Generating Defensive Plays in Basketball Games

Chieh-Yu Chen
National Chiao Tung University
jaycase.cs01@gmail.com

Wenze Lai
National Chiao Tung University
nsknsl.cs05g@g2.nctu.edu.tw

Hsin-Ying Hsieh
National Chiao Tung University
x4ga71@gmail.com

Wen-Hao Zheng
National Chiao Tung University
lobst3rd@gmail.com

Yu-Shuen Wang
National Chiao Tung University
yushuen@cs.nctu.edu.tw

Jung-Hong Chuang
National Chiao Tung University
jhchuang@cs.nctu.edu.tw

ABSTRACT

In this paper, we present a method to generate realistic defensive plays in a basketball game based on the ball and the offensive team's movements. Our system allows players and coaches to simulate how the opposing team will react to a newly developed offensive strategy for evaluating its effectiveness. To achieve the aim, we train on the NBA dataset a conditional generative adversarial network that learns spatio-temporal interactions between players' movements. The network consists of two components: a generator that takes a latent noise vector and the offensive team's trajectories as input to generate defensive team's trajectories; and a discriminator that evaluates the realistic degree of the generated results. Since a basketball game can be easily identified as fake if the ball handler, who is not defended, does not shoot the ball or cut into the restricted area, we add the wide open penalty to the objective function to assist model training. To evaluate the results, we compared the similarity of the real and the generated defensive plays, in terms of the players' movement speed and acceleration, distance to defend ball handlers and non- ball handlers, and the frequency of wide open occurrences. In addition, we conducted a user study with 59 participants for subjective tests. Experimental results show the high fidelity of the generated defensive plays to real data and demonstrate the feasibility of our algorithm.

CCS CONCEPTS

• **Information systems → Multimedia content creation**;

KEYWORDS

Conditional adversarial network, basketball, defensive strategies

ACM Reference Format:
Chieh-Yu Chen, Wenze Lai, Hsin-Ying Hsieh, Wen-Hao Zheng, Yu-Shuen Wang, and Jung-Hong Chuang. 2018. Generating Defensive Plays in Basketball Games. In *2018 ACM Multimedia Conference (MM '18), October 22–26, 2018, Seoul, Republic of Korea.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3240508.3240670

MM '18, October 22–26, 2018, Seoul, Republic of Korea

ACM ISBN 978-1-4503-5665-7/18/10…$15.00
https://doi.org/10.1145/3240508.3240670

1 INTRODUCTION

Analyzing players' performance and behaviors based on statistical and historical data is becoming an effective way for coaches in National Basketball Association (NBA) to develop winning strategies. For example, the analysis can help determine the line-up as well as match-ups in a game and choose the offensive and defensive strategies against the opposing team. It also can help players identify their weaknesses so as to polish their skills. Among the teams in NBA, the Houston Rockets are most well known for using statistical analysis, which has become known as the Moreyball. Recently, more teams, such as the Toronto Raptors, adopted similar methodologies, with their analytics team even creating a system to teach players what they should have done on a play.

Prior research in basketball games is mainly based on analyzing broadcast videos, in addressing challenging problems such as player tracking, perspective effect, dynamic camera motion, and occlusions. Engineers in STATS LLC cooperated with NBA teams and adopted an alternative way to prevent the occurrence of the above mentioned problems when analyzing sport videos. They installed static video cameras at the top of the stadium to track positions of the players and the ball. Some of these data have been released, making sports data analytics a promising research area.

We present a system that can generate realistic defensive plays in a basketball game based on the ball and the offensive team's movements. By using our system, players and coaches can simulate how an opposing team will react to a newly developed offensive strategy for evaluating its effectiveness. To achieve the aim, we train a generative adversarial network (GAN) to learn spatio-temporal interactions between players' movements. The network is composed of a generator and a discriminator. The former takes a latent noise vector and the offensive team's movements as input to generate defensive players' moving trajectories; and the latter evaluates the realistic degree of the generated results. These two components are trained iteratively and alternatively to make the generated trajectories becoming more and more realistic. Observing that a good defensive play must prevent wide open between the ball handler and the basket, we add a penalty term to the objective function to assist training. Note that our convolutional neural network (CNN) has the striking feature of being able to process variable length sequences like recurrent neural networks (RNN). In particular, the network is composed of 1D convolutional layers, with the kernel moving along the temporal dimension to extract features that capture players' relative positions in a time span.

We conducted several experiments to evaluate the models trained with and without the penalty term, and an alternative implementation based on RNN. Figure 2 and the supplemental videos show the results. To measure the realistic degree of the generated plays, we compared the players' moving speed and acceleration, distance to defend ball handlers and non- ball handlers, and frequency of open space occurrences. We also carried out a user study, involving 59 participants for subjective tests. On average, they had a 63% chance of successfully distinguishing between the real and the generated plays, compared to 50% chance by random guessing.

2 RELATED WORKS

Basketball analytic has risen significantly over years and adopted by many teams in NBA. Researchers in statistics and computer science have also started working on new ways to analyze and visualize the game. The works can be classified into several aspects. Many of them are based on statistics and number crunching [5, 14, 18, 19]. They attempted to determine which player is a good defender on specific locations on the court, and how efficient the player is at defending an offensive player. Several works are about predicting if the shot is a miss or make by analyzing basketball and player trajectories [9, 20]. There were also works presented to classify offensive strategies in basketball games [4, 21].

Understanding the player's movement in a basketball game is crucial in practicing and evaluating strategies. Zheng et al.[22] presented a hierarchical policy neural network to generate long-term trajectory of an offensive player based on the other nine player locations on the court. The experiments show that the combination of micro- and macro-policies is able to generate natural moving behaviors of players. In addition to academic researches, Toronto Raptors showcased a private system developed by their analytic team [1], which determines where each individual defensive player should position on the court while playing defense. While the method used in the Toronto Raptors' system is not disclosed, and sometimes the players behave like superman in the introduced videos, we present a generative adversarial network to generate natural moving trajectories of defensive players.

