

Learning Person Trajectory Representations for Team Activity Analysis

by

Nazanin Mehrasa

B.Sc., Amirkabir University of Technology (Tehran Polytechnic), 2015

Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of
Master of Science

in the
School of Computing Science
Faculty of Applied Sciences

© Nazanin Mehrasa 2017
SIMON FRASER UNIVERSITY
Spring 2017

Copyright in this work rests with the author. Please ensure that any reproduction or re-use is done in accordance with the relevant national copyright legislation.

Approval

Name: Nazanin Mehrasa

Degree: Master of Science (Computing Science)

Title: Learning Person Trajectory Representations for Team Activity Analysis

Examining Committee: **Chair:** Jiannan Wang
Assistant Professor

Greg Mori
Senior Supervisor
Professor

Luke Bornn
Co-Senior Supervisor
Assistant Professor
Department of Statistics and Actuarial Science

Mark Drew
Internal Examiner
Professor

Date Defended: April 18, 2017

Abstract

Activity analysis in which multiple people interact across a large space is challenging due to the interplay of individual actions and collective group dynamics. We propose an end-to-end generic approach for learning person trajectory representations for group activity analysis. The learned representations encode rich spatio-temporal dependencies and capture useful motion patterns for recognizing individual events, as well as characteristic group dynamics that can be used to identify groups from their trajectories alone. We develop our deep learning approach in the context of team sports, which provide well-defined sets of events (e.g. pass, shot) and groups of people (teams). We evaluate our model on NBA basketball and NHL hockey games datasets. Analysis of events and team formations using these two sports datasets demonstrate the generality of our approach. Experiments show that our model is capable of (1) capturing strong spatio-temporal cues for recognizing events in hockey dataset (2) capturing distinctive group dynamics for identifying group identity.

Keywords: Trajectory Features; Shared-Compare Trajectory Network; Stacked Trajectory Network; Team Sport Analysis

Dedication

To my lovely family for their love, support and faith.

Acknowledgements

First, I would like to thank my senior supervisor Dr. Greg Mori for his unreserved support, patience, and encouragement during the course of this project. Whenever I got a problem, Greg was always there ready to help, with his immense knowledge, invaluable inspiration, continuing encouragement and great patience. I am really grateful for his great support on my research. In addition, I also would like to thank Greg for demonstrating how to manage every aspect of life as efficiently as possible.

I also would like to thank Yatao Zhong who gave me a great help throughout this project. He is a wonderful friend and a great teammate. Whenever we had a problem, he was always there ready to help. Without his help, we would not be able to finish this project.

I would also like to express my gratitude to Dr. Fred Tung who provided me with deep insights on various aspects of this project. His guidance had a great effect on my research.

I would also like to thank Dr. Luke Bornn, Dr. Mark Drew, and Dr. Jiannan Wang, the members of my supervisory committee, for their helpful suggestions and insightful comments.

I would like to thank my brilliant friends who helped me through this project all around the world, including Sima (she was the "always there" friend), Sajjad, Karoon, Saman, Mehran, Amir, Sepehr, Hossein, Rana, Shokoofe, Kiana, Ramtin, Ehsan, Aydin, and Amirali.

Finally, I would like to thank my mother, father, and sister, who always supported me, whatever path I took. You were always there for me.

Table of Contents

Approval	ii
Abstract	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vi
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Contribution	3
2 Related Work	4
2.1 Deep Learning for Action Recognition	4
2.2 Group Activity Recognition	5
2.3 Trajectory Data Analytics	5
3 Proposed Approach	7
3.1 Shared-Compare Trajectory Network	7
3.2 Stacked Trajectory Network	10
4 Datasets	11
4.1 The SPORTLOGiQ NHL Dataset	11
4.1.1 Data Preprocessing	11
4.2 The STATS SportVU NBA Dataset	12
4.2.1 Data Preprocessing	12
5 Experiments	14
5.1 Event Recognition on the NHL Dataset	14

5.1.1	Task Description	14
5.1.2	Experiment Settings	14
5.1.3	Baselines	15
5.1.4	Train phase	15
5.1.5	Test phase	15
5.1.6	Experiment Results	15
5.2	Team Identification on the NBA Dataset	19
5.2.1	Task Description	19
5.2.2	Experiment Settings	19
5.2.3	Baseline	19
5.2.4	Measurement	19
5.2.5	Results and Analysis	19
6	Conclusion and Future Work	23
6.1	Limitations and Future work	23
6.2	Discussion	24
	Bibliography	25

List of Tables

Table 5.1	Average precision for event recognition on the NHL dataset, unknown key player.	16
Table 5.2	Average precision for event recognition on the NHL dataset, known key player.	16
Table 5.3	Metrics on models with different number of layers. All convolutional layers use a filter size of 3 except the first layer, where the filter size is 5. The number of filters in next layer is double the number in previous layer except the fifth layer (if any), where the number of filters is the same as that in the fourth layer. The number of neurons in fully connected layer is set to 1024.	20
Table 5.4	Metrics on models with different filter sizes. All models in the table use five convolutional layers with no fully connected layer. The filter sizes listed is in a bottom-up order and the number of filters used are 64, 128, 256, 512, 512 (bottom-up order).	20
Table 5.5	Metrics on models with different number of filters. All models in the table use five convolutional layers with no fully connected layer. The base number of filters listed in the table is the number of filters in the first layer. The number of filters in next layer is double the number in previous layer except that the fourth and the fifth layers have the same number of filters.	21
Table 5.6	Stacked trajectory network vs. IDT. Comparison with this baseline show that our model captures valuable collective features for understanding group dynamics	21

List of Figures

Figure 1.1	When activities involve multiple people distributed in space, the relative trajectory patterns of different people can provide valuable cues for activity analysis. We learn rich trajectory representations that encode useful information for recognizing individual events as well as overall group dynamics in the context of team sports. For example this learnt features can be used for recognizing that a hockey player trying to pass the puck t while an opposing player moves to block the pass.	2
Figure 2.1	Our shared-compare trajectory network (Sec.3.1) learns generic individual representations as well as features describing pairwise interactions. Pairs are determined by an ordering based on proximity to a key player. The shared-compare structure is useful for recognizing events such as passes and shots.	6
Figure 3.1	Our stacked trajectory network (Section 3.2) takes trajectories of all group members as input and learns representations that consider all interactions within the group. The stacked network captures overall group dynamics and can be used to identify groups (teams) using member trajectories alone.	9
Figure 4.1	Number of samples per event in the SPORTLOGiQ NHL dataset. .	12
Figure 4.2	Number of possessions of each team in the STATS SportVU NBA dataset.	13
Figure 5.1	Precision-recall curves for each event in the NHL dataset. The first row shows the results of the known key player case while the second row the unknown key player case.	16

Figure 5.2	Top 5 candidates retrieved as shot. Green sequences are true positives while red ones are false positives. The person with a bounding box is the “key” player who is performing the action. We show 8 frames of the 16-frame video clip by sub-sampling. If a frame is black, it means the key player is missing because of failure to detect and track the player.	17
Figure 5.3	Top 5 candidates retrieved as dump in. Green sequences are true positives while red ones are false positives. The person with a bounding box is the “key” player who is performing the action. We show 8 frames of the 16-frame video clip by sub-sampling. If a frame is black, it means the key player is missing because of failure to detect and track the player.	18
Figure 5.4	(a) Visualization of the filters in the first convolutional layer. (b) Confusion matrix based on game-wise prediction. Both figures are created using the 5conv model in Table 5.3.	22

Chapter 1

Introduction

Human activity analysis is a fundamental problem in computer vision and it has a wide range of applications in the context of surveillance and security, health-care monitoring, team-sport analysis and etc. Activities can be divided into two categories: (1) They may involve a single person performing a series of actions to complete a task. (2) they can include multiple people distributed across a large space collectively trying to achieve a shared goal which is referred as the group activity. Group activity analysis is more complex because it has the complexity of both recognizing each individual activity and also capturing the pairwise and global relationship of people in the scene. In this work, we focus on the end-to-end learning of feature representations for analyzing group activities which involve multiple interacting people distributed in space. In the end to end learning procedure, we have trajectories as the input and we will learn trajectory representation and we can have team identity or event label as the output. Our aim is to learn rich feature representations that encode useful information about individual events as well as the overall group dynamics.

