Towards an Optimal Scoring Policy for Simulated Soccer Agents

Jelle Kok, Remco de Boer, Nikos Vlassis, and Frans Groen

Intelligent Autonomous Systems Group, Informatics Institute
Faculty of Science, University of Amsterdam
Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
{jellekok,remdboer,vlassis,groen}@science.uva.nl

Abstract. This paper describes the scoring policy which is used by the agents of the *UvA Trilearn* simulation team. In a given situation this policy enables an agent to determine the best shooting point in the goal, together with an associated probability of scoring when the ball is shot to this point. Our policy is based on an approximate method for learning the relevant statistics of the ball motion which can be regarded as a geometrically constrained continuous-time Markov process.

1 Introduction

The RoboCup Simulation League is based on a soccer simulation system called the *RoboCup Soccer Server* [2]. The *soccer server* provides a realistic multi-agent environment. Various forms of uncertainty have been added into the simulation such as sensor and actuator noise, limited perception and noise in object movement. Since the main purpose of a soccer game is to score goals, it is important for a robotic soccer agent to have a clear policy about whether he should attempt to score in a given situation, and if so, which point in the goal he should aim for. This paper describes the scoring policy that is used by the agents of the *UvA Trilearn* soccer simulation team. In a given situation this policy enables an agent to find the best shooting point in the goal, together with an associated probability of scoring. Using this scoring policy we reached 4th place at *RoboCup-2001* and won the *German Open 2002* with a total score of 110-1 over 10 matches.

2 The Optimal Scoring Problem

The *optimal scoring problem* can be stated as follows: find the point in the goal where the probability of scoring is the highest when the ball is shot to this point in a given situation. A key observation for solving *the optimal scoring problem* is that it can be decomposed into two independent subproblems:

1. Determining the probability that the ball will enter the goal when shot to a specific point in the goal from a given position.
2. Determining the probability of passing the goalkeeper in a given situation.

Since the subproblems are independent, the probability of scoring when shooting at a certain point in the goal is equal to the *product* of the two probabilities.

2.1 Probability that the Ball Enters the Goal

In this section we will show how one can determine the probability that the ball will end up *somewhere* inside the goal when shot at a *specific* goal point. To this end we first need to compute the deviation of the ball from the aiming point. This deviation is caused by the noise which is added to the ball velocity vector in each simulation cycle[2]. Note that the ball motion can be regarded as a *diffusion* process since the position of the ball in each time step diffuses over time. We can treat it as a Markov process because the future development is completely determined by the present state. Finding an analytical solution of the corresponding diffusion process is difficult for two reasons:

- The motion noise which is added by the server is by construction non-white since it depends on the speed of the ball in the previous cycle.
- The process is geometrically constrained since the ball must enter the goal.

This makes an analytical computation of the process statistics non trivial. Therefore, we propose to *learn* the ball motion statistics from experiments. This gives an approximate solution to the problem of computing the statistical properties of a geometrically constrained continuous-time Markov process and we believe that this relatively simple alternative, can be useful in other applications as well[3].

Our solution thus amounts to estimating the *cumulative* noise directly from experiments. To this end, we calculated the deviation of the ball perpendicular to the shooting direction as a function of the distance that the ball had traveled. This function was learned by repeating an experiment in which the ball was placed at even distances between 0 and 32 metres in front of the center of the goal (zero y-coordinate) with a player directly behind it. The player then shot the ball 1,000 times from each distance with maximum power towards the center of the goal. For each instance we recorded the y-coordinate where the ball entered the goal and used these values to compute the sample standard deviation σ of the ball. To a good approximation this function could be represented by

$$\sigma(d) = -1.88 \cdot \ln(1 - d/45) \tag{1}$$

The next step is to compute the distribution of the ball when it reaches the goal line. For this we use a fundamental result from probability theory called the *Central Limit Theorem*. This theorem states that under certain conditions the distribution of the sum of N random variables x_i will be Gaussian as N goes to infinity [1]. This Gaussian g must have a zero mean and a standard deviation $\sigma = \sigma(d)$ from (1). Using this model, we can compute the probability that the ball will end up inside the goal when shot from an arbitrary position on the field perpendicularly to the goal line. This probability equals the area that lies

[2] In the current version of the *soccer server* the ball velocity in cycle $t+1$ is equal to $(v_x^{t+1}, v_y^{t+1}) = 0.94 \cdot (v_x^t, v_y^t) + (\tilde{r}_1, \tilde{r}_2)$ where $\tilde{r}_1$ and $\tilde{r}_2$ are random numbers uniformly distributed over the interval $[-\text{rmax} .. \text{rmax}]$ with $\text{rmax} = 0.05 \cdot \|(v_x^t, v_y^t)\|$.

