

Predicting Behaviors of Basketball Players from First Person Videos

Shan Su¹, Jung Pyo Hong², Jianbo Shi¹, Hyun Soo Park³

¹University of Pennsylvania, ²KAIST, ³University of Minnesota

{sushan,jshi}@seas.upenn.edu jphong@ai.kaist.ac.kr hspark@umn.edu

Abstract

This paper presents a method to predict the future movements (location and gaze direction) of basketball players as a whole from their first person videos. The predicted behaviors reflect an individual physical space that affords to take the next actions while conforming to social behaviors by engaging to joint attention. Our key innovation is to use the 3D reconstruction of multiple first person cameras to automatically annotate each other's visual semantics of social configurations.

We leverage two learning signals uniquely embedded in first person videos. Individually, a first person video records the visual semantics of a spatial and social layout around a person that allows associating with past similar situations. Collectively, first person videos follow joint attention that can link the individuals to a group. We learn the egocentric visual semantics of group movements using a Siamese neural network to retrieve future trajectories. We consolidate the retrieved trajectories from all players by maximizing a measure of social compatibility—the gaze alignment towards joint attention predicted by their social formation, where the dynamics of joint attention is learned by a longterm recurrent convolutional network. This allows us to characterize which social configuration is more plausible and predict future group trajectories.

1. Introduction

We *physically* interact with people around us while *men-tally* engaging with them via joint attention. For example, you as an audience in a concert are locally affected by the people around you and are globally connected to the people on the other side of the stage by sharing joint attention. While the physical connection delineates the proximal space around us, the mental connection encodes the group's intent in a way that facilitates communications, role playing, and group task accomplishment. These connections provide social cues to further reason about the spatial and temporal extent of the social behaviors, which is a key de-



Figure 1. We predict a group trajectory of basketball players from first person videos. The red is the ground truth and blue is the predicted trajectories with gaze direction.

sign factor for an artificial intelligence of social robots.

However, such social cues are rather ambiguous, subtle, and situation dependent, which is challenging to be computationally learned by third person computer vision systems [4, 25, 26, 32] due to their limited expressibility: it is necessary to tap into what we actually see. In this paper, we propose to use first person cameras collectively to decode the social cues and to further predict their future social behaviors.

What visual information makes us to stay connected to people, physically and mentally? We conjecture that two unique signals recorded in first person videos can describe the connections. (1) Individually, a first person video encodes the egocentric visual semantics that provides a social and spatial context to take the next action. (2) Collectively, first person videos follow joint attention spatially arranged by social formation [24, 38], e.g., audiences dynamically change their social formation to secure visibility, which links the individuals to a group. As a proof-ofconcept, we integrate these two learning signals to predict the movement (location and gaze directions) of basketball players, one of most complex forms of social interactions, from their first person videos (Figure 1).

Our method takes an input, the first person videos of basketball players and outputs a set of plausible future trajectories. We learn an egocentric visual representation to recognize similar social and spatial configurations, e.g., which makes us to move, using a Siamese neural network. This representation is used to retrieve a set of future trajectories per player. We find a plausible group trajectory set from the retrieved trajectories of all players by maximizing a measure of social compatibility—the gaze alignment towards joint attention predicted by their social formation via a generalized Dijkstra algorithm. The dynamics of joint attention is learned by a long-term recurrent convolutional network (LRCN) based on social formation features that encode locations and velocities of the players. Note that we predict not only the future locations but also their gaze directions and joint attention.

Our key innovation is leveraging 3D reconstruction of multiple first person videos to automatically annotate each other's visual semantics of social configurations. This labels the location, orientation, and velocity of other players in pixels, precisely (reprojection error is often less than 0.5 pixel). This makes learning visual social signals on a large scale possible, which provides a richer context of the interactions comparing to third person social activity predictions [4, 25, 26, 32].

A challenge of using first person cameras is that they often produce highly jittery, blurry, and narrow view, unlike third person videos captured from mostly static and omniscient views. We virtually stabilize first person images by applying cylindrical projection, and directly learn visual semantics of social configurations from the images via a convolutional neural network. To resolve a limited visibility issue, we consolidate first person images of all players, which substantially extend visible space via 3D registration.

The first person videos have been increasingly adapted to record professional sports such as basketball, soccer, handball, ice hockey, and American football [1]. Our work provides a computational tool to measure team performance and train players based on how they interact with others based on what they see. Beyond sports, decoding such social sensorimotor behaviors can be used to further explain how social cues are encoded in the human mirror neural system [43]. Also this social intelligence system can apply to content generation for social virtual/augmented reality [2], human-robot interactions, and collaborative education.

