IAS technical report series, nr. IAS-UVA-01-06

Towards an optimal scoring policy for simulated soccer agents

Jelle Kok, Remco de Boer and Nikos Vlassis
Computer Science Institute
Faculty of Science
University of Amsterdam
The Netherlands

This paper describes the implementation of a scoring policy that was used by the agents of the *UvA Trilearn 2001* soccer simulation team during the RoboCup-2001 robotic soccer world championship. In a given situation this policy enables agents to determine the best shooting point in the goal, together with an associated probability of scoring when the ball is shot to this point. It turns out that this probability depends both on the position and the angle of the ball with respect to the goal, and also the position of the goalkeeper relative to the striker. We describe the underlying statistical framework for computing these probabilities, show results, and briefly discuss possible extensions.

Contents

Intelligent Autonomous Systems
Computer Science Institute
Faculty of Science
University of Amsterdam
Kruislaan 403, 1098 SJ Amsterdam
The Netherlands

tel: +31 20 525 7461
fax: +31 20 525 7490

http://www.science.uva.nl/research/ias/

Corresponding author:

Jelle Kok
jellekok@science.uva.nl

1 Introduction

The Robot World Cup (RoboCup) Initiative is an attempt to foster AI and intelligent robotics research by providing a standard problem where a wide range of technologies can be integrated and examined [2]. The goal of RoboCup is to have a team of fully autonomous humanoid robot players, which by 2050 can win a soccer game against the winner of the most recent world cup for human players. In order to achieve this goal, the RoboCup organization has introduced several leagues which each focus on different abstraction levels of the problem. One of these leagues is the *RoboCup Simulation League*. This league is based on a soccer simulation system called the *RoboCup Soccer Server* [1]. The soccer server provides a realistic multi-agent environment in which everything happens in real time. Various forms of uncertainty have been added into the simulation such as sensor and actuator noise, limited perception and noise in object movement. One of the advantages of the soccer server is the abstraction made, which relieves researchers from having to handle robot problems such as object recognition and movement. This abstraction makes it possible to focus on higher level concepts such as learning, opponent modelling and strategic reasoning.

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. In the remainder of this paper we describe the implementation of a scoring policy that was used by the agents of the *UvA Trilearn 2001* soccer simulation team during the *RoboCup-2001* robotic soccer world championship. In a given situation this policy enables agents to determine the best shooting point in the goal together with an associated probability of scoring when the ball is shot to this point. This scoring policy was one of the aspects of our team, which reached fourth place at the *RoboCup-2001* world cup.

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. This is not a straightforward problem. The reason for this is that the total number of possible situations is extremely large and that different variables can be decisive for different situations. Furthermore, the problem depends on many uncertain factors. For example, the noise in the ball motion can never be exactly predicted and will be different for different distances that the ball travels. For finding the optimal scoring point by iterating over all possible points in the goal, one will thus have to take many different functions into account, since the distance from the shooting position to the scoring point will be different for each point in the goal. On top of this, the behavior of the opponent goalkeeper cannot be easily predicted but is an important factor for solving the problem. As a result, no simple analytical solution to the problem exists and one has to look for different methods.

A key observation for finding the solution to the optimal scoring problem is that it can be decomposed into two independent subproblems:

1. Determine the probability that the ball will enter the goal when shot to a specific point in the goal from a given position.

2. Determine the probability of passing the goalkeeper in a given situation.

Since the two subproblems are independent, the probability of scoring when shooting at a certain point in the goal is equal to the *product* of these two probabilities. To find the solution to the scoring problem, we thus have to find the solutions to these two subproblems and combine them.

Before discussing the two subproblems, we mention a number of simplifying assumptions that we have made in our solution to the scoring problem. Firstly, we assume that the ball is always shot with maximum power giving it an initial velocity of 2.7 (distance units per simulator cycle), which is the common case in practice. Secondly, we have chosen to neglect the possibility that other players besides the goalkeeper are blocking the path to the goal. We have done this because the goalkeeper's superior interceptive capabilities make passing him the main objective. Finally, in our experiments the ball was always shot from distances smaller than 32 from the target point in the goal (approximately 1/3 of the total field length). Since the distance that the ball will travel when it is shot with maximum power equals 45 (neglecting the movement noise), it will never be the case that the ball comes to a halt before it has reached the goal. It is fairly straightforward to relax these assumptions and extend the method appropriately.

