

Group Activity Recognition in Basketball Tracking Data - Neural Embeddings in Team Sports (NETS)

Sandro Hauri^{a,*} and Slobodan Vucetic^a

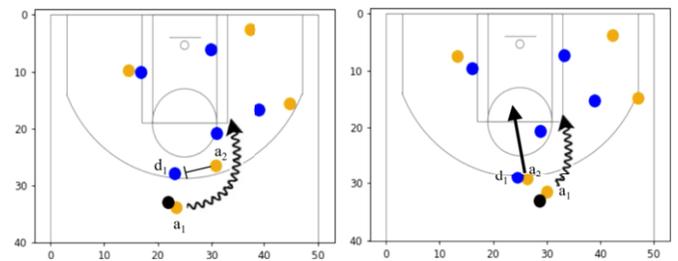
^aTemple University

Abstract. Like many team sports, basketball involves two groups of players who engage in collaborative and adversarial activities to win a game. Players and teams are executing various complex strategies to gain an advantage over their opponents. Defining, identifying, and analyzing different types of activities is an important task in sports analytics, as it can lead to better strategies and decisions by the players and coaching staff. The objective of this paper is to automatically recognize basketball group activities from tracking data representing locations of players and the ball during a game. We propose a novel deep learning approach for group activity recognition (GAR) in team sports called NETS. To efficiently model the player relations in team sports, we combined a Transformer-based architecture with LSTM embedding, and a team-wise pooling layer to recognize the group activity. Training such a neural network generally requires a large amount of annotated data, which incurs high labeling cost. To alleviate this problem, we pretrain the neural network on a self-supervised trajectory prediction task and fine-tune it using a mix of strong and weak labels. We used a large tracking data set from 632 NBA games to evaluate our approach. The results show that NETS is capable of learning group activities with high accuracy, and that self- and weak-supervised training in NETS have a positive impact on GAR accuracy.

1 Introduction

With the recent advances in sensing technology, there is an unprecedented amount of tracking data available for sports analytics, such as Hawk-eye in tennis [25], various tracking systems in soccer (Chyron-Hego, Stats LLC, SciSports), and SPORTLOGiQ in ice hockey. The National Basketball Association (NBA) mandated the installation of a tracking system based on computer vision called SportVU [31] in their sports arenas to collect data about the movements of players and the ball at 25Hz. This data is shared with all NBA teams to ensure equity, with an implicit understanding that the teams with the best ability to benefit from this data will not only improve their own competitiveness, but will advance innovation in the way basketball is played.

Single player actions in basketball, such as shot-taking, can be recognized with rules programmed by an expert. In contrast, group activities involve multiple players that perform a coordinated action dependent on their combined movements over time, which is much more challenging to describe with handwritten rules. As an example, let us consider a commonly observed tactic in basketball called pick-and-roll. For this tactic, two offensive players are collaborating



(a) Start of extracted play sequence. (b) The attacker a_2 blocks the path of the defender. The ball handler a_1 keeps possession.

Figure 1: A classic pick-and-roll. Attackers are displayed in gold, defenders are shown in blue and the ball in black. The roll man a_2 helps the ball handler a_1 by blocking the defender d_1 .

to block the best path of a defender (see Figure 1). To identify this action, we have to understand key concepts of basketball, such as ball possession, defensive assignments, and cooperation of offensive players. To pose a further challenge, there is a closely related tactic called handoff, where two offensive players cross paths and transfer the ball possession with a short pass. Separating the two types of activities can be a challenge, because the setup of a pick-and-roll and a handoff is very similar.

Manually labeling play sequences is time-consuming, because it requires watching all the games with a lot of focus while paying attention to multiple players simultaneously. There are also edge cases that some experts would include and others would not. As a consequence, existing manually labeled datasets of pick-and-rolls only include roughly 1,000 play sequences [24, 21], which is not enough data to train powerful neural network (NN) models. One option to supply more labels is to write rule-based code to recognize actions, yielding weak-labels. As we will elaborate in Section 4.2 and show in the results in Section 6.4, this is not a simple process and even after a lot of effort the quality of handwritten code is inferior to manual labeling. Since rule-based weak-labeling can quickly provide labels for an entire data set, our hypothesis is that those weak-labels can be exploited during training of a deep learning model. Additionally, we propose to use self-supervised pretraining with trajectory predictions. Our hypothesis is that an NN trained to predict trajectories inherently learns about the behavior of players and groups of players. Thus, we expect that an NN pretrained in this way could be fine-tuned for Group Activity Recognition (GAR). In our approach, we combine pretraining on trajectory prediction task, fine-tuning using weak labels and further fine-tuning using manual labels sequentially.

* Corresponding Author. Email: sandro.hauri@temple.edu

We propose the use of a transformer architecture [35] that creates context aware embeddings for each player, which are gathered through a team-wise pooling layer and call our model Neural Embeddings for Team Sports (NETS). Our model uses Long Short-Term Memory (LSTM) [13] to embed the time series to allow the modeling of the temporal aspect of the data. The team-wise pooling layer allows creating play embeddings that are permutation-invariant to the input order of players within a team.

Our main contributions are: 1) A multistage training approach to pretrain on self-supervised data, with subsequent fine-tuning on weak-labels and manual labels to learn complex GAR. 2) A novel Transformer-based NN architecture to create predictions and embeddings that are invariant to player permutation within a team. 3) Extensive experimental evaluation on large basketball tracking data.

2 Related Work

GAR has gained significant interest from computer vision research with practical applications such as social role understanding, security surveillance, and sports analysis [37]. While some researchers study GAR based on tracking information in pedestrian data [44], surveillance video [23], or exercising activity [16], not many works in sports analytics rely only on tracking data [30]. Group activities have been studied in volleyball with the team actions spike, set, pass, and winpoint with a popular benchmark dataset based on short video segments [14]. Various deep learning approaches improved classification accuracy on this dataset, using state-of-the-art image recognition tools to solve this task [39, 38, 7, 40]. GAR was used in soccer analytics to recognize pass, reception, and shot in a given play sequence [28], where the authors use a combination of tracking data and the original video to solve the problem. We believe that using only tracking data in team sports can be advantageous because it reduces the amount of data significantly, making it possible to train on much larger data. Group activities in sports happen in a standardized environment, where the geometry of the court can be learned without the need for the image context.