Contribution To our best knowledge, this is the first paper that predicts long-term activities from a collection of first person videos. The core technical contributions include (a) learning egocentric visual semantics to recognize social and spatial configurations, (b) using a measure of social compatibility to identify plausibility of social behaviors, (c) formulating the trajectory selection process using a dynamic programming, and (d) learning the dynamics of joint attention via LSTM. We demonstrate the predictive validity of our algorithm in real world basketball datasets by comparing with third person prediction systems.

2. Related Work

Our work integrates two core vision tasks: 1) egocentric social perception: identifying social and spatial configuration, e.g., where I am, who I interact with, and how far they are, and 2) long term social behavior prediction: recognizing a plausible collective behaviors where we use joint attention as a social cue.

Unlike third person vision systems operating in social scenes [10, 11, 13, 28, 41, 44, 45, 49], a first person camera provides in-situ measurements of social interactions from an insider's perspective. This unique property allows a camera to record two sources of information simultaneously. (1) The 3D camera pose reconstructed by structure from motion approximates the gaze orientation, and the intersection of the gaze directions is the location of joint attention [36, 37]. (2) The visual semantics (depth, edge, and surface) of first person images encodes what is socially salient. Faces have been used to recognize a group of people [15] and build visual words to describe joint attention [40]. Subtle reciprocal behaviors can also be recognized [51]. Such visual information from first person cameras has been used for social video editing [8], video summarization [30], human-robot interactions [16], and studying autistic behaviors for children [42].

How are my behaviors affected by others? This question has been a central theme in social psychology [7] and neuroscience, e.g., mirror neuron [43], and their models inspire computational algorithms for multi-agent motion planning in robotics [9, 12, 29] and graphics [23, 34, 39]. A notable model is Helbing's social force model [19] that explains crowd movements as a collection of physical interactions between social agents. This model is used to track a crowd [6] and recognize abnormal behaviors [33].

A group as a whole naturally creates a distinctive geometry of social formation that accommodates its social activity, e.g., a street busker's performance surrounded by a crowd with a half circular formation. Therefore, the formation can be a key indicator to classify the type of social configurations that influence individual behaviors with respect to the group. For instance, Kendon's F-formation theory [24] characterizes the spatial arrangements of a social group, that can be used to identify social interactions in an image [13], and its validity is empirically proven using a large social interaction dataset [38]. In dynamic social scenes, the formation enables re-identifying a group of people in a crowd from non-overlapping camera views [5], and the progression of formation change can be learned via inverse reinforcement learning [32] and discriminative analysis (LSTM) [4].

Note that most prior methods in predicting social behaviors rely on the third person measurements which have a limited access to how we perceive the social configurations. We leverage the visual social semantics embedded in first person cameras, which allows us to directly predict a plausible future group trajectory. This also enables predicting not only people's dynamic locations but also their attention, which have not been explored in prior studies.



Figure 2. (a) We model the space around a person using a cylinder that (b) stabilizes a first person image, $\Theta(\mathbf{I})$. The location and orientation of other players in the image are fully automatically labeled using 3D reconstruction. (c) We retrieve egocentric trajectories by associating the visual semantics of social configuration. (d) The retrieved trajectories are projected onto the first person image.

3. First Person Social Behavior Prediction

We predict a group future trajectory (location and gaze direction) up to 5 seconds given their first person videos. We use the 3D pose of a first person camera as a proxy of the head location, c, and orientation (gaze direction)¹, r where \mathbf{c} and \mathbf{r} are camera optical center and the z axis of the camera rotation (optical axis) in the camera projection matrix, respectively. The camera projection matrices for all players are computed by structure from motion. We represent all variables in 2D by projecting 3D camera pose and joint attention on the 2D basketball court (50 ft.×94 ft.) as shown in Figure 3(b): player's location $\mathbf{x} = \mathbf{c}_{1:2}$, gaze direction $\mathbf{g} = \mathbf{r}_{1:2} / \|\mathbf{r}_{1:2}\| \in \mathbb{S}$, velocity \mathbf{v} , joint attention $\mathbf{s} \in \mathbb{R}^2$ where $\mathbf{r}_{1:2}$ is the first two elements of \mathbf{r} assuming that the coordinate system is aligned with the basketball court origin, i.e., $\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$ is the surface normal of the court. The ground truth joint attention is computed by triangulating the gaze directions of players [37, 38]. For each player, a first person image, I is associated with the gaze, (x, g).

Our method is composed of two parts: 1) egocentric trajectory retrieval per player and 2) a group trajectory selection using a measure of social compatibility. For each player, we recognize images that have similar social and spatial configurations and retrieve a set of N future trajectories (location and orientation) in Section 3.1. This generates nN trajectories for n players, and we find a plausible group trajectory set that maximizes a measure of social compatibility (Section 3.2) while localizing joint attention using LRCN (Section 3.3).

