi, Xiaofeng Ren, and James M. Rehg. Learning to recognize objects in egocentric activities. In *CVPR*, pages 3281–3288. IEEE Computer Society, 2011. 3
- [13] Ilaria Gori, J. K. Aggarwal, and Michael S. Ryoo. Building unified human descriptors for multi-type activity recognition. *CoRR*, abs/1507.02558, 2015. 3
- [14] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004. 2, 3
- [15] Mostafa S. Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori. A hierarchical deep temporal model for group activity recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [16] 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
- [17] Kihwan Kim, Dongryeol Lee, and Irfan Essa. Detecting regions of interest in dynamic scenes with camera motions. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, June 2012. 3
- [18] Kris Makoto Kitani, Takahiro Okabe, Yoichi Sato, and Akihiro Sugimoto. Fast unsupervised ego-action learning for first-person sports videos. In *CVPR*, 2011. 3
- [19] Hoang Le, Peter Carr, Yisong Yue, and Iain Matthews. Data-driven ghosting using deep imitation learning. In *MITSSAC*, 2014. 2
- [20] Yong Jae Lee and Kristen Grauman. Predicting important objects for egocentric video summarization. *IJCV*, 2015. 3
- [21] Zheng Lu and Kristen Grauman. Story-driven summarization for egocentric video. In *CVPR*, 2013. 3
- [22] Patrick Lucey, Alina Bialkowski, Peter Carr, Stuart Morgan, Iain Matthews, and Yaser Sheikh. Representing and discovering adversarial team behaviors using player roles. *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 00:2706–2713, 2013. 2
- [23] Patrick Lucey, Alina Bialkowski, Peter Carr, Yisong Yue, and Iain Matthews. How to get an open shot: Analyzing team movement in basketball using tracking data. In *MITSSAC*, 2014. 2
- [24] Andrii Maksai, Xinchao Wang, and Pascal Fua. What players do with the ball: A physically constrained interaction modeling. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [25] Andrew Miller and Luke Bornn. Possession sketches: Mapping nba strategies. In *MITSSAC*, 2016. 2
- [26] Kris Kitani Minghuang Ma. Going deeper into first-person activity recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [27] Hyun Soo Park, Jyh-Jing Hwang, Yedong Niu, and Jianbo Shi. Egocentric future localization. In *CVPR*, 2016. 3
- [28] Hyun Soo Park, Jyh-Jing Hwang, and Jianbo Shi. Force from motion: Decoding physical sensation from a first person video. In *CVPR*, 2016. 3
- [29] Hyun Soo Park, Eakta Jain, and Yaser Sheikh. 3d gaze concurrences from head-mounted cameras. In *NIPS*, 2012. 3
- [30] Hyun Soo Park and Jianbo Shi. Social saliency prediction. In *CVPR*, 2015. 3
- [31] Vignesh Ramanathan, Jonathan Huang, Sami Abu-El-Haija, Alexander Gorban, Kevin Murphy, and Li Fei-Fei. Detecting events and key actors in multi-person videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, USA, June 2016. 3
- [32] Xiaofeng Ren and Chunhui Gu. Figure-ground segmentation improves handled object recognition in egocentric video. In *CVPR*, 2010. 3
- [33] 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. 3
- [34] 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. 5
- [35] Shan Su, Jung Pyo Hong, Jianbo Shi, and Hyun Soo Park. Predicting behaviors of basketball players from first person videos. In *CVPR*, 2017. 2, 3
- [36] Eran Swears, Anthony Hoogs, Qiang Ji, and Kim Boyer. Complex activity recognition using granger constrained dbn (gcdbn) in sports and surveillance video. 06 2014. 3
- [37] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, ICCV '15, pages 4489–4497, Washington, DC, USA, 2015. IEEE Computer Society. 5
- [38] Xinyu Wei, Patrick Lucey, Stuart Morgan, Machar Reid, and Sridha Sridharan. The thin edge of the wedge: Accurately predicting shot outcomes in tennis using style and context priors. In *MITSSAC*, 2016. 2
- [39] Yisong Yue, Patrick Lucey, Peter Carr, Alina Bialkowski, and Iain Matthews. Learning fine-grained spatial models for dynamic sports play prediction. *2014 IEEE International Conference on Data Mining (ICDM)*, 00:670–679, 2014. 2