Human trajectory prediction is an active field of research with many applications, such as navigation of autonomous vehicles, planning of transportation systems, traffic operations, and many more [20]. Modeling the implicit social interactions from trajectory data is a challenging task with a long history of research from social force models [12], locally optimal collision avoidance [33], to the current state-of-the-art of modeling human interactions using deep neural networks [9, 15, 19, 36]. Trajectory prediction on NBA data has found much interest in the computer vision community, with notable examples using Variational AutoEncoders (VAE) to model the players' movements [32, 42, 43]. However, VAEs are notoriously difficult to train [4], and it can be difficult to tune the parameters to generate realistic samples [8]. Another deep learning method to predict NBA trajectories uses an LSTM-base approach [10].

We propose the use of self-supervised learning to help training GAR models. Self-supervised learning defines pretraining tasks that can be learned without manual annotation, but require understanding of the data to be solved [41]. Recent research has shown that self-supervised learning is beneficial in datasets with limited labeled data for many different applications such as natural language processing [6], computer vision [3, 11], robotics [17, 29, 5], or protein modeling [26, 27]. Some research has tried to explain why self-supervised pretraining is particularly beneficial to train NNs. Using the lottery ticket hypothesis for self-supervised training in computer vision [2] suggests that pretraining finds a more compact representation of the complex input data, which is beneficial for downstream tasks.

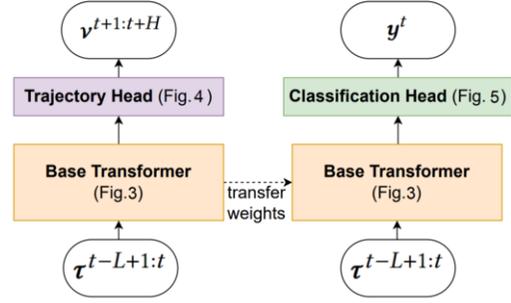


Figure 2: Illustration of the NETS architecture.

3 Methodology

In this chapter, we first provide mathematical definitions necessary to explain the NN architecture and provide a detailed description and motivation for the NN architecture.

3.1 Problem Setting

The goal of this work is to train a NN to recognize group activities during basketball games, given a set of trajectories of the players and the basketball on the court.

Notation

Since 2013, every NBA arena has a camera system to track the players (5 from each team) and the ball during basketball games. The system observes the 11 tracked objects by their locations in an x-y plane, where the x-axis goes across the length of the court and the y-axis goes from sideline to sideline. We pre-process the data, such that the offensive team attacks along the y-axis. We define the location of object o at time step t as $\ell_o^t = [x_o^t, y_o^t]$, with $o \in \mathcal{O} = \{B, A_1, \dots, A_5, D_1, \dots, D_5\}$ where \mathcal{O} denotes the set of tracked objects, namely the ball B , each of the 5 attackers $\{A_1, \dots, A_5\}$, and each of the 5 defenders $\{D_1, \dots, D_5\}$. Using an ordered sequence of L time frames, the trajectory of a tracked object o can be expressed as $\tau_o^{t-L+1:t} = [\ell_o^{t-L+1}, \dots, \ell_o^t]$, with equally spaced time steps at an interval of Δ_t .

Trajectory Prediction Task

We formulate the self-supervised task of predicting future trajectories for every tracked object o , represented by the vector $\tau_o^{t+1:t+H} = [\ell_o^{t+1}, \dots, \ell_o^{t+H}]$, where H is the number of future time steps (or prediction horizon) for which we predict the trajectory.

We denote as $\tau^{t:t'}$ the list of trajectories for all objects, i.e. $\tau^{t:t'} = [\tau_B^{t:t'}, \tau_{A_1}^{t:t'}, \dots, \tau_{D_5}^{t:t'}]$. Through a simple conversion, we can calculate velocity vector $\nu^{t:t'}$ from the trajectory $\tau^{t:t'}$. Although these two vectors are mathematically interchangeable, previous work on NBA data has shown that using velocity vectors as target to deep learning models is advantageous [10]. This leads to the trajectory prediction task, where the objective is to predict $\nu^{t+1:t+H}$ from $\tau^{t-L+1:t}$. This is a self-supervised learning task, because future trajectories can be automatically extracted from the historical trajectory data and stored in dataset $\mathcal{D} = \{(\tau^{t-L+1:t}, \nu^{t+1:t+H}) | t = L, \dots, T-H\}$, where T is the total number of time steps.

Group Activity Recognition (GAR)

The objective of GAR is to predict group activity type $y^t \in 1, \dots, K$ for each play sequence $\tau^{t:t'}$, where K is the number of activity types.

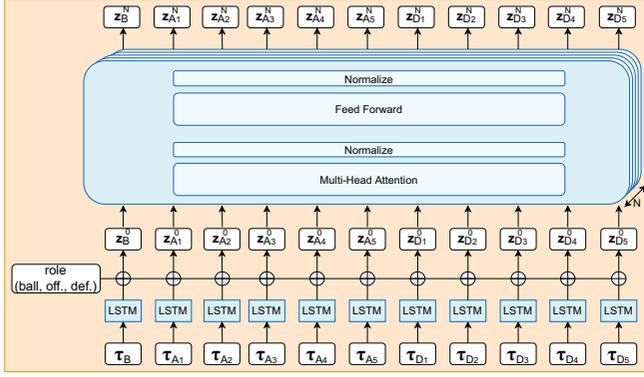


Figure 3: Base Transformer to generate embeddings. Trainable functions are in blue boxes. \oplus stands for concatenation, meaning that one vector is appended to the other.

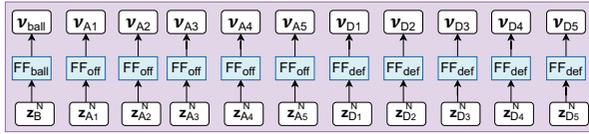


Figure 4: Trajectory head.