Generative adversarial network is an unsupervised learning that maps a latent noise vector z to an output $G(z)$, which is widely used in content generation [7, 11]. The network is composed of a generator and a discriminator. The former is trained to generate samples that can fool the latter; whereas the latter is trained to differentiate the real and the generated samples. Because training a GAN is difficult [1], several strategies were presented. Pathak et al. [17] integrated the L2 loss into the training process of a GAN, and verified its effectiveness. Arjovsky et al. [2] and Gulragini et al. [8] presented Wasserstein GAN, which defines the loss by measuring the earth moving distance between two probability distributions. Compared to traditional GAN, Wasserstein GAN improves stability, captures continuity of distribution distance, and adds meaningful training gradients. Meanwhile, to control the output $G(z)$, Mirza and Osindero [16] introduced conditional GAN that is fed by not only a latent noise vector z but also a condition y so that the output becomes $G(z|y)$. The presented work adopts both the ideas of the Wasserstein GAN and the conditional GAN. It considers the ball and the offensive team's movements as condition when generating defensive players' movements in a basketball game.

[1]Lights, Cameras, Revolution. http://www.grantland.com/features/the-toronto-raptors-sportvu-cameras-nba-analytical-revolution/

3 DEFENSIVE PLAYS GENERATION

3.1 Datasets

The training dataset is from SportVU, in which the ball $\mathcal{R}^3$ and the players' positions $\mathcal{R}^2$ on the court are tracked at 25 frames per second. We parse the NBA play by play text [2] and then partition the whole basketball game into independent plays. Each play starts when the offensive team brings the ball across or inbound from half-court, and ends when an offensive player takes a shot either missed or made. In addition, we down sample the trajectories to 5 frames per second, and partition the plays into two subsets for training (90%) and testing (10%).

3.2 Network architecture

Given the ball and the offensive team's movements, we train a conditional GAN to generate defensive players' movements in a basketball game. We choose GAN rather than supervised learning, for example, input an offensive play and then output the corresponding defensive play, because an offensive play can be defended in different ways. Figure 1 shows the network, which consists of a generator and a discriminator (critic network). The inputs of the generator are a latent noise vector z and a condition vector y, which encodes the ball and the offensive team's player positions in t frames. The generator first linearly projects the two vectors to feature maps that have the same shape, and then adds the maps together. This design allows the noise latent vector z to have global influence on the entire trajectory, rather than players positions in individual frames. The generator then passes the combined feature map through residual blocks [10] and a convolutional layer to generate defensive players' trajectories $G(z|y)$.

The discriminator takes the ball, offensive and defensive players' trajectories as inputs and determines the realistic degree of the basketball game (players movement). The condition y is concatenated with the corresponding defensive players' trajectories x in the dataset and the generated trajectories $G(z|y)$ to form real and fake pairs, respectively. We feed the two pairs into the critic network, which has a similar structure compared to the generator, for training. The output of this critic network is a score that measures the similarity of a real and a fake plays. Specifically, we minimize the objective function

$$\begin{aligned} L(G,C) = & E_{y\sim P_{data}(Y),z\sim P_z(Z)}[C(G(z|y)|y)] \\ & -E_{x,y\sim P_{data}(X,Y)}[C(x|y)] \\ & +\lambda \times E_{\hat{x}\in P_{data}(\hat{X})}E[(\|\nabla_{\hat{x}}C(\hat{x})\|_2-1)^2], \end{aligned} \quad (1)$$

where G and C are the generator and the critic networks, respectively, y is the condition, z is the latent noise vector, $\hat{X}$ is the interpolation of the generated distribution and the real dataset, and λ is the weight for the regularization term. We set $\lambda = 10$ in our experiments.

[2]An NBA play by play text describes the game clock, the type of events, such as shot made or missed, free throw, block, and turnover, and the players involved in the play.

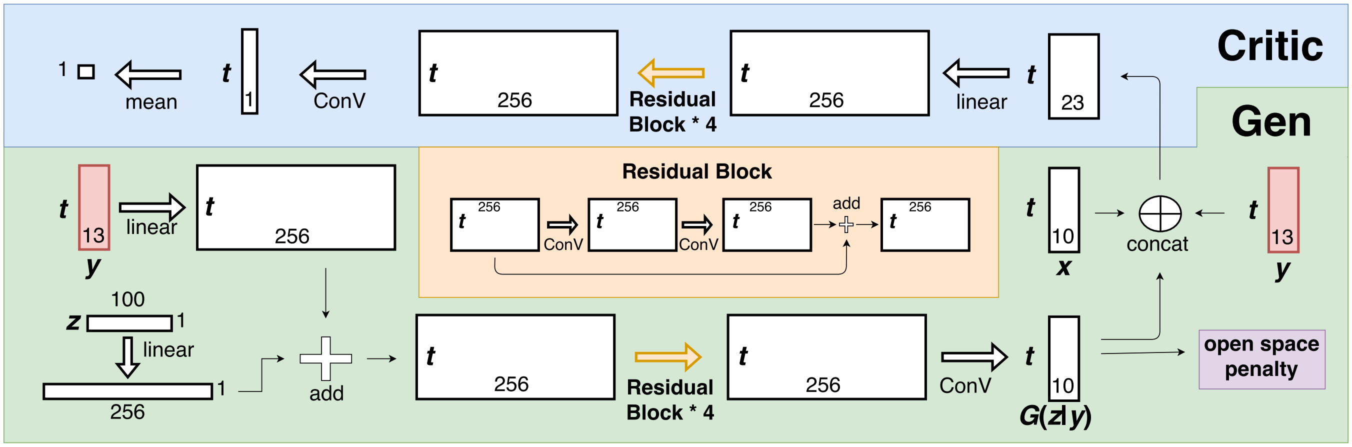

Figure 1: The architecture of the presented network. Each block represents a single feature vector/map with text providing its resolution. All the convolutional layers are one-dimensional in the temporal domain with a kernel size of 5. In addition, each convolutional layer is followed by a normalization layer [3] and a leaky ReLU activation function ($\alpha = 0.2$).