[3] e.g Brownian motion problems. The term *Brownian motion* is used to describe the movement of a particle in a liquid, subjected to collisions and other forces.

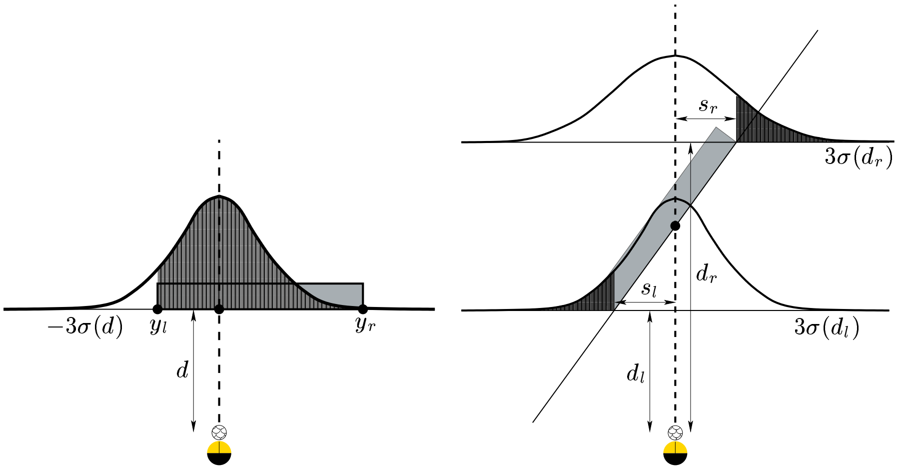

(a) Shooting straight to the goal. (b) Shooting at an angle to the goal.

Fig. 1. Two situations of shooting to the goal (light gray) together with distributions.

under the respective Gaussian density function in between the two goalposts as is shown in Figure 1(a). When the y-coordinates of the left and right goalposts are denoted by y_l and y_r with $y_l < y_r$, this can be computed as follows:

$$\mathrm{P(goal)} = \int_{-\infty}^{y_r} g(y;\sigma)\, dy - \int_{-\infty}^{y_l} g(y;\sigma)\, dy = G(y_r;\sigma) - G(y_l;\sigma) \tag{2}$$

where $G(y;\sigma)$ equals the cumulative distribution function of the Gaussian $g(y;\sigma)$.

Finally, we have to compute the probability that the ball enters the goal when shot at an angle to the goal line (see Figure 1(b)). This case is somewhat more involved than the previous one due to the fact that the noise can cause the ball to travel different distances before it reaches the goal. Since different traveled distances imply different deviations according to (1), the ball distribution along the goal line is no longer Gaussian and this makes an exact calculation of the total probability difficult. However, the key observation is that we want to compute probability *masses* and that for equal masses the particular shape of the distribution that produces these masses is irrelevant and we can compute the probability mass from the identity

$$\mathrm{P(goal)} = 1 - \mathrm{P(not\ goal)} = 1 \text{ - } \mathrm{P(out\ from\ left)} \text{ - } \mathrm{P(out\ from\ right)} \tag{3}$$

where P(not goal) denotes the probability that the ball misses the goal, either going out from the left or right goalpost. This probability mass is easier to compute, to a good approximation, from the tails of the Gaussian distributions corresponding to the goalposts. This is shown in Figure 1(b): when the ball reaches the left goalpost it has *effectively* traveled distance d_l and its corresponding distribution perpendicular to the shooting line is Gaussian with deviation $\sigma(d_l)$ from (1). The probability that the ball misses the goal going out from the left goalpost is approximately[4] equal to the shaded area on the left in Figure 1(b):

$$\text{P(out from left)} \approx \int_{-\infty}^{-s_l} g(y; \sigma(d_l))\, dy \tag{4}$$

where the integration runs up to $-s_l$ which denotes the (negative) shortest distance from the left goalpost to the shooting line. The situation that the ball misses the goal going out from the right post is analogous.