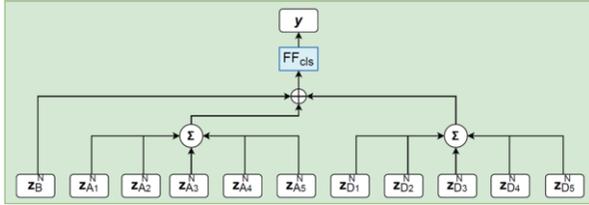


Figure 5: Classification head. Σ stands for summation.

Similar to the trajectory prediction task, we can generate a labeled data set $\mathcal{D} = \{(\tau^{t-L+1:t}, \mathbf{y}^t) | t = L, \dots, T\}$, where each play sequence of L time frames is matched to a label indicating the action type. The main difference compared to the trajectory prediction is that labels need to be provided externally.

3.2 Framework

A standard approach to train an NN for GAR, would be to use supervised learning on ground truth labels. However, as mentioned earlier, manual labeling is time-consuming and expensive. We hypothesize that we can effectively train a powerful model through pretraining on the trajectory prediction task and fine-tune it using a large number of lower-quality weak labels, followed by fine-tuning using a small number of high-quality manual labels. To enable the pretraining and fine-tuning process, we utilize a modular NN architecture, which is described in Figure 2. For the trajectory prediction task, the trajectories $\tau^{t-L+1:t}$ are input to the base model (we specify the chosen model in the next section). The output of the transformer is then used as input to a trajectory prediction head for pretraining. The GAR task uses the same base model, but uses a classification head for activity classification. This modular approach allows the fine-tuning of weights after the pretraining of the trajectory prediction. Our hypothesis is that a majority of the information about the input data will be

modeled in the Base Transformer, which will enable better performance on the second task.

3.3 Neural Network Architecture

In this section, we will explain the detailed implementation of the base transformer, the trajectory head and the classification head. Modeling group interactions in team sports is a challenging task, because the data involves a spatio-temporal aspect for each player as well as interactions between the players. We solve this issue by first modeling the spatio-temporal information for each player and the ball, and then modeling the interactions between the players with an attention mechanism.

3.3.1 Transformer Encoder

The input data consists of all 11 objects $\tau_o^{t-L+1:t}$ of a given play sequence. Because these input vectors represent a time series for each tracked object o , we use a Long Short-Term Memory (LSTM) [13] layer to embed these trajectories into vectors (see Figure 3), to fully exploit the temporal information of the input data.

The 11 objects have different properties, since the ball behaves differently than an attacker. We encode this role information as shown in Figure 3 and choose a one-hot positional encoding to differentiate each object class (i.e. the ball, offensive players, and defensive players). We concatenate these 3-dimensional vectors to the output of the LSTM layer, creating the input to the first attention layer \mathbf{z}_o^0 . Although it would be possible to create embeddings for each player separately, previous work shows that adding player embeddings improves results only marginally while requiring a large feature engineering effort [30].

To generate context-aware player embeddings, given the previously described input embeddings, we use a Transformer encoder with a multi-head self-attention mechanism [35]. Transformer encoders consist of multiple attention-based layers. Each layer learns to adjust the object representations in relation to other objects, where the objects in our application are the ball and the players. More formally, a Transformer processes the input $\mathbf{z}_o^l \in \mathbb{R}^{d_v}$ at layer l to an output embedding $\mathbf{z}_o^{l+1} \in \mathbb{R}^{d_v}$, with input and output dimension d_v . The inputs are transformed into three matrices: query Q , key K , and value V , where each matrix represents the stacked input embeddings \mathbf{z}_o^l . These matrices are then transformed with trainable matrices $W_i^Q \in \mathbb{R}^{d_g \times d_v}$, $W_i^K \in \mathbb{R}^{d_g \times d_v}$, $W_i^V \in \mathbb{R}^{d_g \times d_v}$ and $W_i^O \in \mathbb{R}^{h d_v \times d_g}$, where d_g is a model hyperparameter and h is the number of self-attention heads. The multi-head self-attention function includes the residual connection and is calculated as

$$\mathbf{z}_o^{l+1} = LN(FF(LN(Att(Q, K, V)))) + \mathbf{z}_o^l \quad (1)$$

$$Att(Q, K, V) = Concat(head_1, \dots, head_h)W^O \quad (2)$$

$$head_i = softmax\left(\frac{(QW_i^Q)(KW_i^K)^T}{\sqrt{d_k}}\right)W_i^V, \quad (3)$$

where LN stands for layer normalization [1] and FF is a fully connected feedforward network (see Figure 3). We use the common practice of simplifying the hyperparameters by setting $d_g = d_v = d_k$, which we denote as hidden dimension d_h from here on.

The transformer contains a stack of N identical attention layers. Through this attention mechanism, the NN generates embeddings \mathbf{z}_o^N for each tracked object o , taking into account the information of all other players and the ball.

3.3.2 Trajectory Prediction Head

After generating the embeddings for each tracked object, we can predict the future trajectories by using a FeedForward neural network (FF) to generate one output vector ν_o^t per tracked object (see Figure 4). Because the physical properties of motion for the ball are very different from humans and because the behavior of an offensive player is different from the behavior of a defensive player, we use separate FFs for each role to generate these trajectories. We denote FFs that share the same weights with the same subscripts in Figure 4.

3.3.3 Classification Head

To predict labels for GAR, we consider the context-aware transformer embeddings of the players and the ball from the last transformer layer. For GAR, the input order of the players within a team should not impact the assigned label, and our model should be permutation invariant for each team. We use a team-level pooling layer which sums up all the embeddings of the players that belong to the offense and defense, respectively (see Figure 5). More formally, we can write the output values as

$$\mathbf{y} = \text{softmax} \left(FF \left(\mathbf{z}_B^N \oplus \sum_{i=1}^5 \mathbf{z}_{A_i}^N \oplus \sum_{i=1}^5 \mathbf{z}_{D_i}^N \right) \right), \quad (4)$$

where \oplus stands for concatenation. $\mathbf{y} \in \mathbb{R}^K$ is the output vector containing the probability for each of the K predicted class.

