Where to Score in AI World Cup Football?

Marius Birkhoff[1], Kushal Prakash[2], and Rafael Bidarra[3]

[1,2,3]Delft University of Technology

Juli 1, 2021

Abstract

To push the boundaries of technology, the world cup football for robots, RoboCup, is organized on a yearly basis since 1997. To push the boundaries of artificial intelligence, a simulated version of the RoboCup, AI World Cup Football, is arranged yearly from 2017. This requires skillful attackers, defenders and goalkeeper. A large part of having a competent goalkeeper, is having a well-anticipating goalkeeper. This research will improve the goalkeeper's level of anticipation by finding its own weak spots.

By training an attacker against this goalkeeper, the weak spots can be determined. Based on this data the goalkeeper can estimate in which region of the goal a real opponent is likely to score and decrease that scoring chance by making a move. Heuristics and a deep neural network are used to estimate how likely the attacker is to make a goal.

The results show that the attacker that uses the deep neural network has a chance of scoring a goal which is about 10 to 15 percentage points higher than a random shooting attacker. The deep neural network attacker is however outperformed by the heuristic based player by about 15 to 20 percentage points.

Taking that into account, embedding the heuristic based attacker into the goalkeeper gives the goalkeeper a good sense of what areas of the goal it should cover to decrease the likelihood of conceding a goal.

1 Introduction

'AI is taking over our jobs!' True or not, it is a phrase that can be heard on a regular basis. However, that is not everything. AI is also taking over our sports. Since 1997 the world cup football for robots RoboCup is organized yearly [1]. To make this more accessible the Korea Advanced Institute of Science and Technology (KAIST) started a digital AI world cup football, without the need of physical robots [2]. It is a simulation with two teams of five agents playing against each other. The teams exist of two attackers, two defenders and a goalkeeper. The rules are different from real-life football, e.g. every player can only run a certain distance per game and players will be sent off for creating deadlocks. Prakash et al. have set up a framework to act as a starting point for the journey towards the world cup [3]. The framework enables loosely coupled development of various aspects. One of those is goalkeeping. Keeping clean sheets as a goalkeeper comes down to two factors: anticipation and reaction. This particular research is interested in enhancing the anticipation of the goalkeeper.

One way to achieve this, is by putting us in the attacking opponent's position, anticipating on what he is most likely to do and acting accordingly.

Thus, the main question this research aims to answer is *What are the best scoring regions of the goal given the position of the goalkeeper?* To answer this question it will be broken down into three subquestions.

(1) What is the chance of scoring a goal when aiming at the best scoring region according to the deep learning network compared to simple heuristics? Given a division of the goal in rows and columns, i.e. a grid, a deep learning network should be able to derive the best region to shoot at. Where 'best region' means the region with the highest probability of scoring a goal. The assessment of the scoring regions can also be done in a heuristic based fashion. The performance of the two different attackers can be assessed by measuring how many goals are scored when constantly picking the 'best regions' according to the respective players.

(2) What is the impact of changing subdivisions of the goal (i.e. altering the grid of regions to aim at)? As mentioned in (1), the goal will be subdivided into rows and columns. That is necessary to keep the problem discrete. This makes the second subquestion arise; does the number of cells in the grid have an impact on the performance?

(3) How much will the chance of scoring a goal increase by training against a better goalkeeper? As this research aims to train a deep learning player to assess the probability of scoring a goal for certain regions, it will be dependent on the skill level of the opponent, the goalkeeper. Therefore it will be interesting to find out if the attacker performs better if it is trained against a more skilled goalkeeper.

Once these three subquestions are answered, the results can be used to constantly inform the goalkeeper about where the opponent should aim to achieve the highest probability of scoring a goal. Based upon this information, the goalkeeper can then make a well considered choice on what action to take.

2 Related work

At the moment of writing, not much research towards the AI World Cup Football has been published. A valid reason for this could be that AI World Cup Football is a competition and not many teams are willing to reveal their tactic or technique.

According to Mozgovoy et al. using team games for AI benchmarking will remain relevant, because multi-agent collaboration is key in these games and is not present in many other types of games [4]. This is reflected in the research that was done by Kim et al. who describe in which manner deep reinforcement learning can improve the level of the complete team in the AI World Cup Football [5]. Using a two-stage training all players are first trained individually and then trained collectively using reinforcement strategies. The results of Kim et al. show that this approach produces better teamwork and better individual work than the traditional centralized training in decentralized execution approach for training multi-agent systems, which was used as a baseline. The difference between the study of Kim et al. and this one is present, considering this research is not necessarily aiming at better collaboration in the multi-agent system, but at improving one player.