3.3 Wide open penalty

Wide open is a situation that an offensive player is not defended. If the player is a ball handler, he can drive into the restricted area or shoot the ball to make scores. Although the network trained by minimizing Equation 1 can produce fairly good results in many cases. we observe that in some frames the ball handler does not attempt to make scores even though there is a wide open nearby. Since the offensive players' movements are real, there must be a problem in the generated defensive plays. Therefore, we add a penalty term to penalize the open space of $G(z|y)$. That is,

$$\beta \cdot \left|W(x) - W(G(z|y))\right|,$$

where $W(\bullet) = ((1+\theta) \times (1+|a-b|))$, x is a randomly selected real defensive play, b is the ball position in 2D, a is the front defender's position closest to the ball handler, θ is the angle formed by $(a-b)$ and the edge from the ball to the basket, and $\beta = |C(G(z|y)|y)|$ is the weight used to balance the strength of this wide open penalty and the score given by the critic.

3.4 Implementation details

We train the presented network by using the Adam optimizer [13] with a learning rate= 1e-4 on a single NVIDIA GeForce 1080 ti. The parameters in the network are initialized by using Xavier [6]. We set the batch size to be 128 when training the network. Although the network can process basketball plays with variable length, we train the network with a fixed number ($n = 100$) of frames to enable parallelism. But we emphasize that n can be different in testing. To keep the critic C close to optimum during the training process, in the first 10 epochs, we train the generator only 1 iteration per epoch. After that, we train the critic 5 iterations per generator iteration in general. But for every 10 epochs, we again train the generator 1 iteration in the epoch. The training process stops at 828K iterations because the loss is unable to decrease and the system starts overfitting. We refer readers to [2] for further details.

4 RESULTS AND DISCUSSIONS

We tested the presented network on many offensive conditions. For each condition, we generated 100 sets of defensive plays and then selected the set with the highest score given by the critic. Overall, the generated players' trajectories are visually realistic because of natural moving behaviors such as curving and making sharp turns. In addition, all players in the defensive team strive to prevent the occurrence of wide open. To evaluate the performance of our CNN and the presented wide open penalty, we trained another four models for comparison. They are supervised, CNNONLY, RNNONLY, and RNN+WO, where +WO means *with the wide open penalty*'. The supervised model is similar to the generator in our CNN model, but the input and output are offensive and defensive players' movements, respectively. The RNN model is a bi-directional long short term memory network, in which the number of free parameters was similar to that of the CNN. To approach the best performance of these networks, we stopped the training process when the results in validation sets could not be further improved.

Figure 2 shows an offensive condition, in which the players were executing a pick and roll strategy to earn an open space for the ball handler. If this offensive strategy succeeds, the ball handler can shoot with no defenders standing close to and in front of him. Therefore, the goodness of a defensive strategy can be identified by observing the distance between the ball handler and the closest front defender. Figure 2 (a) and (b) show the real and the generated (by CNN+WO) defensive plays, respectively, under the same offensive condition. In the second and the third columns, the offensive player *A4* is moving toward the defensive player *B5* to execute a pick and roll strategy with *A2* (ball handler). His goal is to set a pick/screen (attempt to separate the defender and the ball handler) on *B5* and allow *A2* more space to create an offensive play. Therefore, in a good defensive play, there should be an additional defensive player moving towards *A2*, following *A4*, in order to get ready to help defend, either by switching defensive assignment or trapping the ball handler by double teaming (Figure 2 (a)). Notice