3.1. First Person Trajectory Retrieval

We behave similarly in similar social situations. The location, velocity, and orientation of other players are recorded in a first person image, **I**, which encodes not only spatial layout, e.g., basket, center line, and background, but also social layout, e.g., where are other players, around the person. In this section, we learn the visual representation of social and spatial configurations from first person images.

We use the 3D reconstruction of first person videos to automatically annotate each other's location and orientation in pixels. We model each player using a cylinder with radius r and height H and project the cylinder onto a first person image, I. The relative gaze direction, Δg is recorded in the label image, M using the HSV color map (Figure 2(b)):

$$\mathbf{M}_{xy} = \begin{cases} (0,0,0) & \text{if } j^* = \emptyset\\ (\Delta \mathbf{g}_{j^*}, 0.9, 0.9) & \text{otherwise} \end{cases}$$

where $j^* = \operatorname*{argmin}_{j} \min_{\lambda} \{\lambda | \mathbf{c} + \lambda \mathbf{r} \in \mathcal{C}_j, \lambda > 0\},$

 $\Delta \mathbf{g} = \angle \mathbf{g}_j - \angle \mathbf{g}$ is the j^{th} relative gaze direction, $C_j = \{\mathbf{c} | \Delta \mathbf{c}_3 < H, \| \Delta \mathbf{c}_{1:2} \| < r, \Delta \mathbf{c} = \mathbf{c}_j - \mathbf{c} \}$ is a set of 3D points in the cylinder of the j^{th} player. This label image directly encodes social configuration around a person.

We stabilize a first person image onto the cylindrical surface² (Figure 2(b)).

$$\Theta(\mathbf{I})_{\theta h} = \mathbf{I}_{xy} \quad \text{where} \quad \left\{ \begin{array}{l} x = \mathbf{r}_x^\mathsf{T} \mathbf{z} / \mathbf{r}^\mathsf{T} \mathbf{z} \\ y = \mathbf{r}_y^\mathsf{T} \mathbf{z} / \mathbf{r}^\mathsf{T} \mathbf{z} \end{array} \right.$$

and $\mathbf{z} = \begin{bmatrix} \cos \theta & \sin \theta & h \end{bmatrix}^{\mathsf{T}}$. \mathbf{r}_x and \mathbf{r}_y is the X and Y axes of the first person camera. The mapping function Θ applies to both first person image I and label image M.

The warped image has three properties that make visual learning effective. 1) Aligned vanishing lines: the head and foot location of the players are dependent solely on the depth given similar height; 2) no perspective distortion: the scale in image linearly proportional to the inverse depth; 3) optical center invariance: the representation is linear in angle where the optical center shift is linear translation in angle.

We learn the visual social semantics using a Siamese neural network. We generate the positive and negative pairs of images based on \mathbf{M} , i.e., positive if $\|\Theta(\mathbf{M}_i) - \Theta(\mathbf{M}_j)\| < \epsilon$ and negative otherwise, where $\Theta(\mathbf{M})$ is the stabilized label image. We minimize the following contrastive loss for training:

$$L_{\text{soc}} = \sum_{(i,j)\in\mathcal{P}} l_{ij} \|\Delta\phi\|^2 + (1 - l_{ij}) \max(0, m^2 - \|\Delta\phi\|^2)$$

¹Optionally, the fixed spatial relationship between camera optical axis and primary gaze direction can be calibrated [37].

²Similar projection has been used to generate a panoramic image [47].



Figure 3. (a) In conjunction with location and orientation prior, we use the visual semantics in a first person image to retrieve trajectories, which shows strong selectivity. GT: ground truth trajectory, Ret: retrieved trajectory. (b) We select K best trajectory sets in a trellis graph using a generalized Dijkstra algorithm. A vertical space represents the retrieved trajectories per player and the path cost is computed by Equation 1.

where l_{ij} is a label indicating positive and negative pairs, $\phi(\Theta(\mathbf{I}))$ is the visual feature of the warped image $\Theta(\mathbf{I})$ learned by a convolutional neural network (CNN). $\Delta \phi = \phi(\Theta(\mathbf{I}_i)) - \phi(\Theta(\mathbf{I}_j))$, \mathcal{P} is the set of pairs, and m is a margin between positive and negative pairs. We use the pre-trained CNN [27] and refine the weights through the training. The paired stabilized images are passed through two identical subnetworks, and L_{soc} is computed by the contrastive loss of FC7 feature of the two images.

We empirically observed that this pairing across all locations inclines to learn the background because a first person image is dominated by background pixels, e.g., the network learns ego-motion rather than social configurations [3, 20, 21, 35]. Instead, we make pairs that are located and oriented in the similar area of the basketball courts, i.e., $\|\mathbf{x}_i - \mathbf{x}_j\| < \epsilon_{\mathbf{x}}$ and $|\angle \mathbf{g}_i - \angle \mathbf{g}_j| < \epsilon_{\mathbf{g}}$. Our learning based approach is beneficial in particular dynamic social scenes that include severe motion blur, illumination and view point changes where standard structure from motion often fails.