A study by Zandsteeg et al. comparing multiple types of shooting mechanisms for the RoboCup, recommends the mechanism using a solenoid [6]. That mechanism is also deemed the most dangerous one and can only be applied in the RoboCup when taking good safety measures. Safety is however not a factor that has to be taken into account in the simulated AI World Cup Football. Therefore the solenoid solution seemed like a good one for precise shooting in the AI World Cup Football. However, the shooting mechanism turned out to be completely built into the game and there is no method to adjust that.

3 Method

All steps are executed within Webots, the simulation engine used by AI World Cup Football, except for training the deep neural network [7].

3.1 Evaluation Strategy

To evaluate the scoring capacity of the attacker (and therefore eventually the anticipation skills of the goalkeeper), the attacker has to repeatedly shoot at the goal from a random position. This is done a hundred times. To make this smooth, all opponents, except the goalkeeper, run towards the sideline to not be disruptive (Figure 1). In the mean time the kickoff is taken. Thereafter the ball possessing attacker moves to a random position on the opponent's half (Figure 2) and fires a shot at goal. After determining whether the shot ended up in the goal or was saved by the goalkeeper, the procedure will be repeated. The score of the method under evaluation is measured as the percentage of shots that ended up in the goal.

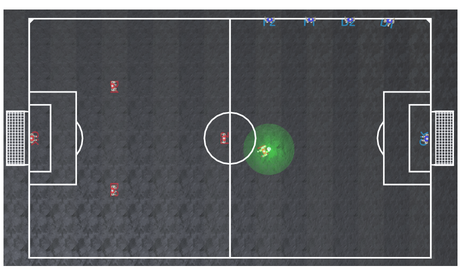

Figure 1: Top view with opponents in blue on the side line.

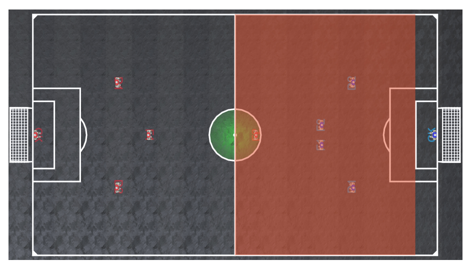

Figure 2: Top view with the red area indicating possible shooting positions.

3.2 Goal Subdivision Algorithm

To end up with a discrete number of targets to aim at, the attacker will subdivide the goal into a grid. The number of desired grid cells can be given as a parameter, which makes answering subquestion 2 possible. Given the number of desired grid cells, the number of rows and columns is its pair of factors such that their sum is minimal:

$$(n_{rows}, n_{columns}) \in \{(x, y);\ x \times y = n_{cells}\},\ s.t.\ x + y\ is\ minimal$$

Figure 3 demonstrates what a 12 cell grid would look like when utilizing the goal subdivision algorithm.

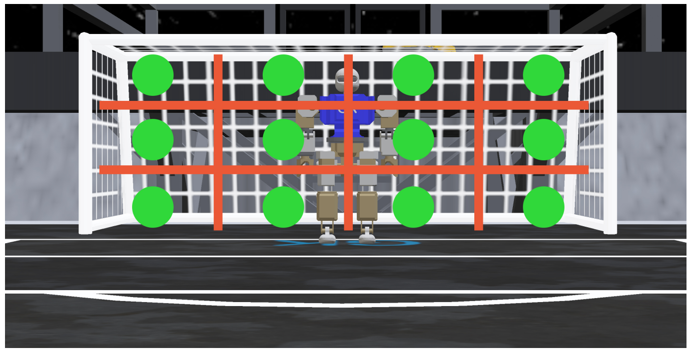

Figure 3: Front view of the goal with all targets (green) for a 12 cell grid (red).

3.3 Heuristic Based Attacker

The heuristic based attacker assigns every target, given by the goal subdivision algorithm, a score. This score is the distance between the goalkeeper and the estimated trajectory of the ball to the target. Put abstractly, it is the Euclidean distance between a point and a line. Additionally, targets receive a bonus for their height. Therefore high shots are always preferred over low shots. An example is given in Figure 4 where the trajectory to a target is drawn in red and the distance past the goalkeeper is displayed as a green line. The score that is given to the particular target of Figure 4 is:

$$\text{Length of green line} * (1 + y_{target})$$

For other targets of the goal, the calculation is the same. With the trivial change that the red line and therefore the green line differ from the given example. During evaluation the attacker shoots at the target with the highest score. By doing this, the heuristic based attacker should always aim at either the top left or the top right corner of the goal.