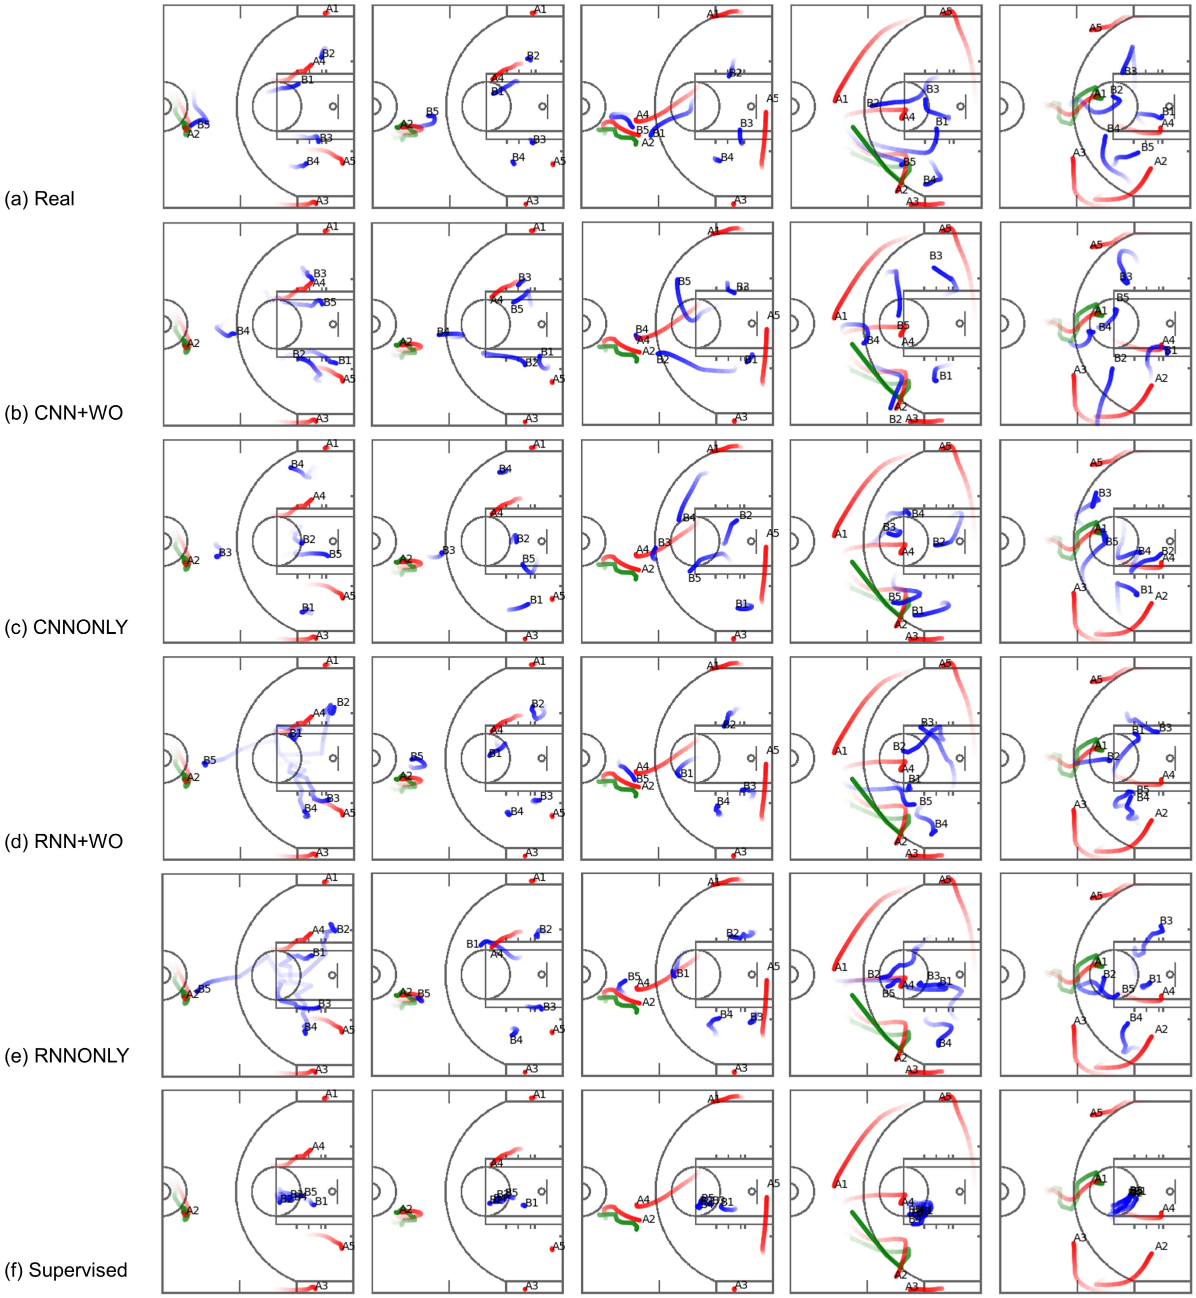

Figure 2: From (a) to (f) are the real and the generated defensive plays conditioned on the same offensive play. We partition the play into five segments according to time (i.e., 20 frames per segment) and show them from left to right to achieve a clear visualization. The ball, offensive team's players, and the defensive team's players are in green, red, and blue, respectively. In addition, the moving direction of each trajectory is from transparent to opaque.

that the defensive plays generated by CNN+WO (Figure 2 (b)) fit the situation we just mentioned, as *B2* quickly comes over to cover *A4*, giving *A2* no open space to shoot.

Figure 2 (c) shows the defensive play generated by CNNONLY. Without the wide open penalty, an open space is created by the offensive players when they execute the pick and roll strategy. Basketball experts would easily notice the unreasonable behavior of *A2* because he does not shoot the ball even though the space has been created. With regard to RNN+WO, as shown in Figure 2 (d), when offensive players *A2* and *A4* execute the pick and roll

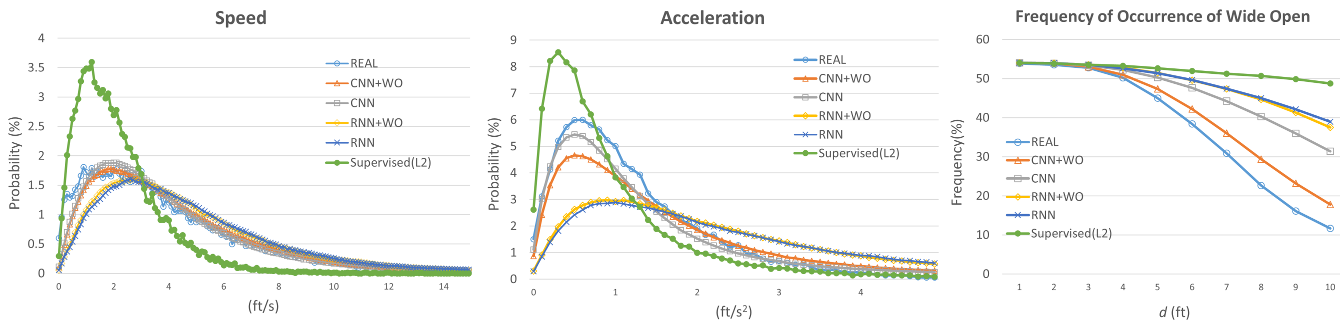

Figure 3: (Left and middle) Frequencies of speed and acceleration of the defensive players' movements. (Right) Frequencies of occurrence of an open space under the condition of d feet. In this chart, we only analyze ball handlers.