We use a softmax activation function on the last layer so that the output can be interpreted as probabilities and trained by minimizing the Negative Log-Likelihood (NLL) loss for K classes:

$$\mathcal{L}_{\text{NLL}} = -\frac{1}{|\mathcal{D}|} \sum_{\mathcal{D}} \sum_{k=1}^K \alpha_k \mathbf{y}_k \ln(\hat{\mathbf{y}}_k), \quad (5)$$

where α_k stands for class adjustments that balance the unequal class distribution [18]. \mathbf{y}_k is the one-hot representation of ground truth, and $\hat{\mathbf{y}}_k$ is the predicted probability that the play belongs to class k .

4 Experimental Setup

In this section, we describe the data set and the design of our labeling approach, which consists of weak-labeling and manual labeling. Note that the dataset with the full trajectories for 632 NBA games is publicly available, and we will provide a full dataset of our weak-labeled play segments.

4.1 Data Set

We used publicly available movement data collected from 632 NBA games during the 2015-2016 season¹. To avoid including play sequences without significant action (e.g. time-outs or free throws), we segmented the data into possessions, which start when the shot clock resets and end on the next reset. We only kept the part of the possession when all 10 players are in the offensive half of the court to focus on significant offensive tactics. Possessions shorter than 3s were discarded, resulting in 113,760 possessions. This amounts to 1.1 million seconds of game play, where locations of the players and the ball are captured every 0.04s. We downsampled the data by a factor of 3 to

¹ <https://github.com/sealneward/nba-movement-data>, last accessed July 2023; we are not associated with the data creator in any way.

reduce computational cost (similar to [42]) and obtain a sampling rate of $\Delta_t = 0.12s$. For all models evaluated in the experimental section, we used $L = 10$ time steps as input, corresponding to 1.2s of game play. We constructed a data set by splitting the possessions into 1.2s long non-overlapping segments, resulting in 869,905 play sequences. For each of those play sequences, we also included the future horizon of the next H time steps to allow training on a trajectory prediction task, where we experimented with varying values of $H = 10$, $H = 20$, and $H = 40$.

4.2 Labeling Group Activities in Basketball

We generated labels for pick-and-rolls and handoffs, two commonly used tactics by NBA teams. Plays that do not fall into these two categories were labeled as "other", resulting in $K = 3$ classes. Generating such labels manually is very time-consuming because it not only requires watching entire basketball games, but also high concentration to follow multiple players at the same time. There are also many edge cases that require multiple viewings to decide on the correct label. Here, we provide a short form explanation of the weak-labeling rules for pick-and-rolls and handoffs, and we explain the weak-labeling process in more detail in the Supplementary Material.

4.2.1 Weak-Labeling Pick-and-rolls

Pick-and-roll is an offensive tactic in which the attacking team tries to block the defender guarding the ball handler. Another attacker (the so-called roll-man) helps the ball handler by standing in the way of the defender (see Figure 1). This creates a difficult situation for the defender, who has to either run around the roll-man to keep guarding the ball handler or switch the assignment and guard the roll-man, which leads to a possible mismatch. In short, the rule-based approach matches each defender to an offensive player and then identifies situations in which the roll-man is very close to the ball handler's defender (see details in Supplementary Material).

4.2.2 Weak-Labeling Handoffs

A handoff is a different offensive tactic in which two attackers cross paths, and the ball is handed off when the players are close to each other. The action can also involve a very short pass. This tactic allows one to give the ball to another player with a low risk of losing it, but it can also be used as an opportunity to stand in the defender's way. In this sense, it is closely related to a pick-and-roll, but the ball possession changes during the action. The rule-based approach to identifying a handoff tries to identify a change in possession between two players, with a small distance between the players and a short time of ball transition (see details in Supplementary Material).

The weak-labeling process produced 45,802 "pick-and-rolls", 15,251 "handoffs", and 808,852 "other" play sequences.

4.2.3 Manual Labeling

We assigned manual labels for a total of 1,800 play segments. Because the data set is highly unbalanced according to weak-label distribution, we sampled play sequences equally from all 3 weak-labeled subgroups. To assign the labels, we generated video representations of 1.2s in length for each play segment. We carefully watched each video, replaying and freezing it until we were confident with its label. We sampled and labeled one video from each of the 3 weak labels until we found 600 pick-and-rolls manually labeled samples

for each class. We used 50% of these manual labels for testing and 50% for fine-tuning our NETS model.

5 Trajectory Prediction Problem

We used the self-supervised trajectory prediction problem as a pre-training step for GAR. This task itself has found recent interest in the computer vision community [32, 42, 43, 10]. In this section, we will show that our Transformer architecture can achieve state-of-the-art performance in trajectory prediction.

5.1 Experimental Design

We trained our models using pytorch. We used Adam optimizer with an initial learning rate of $5 \cdot 10^{-5}$ to reduce the mean squared error of predicted velocities

$$\mathcal{L}_{\text{MSE}}(\nu, \hat{\nu}) = \frac{1}{2H} \|\nu - \hat{\nu}\|_2^2. \quad (6)$$

We set all hidden dimension to $d_h = 256$, and used a Transformer with $L = 8$ layers and $h = 64$ prediction heads (also see 6.3). We used a ReLU activation function between all layers except the output. We trained in batches of 512 samples. Early stopping was applied after 50 epochs if the validation error did not improve, which typically took around 400 epochs, corresponding to about 30 hours of training on a single GPU.

To evaluate the trajectory prediction task, different prediction horizons $H = 10$, $H = 20$ and $H = 40$ were tested. From the predicted velocities, we computed the trajectories for the players and evaluated different approaches using two standard metrics [9, 22] in trajectory prediction: average displacement error (ADE) and final displacement error (FDE). FDE is the expected Euclidean distance between the predicted final location and the true final location of the tracked object after H time steps. ADE describes the expected average Euclidean distance between the predicted and the true trajectory at every predicted time step.

5.2 Baselines

Since the trajectory prediction task has been studied in literature before, we compared our NETS architecture with several strong baselines.

LSTM The baseline LSTM refers to an approach using a 2-layer LSTM, with 128 hidden nodes in each layer as an input encoder, followed by a 2-layer FF network, also with 128 hidden nodes. We used the same loss function and training setup as for our NETS experiments.