We observe that a sequence of a person’s movement across space over time (called trajectory), can provide valuable information such as the person’s goal or intention. When activities involve multiple interacting people, the relative trajectory patterns of different people can also provide important interaction cues. For example, consider the snapshot of time shown in Fig. 1.1, which shows a hockey player trying to pass the puck to a teammate while an opposing player moves to block the pass.

A large volume of related works have focused on visual features for recognizing individual actions. These are typically built from challenging unconstrained Internet video datasets such as UCF Sports [23], UCF-101 [27], HMDB-51 [16], and Sports-1M [13]. However, it is hard to determine what a crowd of players is doing by simply inspecting the pixels inside the bounding boxes of individuals in a video clip. It is helpful to analyze the relative positions of people over time as well, so as to understand the spatio-temporal context and then understand the behavior of each person. We will show that learned person trajectory

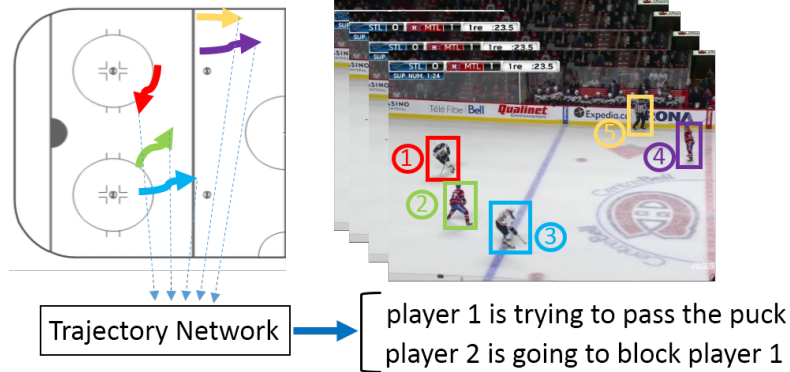


Figure 1.1: When activities involve multiple people distributed in space, the relative trajectory patterns of different people can provide valuable cues for activity analysis. We learn rich trajectory representations that encode useful information for recognizing individual events as well as overall group dynamics in the context of team sports. For example this learnt features can be used for recognizing that a hockey player trying to pass the puck while an opposing player moves to block the pass.

representations can encode useful individual and collective movement patterns for analyzing group actions in the video.

We develop our deep learning approach in the context of team sports. Team sports are a useful testbed for complex group activity analysis as they provide: (1) a well-defined set of events, such as pass, dump-in, and shot; (2) well-defined groups (teams) with distinct dynamics due to differences in strategies and tactics; (3) detailed annotations from sports analytics professionals. At the same time, sports video analysis presents numerous challenges. Players move quickly and often frames are blurred due to this rapid movement. Thus, the input video clips do not always carry the rich visual information as expected. Moreover, sports video, especially team sports, contains many player interactions. Interpreting those interactions can help understand their activities as a group, but the representations used to decode such interactions remains an open challenge.

A body of literature focuses on group activity and human interaction [1, 2, 4, 5, 11, 14, 17, 18, 22], some of which incorporate spatial information of individuals. However, these representations tend to be hand-crafted and do not sufficiently encode the rich information of individual movements and their interactions over time.

We propose a method to analyze team activities by using player locations in world coordinates over time. Analysis of trajectory data is not a new area of research. Many works have proposed computational methods to analyze player behaviors using trajectories. We refer readers to [9] for a recent survey on this field of research. However, to the best of our knowledge, we are the first to learn features describing trajectories with convolutional neural networks and the first to combine trajectory representations of multiple interacting actors in a single framework for group activity analysis.

Our main contribution is a novel deep learning approach for group behavior analysis, which encodes spatio-temporal dependencies from person trajectories. To demonstrate the generality of our model, we analyze events and team formations using NHL (hockey) and NBA (basketball) datasets. We show that our learned representations capture useful information for recognizing events in team sports video. Our learned representations also capture distinctive group dynamics. We show that the collective trajectory patterns of a group can provide strong cues for identifying the group.

1.1 Contribution

This work has been done in collaboration with Yatao Zhong, Dr. Frederick Tung, Dr. Luke Bornn, and Dr. Greg Mori. In addition to this project, I also collaborated on a video segmentation project [15] which was accepted to IAPR International Conference on Pattern Recognition (ICPR), 2016.

Chapter 2

Related Work

The existing literature on analyzing human activities is extensive. Thorough surveys of earlier work include Gavrilu [8] and Weinland et al. [33]. Below, we review closely related work in activity recognition, including individual actions, group multi-person activities, and trajectory analysis.

2.1 Deep Learning for Action Recognition

Recently, deep learning has been brought to bear on the problem of action recognition. Approaches for video-based action recognition include the two-stream network of Simonyan and Zisserman [26], which fuse motion and appearance feature branches into a single network. It has two recognition streams. One ConvNet for capturing spatial information and one other ConvNet for capturing temporal feature. The obtained information of these two streams is mixed by late fusion. Feichtenhofer et al. [7] investigate different methods of fusing motion and appearance information. Karpathy et al. [13] did extensive experiments on when and how to fuse information extracted from video frames to extend convolutional neural networks to video classification. three time information fusion methods has been investigated in this paper slow fusion, early fusion, and late fusion. Veeriah, et al. [30] proposed a differentiable recurrent neural network which highlights the change caused by the salient motions between the successive frames. Donahue et al. [6] extract features from each frame and encode temporal information using a recurrent neural net (LSTM [10]) for action recognition. Their proposed method is deep in both spatial and temporal dimension. Tran et al. [29] extended traditional 2D convolution to the 3D case, where filters are applied to the spatial dimensions and temporal dimension simultaneously. Sun et al. [28] introduce a factorized spatio-temporal convolutional network which factorize the 3D convolution kernel as a sequence of 2D spatial kernels followed by 1D temporal kernels. These approaches have surpassed the performance of the best hand-crafted features, based on dense trajectories and motion features [31].

2.2 Group Activity Recognition

Group activity recognition examines classifying the behaviour of multiple, interacting people. Effective models typically consider both individual actions and person-level interactions within the group. Previous works use hand-crafted features and model interactions with graphical models. Choi et al. [3] build hand-crafted descriptors of relative human poses. Their proposed spatio-temporal local descriptor(STL) is a histogram which takes into account the number of people and their poses. Lan et al. [18] and Amer et al. [1] utilize hierarchical models to understand collective activity among a group of people at different levels, ranging from atomic individual action to group activity in the scene. The concept of social roles performed by people during interactions has also been studied [17, 22]. All of these methods use hand-crafted representations of inter-person relationships.