Model	Speed (ft/s)		Acc (ft/s^2)	
	Mean	SD	Mean	SD
Real	3.85	2.85	1.28	1.02
CNN+WO	4.07	2.82	2.00	2.27
CNNONLY	3.90	2.76	1.95	2.67
RNN+WO	4.58	2.94	2.94	3.01
RNNONLY	4.68	2.96	3.16	3.35
Supervised	2.15	1.49	0.99	0.97

Table 1: Mean and standard deviation of defensive players' speed and acceleration.

Model	BH		nBH	
	Mean	SD	Mean	SD
Real	8.866	3.716	12.044	5.576
CNN+WO	9.450	3.819	15.093	6.015
CNNONLY	12.163	4.931	15.101	6.024
RNN+WO	14.504	5.742	17.074	7.106
RNNONLY	14.871	5.753	17.104	7.273
Supervised	19.699	5.979	24.308	8.341

Table 2: Mean and standard deviation of distance to defend an offensive player. The distances to defend ball handlers (BH) and non ball handlers (nBH) are measured.

play to create an open space, the defender *B5* gets no help, and has to fight over the pick himself to defend, although wide open penalty is considered in this network. We also found a similar situation when the wide open penalty is not considered (Figure 2 (e)). Finally, the defensive play generated by the supervised model is the least realistic since all defensive players gather at around the restricted area and move slowly in a game. The problem occurs because an offensive play can be defended in different ways. Under this circumstance, the supervised model tends to fit player positions in a regression manner. Please refer to our supplemental videos (*Results.mp4*, *CNN_VS_RNN.mp4* and *comp_penalty.mp4*) for comparison because players' movements in a basketball game are difficult to visualize by pictorial representation. We also show the defensive plays generated by the network trained at different iterations in *comp_iterations.mp4*.

4.1 Objective evaluation

In NBA, a player's defensive performance is evaluated based on statistic, such as rebounds per game and blocks per game. However, the statistic can only be obtained when games are played. Therefore, we evaluate the generated defensive plays by comparing with real plays in terms of 1) the degree of realism of players' movements, and 2) the fulfillment of defense objectives.

Degree of realism. We compare the real and the generated movements in terms of speed and acceleration, and check if the movements are similar. This evaluation is presented because the generated defensive plays are used to forecast how the opposing team will react when an offensive strategy is applied. If the players'

Model	Speed	Acc	Dis (BH)	Dis (nBH)
CNN+WO	0.51	1.22	0.017	0.051
CNNONLY	0.50	0.43	0.050	0.054
RNN+WO	0.67	3.02	0.083	0.092
RNNONLY	0.74	3.13	0.082	0.093
Supervised	1.84	2.82	0.171	0.234

Table 3: We measured the similarity between the real and the generated defensive plays by the Hausdorff distance, in terms of players' speed and acceleration, and the distance to defend ball handlers and non- ball handlers.

movements are unrealistic, the generated plays would be meaningless. Figure 3 left and middle show the statistic of speed and acceleration. The horizontal and vertical axes indicate speed/acceleration and frequency, respectively. We also showed the means and the standard deviations in Table 1. As indicated by the statistic, although players' movements can be fast, they run slowly in most of the time because they have to save energy. In addition, the lines indicate that players' movements generated by the CNNs are more similar than the RNNs and the supervised model to the real players' movements. To prevent visual misleading, we verified the result by the Hausdorff distance (Table 3) between the statistic of the real and the generated players' movements.

Fulfillment of defense objectives. We attempt to understand whether the generated strategies can effectively minimize the opportunity of making scores by the opposing team. Since players