SocialGAN[9] This model uses an LSTM-based generator, coupled with a social pooling layer, to account for nearby actors and a discriminator that learns to distinguish between actual and simulated trajectories. We used the same hyperparameter settings as in [9]. It is important to note that SocialGAN was designed for pedestrian prediction, where all objects are embedded the same way in a shared layer. Therefore, adding the ball location would require significant modification, which is why this model does not consider the ball’s location in our experiments.

M. VRNN [42] This model uses a hierarchical approach in which a Variational RNN is trained on programmatic weak supervision to first predict a location that a player wants to reach. Then it uses a second identical layer to predict a trajectory that the player will take to reach it. It is trained like a Variational Autoencoder, aiming to

Table 1: Comparison of various models in terms of error metrics ADE and FDE (in feet) for prediction horizons $H = 10$ (1.2 seconds), $H = 20$ (2.4 seconds), and $H = 40$ (4.8 seconds).

Method	$H = 10$		$H = 20$		$H = 40$	
	ADE	FDE	ADE	FDE	ADE	FDE
LSTM [10]	1.61	2.98	3.43	6.91	6.59	11.97
M. VRNN [42]	1.70	3.43	4.46	8.66	8.48	14.98
Social GAN [9]	1.25	2.75	3.09	6.67	6.47	12.35
NETS _{no LSTM}	1.18	2.51	2.98	6.42	6.31	11.84
NETS	1.08	2.34	2.78	5.87	5.70	10.88

reproduce realistic behavior of offensive players. We used the same hyperparameter settings as in [9]. Because this baseline only predicts the trajectories of offensive players, we compared the performance of all models for the offensive players only.

NETS_{no LSTM} This is an ablation of our NETS model, with the only difference being that the input embedding is a 2-layer feedforward network with 256 nodes each, instead of the LSTM-layer in NETS.

5.3 Results

We compared the performance of NETS on the trajectory prediction task to several strong baselines. Results for various prediction horizons H are shown in Table 1. NETS_{no LSTM} and NETS, which are based on transformers, outperformed the other baselines, showing improvements on different prediction horizons. The results held for both ADE and FDE, suggesting that the improvement can be observed along the entirety of a player’s path. MACRO VRNNs trajectories are designed to resemble the behavior of basketball players, but can be far off from the ground truth. SocialGAN produced reasonably accurate predictions, but was inferior to other approaches that were specifically designed for this basketball dataset.

Another interesting finding is that NETS_{no LSTM} performed worse than our full NETS model. We believe that the LSTM input embedding allows the model to more easily extract temporally dependent information in the input data, therefore explaining a reduced accuracy of a model that has to learn temporal patterns from unstructured input data.

6 Group Activity Recognition (GAR)

In this section, we report the performance of our proposed approach on GAR. We first evaluated the ability of these models to classify a large weak-labeled dataset. We were particularly interested to investigate if the self-supervised task of trajectory prediction can improve the accuracy of GAR. Then we evaluated our NETS model on manually labeled data to determine if a deep learning method can outperform the rule-based labeling approach.

6.1 Experimental Design

Due to our modular architecture design, the NN architecture was the same as for the trajectory prediction task, except for the last prediction layer (see Section 3.3). We used an 80 / 10 / 10 train-, validation-, and test-split on all of these play sequences, and we applied early stopping after 50 epochs. To balance the training set distribution of 45,802 “pick-and-rolls”, 15,251 “handoffs”, and 808,852 “other” play sequences, we downsampled other plays to 45,802 (same as pick-and-roll), we used weighting factors α_k of 0.77, 2.34, and 0.77 for pick-and-rolls, handoffs, and other, respectively.

We note that on the test set the group activity labels were dominated by play sequences labeled as ‘other’ and a trivial classifier

predicting that class would achieve 93.0% test accuracy. Therefore, instead of reporting accuracy, we report the multi-class F1 score to compare the performance of the models. We calculated the confusion matrix M , where M_{ij} represents the numbers of plays with ground truth labels i classified as j . We calculated F1 scores by converting the 3-class classification problem into three binary classification problems, yielding the one vs. all F1-scores

$$\text{precision}_i = \frac{M_{ii}}{\sum_j M_{ji}}; \quad \text{recall}_i = \frac{M_{ii}}{\sum_j M_{ij}} \quad (7)$$

$$\text{F1}_i = 2 \cdot \frac{\text{precision}_i \cdot \text{recall}_i}{\text{precision}_i + \text{recall}_i}. \quad (8)$$

6.2 Baselines

Although GAR has been studied in different sports (see Section 2), they rely on the original video of the plays, whereas the dataset studied in this paper only relies on trajectory data and the original videos are not available. Due to the lack of baselines using deep learning, we used standard shallow baselines, consisting of logistic regression **LReg**, random forest classifier **RForest**, and gradient boosting classifier **GBoost**. To establish a baseline, we trained 3 shallow models using the popular *sklearn* implementation. The input to these shallow models were the play sequences $\tau_o^{t-L+1:t}$ consisting of 220 features. We found the optimal parameter settings through grid search.

To add a deep-learning baseline, we compared our NETS architecture to a model inspired by DeepHoops [30], which is an **LSTM**-based neural network with the same settings described in Section 5.2. To allow a fair comparison to our NETS architecture, we used the same training procedure of pretraining the model using the trajectory prediction task and then changing the prediction head to enable play classification.

6.3 Evaluation on Weak-Labels

In Table 2 we show F1 scores on the test-set of the weak labeled data. Our NETS model outperformed shallow approaches by a large margin. The results show that the classification of group activities even if labels are generated with a rule-based approach is a difficult problem for shallow models such as LReg, RForest, and GBoost, as indicated by the relatively low F1 scores of these traditional machine learning algorithms. Interestingly, we did not find that the LSTM baseline outperforms GBoost on the GAR problem.

Table 2: Classification test performance compared to baselines. Tested on 4,581 pick-and-rolls (p&r), 1,525 handoffs and 80,884 other plays.

