UT Austin Villa: Deep Learning for Passing Strategy

Patrick MacAlpine, Brahma Pavse, Faraz Torabi, and Peter Stone

Department of Computer Science, The University of Texas at Austin
{patmac,brahmasp,faraztrb,pstone}@cs.utexas.edu

This presentation discusses work by the UT Austin Villa team toward learning where to kick the ball for passing. Currently the UT Austin Villa team decides where to kick the ball for a pass by using a hand-coded value function that assigns a score or value to every possible location to kick the ball to from the ball's current location [1]. Each potential kick location is given a score according to Equation 1, and the location with the highest score is chosen as the location to kick the ball to. Equation 1 rewards kicks for moving the ball toward the opponent's goal, penalizes kicks that have the ball end up near opponents, and also rewards kicks for landing near a teammate. All distances in Equation 1 are measured in meters.

$$\texttt{score}(\mathit{target}) = \begin{aligned}[t] & -\|\mathit{opponentGoal} - \mathit{target}\| \\ & \forall \mathit{opp} \in \mathit{Opponents}, -\max(25 - \|\mathit{opp} - \mathit{target}\|^2, 0) \\ & + \max(10 - \|\mathit{closestTeammateToTarget} - \mathit{target}\|, 0) \end{aligned} \tag{1}$$

Figure 1 shows a graphical representation of the values Equation 1 assigns to potential locations to kick the ball to.

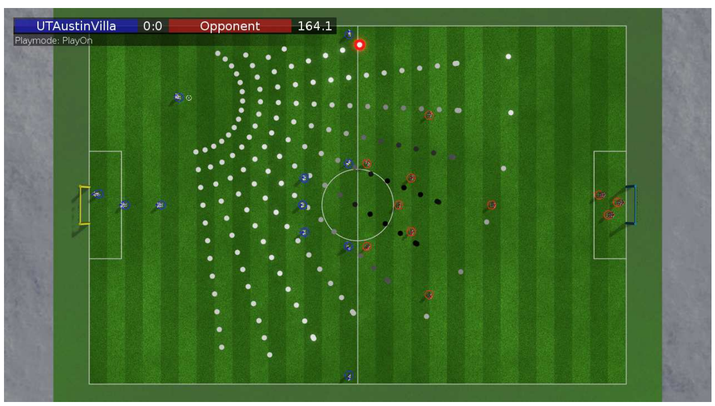

Figure 1. Potential kick target locations with lighter circles having a higher score. The highest score location is highlighted in red.

Instead of using a hand-coded scoring function (Equation 1), UT Austin Villa has now trained a deep neural network with game data to determine values for kicking the ball to different locations on the field. Training such a network involves the following three steps:

1. Play games and record scenarios where players kick the ball. Each scenarios is represented by the locations of all the players and ball on the field as well as all the potential discrete locations that the player on the ball can kick the ball to.
2. Determine the value for each potential kick location for each scenario by restoring the game state of the scenario and kicking the ball to each potential kick location ten times. The percentage of times a team is able to score a goal within 20 seconds of kicking the ball is recorded as the value for kicking the ball to each potential location.
3. Train a neural network using backprop to represent the value for each kick location using the data from the previous step. The locations of all the players, the ball, and location to kick the ball are given as input to the network, and the output of the network is the estimated value of kicking the ball to the given location.

Data from playing 1000 games against opponents shows some improvement in performance (average goal difference) when using the trained deep neural network instead of the hand-coded function for determining the value of kicking the ball to different locations.

1. P. MacAlpine, J. Hanna, J. Liang, and P. Stone. UT Austin Villa: RoboCup 2015 3D simulation league competition and technical challenges champions. In L. Almeida, J. Ji, G. Steinbauer, and S. Luke, editors, *RoboCup-2015: Robot Soccer World Cup XIX*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2016.