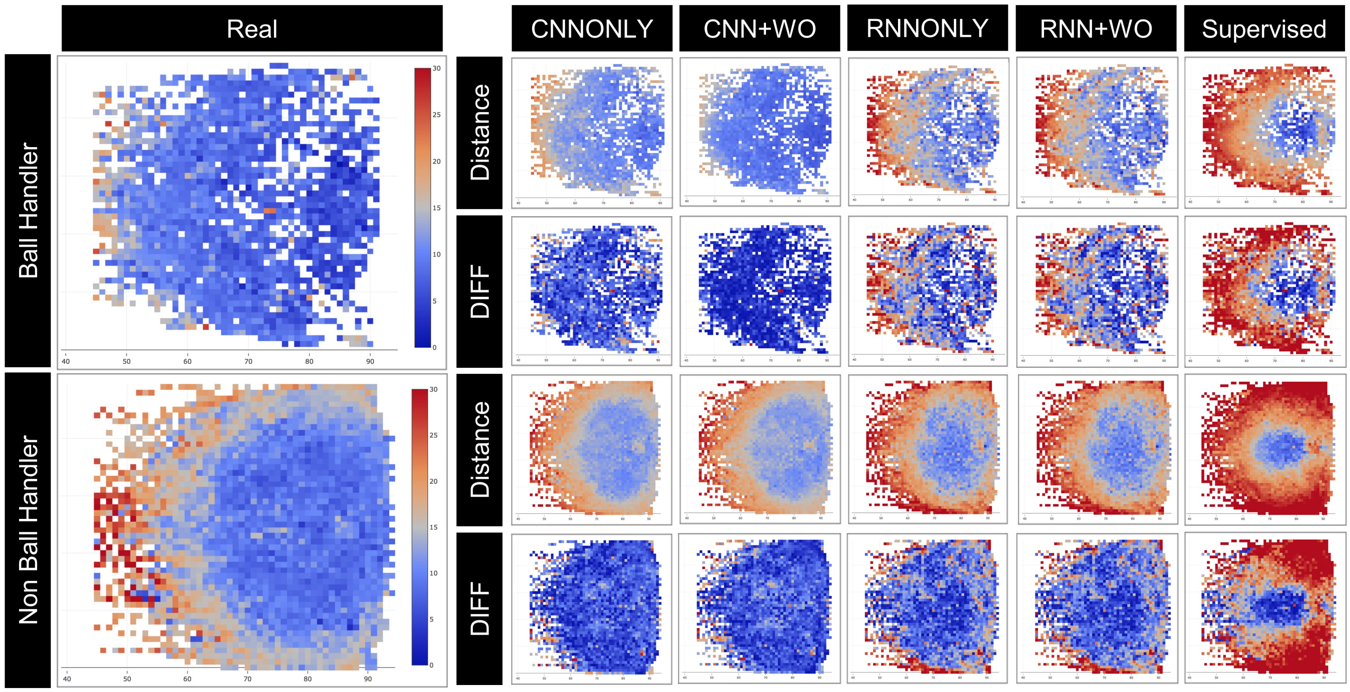

Figure 4: We visualize the mean distance to defend an offensive player with respect to the offensive player's position. The color from blue to red indicates the distance from short to long; whereas white means the position has no data. Because of different behaviors of ball handlers and non- ball handlers on the court, we analyze the distances of these two types of players separately. We also visualize the difference of distance for comparison.

would quickly run out of energy if they keep playing man to man defense, expecting a zero occurrence of wide open is problematic. Therefore, considering that the real data were collected from the best players worldwide, we compare the frequency of the occurrence and the positions of wide opens between the real and the generated plays. In other words, the occurrence of wide open is allowed if it is far away from the basket or to non-ball handlers.

For each position on the basketball court, if there is an offensive player, we measured the mean distance between the player and his closest front defender. The formula used to compute the distance is shown in Equation 2. Note that the defender who stands at the back of an offensive player will have a large distance by this formula. Since a ball handler can shoot the ball directly if there is an open space nearby, whereas the remaining offensive players cannot, we analyzed the defensive behaviors for these two types of players separately. Table 2 shows the mean distances and the standard deviations. We also show the heat maps in Figure 4 to convey the mean defensive distance at each position. The heat maps reveal that, ball handlers are always closely defended; but non- ball handlers are not if they are far away from the restricted area. This phenomenon is reasonable because a non- ball handler has to obtain the ball before he shoot, and making a three point shot is more difficult than making a two point shot. In addition, the heat maps indicate that real defenders perform the best; and the wide open penalty can effectively reduce the distance between the defender and the ball handler. To prevent misleading, we visualize the deviation of the defensive distances between the real and the generated plays in Figure 4. We also verify the similarity by the Hausdorff distance and show the results in Table 3.

Wide opens in the real basketball plays seldom occur because the dataset was recorded from NBA players. Accordingly, we determine the frequency of open space occurrences in the real and generated defensive plays. Suppose that a ball handler can shoot if the closest front defender is d feet away from him. We determine the frequency of open space occurrences under different conditions d. Figure 3 right shows the result. It is not surprising that the frequency of occurrence in the real plays is the lowest since NBA players are world class professional. The defensive plays generated by our CNN+WO is the second best.

Summary. We compared the similarity of the real and the generated defensive plays based on the players' moving speed and acceleration, the distance to defend offensive players, and the frequency of wide open occurrences. An interesting finding is that, defensive players in the real plays run the slowest but perform the best, as indicated in the statistic shown in Tables 1 and 2, and Figure 3. In other words, the real players adopted a smarter strategy than the generated players in defending the offensive team. Among the networks, the performance of the CNN GAN is much better than that of the RNN GAN and the supervised CNN. In addition, the wide open penalty overall has positive effects to the results. The mean distance to ball handlers becomes small and the frequency of open space occurrences are greatly reduced. Although the non-ball

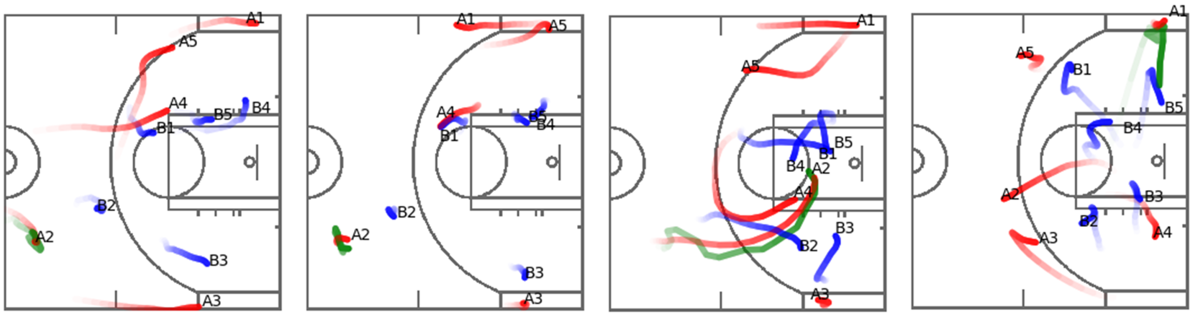

Figure 5: Although defensive players B4 and B5 seem to wonder aimlessly, the play generated by our system was typically a zone defense strategy. Since the participants in PRO identified the strategy, they considered the generated play to be real. From left to right are the four consecutive segments of the game play, in which the offensive players, defensive players, and the ball are visualized in red, blue, and green, respectively. The moving direction of each trajectory is from transparent to opaque.