Another line of work introduces structures into deep learning frameworks by integrating neural networks and graphical models in a unified framework [24, 25, 34]. For example, Deng et al. [5, 4] apply deep structured models to collective activity recognition, learning dependencies between the actions of people in a scene. However, these works do not consider spatio-temporal relationships between participants, which we believe would provide strong indication about how a group activity is formulated. Thus, we propose a model to incorporate spatial information by learning the dynamics of trajectories of each participant as well as their relative movements.

The incorporation of track-level features as extra cues for interaction modeling was done by Choi et al. [2] and Khamis et al. [14]. Recent work has developed more sophisticated deep temporal models for activity analysis. Ramanathan et al. [21] utilize attention models to focus on key players in sports activities. Ibrahim et al. [11] build hierarchical LSTMs to model multiple interacting people over time. In contrast with this, our work learns trajectory features directly from human position inputs.

2.3 Trajectory Data Analytics

There exists significant literature on trajectory analysis focusing on team sports, such as basketball, soccer, and hockey. Applications within sports analytics include analyzing player and team performance, and mining underlying patterns that lead to certain results. Work in this field has included various statistical models to capture the spatio-temporal dynamics in player trajectories. We refer readers to a recent survey [9] on detailed team sports analysis with trajectory data.

Classic examples in the vision literature include Intille and Bobick [12] who analyzed American football plays based on trajectory inputs. Médioni et al. [20] utilized relative positioning between key elements in a scene, such as vehicles and checkpoints, to recognize activities.

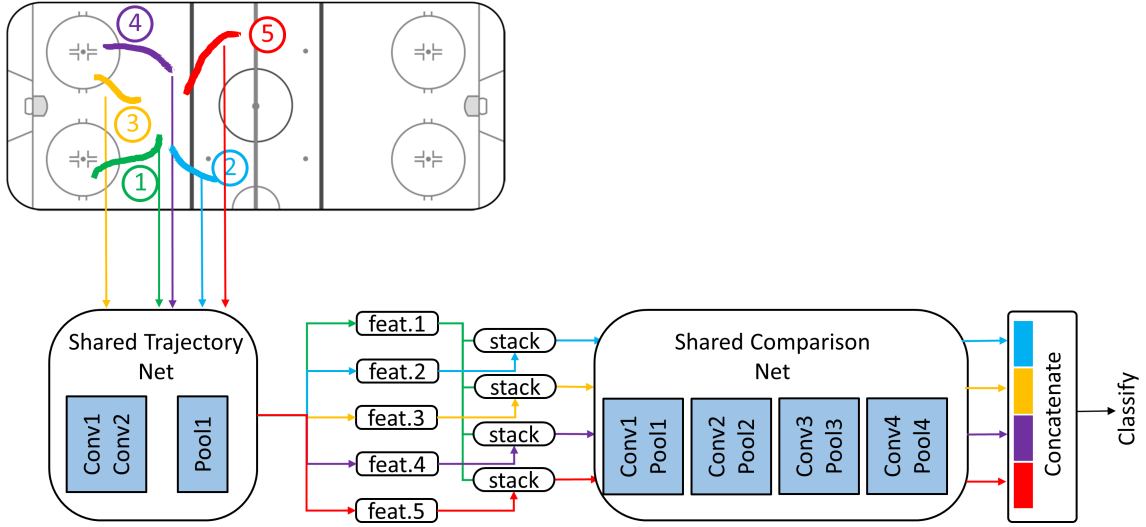


Figure 2.1: Our shared-compare trajectory network (Sec.3.1) learns generic individual representations as well as features describing pairwise interactions. Pairs are determined by an ordering based on proximity to a key player. The shared-compare structure is useful for recognizing events such as passes and shots.

In the sports context, Lucey et al. [19] use a basis representation of person trajectories that utilizes roles. The work that is most similar to ours is Wang and Zemel [32], which uses a traditional CNN plus RNN to classify NBA offensive patterns. Person trajectories are converted to an image representation for input to the CNN. However, an image representation of trajectories is not an effective coding mechanism in that time is ignored, and further most pixels are zeros, making this representation redundant. We instead propose to build a neural network directly on top of raw trajectories.

Chapter 3

Proposed Approach

We aim to learn trajectory representations that encode rich spatio-temporal patterns for recognizing individual events as well as characteristic group dynamics. Our proposed trajectory networks use layers of 1D temporal convolutions over person location inputs in world coordinates. We present two instantiations of this idea.

1. The *shared* trajectory network takes a single person’s trajectory as input. Since the weights of the network are shared across all people, the network learns a generic representation for any individual. The resulting individual representations are passed through a *compare* network, which takes pairs of people as input, with pairs determined by an ordering based on proximity to a “key” player. We show that the shared-compare network structure is useful for recognizing events such as passes and shots.
2. The *stacked* trajectory network takes as input the stacked trajectories of all group members. This network learns representations that take into account all interactions among group members, and is effective at capturing overall group dynamics. We show that the stacked network structure can accurately predict group (team) identities by their member trajectories alone.

3.1 Shared-Compare Trajectory Network

The shared-compare network structure is illustrated in Fig. 2.1. The input to the shared network is a sequence of person world coordinates in the form of (x_t, y_t) , where t is the frame number, associated with a single person. These inputs are obtained via state of the art tracking and camera calibration systems, which provide reasonably accurate, though sometimes noisy, data. To learn the space-time variations in person trajectories, we propose to use layers of 1D convolutions. The network is “shared” in that the same 1D convolutions are applied to each person.

Recall that a person trajectory is essentially a continuous signal. We propose a direct way of interpreting a trajectory. A 2D trajectory in world coordinates (e.g. player position in court or rink coordinates) has two separate continuous signals, one for the x series and one for the y series. We can split the input $[(x_1, y_1), (x_2, y_2), \dots, (x_T, y_T)]$ into two sequences $[x_1, x_2, \dots, x_T]$ and $[y_1, y_2, \dots, y_T]$, each being a 1D continuous signal. In our approach we treat these two sequences as two channels. We build a convolutional neural network on top of these inputs, with 1D convolution operating on each input. By stacking layers of 1D convolution, we can learn combinations of x and y movements that are indicative of particular action classes.

In detail, let $X \in \mathbb{R}^{N \times T}$ denote the input, $F \in \mathbb{R}^{N \times W \times M}$ denote the filters in a convolutional layer and $O \in \mathbb{R}^{M \times T}$ denote the output, where N is the number of input channels, T is the length of input sequence, W is the filter size and M is the number of filters. To model the behaviour of a convolutional layer¹, we perform the basic operation as follows:

$$O_{k,t} = \sigma\left(\sum_{i=1}^N \sum_{j=1}^W X_{i,t+j-1} F_{i,j,k}\right). \quad (3.1)$$

In the above formula, $\sigma(\cdot)$ can be any activation function. In our case, we choose ReLU for all activations. Two convolutional layers are followed by a max pooling layer to make the model shift-invariant and help reduce the dimension of the output.