2.2 The Probability of Passing the Goalkeeper

The second subproblem can be stated as follows: given a shooting point in the goal, determine the probability that the goalkeeper intercepts the ball before it reaches the goal line. We propose an empirical method for learning this probability from examples of successful and unsuccessful scoring attempts. In our experiments we have used the goalkeeper of *RoboCup-2000* winner *FC Portugal 2000*, since it appeared to be one of the best available goalkeepers at that time.

To cast the problem into a proper mathematical framework, we note that ball interception can be regarded as a two-class classification problem: given the shooting point in the goal together with the positions of the goalkeeper and the ball (input feature vector), predict which class (intercepting or not) is most probable. Moreover, we are interested in the *posterior* probability associated with the prediction of each class. We performed an experiment in which a player repeatedly shot the ball from a fixed position straight to the goal, while the goalkeeper was placed randomly in different positions relative to the ball. A data set was formed by recording 10,000 scoring attempts, together with a boolean indicating whether the goalkeeper had intercepted the ball. An analysis of the resulting data revealed that the relevant features for classification were the following[5]:

- The absolute angle a between the goalkeeper and the shooting point
- The distance d between the ball and the goalkeeper

These two values form a two-dimensional feature vector on which the classification has been based. The recorded data set is depicted in Figure 2(a) which shows that there is an almost linear discriminant function between the two classes. We determined this discriminant function via linear regression on the boolean class indicator. This procedure is known to give the optimal *Fisher's Linear Discriminant* which has the property that it maximizes the ratio of the *between-group*

[4] There is a small probability that the ball ends up to the right of the left goalpost after travelling an 'illegal' trajectory outside the field. The ball thus actually went out from the left in this case but we neglect this probability in (4).

[5] Principled methods for automatic feature extraction are described in [3].

variance and the *within-group variance* for the two classes [3]. The resulting discriminant function is characterized by the equation

$$u = (a - 26.1) * 0.043 + (d - 9.0) * 0.09 - 0.2 \tag{5}$$

for distance values d between 3 and 15. This can be interpreted as follows: for a new angle-distance pair (a, d), the sample mean $(26.1, 9.0)$ is subtracted after which the inner product with the vector $(0.043, 0.09)$ is computed. The resulting vector is perpendicular to the discriminant boundary which is shifted by the offset -0.2. The pairs (a, d) for which (5) equals zero form the boundary between the two classes. This is plotted by a dotted line in Figure 2(a).

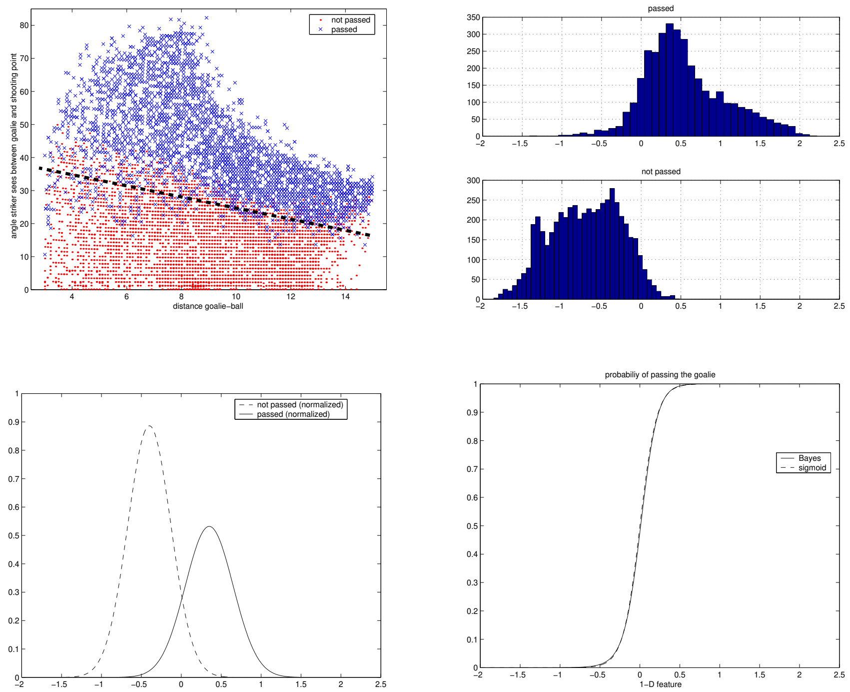

Fig. 2. Data set for goalkeeper interception experiment together with derived statistics.