Based on the learned feature of the target image, $\phi(\Theta(\mathbf{I}_{tar}))$, we retrieve N 2D trajectories, $\mathcal{T} = {\mathbf{T} | \epsilon > \| \phi(\Theta(\mathbf{I}_{tar})) - \phi(\Theta(\mathbf{I})) \|}$ where $\mathbf{T} = {\mathbf{x}^t, \mathbf{g}^t}_{t=1}^T$ is a trajectory (location and gaze direction) of each player and ϵ is the feature decision boundary learned by the neural network. Similar to the training phase, we restrict the training data samples based on location and orientation. In practice, we cluster the trajectories, \mathbf{T} using Medoidshift [46] to identify topologically distinctive trajectories that are projected onto the first person image in Figure 2(d). Note that our first person trajectory retrieval is highly selective as shown in Figure 3(a).

The retrieved trajectories have three properties: 1) they discover egocentric physical space to move based on social configurations; 2) they include diverse topological structure, i.e., different trajectories may be plausible given a social configuration; and 3) they reflect spatial layout.

Why learning? We exploit visual learning based retrieval

because 1) it can encode subtle pose signals that describe where they will move next; and 2) large reprojection error caused by localization error can be handled.

3.2. Group Trajectory via Social Compatibility

There exist N^n possible combinations of group trajectories where n is the number of players. The trajectories are retrieved independently, and not all combinations are *socially* plausible. In this section, we recognize the plausible trajectory combinations using a measure of social compatibility—the gaze alignment towards joint attention predicted by social formation. Note that we consolidate all retrieved egocentric trajectories by registering them into the basketball court.

There are two ways of computing joint attention in a static social scene: 1) geometrically finding the intersection of gaze directions [37] and 2) statistically learning the characteristics of the social formation, which does not require knowing gaze directions [38]. Note that we denote the geometrically computed joint attention as s to differentiate with the statistically estimated joint attention \hat{s} .

Ideally, these two locations of joint attention must agree, and we define a measure of social compatibility based on the alignment between two joint attentions:

$$\eta = \frac{1}{n} \sum_{\mathbf{x} \in \mathcal{X}} \frac{\left(\hat{\mathbf{s}} - \mathbf{x}_{i}\right)^{\mathsf{T}} \mathbf{g}_{i}}{\|\hat{\mathbf{s}} - \mathbf{x}_{i}\|}$$

where $\mathcal{X} = \{\mathbf{x}_i, \mathbf{g}_i\}_{i=1}^n$ is a set of player locations and gaze directions. The social compatibility measures how the gaze directions are geometrically aligned with statistically computed joint attention, and it characterizes which social formation and corresponding gaze directions are socially plausible. Note that $\hat{\mathbf{s}}$ is a function of $\{\mathbf{x}_i\}_{i=1}^n$.

We integrate the social compatibility over time to evalu-



(a) Joint attention prediction

(b) Dynamics of joint attention

Figure 4. (a) We predict joint attention using a social formation feature that encodes the player's location and instantaneous velocity (bottom right). Top row: the predicted joint attention is projected onto each first person image. (b) The social formation features are used to learn the dynamics of joint attention using LRCN [14].

ate a group trajectory set:

$$\eta = \frac{1}{nT} \sum_{i=1}^{n} \eta_t \left(\{ \hat{\mathbf{s}}_t \}_{t=1}^T, \mathbf{T}_i \right),$$

where η_t is the accumulated measure of social compatibility over T time instances.

Group Trajectory Selection Among the nN retrieved trajectories from all players, $\{\mathcal{T}_i\}_{i=1}^n$, we find a group trajectory set that maximizes the measure of social compatibility:

$$\underset{\{p_i\}_{i=1}^n}{\operatorname{argmin}} \quad -\frac{1}{nT} \sum_{i=1}^n \eta_t \left(\{\hat{\mathbf{s}}_t\}_{t=1}^T, \mathbf{T}_{p_i}\right), \tag{1}$$

where $\{p_i\}_{i=1}^n$ is an index set for the retrieved trajectory of each player.

Solving Equation (1) by the exhaustive search is computationally prohibited, $\mathcal{O}(N^n)$. A stochastic search such as Monte Carlo simulations does not apply due to the low probability to choose a correct model. Instead, we employ the generalized Dijkstra algorithm, or Yen's algorithm [50] to efficiently find the K best trajectory sets.

We construct a trellis graph where the vertical slice represents a set of the retrieved trajectories per players, \mathcal{T}_i , i.e., each node is a trajectory and an edge indicates the trajectory selection as shown in Figure 3. A path along the trellis graph determines the selected trajectory set where the path cost is defined in Equation (1). Despite the greedy search due to a nonlinear prediction of joint attention (Section 3.3), in practice, the algorithm finds "good" solutions that have high social compatibility. We predict the group behaviors using the selected trajectory set.