Let $Z \in \mathbb{R}^{M \times \lceil \frac{T}{S} \rceil}$ be the output of max pooling, where S is the step size in the pooling operation, then we have

$$Z_{k,t} = \max_{1 \leq j \leq S} O_{k, (t-1) \cdot S + j}. \quad (3.2)$$

To build a network with stacked convolutional and max pooling layers, we use the output Z^{l-1} at layer $l-1$ as the input X^l at layer l :

$$X^l = Z^{l-1}. \quad (3.3)$$

We repeat the process described in Eq. 3.1 and Eq. 3.2 for a number of layers. To obtain the output of the shared trajectory network, we flatten the output of the last layer.

The outputs of the shared network are then paired and passed to the compare network. The compare network aims to learn representations for the relative motion patterns of pairs of people. The comparison of trajectory pairs allows the network to learn interaction cues that can be useful for recognizing events or actions (e.g. pass, shot).

Pairs are formed relative to a "key" person according to a pre-defined rule (we discuss the case of an unknown key person later). Denote the key person as $P^{(1)}$ and the other people as $P^{(2)}, P^{(3)}, \dots, P^{(N_p)}$. N_p is the number of players on the rink. Given a pair

¹We choose a step size of 1 when doing convolution and pad zeros to the input if the domain of a filter goes outside of the input.

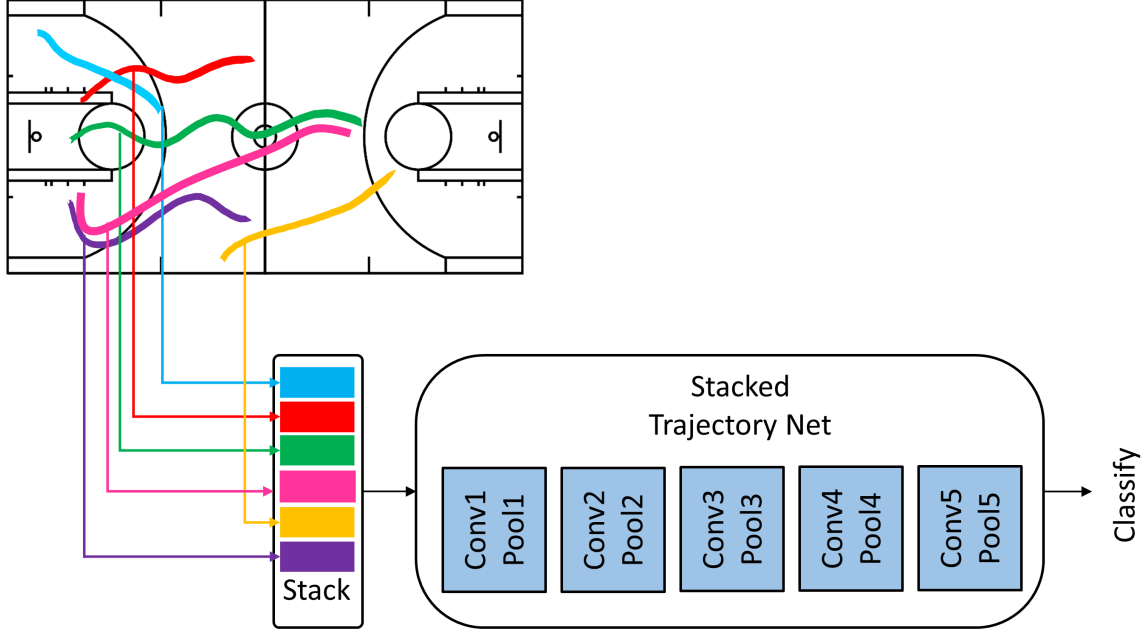


Figure 3.1: Our stacked trajectory network (Section 3.2) takes trajectories of all group members as input and learns representations that consider all interactions within the group. The stacked network captures overall group dynamics and can be used to identify groups (teams) using member trajectories alone.

$(P^{(1)}, P^{(i)}), 2 \leq i \leq N_p$, the individual trajectory representations of $P^{(1)}$ and $P^{(i)}$ (computed using the shared network) are input to a compare network consisting of several convolutional and max pooling layers. The output feature of the compare network can then be used as a learned feature representation of the trajectory pair. We apply the compare network to all pairs $(P^{(1)}, P^{(i)}), 2 \leq i \leq N_p$ relative to the key person, and concatenate the output features $Z^{(i)}$ to obtain the final feature descriptor of the group of people, $[Z^{(1)}, Z^{(2)}, \dots, Z^{(N_p)}]$. Next, a fully connected layer and softmax can be applied depending on the activity analysis task:

$$Z_e = \sigma([Z^{(1)}, Z^{(2)}, \dots, Z^{(N_p)}]W_e) \quad (3.4)$$

where $\sigma(\cdot)$ denotes softmax normalization, F_e is the output feature size, $W_e \in \mathbb{R}^{(F \cdot N_p) \times F_e}$ are the weights of the fully connected layer, and $Z_e \in \mathbb{R}^{F_e}$ is the prediction vector which can be fed into a loss function for end-to-end training

Note that when we perform the concatenation, we implicitly enforce an ordering among the group of people. Arbitrarily enforcing such an order is problematic. We set a consistent ordering by spatial proximity to the key person: $P^{(1)}$ is the key person, $P^{(2)}$ is the person closest to the key person, $P^{(3)}$ is the next closest person to the key person, and so on.

The key person is often available in professional annotations for sports analytics, and typically corresponds to the player with possession of the ball or puck. However, the key

person in a general non-sports setting may not be as well-defined. In addition, the key person may be unknown, or difficult to determine automatically, in general broadcast sports videos. When the key person is not provided, we adopt an average pooling strategy to determine Z_e . For each person in the scene, we fix that person as $P^{(1)}$, rank the other people accordingly based on proximity, and compute the network output. This produces N^p prediction vectors $\{Z_e^{(1)}, Z_e^{(2)}, \dots, Z_e^{(N^p)}\}$. We then apply element-wise average pooling over the N^p prediction vectors to obtain the final prediction vector Z^e . Results for both known and unknown key player settings are presented in the experiments.

3.2 Stacked Trajectory Network

The stacked trajectory network structure is illustrated in Fig. 3.1. The input to this network is a stack of (x_t, y_t) sequences for all members of the group (e.g. all members of the same team) and the ball. For example, suppose a group consists of N_p people. Then the sequences $[(x_1^{(i)}, y_1^{(i)}), (x_2^{(i)}, y_2^{(i)}), \dots, (x_T^{(i)}, y_T^{(i)})]$ for each person i and ball are stacked to form a $(C \cdot (N_p + 1)) \times T$ dimensional input ($C = 2$ for two channels x and y). However, after stacking, its actual number of channels is $C \cdot (N_p + 1)$. Similar to Section 3.1, a convolutional neural network consisting of layers of 1D convolutions is trained on this input. We set a consistent ordering by spatial proximity to the ball. The output of the final layer is flattened and input to a fully connected layer and softmax for task-specific end-to-end training.

By taking into account all combinations of interactions within the group, the stacked trajectory network learns overall group dynamics instead of individual representations.

Chapter 4

Datasets

We conduct experiments on two datasets. The first includes player trajectories and visuals from broadcast video footage of NHL hockey games. The second consists of player trajectories extracted from an external tracking system recording player positions in NBA basketball games.