Method	p&r F1-score	handoff F1-score	other F1-score
LReg	0.188	0.099	0.755
RForest	0.329	0.261	0.810
GBoost	0.398	0.443	0.915
LSTM	0.360	0.490	0.905
NETS	0.856	0.768	0.988

We performed an ablation study of the NETS models to evaluate our architecture design choices: We created a model ablation by removing the LSTM-embedding (see Figure 3), in which the inputs were embedded with a 2-layer feedforward neural network. We also examined the influence of the team-wise pooling layer depicted in Figure 5 by removing the summation and instead concatenating all representations \mathbf{z}_i^N .

Table 3: Classification test performance of NETS with different architecture ablations. Tested on 4,581 pick-and-rolls (p&r), 1,525 handoffs and 80,884 other plays.

pretrain	LSTM embed.	pooling	p&r F1-score	handoff F1-score	other F1-score
✗	✗	✓	0.705	0.667	0.973
✗	✓	✗	0.777	0.644	0.980
✗	✓	✓	0.802	0.675	0.982
✓	✗	✓	0.803	0.731	0.983
✓	✓	✗	0.829	0.718	0.985
✓	✓	✓	0.856	0.768	0.988

All the models in Table 3 were ablations of NETS and were based on a Transformer architecture. We see that all these models performed better than the shallow models or the LSTM-based model presented in Table 2, indicating that our domain-adapted Transformer architecture is well suited for this problem. All models that were pre-trained showed significant improvement compared to the same models without pretraining. As in the trajectory prediction task, we noted significant improvement when utilizing an LSTM-embedding at the input of the base Transformer instead of dense embedding layers. We also observed an improvement when using the team-pooling layer at the output. Furthermore, we observed that training GAR without pre-training took around 500 epochs when we trained from scratch but only roughly 200 epochs when we started with a pretrained model.

6.4 Evaluation on Manual Labels

Table 4: Confusion matrix of weak-labels tested on manually labeled data, consisting of 300 pick-and-rolls (p&r), 300 handoffs and 300 other plays.

	ground truth		
	p&r	handoff	other
predicted p&r	282	43	24
predicted handoff	5	253	21
predicted other	13	4	260

Table 5: Accuracies of four variants of NETS. "weak-labels" refers to concordance between weak- and manual labels. Tested on manually labeled data, consisting of 300 pick-and-rolls (p&r), 300 handoffs and 300 other plays.

method	p&r F1-score	handoff F1-score	other F1-score
1) weak-labels	0.869	0.874	0.893
2) NETS	0.915	0.863	0.908
3) NETS, finetune on ml* only	0.784	0.813	0.844
4) NETS, validate on ml*	0.932	0.930	0.908
5) NETS, finetune on ml*	0.951	0.938	0.902

* ml = manual labels

As described in Section 4.2, we manually labeled 600 pick-and-rolls, 600 handoffs and 600 other plays, which we split into 50% train- and 50% test-set.

To get an insight into the performance of our rule-based weak-labeling approach, we show the confusion table of the weak-labels compared to the assigned ground truth from manual labeling in Table 4. While the weak-labeling rules tended to extract many correctly labeled play sequences, both rules for pick-and-rolls and handoff included 24 and 21 play sequences that were labeled as other by manual labeling. Furthermore, there were relatively frequent misclassification of handoffs as pick-and-rolls.

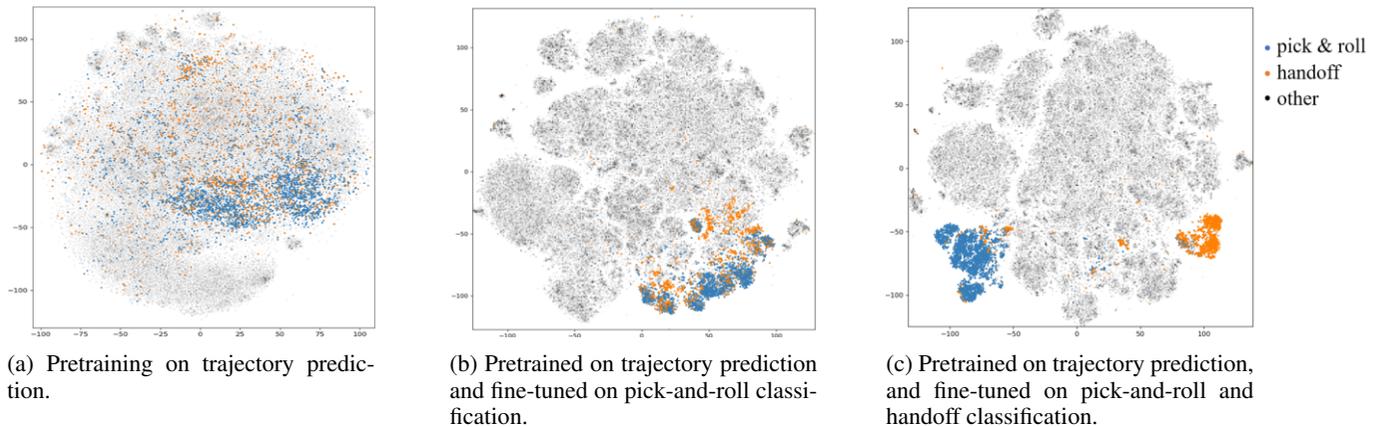


Figure 6: t-SNE of embeddings of - pick-and-rolls (blue), - handoffs (orange) and - random plays (black).

In the previous subsection, we showed that our NETS architecture can predict the weak-labels with high accuracy. Next, we evaluated the performance on manual labels. We also examined the usefulness of weak and manual labels when testing on these manual labels.

We evaluated the following approaches: 1) Using the rule-based weak-labeling approach to classify the group activities, which requires human expertise to define the rules and is not a deep learning approach. 2) Pretrain the NETS model on a trajectory prediction task and fine-tune the model using a large amount of weak-labels. 3) Pretrain the NETS model on a trajectory prediction task and fine-tune the model only using 900 manual labels. 4) Pretrain the NETS model on a trajectory prediction task, use the weak-labels in the training set and the manual labels in the validation set for early stopping. 5) Pretrain the NETS model on a trajectory prediction task, fine-tune it first on the weak-labels and then fine-tune again on the manual labels.