3.3. Joint Attention Dynamics

Equation (1) requires joint attention prediction, \hat{s} . In this section, we learn the dynamics of joint attention with respect to social formation using LRCN [14].

As an input of the network, we generate a formation feature image, $\Phi(\mathcal{X})$ that encodes the occupancy and instantaneous velocity, v, of the players in a discretized basketball court. The HSV value of the formation feature image is set to:

$$\begin{split} \Phi_{ij}(\mathcal{X}) &= \begin{cases} (\angle \bar{\mathbf{v}}_{ij}, \|\bar{\mathbf{v}}_{ij}\|, 0.9) & \text{if } |\mathcal{C}_{ij}| > 0\\ (0, 0, 0) & \text{if } |\mathcal{C}_{ij}| = 0 \end{cases} \\ \text{where } \bar{\mathbf{v}}_{ij} &= \frac{1}{|\mathcal{C}_{ij}|} \sum_{\mathbf{v}_k \in \mathcal{C}_{ij}} \mathbf{v}_k, \end{split}$$

 $C_{ij} = \{\mathbf{v} | \mathbf{x} \in C_{ij}\}$ and the (i, j) cell of the court. $\Phi(\mathcal{X})$ is illustrated in the bottom right of Figure 4(a). Note that unlike social dipole moment [38], this representation is independent on the location of center of mass and joint attention, which is robust to missing data.

We use LRCN with a few minor modifications to learn the dynamics of joint attention. We minimize the following joint attention error:

$$L_{\text{LSTM}} = \sum_{t=1}^{T} \|\mathbf{s}_t - \hat{\mathbf{s}}_t\|^2,$$

where the $\hat{\mathbf{s}}_{t+1}$ is recursively computed by

$$\hat{\mathbf{s}}_{t+1} = f(\hat{\mathbf{s}}_t, \Phi(\mathcal{X}_t); \mathbf{w}_{\text{CNN}}, \mathbf{w}_{\text{LSTM}}).$$
(2)

f is the dynamics parametrized by the weights of a convolutional neural network, \mathbf{w}_{CNN} , and a long short-term memory unit, \mathbf{w}_{LSTM} as shown in Figure 4(b). We initialize \mathbf{w}_{CNN} based on pre-trained model [27] separately with further refinement by regressing the static location of joint attention from social formation, $\mathbf{s} = g(\Phi(\mathbf{X}); \mathbf{w}_{\text{CNN}})$.

4. Basketball Dataset Analysis

We use the first person basketball video data collected by the university team at Northwestern Polytechnical University in China [8, 38]. The dataset includes 10.5 hours of basketball games. We take two steps for reconstruction: (1) reference reconstruction: we subsample images from each player to reconstruct the reference 3D points and cameras



Figure 5. (a) Players consistently engage joint attention while playing. They look at the joint attention more than 60% of their play. (b) Role is a key factor to determine social formations. We illustrate distributions of the Center and Wing players given the ball holder's location. (c) The role of a player is a strong prior to predict other players, e.g., two Centers from different teams often move together to block each other. PF: Power Forward, C: Center, PG: Point guard, SF: Small Forward, SG: Shooting guard.

 $(\sim3,000 \text{ images})$ using structure from motion [18]; and (2) camera registration: we register each image into the reference reconstruction coordinate system using a camera resectioning algorithm [31] with local bundle adjustment up to 500 consecutive images.

Figure 5(a) illustrates a normalized angle histogram of joint attention engagement. This indicates that the players consistently align their gaze directions to joint attention (< 40 degree): 83%, 65%, and 48% of their play at 0 m/s, 1.5 m/s, and 3 m/s speed, respectively. As the speed gets faster, the player's gaze direction tends to deviate from the joint attention: it often follows the fast motion, which forms behind the person (180 degrees) at high speed.

Player's role is a key factor to characterize social formations. Figure 5(b) illustrates a spatial distribution of players based on their role, given joint attention. For instance, when the Center possesses a ball, Power Forward and Center are likely located near the basket area for blocking and rebound. When a Point Guard possesses the ball, players tend to be distributed widely to create space to receive the ball. Also the role is a strong predictor of the play as similar roles in different teams enforces them to move together. Figure 5(c) shows that a strong correlation of roles in different teams.

5. Result

We evaluate our social behavior prediction on two basketball datasets, which are captured at NPU and CMU, respectively. The two datasets significantly differ from each other in terms of at least 5 aspects. (1) Location: the visual appearance of the court does not match; (2) Lighting: illumination and reflectance from floor are different; (3) Level of players' skill: NPU players are semi-professionals while CMU players range from novice players to experts; (4) Strategy: no team coach directs the CMU games unlike NPU; (5) Jersey: CMU players wear team jersey unlike NPU.