4.1 The SPORTLOGiQ NHL Dataset

The SPORTLOGiQ NHL dataset includes both video and trajectory data. Unlike the NBA dataset where person trajectories are obtained from a multi-camera system, the player positions in the NHL dataset are estimated using a homography, which maps a pixel in image coordinates to a point in world coordinates. State of the art algorithms are used to automatically detect and track players in raw broadcast videos. If we have the bottom-middle point of a player bounding box, we can map this point to world coordinates with a homography matrix, hence acquiring the player position. The NHL dataset has detailed event annotation for each frame, each event being categorized into a super class and a fine-grained class. In our experiment, we use 8 games with 6 super classes: pass, dump out, dump in, shot, carry and puck protection. Fig. 4.1 shows the fraction of each event in the 8-game dataset. We will describe how we handle the event imbalance in Sec. 5.1.

4.1.1 Data Preprocessing

In a hockey game, typically there are 4 on-ice officials and 12 players (6 on each team). Thus, there can be at most 16 persons on the rink at the same time. In the following we do not make any distinction between officials and players and we use "player" to refer to all people on the rink. Because the dataset is created from NHL broadcast videos where not all players are visible in each frame, we need to set a threshold N_p so that our model can handle a fixed number of players. If the number of players available in a frame is less than N_p , we pad with zeros the part where players are unavailable. Each training sample consists

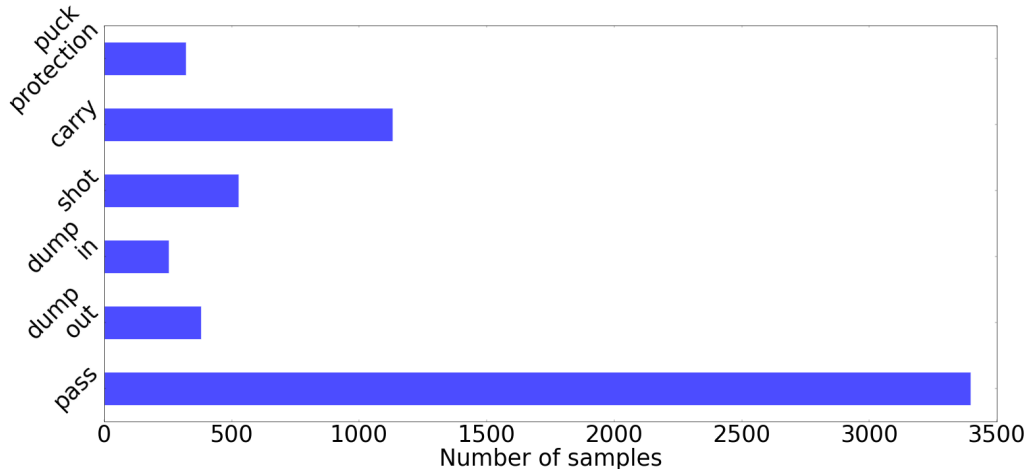


Figure 4.1: Number of samples per event in the SPORTLOGiQ NHL dataset.

of data from N_p players. The data of each player includes a T -frame video clip (cropped from raw video using bounding boxes) and the corresponding T -frame trajectory estimated from this video clip. Note that our model supports variable-length input. If in some frames a player is not available, we set the data in these frames to zeros. In our experiments, N_p is set to 5 and video frame size is set to 96×96 . We set T to 16 by first locating the center frame where an event happens and then cropping 7 frames before the center frame plus 8 frames after it. If the center frame of a certain event happens to be close to that of another event within 15 frames, we drop this sample.

4.2 The STATS SportVU NBA Dataset

The STATS SportVU NBA dataset consists of real-time positions of players and the ball in 2D world coordinates captured by a six-camera system at a frame rate of 25 Hz. Each frame has complete annotations of the events happening in the frame, such as dribble, possession, shot, pass and rebound. The dataset we use has 1076 games during the 2013–2014 NBA season with approximately 10^6 frames in each game. We will use this dataset for evaluating the ability of our model to capture characteristic group dynamics. In particular, we attempt to predict the identity of a team using only the trajectories of its players during a game.

4.2.1 Data Preprocessing

We extract 137176 possessions from the 1076 games for experiments. Each possession starts with an offensive team having possession of the ball and ends with a shot. We fix possession length to 200 frames. If a possession is longer than 200 frames, we crop it starting from the last frame and count the number of frames backward until it reaches 200. If a possession is shorter than 200 frames, we pad zeros to it. Originally there are 25 frames per second, but

we sample only half of the frames in a second, so the sampled 200 frames actually represent a 16 second long sequence. There are in total 30 NBA teams. Fig. 4.2 shows the number of possessions we extracted from each team in the dataset. We can see that this is a relatively balanced dataset, each team having a similar number of samples for experiments.

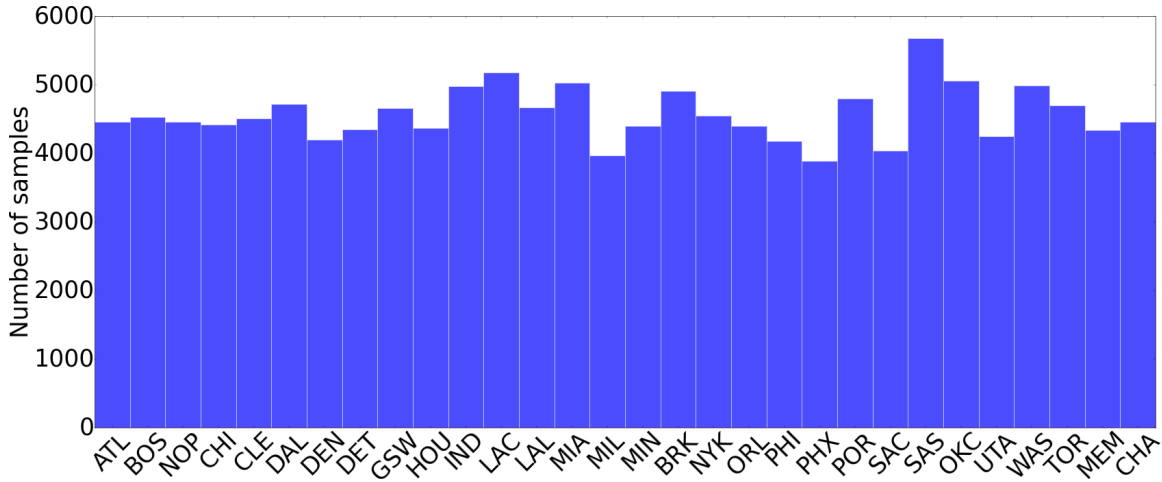


Figure 4.2: Number of possessions of each team in the STATS SportVU NBA dataset.

Chapter 5

Experiments

We conduct our experiments on the NHL and NBA datasets described above. To demonstrate that our shared-compare network is capable of learning temporal dynamics of person trajectories and group interactions, we perform event recognition using the NHL dataset. We then show that the stacked trajectory network captures overall group dynamics by performing team identity prediction on the NBA dataset, using only player trajectories.