Group	Q1_1	Q1_2	Q1_3	Q1_4	Q1_5	Q1_6	Q2_1	Q2_2	Q2_3	Q2_4	Q2_5	Q2_6
PRO	0.47	0.12	0.59	0.82	0.88	0.76	0.71	0.94	0.41	0.94	0.71	0.65
FAN	0.30	0.35	0.70	0.78	0.70	0.74	0.65	0.78	0.43	0.78	0.78	0.57
ORD	0.47	0.53	0.68	0.79	0.58	0.58	0.53	0.63	0.58	0.53	0.42	0.79

Table 4: The correct rates of questions answered by the participants in PRO, FAN, and ORD, respectively. An interesting finding was that, the participants in PRO and FAN could easily identify or be fooled by the generated defensive plays because of the relatively high and low rates. In contrast, the correct rates of the participants in ORD did not vary considerably.

handlers are loosely defended caused by the penalty, the influence to realism is not considerable because they do not have a ball.

The evaluation presented in this paper is mainly achieved by comparing statistics. Another possible way to evaluate a model is letting the models to compete in a simulated environment. Therefore, we are thinking about training a model that can generate offensive strategies, and let the offensive and defensive models to play basketballs. In other words, the effectiveness of offensive and defensive strategies can be simply evaluated by the scores. We will work on this direction in the near future.

4.2 User study

Procedure. We conducted a user study to obtain subjective measurements of the defensive plays generated by CNN+WO. Specifically, we randomly selected 12 offensive conditions, which were about 15-20 seconds, from the testing dataset, and then merged them with real or the generated defensive plays. We applied the plays to create a questionnaire with two sessions. In the first session, there were 6 questions, in which the participants were shown a game play and were asked to judge whether the play was real or fake. Half of the defensive plays in the questions were real and the other half were generated by CNN+WO. In the second session, there were 6 questions, too. We compared our generated defensive plays with real data. The participants were shown two game plays (top and bottom in random order) that had the same offensive condition, and were asked to select which play was real.

Results. 59 participants joined our user study. Among the participants, 17 of them were players on the varsity/department team, 23 of them were NBA fans, and 19 of them were ordinary people but familiar with basketball rules. The participants in the three groups were denoted by PRO, FAN, and ORD, respectively. Ideally, the correct rate in the conducted user study would be 0.5 if the real and the generated plays were indistinguishable, because there were only two choices in each question. Table 4 shows the correct rates of the questions answered by the participants. As indicated, the participants could more or less distinguish the real and our generated defensive plays, since the correct rate on average (session 1: 0.60, session 2: 0.66) was higher than the rate of random guess (i.e., 0.5).

In the first session, the participants in PRO (M=0.61, SD=0.28) did not outperform the participants in FAN (M=0.59, SD=0.21) and ORD (M=0.61, SD=0.11). Particularly, most of the participants in PRO was fooled by the second question, in which the defensive play generated by our network was a 2-3 zone defense. We visualize the game play in Figure 5. That was a defensive formation, where defenders were designated to defend certain areas of the court that makes it harder to pass the ball into the restricted area. Compared to a one-to-one defense, where the defenders follow an assigned offensive player around the court, players can save energy by playing a zone defense. Since the participants in PRO could identify such formations, they considered that the generated defensive play was real. However, the participants in FAN and ORD could not identify the formation. They simply thought that the defensive play was fake because two players seem to wonder aimlessly in the restricted area. In addition to the correct rate, we found that 73% and 47% of the real and the generated plays were identified as real, respectively.

In the second session, the participants in PRO (M=0.73, SD=0.20) could answer a lot more questions correctly than the participants in FAN (M=0.67, SD=0.14) and ORD (M=0.58, SD=0.12). We also

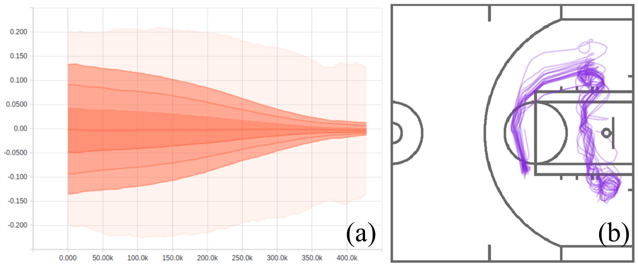

Figure 6: (a) The distributions of weights in the linear projection layer for reshaping the latent noise vector z. The horizontal and vertical coordinates indicate the number of iterations and weight values, respectively. Let μ and σ be the mean and the standard deviation of the weights, respectively. The distributions $\mu \pm 0.5\sigma$, $\mu \pm \sigma$, $\mu \pm 1.5\sigma$ are visualized from the most saturated to the least saturated orange, respectively. (b) The trajectories of a defensive player generated by 20 different latent vectors z. As indicated, the weights become smaller as the number of iteration increases, and the diversity of player trajectories falls away.

observed that the correct rates of the three groups were very different, and the rates could reflect their familiarity with basketball. In other words, when both the real and the generated plays were compared side by side, the participants could make better decisions if they knew more about basketball.

Summary. We have subjectively evaluated our system by conducting a user study. The mean correct rate of the questions indicated that the defensive plays generated by our system were realistic. Particularly, our system could fool participants in ORD because of the low correct rate and the low standard deviation of the rate. Although several plays generated by our system could be easily identified as fake by the participants in FAN and PRO, sometimes these participants made mistakes because the generated plays contain strategies. The generated defensive players were not just chasing the offensive players, but defending the players with a strategy.