We use AlexNet [27] to train the Siamese network with Caffe [22]. 240k image pairs are generated from first person



Figure 6. We compare our predicted joint attention with 7 baseline algorithms which it shows 7m error after 5 seconds.

images of players where the pairs are selected within similar location ($\epsilon_x < 3m$) and orientation ($\epsilon_g < 45$ degrees) in the basketball court. Due to the location and orientation prior, the network can be efficiently trained with strong generalization power (98.7% testing accuracy). For training the dynamics of joint attention, we concatenate the AlexNet FC7 layer with LSTM through Theano [48]. We generate 85k sequences of joint attention and corresponding social formation feature (210×410). The testing average error over 5 seconds is 3.12m.

5.1. Quantitative Evaluation

We evaluate our prediction in three categories: joint attention, missing trajectory, and social trajectories.

Joint attention prediction We compare our method with 7 baseline algorithms for predicting 5 seconds. A) Zero velocity (ZV) and linear constant velocity (LV) extrapolate the location of joint attention by taking into account instantaneous velocity; B) Center of mass (COM) and center of circumcircle (CC) are geometrically computed based on the locations of players; C) Social dipole moment [38] (SDM) is used to learn a binary classifier (AdaBoost) to recognize the location of joint attention; D) Our social formation fea-



Figure 7. (a) and (b): We evaluate our algorithm by comparing with 7 baseline algorithms including state-of-the-art Social LSTM, and our method consistently outperforms other methods on missing data prediction task. In particular, our method shows a strong predictive power on gaze direction. (c) and (d): We predict a group trajectory set and compare with Vanilla LSTM and Social LSTM while no comparable algorithm exists for graze prediction.

ture image (SFI), $\Phi(\mathcal{X})$, with LSTM is used to predict joint attention using a convolutional neural network [27]. We train the network to minimize the Euclidean loss of $||\mathbf{s} - \hat{\mathbf{s}}||$; E) A Bayesian filtering (BF) is applied for temporal smoothing by learning a stochastic dynamics of joint.

Figure 5 illustrates the predictive validity where our method outperforms all baseline algorithms. In particular, it shows a strong predictive power up to 4 seconds with 5 m error in a highly dynamic scene. The error in LV and ZV indicates the nature of dynamics of the basketball game. COM, CC, SDM, and SFI are time independent predictors where COM shows the most consistent and strongest prediction. This is caused by the fact that social formations in basketball data are often distributed near the basket area where the center of mass of players is likely located.

Missing trajectory prediction We apply our method for missing trajectory prediction. We leave out a trajectory and predict its behaviors using social compatibility.

We compare our method with 7 baseline algorithms. A) We use a kinematic prior to predict a trajectory: location (Loc), orientation (Ori), velocity (Vel), and their combinations. B) We compare with state-of-the-art third person prediction systems based on Vanilla LSTM [17] and Social LSTM [4]. We use the occupancy based Social LSTM which applies pooling based on social proximity. C) We compare with first person prediction based solely on visual features (Img) (no kinematic knowledge). The visual features are learned by our Siamese network. Note that we compare not only future locations but also gaze directions except for Vanilla and Social LSTMs where gaze prediction is not possible with their trivial extension.

Figure 8(d) indicates that orientation or velocity is a strong prior to predict future while our method produces more selective trajectories due to the social compatibility measure. Vanilla LSTM produces unconvincing results due to its limited expressibility on social interactions and Social LSTM shows drifts because the behaviors of basketball players are often affected by long range team players. Notably a first person image based method without kinematic knowledge (Img) performs poorly, which indicates visual information alone can be ambiguous.

Our method outperforms all baseline algorithms. In par-

ticular, our method shows strong predictive power on gaze direction driven by joint attention.

Social trajectory prediction We focus on comparing with third person approaches: Vanilla LSTM and Social LSTM. Note that both LSTMs require longer observation time (10 seconds) to predict 5 seconds while our first person based method needs 0.5 second (instantaneous velocity).

Note that Vanilla LSTM behaves similarly to the missing data prediction as it has no consideration on social behaviors. Our method produces the error range, 6 m and 60 degree error after 5 seconds as shown in Figure 7.

We also characterize the prediction error based on player's role summarized in Table 1. This error indicates that the predictive power differs by the roles, e.g., predicting Shooting Guard's behaviors is relatively more difficult than Centers because of diverse interactions across the court.

5.2. Qualitative Evaluation

We apply our method to predict players future behaviors in diverse basketball scenarios. Figure 8 shows trajectory and joint attention predictions. We also show the retrieved sequences that have similar social configuration to reason about predictions.