The results in Table 5 show the F1 scores for the rule-based weak-labels calculated from the confusion matrix (Table 4). We observed a slight improvement when training NETS on the weak labels compared to the weak labels themselves, which we hypothesize stems from the ability of the neural network to generalize better than the rules alone. In contrast, fine-tuning only on manual labels performed worse than the rule-based labeling, indicating that we did not have enough manual labels to train a large NN. Using manual labels in the validations set further improved the accuracy, implying that NETS overfit the weak labels without access to manual labels. Using a sequential fine-tuning paradigm resulted in the highest accuracy (row 5). When comparing row 5) with the weak-label F1 scores in row 1), we can observe that the accuracy on pick-and-rolls increased from 0.869 to 0.951, for handoffs it increased from 0.874 to 0.938, and for others from 0.893 to 0.902.

7 Evaluation of Representations

In this section, our goal is to show that the learned embeddings got increasingly better with each learning step. To this end, we perform a qualitative analysis where we transformed the higher dimensional embeddings into 2D space to allow for visual inspection by using t-SNE [34] on the test set. We then visually examined how tightly pick-and-rolls and handoffs were packed together. Scatter plots shown in Figure 6 illustrate the difference between internal representations of the three NETS models.

Figure 6a shows the embeddings generated from a model trained on the trajectory prediction task. The embeddings of pick-and-roll plays and handoffs did not seem to be clustered together and were

scattered within other plays. Some clusters were forming, but they were not very distinct.

Figure 6b shows that pick-and-rolls are clearly separated from other plays. More interestingly, handoffs are often clustered together and are mostly placed between pick-and-rolls and other plays, correctly indicating that there are some similarities between the plays. It also shows some unwanted behavior, since many handoffs are clearly placed among pick-and-rolls, indicating that it is difficult to distinguish the two types of play from each other.

Figure 6c shows the embeddings created after fine-tuning on both pick-and-rolls and handoffs. The two types of strategies were now clearly separated, with only few exceptions. Furthermore, we can see more distinct clusters forming among the other plays, indicating that the embeddings were of higher quality and more capable of distinguishing different types of strategies, even for those that have no labels.

8 Conclusion and Future Work

We have shown that our NETS model with a specifically designed transformer architecture to address common challenges in sports analytics can classify group activities with high accuracy. Pretraining the model on a self-supervised trajectory prediction task significantly improved the model performance on downstream tasks. In future work, we aim to utilize the generated play embeddings to discover strategic differences employed by different teams and generate a tool that can help automatically analyze a game.

To allow reproducibility, we share our code online at www.github.com/S-Hauri/NETS.