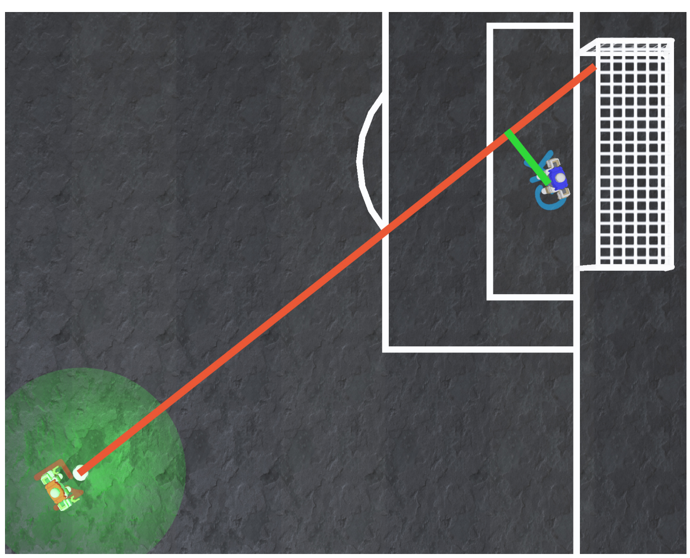

Figure 4: Top view of the distance past the goalkeeper

3.4 Deep Learning Attacker

The deep learning attacker has to be trained. By first generating a dataset and then training a neural network using the dataset, training becomes more versatile; multiple types of networks can be built for experimentation based on the dataset.

Dataset Generation

Generation of the dataset is similar to the evaluation strategy described in Section 3.1. However, attacker 2 does not only move to a random position on the pitch, it also fires the shots at random targets. After a shot is fired, a row is added to the dataset, containing:

- x, y, z coordinates of the player
- angle of the player
- x, y, z coordinates of the ball
- x, y, z coordinates of the goalkeeper
- angle of the goalkeeper
- x, y, z coordinates of the target
- reward, 1 for a goal; 0 for a miss

Deep Neural Network Training

This dataset was used to train a deep neural network outside Webots by making use of PyTorch and PyTorch Lightning [8, 9]. The deep neural network has 8 input values and 2 output values. Section 4 explains this setup in greater detail.

Target Choice

The evaluation strategy requires the attacker to shoot from a certain position on the pitch. The deep neural network attacker will make its choice by forwarding all the targets through the pretrained network and aiming at the target with the highest score and certainty.

4 Dive into the deep neural network

As described in section 3.4 the deep neural network that was used has a vector of 8 elements as input and a vector of 2 elements as output. This setup contains two non-trivial parts. Firstly, giving scores to a regions of the goal can seemingly be achieved by 1 output value. Secondly, section 3.4 described the generation of 14 input and 1 output (the reward) values, which does not comply with the stated input vector of 8 elements. This section addresses why the deep neural network has this structure.

4.1 Two output values

The network was was designed to one-hot-encode the reward, turning a regression problem into a classification problem. The two labels to classify are *goal* and *no goal*. Since the two output values will never exactly be ones and zeroes in practice, all scoring regions can actually be evaluated and compared to each other. To select the target, the following formula is applied:

$$t_{selected} = t \in \ targets \ s.t. \ \hat{y}_t^{goal} - \hat{y}_t^{no\ goal} = max_{t_n \in targets}(\hat{y}_{t_n}^{goal} - \hat{y}_{t_n}^{no\ goal})$$

Where,

$targets$ is the set of all targets/regions of the goal to shoot at

$\hat{y}_t^{goal}$ is the first output of the neural network for target t

how likely a shot at this target is to become a goal

$\hat{y}_t^{no\ goal}$ is the second output of the neural network for target t

how likely a shot at this target is to not become a goal

4.2 Eight input values

Not all 14 input values that were collected during the generation of the dataset were actually used to train the neural network. The position (x-, y- and z-coordinate and angle) of the player was left out completely, because this has a high correlation with the position of the ball. The z-coordinate of the ball was also not taken into account, since the ball must be on ground-level to be able to shoot and therefore this value is the same in any scenario. This also holds for the z-coordinate of the goalkeeper; before shooting the goalkeeper always has the same z-coordinate. Thus, out of the 14 collected values, only 8 were part of the training sessions and 6 were disregarded.