6. Summary

We present a method to predict the future location and gaze direction of basketball players from their first person videos. 3D reconstruction of multiple first person videos provides the automatic supervision for learning visual social semantics. We use the learned representation to retrieve trajectories per player. We evaluate the plausibility of each group trajectory using social compatibility. We select K best group trajectories using a generalized Dijkstra's algorithm. We demonstrate that our first person based method is effective, outperforming state-of-the-art social activity prediction systems that use third person views.

7. Acknowledgement

This work is partially supported by the National Science Foundation (IIS 1651389) and Facebook/Oculus gift.

	Power Forward			Point Guard			Small Forward			Shooting Guard			Center		
	1sec	3 sec	5 sec	1sec	3 sec	5 sec	1sec	3 sec	5 sec	1sec	3 sec	5 sec	1sec	3 sec	5 sec
Ours	0.50	0.60	0.40	1.59	3.76	5.71	0.64	0.25	2.41	1.51	4.95	7.79	1.39	0.65	1.70
Vanilla LSTM	6.50	10.86	13.77	6.85	12.86	8.30	3.52	3.81	9.05	1.98	15.43	7.69	9.35	5.54	12.23
Social LSTM	2.98	3.02	3.59	6.55	11.49	12.32	0.53	5.35	7.84	1.99	10.19	2.43	5.14	1.60	7.36

Table 1. Trajectory prediction error based on player's role



(a) Taking-turn

 Trajectory prediction
 Image: Sequence

 Joint attention prediction
 Image: Sequence

(b) Attack







(d) Drive-in

Figure 8. We evaluate our algorithm qualitatively in diverse scenarios (taking-turn, attack, drive-in, and shot) on both datasets. The first column and top row: a comparison between the predicted trajectories with gaze directions in blue with ground truth trajectory in red up to 5 seconds. First column and bottom row: a comparison between the predicted joint attention in green with the ground truth joint attention in orange. Transparency encodes time. Second column: a comparison between a target sequence (top row) and the retrieved sequence (bottom row). We also show the retrieved sequences to reason about our prediction. The retrieved sequence has similar social configuration as time evolves. The predicted trajectories and joint attention are projected onto the target sequence to validate the prediction. The joint attention agrees with scene activities. The blank space is missing data where structure from motion fails.

References

- [1] http://firstvision.tv/.
- [2] https://techcrunch.com/2016/10/06/facebook-social-vr/.
- [3] P. Agrawal, J. Carreira, and J. Malik. Learning to see by moving. In *ICCV*, 2015.
- [4] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *CVPR*, 2016.
- [5] A. Alahi, V. Ramanathan, and L. Fei-Fei. Socially-aware large-scale crowd forecasting. In *CVPR*, 2014.
- [6] S. Ali and M. Shah. Floor fields for tracking in high density crowd scenes. In ECCV, 2008.
- [7] G. W. Allport. The historical background of social psychology. *McGraw Hill*, 1985.
- [8] I. Arev, H. S. Park, Y. Sheikh, J. K. Hodgins, and A. Shamir. Automatic editing of footage from multiple social cameras. *SIGGRAPH*, 2014.
- [9] S. Bhattacharya, M. Likhachev, and V. Kumar. Multi-agent path planning with multiple tasks and distance constraints. In *ICRA*, 2010.
- [10] I. Chakraborty, H. Cheng, and O. Javed. 3D visual proxemics: Recognizing human interactions in 3D from a single image. In *CVPR*, 2013.
- [11] W. Choi, Y. Chao, C. Pantofaru, and S. Savarese. Discovering groups of people in images. In ECCV, 2014.
- [12] L. Cohen, T. Uras, S. Kumar, H. Xu, N. Ayanian, and S. Koenig. Improved solvers for bounded-suboptimal multiagent path finding. In *IJCAI*, 2016.
- [13] M. Cristani, L. Bazzani, G. Paggetti, A. Fossati, D. Tosato, A. D. Bue, G. Menegaz, and V. Murino. Social interaction discovery by statistical analysis of F-formations. In *BMVC*, 2011.
- [14] J. Donahue, L. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.
- [15] A. Fathi, J. K. Hodgins, and J. M. Rehg. Social interaction: A first-person perspective. In CVPR, 2012.
- [16] I. Gori, J. K. Aggarwal, L. Matthies, and M. S. Ryoo. Multitype activity recognition in robot-centric scenarios. In *ICRA*, 2016.
- [17] K. Greff, R. K. Srivastava, J. Koutnik, B. R. Steunebrink, and J. Schmidhuber. Lstm: A search space odyssey. *arXiv*:1503.04069, 2015.
- [18] R. Hartley and A. Zisserman. Multiple View Geometry in Computer Vision. Cambridge University Press, second edition, 2004.
- [19] D. Helbing and P. Molnár. Social force model for pedestrian dynamics. *Physics Review*, 1995.
- [20] D. Jayaraman and K. Grauman. Learning image representations tied to ego-motion. In *ICCV*, 2015.
- [21] D. Jayaraman and K. Grauman. Look-ahead before you leap: End-to-end active recognition by forecasting the effect of motion. In *ECCV*, 2016.
- [22] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv*:1408.5093, 2014.