9 Acknowledgement

Research was in part sponsored by the DEVCOM Analysis Center and was accomplished under Cooperative Agreement Number W911NF-22-2-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. We thank Zach Cameron and Harith Siddiqui for their contribution in labeling and data analysis in this work.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton, ‘Layer normalization’, *arXiv preprint arXiv:1607.06450*, (2016). 3
- [2] T. Chen, J. Frankle, S. Chang, S. Liu, Y. Zhang, M. Carbin, and Z. Wang, ‘The lottery tickets hypothesis for supervised and self-supervised pre-training in computer vision models’, *arXiv preprint arXiv:2012.06908*, (2020). 2
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton, ‘A simple framework for contrastive learning of visual representations’, *arXiv preprint arXiv:2002.05709*, (2020). 2
- [4] Bin Dai, Ziyu Wang, and David Wipf, ‘The usual suspects? reassessing blame for vae posterior collapse’, in *International Conference on Machine Learning*, pp. 2313–2322. PMLR, (2020). 2
- [5] X. Deng, Y. Xiang, A. Mousavian, C. Eppner, T. Bretl, and D. Fox, ‘Self-supervised 6d object pose estimation for robot manipulation’, in *2020 IEEE ICRA*, pp. 3665–3671. IEEE, (2020). 2
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, ‘Bert: Pre-training of deep bidirectional transformers for language understanding’, *arXiv preprint arXiv:1810.04805*, (2018). 2
- [7] Haodong Duan, Yue Zhao, Kai Chen, Dian Shao, Dahua Lin, and Bo Dai, ‘Revisiting skeleton-based action recognition’, *arXiv preprint arXiv:2104.13586*, (2021). 2
- [8] Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, and Lawrence Carin, ‘Cyclical annealing schedule: A simple approach to mitigating kl vanishing’, *arXiv:1903.10145*, (2019). 2
- [9] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi, ‘Social gan: Socially acceptable trajectories with generative adversarial networks’, in *Proceedings of the IEEE CVPR*, pp. 2255–2264, (2018). 2, 5
- [10] Sandro Hauri, Nemanja Djuric, Vladan Radosavljevic, and Slobodan Vucetic, ‘Multi-modal trajectory prediction of nba players’, in *Proceedings of the IEEE/CVF WACV*, pp. 1640–1649, (2020). 2, 5
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick, ‘Momentum contrast for unsupervised visual representation learning’, in *Proceedings of the IEEE/CVF CVPR*, pp. 9729–9738, (2020). 2
- [12] Dirk Helbing and Peter Molnar, ‘Social force model for pedestrian dynamics’, *Physical review E*, **51**(5), 4282, (1995). 2
- [13] Sepp Hochreiter and Jürgen Schmidhuber, ‘Long short-term memory’, *Neural computation*, **9**(8), 1735–1780, (1997). 2, 3
- [14] Mostafa S Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori, ‘A hierarchical deep temporal model for group activity recognition’, in *Proceedings of the IEEE CVPR*, pp. 1971–1980, (2016). 2
- [15] Boris Ivanovic and Marco Pavone, ‘The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs’, in *Proceedings of the IEEE/CVF ICCV*, pp. 2375–2384, (2019). 2
- [16] Neziha Jaouedi, Nouredine Boujnah, and Med Salim Bouhlel, ‘A new hybrid deep learning model for human action recognition’, *Journal of King Saud University-Computer and Information Sciences*, **32**(4), 447–453, (2020). 2
- [17] Gregory Kahn, Adam Villafior, Bosen Ding, Pieter Abbeel, and Sergey Levine, ‘Self-supervised deep reinforcement learning with generalized computation graphs for robot navigation’, in *2018 IEEE ICRA*, pp. 1–8. IEEE, (2018). 2
- [18] Gary King and Langche Zeng, ‘Logistic regression in rare events data’, *Political analysis*, **9**(2), 137–163, (2001). 4
- [19] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezaatofghi, and Silvio Savarese, ‘Social-bigat: Multimodal trajectory forecasting using bicycle-gan and graph attention networks’, *NeurIPS*, **32**, (2019). 2
- [20] Parth Kothari, Sven Kreiss, and Alexandre Alahi, ‘Human trajectory forecasting in crowds: A deep learning perspective’, *arXiv preprint arXiv:2007.03639*, (2020). 2
- [21] Christos Koutsouridis, Georgios Karamousalidis, and Christos Galatzoulas, ‘The efficacy of “high pick and roll” in relation to the defence’s reaction and its effect on the result of the game’, *International Journal of Performance Analysis in Sport*, **18**(4), 554–567, (2018). 1
- [22] Sajjad Mozaffari, Omar Y Al-Jarrah, Mehrdad Dianati, Paul Jennings, and Alexandros Mouzakitis, ‘Deep learning-based vehicle behavior prediction for autonomous driving applications: A review’, *IEEE Transactions on Intelligent Transportation Systems*, (2020). 5
- [23] Bingbing Ni, Shuicheng Yan, and Ashraf Kassim, ‘Recognizing human group activities with localized causalities’, in *2009 IEEE CVPR*, pp. 1470–1477. IEEE, (2009). 2
- [24] H Nunes, X Iglesias, G Daza, A Iruiria, T Caparrós, MT Anguera, et al., ‘The influence of pick and roll in attacking play in top-level basketball.’, *Cuadernos de Psicología del Deporte*, **16**(1), 129–142, (2016). 1
- [25] NEIL Owens, C Harris, and C Stennett, ‘Hawk-eye tennis system’, in *2003 International Conference on Visual Information Engineering VIE 2003*, pp. 182–185. IET, (2003). 1
- [26] Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Peter Chen, John Canny, Pieter Abbeel, and Yun Song, ‘Evaluating protein transfer learning with tape’, in *NeurIPS*, pp. 9689–9701, (2019). 2
- [27] Alexander Rives, Siddharth Goyal, Joshua Meier, Demi Guo, Mylle Ott, C Lawrence Zitnick, Jerry Ma, and Rob Fergus, ‘Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences’, *bioRxiv*, 622803, (2019). 2
- [28] Ryan Sanford, Siavash Gorji, Luiz G. Hafemann, Bahareh Pourbabae, and Mehrsan Javan, ‘Group activity detection from trajectory and video data in soccer’, in *Proceedings of the IEEE/CVF CVPR (CVPR) Workshops*, (June 2020). 2
- [29] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain, ‘Time-contrastive networks: Self-supervised learning from video’, in *2018 IEEE ICRA*, pp. 1134–1141. IEEE, (2018). 2
- [30] Anthony Sicilia, Konstantinos Pelechris, and Kirk Goldsberry, ‘Deep-hoops: Evaluating micro-actions in basketball using deep feature representations of spatio-temporal data’, in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2096–2104, (2019). 2, 3, 6
- [31] LLC Stats, ‘Stats sportvu basketball player tracking’, *SportVU website*, available at: <https://www.stats.com/sportvu-basketball/>, last accessed: Feb, 12, (2019). 1
- [32] C Sun, P Karlsson, J Wu, J B Tenenbaum, and K Murphy, ‘Stochastic prediction of multi-agent interactions from partial observations’, *arXiv preprint arXiv:1902.09641*, (2019). 2, 5
- [33] Jur Van den Berg, Ming Lin, and Dinesh Manocha, ‘Reciprocal velocity obstacles for real-time multi-agent navigation’, in *2008 IEEE international conference on robotics and automation*, pp. 1928–1935. Ieee, (2008). 2
- [34] Laurens Van der Maaten and Geoffrey Hinton, ‘Visualizing data using t-sne.’, *Journal of machine learning research*, **9**(11), (2008). 7
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, ‘Attention is all you need’, *NeurIPS*, **30**, 5998–6008, (2017). 2, 3
- [36] Anirudh Vemula, Katharina Muelling, and Jean Oh, ‘Social attention: Modeling attention in human crowds’, in *2018 IEEE ICRA*, pp. 4601–4607. IEEE, (2018). 2
- [37] Li-Fang Wu, Qi Wang, Meng Jian, Yu Qiao, and Bo-Xuan Zhao, ‘A comprehensive review of group activity recognition in videos’, *Int. Journal of Automation and Computing*, **18**(3), 334–350, (2021). 2
- [38] Dezhong Xu, Heng Fu, Lifang Wu, Meng Jian, Dong Wang, and Xu Liu, ‘Group activity recognition by using effective multiple modality relation representation with temporal-spatial attention’, *IEEE Access*, **8**, 65689–65698, (2020). 2
- [39] Hangjie Yuan, Dong Ni, and Mang Wang, ‘Spatio-temporal dynamic inference network for group activity recognition’, in *Proceedings of the IEEE/CVF ICCV*, pp. 7476–7485, (2021). 2
- [40] F. Zappardino, T. Uricchio, L. Seidenari, and A. Del Bimbo, ‘Learning group activities from skeletons without individual action labels’, in *2020 25th ICPR*, pp. 10412–10417. IEEE, (2021). 2
- [41] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer, ‘S4l: Self-supervised semi-supervised learning’, in *Proceedings of the IEEE/CVF ICCV*, pp. 1476–1485, (2019). 2
- [42] Eric Zhan, Stephan Zheng, Yisong Yue, Long Sha, and Patrick Lucey, ‘Generating multi-agent trajectories using programmatic weak supervision’, in *ICLR*, (2018). 2, 4, 5
- [43] S. Zheng, Y. Yue, and J. Hobbs, ‘Generating long-term trajectories using deep hierarchical networks’, *NeurIPS*, **29**, (2016). 2, 5
- [44] Yue Zhou, Bingbing Ni, Shuicheng Yan, and Thomas S Huang, ‘Recognizing pair-activities by causality analysis’, *ACM Transactions on Intelligent Systems and Technology (TIST)*, **2**(1), 1–20, (2011). 2