4.3 Diversity of the generated trajectories

We evaluated the diversity of the results generated by our network. Specifically, given the same condition y but different latent noise vectors z, we would like to know whether the trajectories $G(z|y)$ were different. Figure 6 (b) shows the trajectories of a player generated by various latent vectors z. As indicated, although the trajectories were different, they were visually similar. To figure out whether the latent vectors z had sufficient influence to the results, we visualized the weights that were used to linearly project vector z to a feature map. Figure 6 (a) shows that the weights gradually degenerated to zero when the number of training iteration increased. In other words, the network considered that noise vectors were not helpful in generating realistic defensive plays. Similar phenomenon also appears in another applications such as image to image translation [12] and video prediction [15].

4.4 Attributes of the dataset

The SportVU datasets contain player positions in each frame. Although the offensive and defensive players can be distinguished, positions of the players are unknown. The generator has no knowledge about point guard, forward, and center when it generates defensive plays. Experimentally, we did not find a generated player whom plays two roles in a game. If a player plays as a point guard, he would not play as a center. We reason that the training dataset contains no (or very few) such examples. However, we believe that addition information, such as profiles of the players, would be beneficial to training the network. In our current system, we assume that all players have an identical weight and height, and can run equally fast on the court. If the assumption is not the case, the generated defensive plays can be less reliable.

4.5 Limitations

Although the defensive plays generated by the presented network look realistic in most of the situations, abnormal player movements may occur occasionally. For example, defensive players may switch positions back and forth or stay very close to each other. Even participants who were not familiar with basketball games can notice the unnatural behaviors. We show the failure examples in our supplemental results (*failure example.mp4*). Therefore, one of our future goals is to improve the result quality by preventing the occurrence of such abnormal movements. Another limitation is the way to input the offensive condition. Users would prefer specifying the offensive condition by sketching rather than drawing precise moving trajectories on a tactical board. We plan to provide an intuitive graphical interface for users in the near future.

5 CONCLUSIONS AND FUTURE WORKS

We have presented a generative adversarial network to simulate defensive plays corresponding to the ball and the movements of players on the offensive team. The generated results are visually realistic, which have been verified by the conducted user study. Although players and coaches can simulate how the opposing team will react to an offensive strategy by using our system, there is still space for improvement. For example, the system does not consider profiles, skillsets, and performances of the players, when it simulates defensive plays. While the way to defend a tall and a short players could be different, considering additional information is needed. In addition, the simulation can only mimic players' movements from the historical data. Considering that the reinforcement learning allows players in the two teams to play games, we plan to apply the technique to simulate basketball games according to players' abilities. We also attempt to explore the possibility in inventing new basketball strategies by this new technique.

ACKNOWLEDGMENTS

We thank anonymous reviewers for their insightful comments and suggestions. We are also grateful to Prof. Wen-Hsiao Peng for the valuable discussions, and all the participants who joined the user study. This work is partially supported by the Ministry of Science and Technology, Taiwan, under Grant No. 104-2221-E-009 -041 -MY3 and 105-2221-E-009 -135 -MY3.

REFERENCES

[1] Martin Arjovsky and Léon Bottou. 2017. Towards principled methods for training generative adversarial networks. *stat* 1050 (2017), 17.

[2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein gan. *arXiv preprint arXiv:1701.07875* (2017).

[3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *stat* 1050 (2016), 21.

[4] Ching-Hang Chen, Tyng-Luh Liu, Yu-Shuen Wang, Hung-Kuo Chu, Nick C Tang, and Hong-Yuan Mark Liao. 2015. Spatio-Temporal Learning of Basketball Offensive Strategies. In *Proceedings of ACM international conference on Multimedia*. 1123–1126.

[5] Alexander Franks, Andrew Miller, Luke Bornn, and Kirk Goldsberry. 2015. Counterpoints: Advanced defensive metrics for nba basketball. In *Proc. of MIT Sloan Sports Analytics Conference*.

[6] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proc. of the International Conference on Artificial Intelligence and Statistics*. 249–256.

[7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.

[8] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. 2017. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028* (2017).

[9] Mark Harmon, Patrick Lucey, and Diego Klabjan. 2017. Predicting Shot Making in Basketball Learnt from Adversarial Multiagent Trajectories. *arXiv preprint arXiv:1609.04849* (2017).

[10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity mappings in deep residual networks. In *European Conference on Computer Vision*. Springer, 630–645.

[11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2016. Image-to-image translation with conditional adversarial networks. *arXiv preprint arXiv:1611.07004* (2016).

[12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-To-Image Translation With Conditional Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1125–1134.

[13] Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the International Conference for Learning Representations*.

[14] Justin Kubatko, Dean Oliver, Kevin Pelton, and Dan T Rosenbaum. 2007. A starting point for analyzing basketball statistics. *Journal of Quantitative Analysis in Sports* 3, 3 (2007).

[15] Michael Mathieu, Camille Couprie, and Yann LeCun. 2015. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440* (2015).

[16] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).

[17] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2536–2544.

[18] Jaime Sampaio, Eric J Drinkwater, and Nuno M Leite. 2010. Effects of season period, team quality, and playing time on basketball players' game-related statistics. *European Journal of Sport Science* 10, 2 (2010), 141–149.

[19] Jaime Sampaio, Manuel Janeira, Sergio Ibáñez, and Alberto Lorenzo. 2006. Discriminant analysis of game-related statistics between basketball guards, forwards and centres in three professional leagues. *European journal of sport science* 6, 3 (2006), 173–178.

[20] Rajiv Shah and Rob Romijnders. 2016. Applying deep learning to basketball trajectories. *arXiv preprint arXiv:1608.03793* (2016).

[21] Kuan-Chieh Wang and Richard Zemel. 2016. Classifying NBA offensive plays using neural networks. In *Proc. of MIT Sloan Sports Analytics Conference*.

[22] Stephan Zheng, Yisong Yue, and Jennifer Hobbs. 2016. Generating Long-term Trajectories Using Deep Hierarchical Networks. In *Advances in Neural Information Processing Systems*. 1543–1551.