Projecting all the (a_i, d_i) pairs perpendicularly to the discriminant line via (5) gives a set of one-dimensional points u_i that, to a good approximation, describe the two classes. The histogram class distributions of these points are plotted in Figure 2(b). Instead of trying to model these two distributions parametrically, we note that the relevant range for classification is only where the two histograms *overlap*. It is easy to see that the posterior probability of non-interception will be zero for approximately $u \leq -0.5$, one for $u > 0.5$ and will increase smoothly from zero to one in the interval in between. The posterior probability can thus be represented by a sigmoid function. In the region where the class distributions for interception and non-interception overlap, we fit a univariate Gaussian function on each class as shown in Figure 2(c). For each class C this gives us a Gaussian model for the class-conditional density function $\mathrm{P}(u|C)$. With this model we can compute the posterior probability $\mathrm{P}(C|u)$ for a class C using the Bayes rule

$$\mathrm{P}(C|u) = \frac{\mathrm{P}(u|C)\,\mathrm{P}(C)}{\mathrm{P}(u|C)\,\mathrm{P}(C) + \mathrm{P}(u|\bar{C})\,\mathrm{P}(\bar{C})} \tag{6}$$

which is a sigmoid-like function. Since this is a simple two-class classification problem, $\bar{C}$ refers to the 'other' class in this case, while the prior probability $\mathrm{P}(C)$ for a class C is computed as the proportion of points u_i in the data set which belong to C. In Figure 2(d) we have plotted the posterior probability for the non-interception class, together with the sigmoid approximation

$$\mathrm{P}(\text{pass goalkeeper} \mid u) = \frac{1}{1 + \exp(-9.5u)} \tag{7}$$

which allows for an easy implementation.

2.3 Determining the Best Scoring Point

Having computed the probability that the ball will end up inside the goal (3) and the probability that the goalkeeper will not intercept it (7), the assumption of independence gives the total probability of scoring in a given situation as the product of these two values. This total probability is a bell-shaped function which represents the probability that the ball will enter the goal and which has a valley around the position of the goalkeeper. This curve will always have two local maxima which correspond to the left and right starting point of the valley, which can be located using a simple *hill-climbing* algorithm [4]. The highest one denotes the global maximum and is selected as the best scoring point.

3 Implementation and Demonstration

We have incorporated our scoring policy into the agent's decision loop as follows. When the agent has control of the ball, he first checks whether the probability of scoring is larger than a certain threshold[6]. If the total scoring probability exceeds this threshold then the agent tries to score by shooting the ball with maximum power towards the best scoring point. Otherwise he considers different options, such as dribbling, which he performs when the predicted success rate is high enough. However, when all possible alternatives fail and the agent stands at a close distance to the goal, he decides to shoot to the best scoring point anyhow.

Figure 3 shows two successive situations which were taken from a real match played by *UvA Trilearn*. In Figure 3(a) the player with the ball stands to the

[6] 90% in our current implementation.

left of the goal which is covered well by the opponent goalkeeper. The scoring probability curve which corresponds to this situation is shown in Figure 3(c). The total scoring probability (solid line) is very low for all the points on the goal line[7] and the player with the ball thus rightly decides not to shoot to the goal. However, several cycles later the situation is completely different. The right wing attacker now has a high probability of scoring in the right half of the goal.

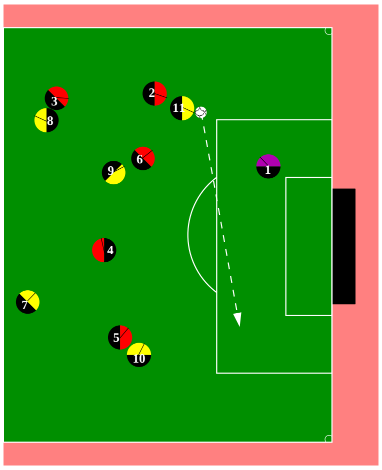

(a) Goalkeeper covers his goal well. A through pass is given.