5.1 Event Recognition on the NHL Dataset

5.1.1 Task Description

The goal is to predict the event label given the trajectories of players on the rink. The events used are pass, dump out, dump in, shot, carry and puck protection. The number of samples of each event in the dataset is shown in Fig. 4.1. This dataset is highly unbalanced with the pass event taking up half of the dataset. To resolve this problem, we minimize a weighted cross-entropy loss function during training. The weighting for each class is in inverse proportion to its frequency in the dataset.

5.1.2 Experiment Settings

We use the exact model shown in Fig. 2.1 for the shared-compare model. The 1D convolutions of the shared trajectory network have 64 and 128 filters with a filter size of 3. The first two layers of compare network have 128 filters with filter size 3, and the last two have 256, 512 filters with filter size 3,2 respectively. All max-pooling layers are performed with a window size of 2 with stride 2. The weights in the loss function for pass, dump out, dump in, shot, carry and puck protection are 0.07, 0.6, 1, 0.4, 0.2 and 0.7 respectively. We use average precision as the evaluation metric.

5.1.3 Baselines

We compare the proposed model with (1) a hand-crafted trajectory feature, IDT [31]; (2) deep learning using video features, C3D [29]. Comparisons with these baselines demonstrate that our learned trajectory features are better than previous trajectory descriptors; and our approach, which *only uses features from person trajectories*, can outperform methods that learn on top of video input.

- **IDT** [31]: We use the same input data as in our method: sequences of x, y coordinates of 5 players. We represent each trajectory using the IDT *Trajectory shape descriptor* which is the normalized displacement vector of the trajectory. An SVM with RBF kernel and 'one vs. rest' mechanism is deployed for multi-class classification.
- **C3D** [29]: We use a C3D model as the visual baseline, either trained from scratch using only the NHL dataset, or fine-tuned from a model pre-trained on Sports-1M. We obtain a short clip for each player by concatenating player bounding boxes across time. C3D takes each player clip as input. The weights of the network are shared across all players. Then the extracted features of players are concatenated using the same ordering as in our approach and passed to a softmax layer for classification.

5.1.4 Train phase

We train all three networks by providing the key player, defined as the player who performs the action. The remaining players are ranked by proximity to the key player as described in Sec. 3.1.

5.1.5 Test phase

For test time evaluation, we consider both the case where the key player is given, and the more general case where the key player is unknown. When the key player is not given, we adopt an average pooling strategy as described in Sec. 3.1. In our experiments, we use 4 games for training, 2 games for validation and 2 games for testing.

5.1.6 Experiment Results

Tables 5.1 and 5.2 show experimental results when the key player is unknown and known, respectively. When the key player is unknown, our learned trajectory representations achieve 13.2% higher mean average precision compared to state-of-the-art hand-crafted trajectory features (IDT). Compared to deep learning using video features, our shared-compare model obtains 8.7% higher mAP when the C3D model is trained from scratch. Compared to the C3D baseline pre-trained on Sports-1M and fine-tuned on the NHL dataset, our shared-compare model obtains 1.7% higher mAP while not requiring auxiliary training data (note

	IDT	C3D	Fine-tuned C3D	Shared-Cmp
pass	72.86%	71.10%	77.45%	78.13%
dump out	13.75%	11.66%	18.15%	22.14%
dump in	6.35%	7.58%	19.04%	26.63%
shot	13.05%	23.37%	38.96%	40.52%
carry	45.66%	64.75%	65.65%	61.10%
puck protection	6.28%	6.50%	7.98%	8.72%
mAP	26.32%	30.83%	37.87%	39.54%

Table 5.1: Average precision for event recognition on the NHL dataset, unknown key player.

	IDT	C3D	Fine-tuned C3D	Shared-Cmp
pass	73.35%	77.30%	84.34%	81.33%
dump out	14.34%	10.17%	17.10%	23.11%
dump in	5.77%	10.25%	24.83%	50.04%
shot	13.07%	34.17%	58.88%	48.51%
carry	47.38%	86.37%	90.10%	85.96%
puck protection	7.28%	11.83%	13.99%	11.54%
mAP	26.86%	38.35%	48.21%	50.08%

Table 5.2: Average precision for event recognition on the NHL dataset, known key player.

that Sports-1M is a large sports dataset). Fig. 5.1 shows precision-recall curves for each event. Our learned trajectory features consistently outperform the hand-crafted IDT trajectory features, and obtain mixed performance compared to the fine-tuned C3D model pre-trained on Sports-1M. When the key player is known, performance increases for all models; the shared-compare model improves by 10.5% mAP. This result suggests that a useful direction for future work is the automatic prediction of the key person in the scene.

For the case where key player is provided, we visualize the top 5 candidates retrieved as shot and dump in in Fig. 5.2 and 5.3. The retrieved events look similar visually, highlighting the benefit of using the spatio-temporal relationships of the players to recognize events.

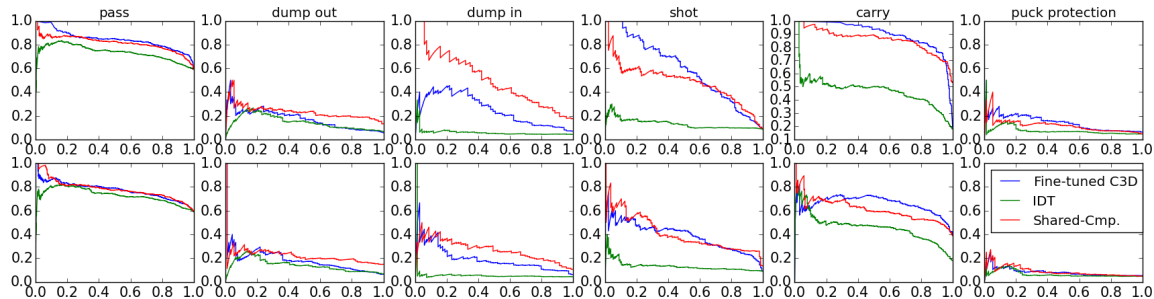


Figure 5.1: Precision-recall curves for each event in the NHL dataset. The first row shows the results of the known key player case while the second row the unknown key player case.