The structure of the deep neural network is given in Table 1. All layers are accompanied by a rectified linear unit (ReLU) as activation function, except the output layer which utilizes a Sigmoid activation function to ensure classification scores between 0 and 1. The Adam optimizer that was used has a learning rate of 1e-3.

Table 1: Deep neural network structure

Layer	**Dimension**
Input	8
Hidden 1	64
Hidden 2	32
Output	2

5 Results

The table below contains the results of multiple setups assessed using the evaluation strategy from 3.1. The results are grouped by two of Oude Elferink's goalkeepers that the attacking player was facing [10]. The scoring percentages against these goalkeepers are given below in Table 2.

Table 2: Scoring percentages of distinct attackers against two goalkeepers

	arc goalkeeper			line goalkeeper		
subdivisions \ *attacker*	**random**	**heuristic**	**neural net**	**random**	**heuristic**	**neural net**
12	38%	61%	45%	39%	44%	45%
25	45%	79%	54%	33%	66%	50%
56	35%	77%	58%	37%	74%	40%
100	37%	71%	58%	29%	71%	50%
mean	38.8%	72.0%	53.8%	34.5%	63.8%	46.3%

The weighted scatterplots below show a heatmap of which target the deep neural network attacker aims at when facing the two goalkeepers and having subdivided the goal in 100 regions. Obviously a large margin at the top of the goal is present. This can be explained by the fact that the bottom of the ball is aimed at the target.

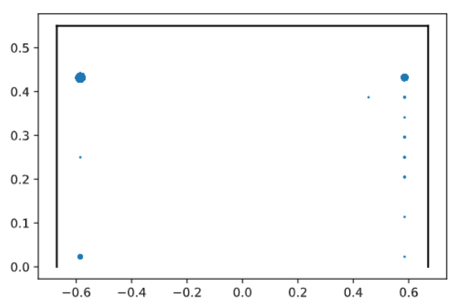

Figure 5: Targeted regions facing the arc goalkeeper (100 samples).

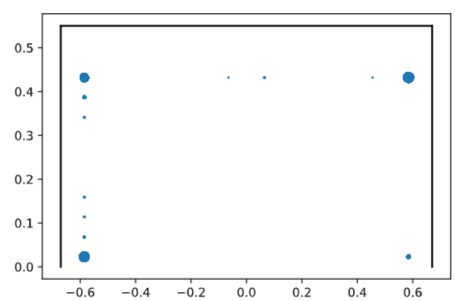

Figure 6: Targeted regions facing the line goalkeeper (250 samples).

6 Responsible research

This research has little to none ethical implications as it only impacts a simulated game. When techniques used in this paper are transferred to physical robots playing against humans, then this has to be reconsidered [11].

Traditionally research is only trusted when others, with the right knowledge, are capable of attaining the same results by redoing the prescribed steps. With the rise of AI, this has become more complicated. A lot of (sensitive) data would have to be published alongside, as Henderson et al. pointed out, hyperparameters, network architectures and random seeds [12]. To mitigate the problem of reproducibility as much as possible, this research attempts to state all important details in section 4. However, it is apparent that these details might not be enough to fully reproduce the results. Partially, this is due to the randomness of the dataset generation. But, when having a large dataset this should not produce major discrepancies. The fact that both the framework and Oude Elferink's goalkeepers are not open source projects unquestionably contributes to the difficulty of reproducing these results [3, 10].

7 Discussion and Future Work

The neural network that was trained is not very stable - the validation loss was much dependent on the training/validation split. Having more training samples could solve this problem. It is expected that the number of goals scored will also increase then. Due to time constraints and difficult training conditions in Webots, it was not possible to generate a larger dataset for this research. It would be interesting to expand the dataset in the future and investigate if the number of scored goals does actually rises and catches up with the heuristic based attacker.

From Figures 5 and 6 it stands out that the deep neural network favours the corners of the goal, especially the upper corners. This is not very surprising given that those regions require both diving and positioning skills of a goalkeeper. However, not all shots are aimed at the top corners, even though that would give better results, looking at the heuristic based player. This is probably a result of the poor training set. Besides that, it is not apparent why the deep neural network decides to fire some shots at the (vertical) middle of the goal near only of the posts. Further investigation could uncover if this is also due to the lack of training samples, asymmetry in the goalkeepers or asymmetry in shooting.