- [23] I. Karamouzas, B. Skinner, and S. J. Guy. Universal power law governing pedestrian interactions. *Physics Review Letter*, 2014.
- [24] A. Kendon. Conducting Interaction: Patterns of Behavior in Focused Encounters. Cambridge University Press, 1990.
- [25] K. Kim, M. Grundmann, A. Shamir, I. Matthews, J. Hodgins, and I. Essa. Motion field to predict play evolution in dynamic sport scenes. In *CVPR*, 2010.
- [26] K. M. Kitani, B. Ziebart, J. D. Bagnell, and M. Hebert. Activity forecasting. In ECCV, 2012.
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [28] T. Lan, Y. Wang, W. Yang, S. Robinovitch, and G. Mori. Discriminative latent models for recognizing contextual group activities. *PAMI*, 2012.
- [29] S. Lee, Y. Diaz-Mercado, and M. Egerstedt. Multi-robot control using time-varying density functions. *TRO*, 2015.
- [30] Y. J. Lee, J. Ghosh, and K. Grauman. Discovering important people and objects for egocentric video summarization. In *CVPR*, 2012.
- [31] V. Lepetit, F. Moreno-Noguer, and P. Fua. Epnp: An accurate o(n) solution to the pnp problem. *IJCV*, 2008.
- [32] W.-C. Ma, D.-A. Huang, N. Lee, and K. M. Kitani. A gametheoretic approach to multi-pedestrian activity forecasting. In arXiv:1604.01431v2, 2016.
- [33] R. Mehran, A. Oyama, and M. Shah. Abnormal crowd behavior detection using social force model. In CVPR, 2009.
- [34] R. Narain, A. Golas, S. Curtis, and M. C. Lin. Aggregate dynamics for dense crowd simulation. SIGGRAPH Asia, 2009.
- [35] H. S. Park, J.-J. Hwang, Y. Niu, and J. Shi. Egocentric future localization. In CVPR, 2016.
- [36] H. S. Park, E. Jain, and Y. Sheikh. Predicting primary gaze behavior using social saliency fields. In *ICCV*, 2013.
- [37] H. S. Park, E. Jain, and Y. Shiekh. 3D social saliency from head-mounted cameras. In *NIPS*, 2012.
- [38] H. S. Park and J. Shi. Social saliency prediction. In CVPR, 2015.
- [39] J. Pettré and M. C. Lin. New generation crowd simulation algorithms. SIGGRAPH, 2014.
- [40] G. Pusiol, L. Soriano, L. Fei-Fei, and M. C. Frank. Discovering the signatures of joint attention in child-caregiver interaction. In *CogSci*, 2014.
- [41] V. Ramanathan, J. Huang, S. Abu-El-Haija, A. Gorban, K. Murphy, and L. Fei-Fei. Detecting events and key actors in multi-person videos. In *CVPR*, 2016.
- [42] J. M. Rehg, G. D. Abowd, A. Rozga, M. Romero, M. A. Clements, S. Sclaroff, I. Essa, O. Y. Ousley, Y. Li, C. Kim, H. Rao, J. C. Kim, L. L. Presti, J. Zhang, D. Lantsman, J. Bidwell, and Z. Ye. Decoding children's social behavior. In *CVPR*, 2013.
- [43] G. Rizzolatti and L. Craighero. The mirror-neuron system. Annual Review of Neuroscience, 2004.
- [44] M. Rodriguez, J. Sivic, I. Laptev, and J.-Y. Audibert. Datadriven crowd analysis in videos. In *ICCV*, 2011.
- [45] F. Setti, O. Lanz, R. Ferrario, V. Murino, and M. Cristani. Multi-scale F-formation discovery for group detection. In *ICIP*, 2013.

- [46] Y. Sheikh, E. Khan, and T. Kanade. Mode-seeking by medoidshifts. In *ICCV*, 2007.
- [47] R. Szeliski and H.-Y. Shum. Creating full view panoramic image mosaics and environment maps. SIGGRAPH, 1997.
- [48] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv:1605.02688*, 2016.
- [49] Y. Yang, S. Baker, A. Kannan, and D. Ramanan. Recognizing proxemics in personal photos. In CVPR, 2012.
- [50] J. Y. Yen. Finding the k-shortest loopless paths in a network. *Management Science*, 1971.
- [51] R. Yonetani, K. M. Kitani, and Y. Sato. Recognizing microactions and reactions from paired egocentric videos. In *CVPR*, 2016.