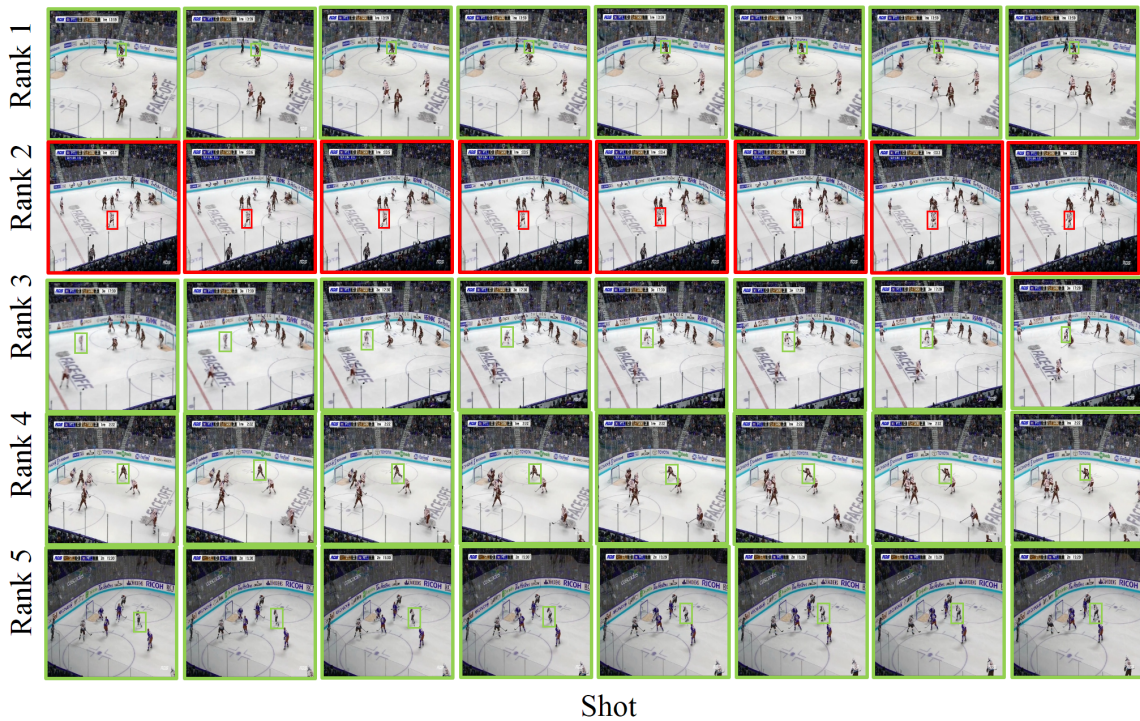


Figure 5.2: Top 5 candidates retrieved as shot. Green sequences are true positives while red ones are false positives. The person with a bounding box is the “key” player who is performing the action. We show 8 frames of the 16-frame video clip by sub-sampling. If a frame is black, it means the key player is missing because of failure to detect and track the player.



Dump in

Figure 5.3: Top 5 candidates retrieved as dump in. Green sequences are true positives while red ones are false positives. The person with a bounding box is the “key” player who is performing the action. We show 8 frames of the 16-frame video clip by sub-sampling. If a frame is black, it means the key player is missing because of failure to detect and track the player.

5.2 Team Identification on the NBA Dataset

5.2.1 Task Description

The goal is to predict the team identity given only its member trajectories. Good performance on this task demonstrates the capacity of our model to capture characteristic group dynamics.

5.2.2 Experiment Settings

We use the exact model shown in Fig. 3.1 and described in Sec. 3.2 for the stacked trajectory network. The x and y coordinates of the ball and 5 players are stacked together, resulting in a 200×12 matrix as input, where 200 is the length of the input sequence and 12 is the number of channels. We use 60% of the 1076 games for training, 20% for validation and 20% for testing.

5.2.3 Baseline

We compare our model with IDT [31]. The same input data is used as in our method: sequences of x, y coordinates of the ball and 5 players. Each trajectory is represented by the IDT *trajectory shape descriptor* which is the normalized displacement vector of the trajectory. Then we concatenate the IDT features of the ball and the players and apply an SVM with RBF kernel and 'one vs. rest' mechanism for classification.

5.2.4 Measurement

We measure the performance of our model according to accuracy and hit-at- k accuracy¹ metrics, both of which are calculated over possessions. However, a single possession can hardly capture the distinctive group dynamics a team might possess. We therefore also aggregate all of a team's possessions in a game and predict the team identity using majority voting.

5.2.5 Results and Analysis

Using the network architecture shown in Fig. 3.1, we are able to achieve 95% accuracy in predicting a team's identity by majority voting over all its possessions in the game, given trajectory information alone. We next explore the architecture of our model by varying the number of convolutional layers, the filter size and the number of filters in each layer. Tables 5.3, 5.4 and 5.5 show the results respectively. From Tables 5.3 and 5.5, we see that by increasing the number of layers and filters, generally we can obtain a more complex model

¹Hit-at- k accuracy: if any one of the top- k predictions equals the ground truth label, we consider it as being correctly classified.

layers	acc	hit@2	hit@3	game acc
2conv	10.68%	18.09%	24.31%	50.00%
3conv	18.86%	28.89%	36.47%	87.05%
4conv	22.34%	33.03%	40.47%	93.41%
5conv	24.78%	35.61%	42.95%	95.91%
5conv+2fc	25.08%	35.83%	42.85%	94.32%

Table 5.3: Metrics on models with different number of layers. All convolutional layers use a filter size of 3 except the first layer, where the filter size is 5. The number of filters in next layer is double the number in previous layer except the fifth layer (if any), where the number of filters is the same as that in the fourth layer. The number of neurons in fully connected layer is set to 1024.

filter sizes	acc	hit@2	hit@3	game acc
3 3 3 2 2	24.24%	35.36%	43.25%	94.10%
5 3 3 3 3	24.78%	35.61%	42.95%	95.91%
7 5 5 3 3	23.12%	33.48%	41.04%	95.45%
9 7 7 5 5	14.13%	23.15%	30.01%	62.05%

Table 5.4: Metrics on models with different filter sizes. All models in the table use five convolutional layers with no fully connected layer. The filter sizes listed is in a bottom-up order and the number of filters used are 64, 128, 256, 512, 512 (bottom-up order).

to achieve better performance. However, as we increase the number of parameters in the model, we reach a limit that prevents us from acquiring further improvement by increasing the model complexity. For example, by adding two fully connected layers after the 5conv model in Table 5.3, we obtain only a slight increase in possession-based accuracy and a drop in game-based accuracy. Also note that in Table 5.4, using small filter sizes generally leads to good results (see the first three models in Table 5.4). If we increase the filter size, we have a large decrease in model performance (see the last model in Table 5.4). Table 5.6 compares our model with the IDT baseline and shows that our model learns distinctive group dynamics better than hand-crafted trajectory features.

Fig. 5.4b shows the confusion matrix created from the 5conv model in Table 5.3. For most teams, our model can correctly predict the team identity when aggregating over all of its possessions in a game. The worst case is the Phoenix Suns (PHX in Fig. 5.4b); the model has only a probability around 65% to classify the Suns correctly, but this is still much better than chance performance.

To see what kind of patterns the model learns over the time dimension, we visualize a small fraction of the filters in the first convolutional layer. In Fig. 5.4a, we show 64 filters learned from the input sequence of x coordinates of the ball. Some of them appear to be “Z” or “S” shaped and some appear to be “M” or “W” shaped. Some of them are similar, so

base # filters	acc	hit@2	hit@3	game acc
16	20.37%	30.71%	38.21%	81.14%
32	23.73%	34.55%	41.85%	92.95%
64	24.78%	35.61%	42.95%	95.91%
128	21.81%	32.10%	39.24%	94.45%

Table 5.5: Metrics on models with different number of filters. All models in the table use five convolutional layers with no fully connected layer. The base number of filters listed in the table is the number of filters in the first layer. The number of filters in next layer is double the number in previous layer except that the fourth and the fifth layers have the same number of filters.

models	acc	game acc
IDT	5.74%	9.10%
Stacked Traj. Net	25.78%	95.91%

Table 5.6: Stacked trajectory network vs. IDT. Comparison with this baseline show that our model captures valuable collective features for understanding group dynamics

there could be redundancy in these filters. These temporal patterns are the building blocks that form discriminative representations to distinguish teams.

Chapter 6

Conclusion and Future Work

We have presented deep neural network models for learning person trajectory representations for group activity analysis. Team sports provide a useful testbed for understanding activities involving groups of people interacting to achieve shared goals. We evaluated our models on two datasets for two group activity analysis tasks: event recognition on NHL hockey games dataset and team classification on NBA basketball games dataset. Experiments show that our learned trajectory representations encode complex spatial-temporal dependencies for recognizing individual events and also capture distinctive group dynamics that can be used to identify teams given only member trajectories.