A weak point of this research is that only the opponent's goalkeeper is taken into account whilst deciding where to aim. In a real match, the four other players could also form obstacles for the shot. Considering their positions and movements can then make a difference. Especially given the fact that the attacker will eventually be implemented as the anticipating part of the goalkeeper, it could then serve as a way to make the goalkeeper take decisions based upon actions of its own defenders.

The obvious next step to take is to actually utilize the trained attacker as the anticipating part of the goalkeeper. Once this is implemented, comparing the newly created goalkeeper with the currently available one without anticipation will give valuable insight into whether the anticipation makes a difference in AI World Cup Football. It could also turn out, that in this simulation reaction is far more important than anticipation.

8 Conclusion

This paper's purpose was to find out *what the best scoring regions of the goal are, given the position of the goalkeeper* in AI World Cup Football. The result can afterwards be used to improve the goalkeeper's anticipation. The research was carried out by creating and evaluating three different parameterized attackers that played against two goalkeepers of different levels.

The research shows that a deep neural network can be used to decide where to aim in AI World Cup Football, because the deep neural network is clearly outperforming a random shooting attacker. Playing against the least skillful goalkeeper the random shooter was beaten by the deep neural network by exactly 15 percentage points. When facing the better goalkeeper, this gap decreased to 11.8 percentage points, which is still significant. Moreover, heatmaps show that it is evident that the deep neural network has truly learned where to aim and which targets not to choose.

The chance of scoring a goal using the deep neural network does not exceed that chance when making use of simple heuristics. Only one out of eight evaluations showed that the deep neural network performed somewhat better (1 percentage point). In fact, in six out of eight evaluations the heuristic based attacker surpassed the deep neural network by 15 or more percentage points.

By tweaking the parameter for the number of possible targets in the goal, it has become clear that the optimal number of grid cells to use is dependent on both the level of the goalkeeper as well as the type of attacker. Nonetheless, for all goalkeepers and attackers hold that either 25 or 56 subdivisions result in the highest scoring chances. Thus, the number of subdivisions does definitely have an impact on the percentage of scoring, but it is incorrect to assume that the higher the number of grid cells the higher the chance of scoring.

The deep neural network was trained against the most skillful goalkeeper. Still, it performed better on the less skillful goalkeeper by 7.5 percentage points. This gap is larger than the difference of the random shooting player which only scores 4.3 percentage points more goals facing the less skillful goalkeeper. Therefore the attacker does perform better against less skillful goalkeepers, when trained on a better goalkeeper, but only by a few percent.

The main takeaway of this paper is that the regions of the goal with the highest probability of scoring are the upper left and upper right corners of the goal. Although it should be taken into account that this can differ when defenders are deployed to block shots.

References

[1] Hiroaki Kitano et al. "RoboCup: A challenge problem for AI". In: *AI magazine* 18.1 (1997), pp. 73–73.

[2] *AI World Cup*. URL: `http://aiworldcup.org`. (accessed: 19.04.2021).

[3] Kushal Prakash and Rafael Bidarra. *Framework for AI World Cup*. Delft University of Technology. Fac. EEMCS, 2021.

[4] Maxim Mozgovoy et al. "Team sports for Game AI benchmarking revisited". In: *International Journal of Computer Games Technology* 2021 (2021).

[5] Taeyoung Kim et al. "Two-stage training algorithm for AI robot soccer". In: *arXiv preprint arXiv:2104.05931* (2021).

[6] CJ Zandsteeg and Ir MJG van de Molengraft. "Design of a robocup shooting mechanism". In: *University of Technology Eindhoven* (2005), p. 44.

[7] Olivier Michel. "Cyberbotics Ltd. Webots™: professional mobile robot simulation". In: *International Journal of Advanced Robotic Systems* 1.1 (2004), p. 5.

[8] *PyTorch*. URL: `https://pytorch.org`. (accessed: 24.04.2021).

[9] *PyTorch Lightning*. URL: `https://www.pytorchlightning.ai`. (accessed: 24.04.2021).

[10] Rens Oude Elferink. *Defensive Positioning Strategies For A Goalkeeper In The AI World Cup*. Tech. rep. CSE3000-G40-2. Fac. EEMCS: Delft University of Technology, 2021, p. 8.

[11] Hans-Dieter Burkhard et al. "The road to RoboCup 2050". In: *IEEE Robotics & Automation Magazine* 9.2 (2002), pp. 31–38.

[12] Peter Henderson et al. "Deep reinforcement learning that matters". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32. 1. 2018.