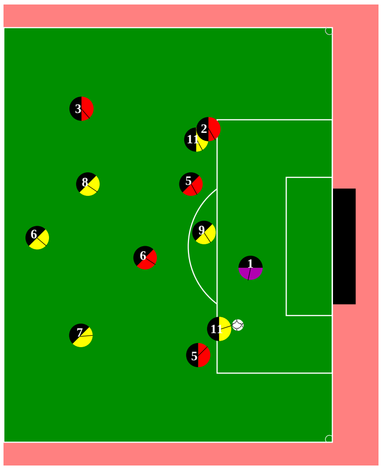

(b) Goalkeeper has been outplayed. Good chance to score.

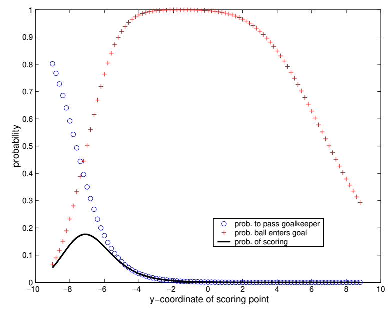

(c) Low scoring probability for all goal points. No shot is attempted.

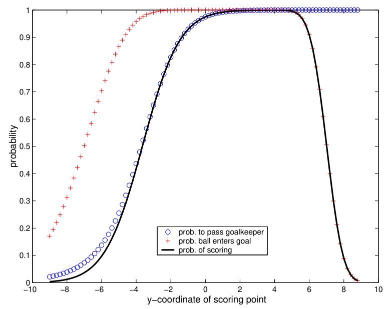

(d) High scoring probability in right half of goal. Shot is attempted.

Fig. 3. Two successive match situations with associated scoring probability curves.

[7] Note that the noise in the ball movement causes a non-zero scoring probability when the ball is shot to a point just outside the goal. In our implementation these points are never selected however, since we only consider points on the goal line for which the single probability of entering the goal is more than 0.7 (=*EnterGoalThr*).

4 Conclusion

In this paper we have described a methodology that allows a simulated soccer agent to determine the probability of scoring when he shoots the ball to a specific point in the goal in a given situation. The single probability that the ball enters the goal (first subproblem) depends on the values of various parameters which control the movement noise of the ball, the size of the goal, etc. The approach is general in the sense that it enables one to 'learn' this probability even when these parameter values change. However, the single probability of passing the goalkeeper (second subproblem) depends on the opponent goalkeeper and different goalkeepers exhibit different behaviors. In our current implementation we have based this probability entirely on the goalkeeper of *FC Portugal*. Since this is a good goalkeeper, the approach is useful against other goalkeepers as well[8].

Ideally however, the probability of passing the goalkeeper should be *adaptive* and the model should incorporate information about the current opponent goalkeeper instead of using that of a particular team. The desired case would be to let the model adapt itself during the game, using little prior information about the current goalkeeper. This is a difficult problem because learning must be based on only a few scoring attempts. It is therefore important to extract the most relevant features and to parametrize the intercepting behavior of the opponent goalkeeper in a compact manner that permits on-line learning.

Acknowledgements

This research is/was supported by PROGRESS, the embedded systems research program of the Dutch organisation for Scientific Research NWO, the Dutch Ministry of Economic Affairs and the Technology Foundation STW.

References

1. W. Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons, Inc., New York, 1957.
2. Ehsan Foroughi, Fredrik Heintz, Spiros Kapetanakis, Kostas Kostiadis, Johan Kummeneje, Itsuki Noda, Oliver Obst, Pat Riley, and Timo Steffens. RoboCup Soccer Server User Manual: for Soccer Server version 7.06 and later, 2001. At http://sourceforge.net/projects/sserver.
3. B.D. Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, Cambridge, U.K., 1996.
4. Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ, 1995.
5. Peter Stone, Patrick Riley, and Manuela Veloso. Defining and Using Ideal Teammate and Opponent Agent Models. In *Twelfth Innovative Applications of AI Conference (IAAI-2000)*, 2000.

[8] A different approach is to base the probability on the *ideal* goalkeeper performance using the known model of the world dynamics [5]. But in this case it is not straightforward to create a continuous probability function for all different situations.