6.1 Limitations and Future work

One of our limitation for activity recognition on NHL dataset is that we limit our framework to know when events happen. As we mentioned in Sec. 4.1, the input samples to our model are 16-frames player trajectory and center frame is when the event happens. In general, we usually don't have the annotation for the time that events happen. Hence, one of our future direction will be the temporal localization of events in video. Then, we can chunk the interval around the time that event is happening and feed it to the network.

In NHL event recognition we classify events into 6 classes of 'pass', 'puck protection', 'dump in', 'dump out' and shot. These 6 classes don't occur at the same time by two players. So, what we are doing is classifying single person action in a group. One future work can be covering events that happen simultaneously and adding one more class of 'others'. Thus, we don't need to focus on predicting single person activity and we can have a label for each player activity in the scene.

In our proposed model for event recognition, we evaluated our model by considering both the case where the key player is given and the more general case which key player is not provided. We observed that when the key player is known, the performance increased compared to the case of the unknown key player. In our case, NHL hockey dataset provides

the annotation for the key player. However, the key person in a general non-sport setting may not be as well-defined. This suggests that a direction for future work is the automatic detection of the key person in the scene.

Another future work can be the use of team membership information. In our experiments on NHL dataset, we didn't make any distinction between teammates, opponents, and referees and We used 'player' to refer to all people on the rink. I believe that use of this information can be helpful. For example, comparing the trajectory of key player with a teammate/opponent trajectory provides more information than the case of comparing key player trajectory with referee trajectory.

In our shared-compare model, the comparison network tries to find the relationship of key player with others. In some cases, there might be some other interesting relationships like triplet relations, the relationship of other players (excluding key player) and etc. Thus, one future work can be considering of these kind of relationships in our model.

6.2 Discussion

In addition to team sports analysis, we believe that our trajectory representation can be useful in many human activity analysis tasks including surveillance and security, health care monitoring, smart homes and etc. For example, in surveillance and security, these trajectory features can be helpful in building a system for monitoring activities of a crowd of people including determining somebody is running, some people are meeting each other, a crowd of people is fleeing toward a location and detect any suspicious or unusual activity that is happening. In general, our proposed models are applicable to any activity recognition task that cares about movements, tracking key points and the interaction and relationship of those key points.

Bibliography

- [1] Mohamed Rabie Amer, Peng Lei, and Sinisa Todorovic. Hirf: Hierarchical random field for collective activity recognition in videos. In *European Conference on Computer Vision (ECCV)*, pages 572–585. Springer, 2014.
- [2] Wongun Choi and Silvio Savarese. A unified framework for multi-target tracking and collective activity recognition. In *European Conference on Computer Vision (ECCV)*, pages 215–230. Springer, 2012.
- [3] Wongun Choi, Khuram Shahid, and Silvio Savarese. What are they doing?: Collective activity classification using spatio-temporal relationship among people. In *Computer Vision Workshops (ICCV Workshops)*, pages 1282–1289. IEEE, 2009.
- [4] Z. Deng, M. Zhai, L. Chen, Y. Liu, S. Muralidharan, M. Roshtkhari, , and G. Mori. Deep structured models for group activity recognition. In *British Machine Vision Conference (BMVC)*, 2015.
- [5] Zhiwei Deng, Arash Vahdat, Hexiang Hu, and Greg Mori. Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [6] Jeffrey 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 *Computer Vision and Pattern Recognition (CVPR)*, pages 2625–2634, 2015.
- [7] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [8] D. M. Gavrilu. The visual analysis of human movement: A survey. *Computer Vision and Image Understanding (CVIU)*, 73(1):82–98, 1999.
- [9] Joachim Gudmundsson and Michael Horton. Spatio-temporal analysis of team sports—a survey. *arXiv preprint arXiv:1602.06994*, 2016.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [11] Moustafa 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.

- [12] Stephen S. Intille and Aaron Bobick. Recognizing planned, multiperson action. *Computer Vision and Image Understanding (CVIU)*, 81:414–445, 2001.
- [13] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1725–1732, 2014.
- [14] Sameh Khamis, Vlad I Morariu, and Larry S Davis. Combining per-frame and per-track cues for multi-person action recognition. In *European Conference on Computer Vision (ECCV)*, pages 116–129. Springer, 2012.
- [15] Mehran Khodabandeh, Srikanth Muralidharan, Arash Vahdat, Nazanin Mehrasa, Eduardo M. Pereira, Shin-ichi Satoh, and Greg Mori. Unsupervised learning of supervoxel embeddings for video segmentation. In *IAPR International Conference on Pattern Recognition (ICPR)*, 2016.
- [16] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. In *The International Conference on Computer Vision (ICCV)*, 2011.
- [17] Tian Lan, Leonid Sigal, and Greg Mori. Social roles in hierarchical models for human activity recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [18] Tian Lan, Yang Wang, Weilong Yang, Stephen Robinovitch, and Greg Mori. Discriminative latent models for recognizing contextual group activities. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 34(8):1549–1562, 2012.
- [19] Patrick Lucey, Alina Bialkowski, G. Peter K. Carr, Stuart Morgan, Iain Matthews, and Yaser Sheikh. Representing and discovering adversarial team behaviors using player roles. In *Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [20] G. Médioni, I. Cohen, F. Brémond, S. Hongeng, and R. Nevatia. Event detection and analysis from video streams. *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 23(8):873–889, 2001.
- [21] 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 *Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, USA, June 2016.
- [22] Vignesh Ramanathan, Bangpeng Yao, and Li Fei-Fei. Social role discovery in human events. In *Computer Vision and Pattern Recognition (CVPR)*, June 2013.
- [23] Mikel D. Rodriguez, Javed Ahmed, and Mubarak Shah. Action mach: a spatio-temporal maximum average correlation height filter for action recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [24] Stephane Ross, Daniel Munoz, Martial Hebert, and J Andrew Bagnell. Learning message-passing inference machines for structured prediction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2737–2744. IEEE, 2011.
- [25] Alexander G Schwing and Raquel Urtasun. Fully connected deep structured networks. *arXiv preprint arXiv:1503.02351*, 2015.

- [26] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems (NIPS)*, pages 568–576, 2014.
- [27] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [28] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E. Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [29] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *International Conference on Computer Vision (ICCV)*, pages 4489–4497. IEEE, 2015.
- [30] Vivek Veeriah, Naifan Zhuang, and Guo-Jun Qi. Differential recurrent neural networks for action recognition. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [31] Heng Wang, Alexander Kläeser, Cordelia Schmid, and Cheng-Lin Liu. Dense trajectories and motion boundary descriptors for action recognition. Research Report RR-8050, August 2012.
- [32] Kuan-Chieh Wang and Richard Zemel. Classifying nba offensive plays using neural networks. In *Sloan Sports Analytics Conference*, 2017.
- [33] Daniel Weinland, Remi Ronfard, and Edmond Boyer. A survey of vision-based methods for action representation, segmentation and recognition. *Computer Vision and Image Understanding (CVIU)*, 115(2):224–241, 2011.
- [34] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1529–1537, 2015.