2.1 Subproblem 1: Probability that the Ball Enters the Goal

If there were no noise in the movement of the ball, it would always enter the goal when shot to a point in the goal from a given position. However, due to the limited goal width and the noise introduced by the server, the ball may miss the goal. We will show how one can determine the probability that the ball will end up *somewhere* inside the goal when shot at a *specific* point.

First we need to compute the deviation of the ball from the aiming point. This deviation is caused by the noise which is added to the velocity vector of the ball in each simulation cycle. Although the underlying noise statistics are known from the server implementation,[1] an analytical computation of the cumulative noise is not trivial: the ball motion can be regarded as a nonstationary Markov process whose statistics can be computed by solving a corresponding Fokker-Planck equation [3]. The complication arises from the fact that the added noise in each cycle depends on the speed of the ball in the previous cycle, making the noise non-white.

A simple alternative, which also allows also for easy adaptation when the server noise parameters change, is to estimate the cumulative noise directly from experiments. To this end, we computed the deviation of the ball perpendicular to the shooting direction as a function of the travelled distance. This function was learned by repeating an experiment in which a player was placed at even distances between 0 and 32 in front of the center of the goal (zero y-coordinate) and shot the ball 1000 times from each distance perpendicularly to the goal line. For each instance we recorded the y-coordinate of the point on the goal line where the ball entered the goal. From these values we computed the sample standard deviation of the ball.[2]

As expected, we saw that the deviation σ of the ball was different for each distance d. We empirically found that, to a good approximation, the standard deviation of the ball perpendicular to the shooting direction was a monotone increasing function

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

for distance values d between 0 and 32 and where $\ln(\cdot)$ is the natural logarithm. This function, together with the recorded deviation values, is plotted in Figure 1. The surprisingly simple formula of σ as a function of d implies that the deviation of the process increases linearly with time.[3] This contrasts with most Brownian motion problems where $\sigma(t) = O(\sqrt{t})$. This difference can mainly be attributed to the non-white motion noise.

The next step is to compute the distribution of the ball when it reaches the goal line. We note that the deviation of the ball is caused by noise that is added in each cycle. Under conditions,

[1] The ball velocity vector (v_x^{t+1}, v_y^{t+1}) in cycle $t+1$ is equal to $0.94 * (v_x^t, v_y^t) + (\tilde{r}_1, \tilde{r}_2)$ where $\tilde{r}_1$ and $\tilde{r}_2$ are random numbers uniformly distributed in [-rmax, rmax], with rmax $= 0.05 * ||(v_x^t, v_y^t)||$.

[2] The sample standard deviation for n zero-mean points x_i is $\sigma = (\frac{1}{n}\sum_{i=1}^{n} x_i^2)^{1/2}$.

[3] This can be easily seen by solving the differential equation of the forward motion of the ball (ignoring the noise) and expressing time as a function of travelled distance.

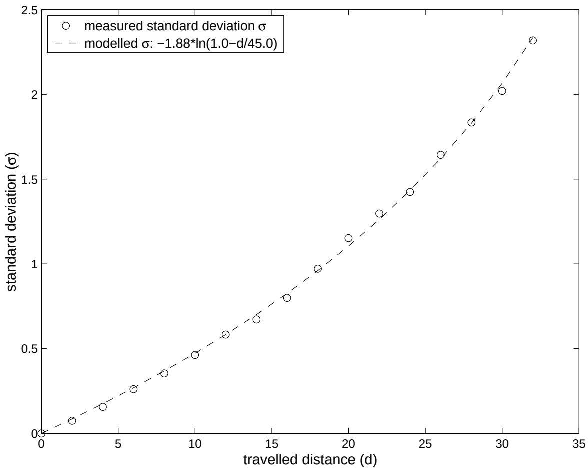

Figure 1: The standard deviation of the ball vs. the travelled distance.

the Central Limit Theorem guarantees that the cumulative noise will be approximately Gaussian [3]. Moreover, this Gaussian must have zero mean and standard deviation $\sigma = \sigma(d)$ from (1)

$$g(y;\sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{y^2}{2\sigma^2}\right) \tag{2}$$

Having such a 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 under the respective Gaussian density function and in between the two goalposts (see Figure 2(a)). When the y-coordinates of the goalposts are denoted by y_1 and y_2 with $y_1 < y_2$, this can be computed as

$$\mathrm{P}\{\text{goal}\} = \int_{y_1}^{y_2} g(y;\sigma)\,dy = \int_{-\infty}^{y_2} g(y;\sigma)\,dy - \int_{-\infty}^{y_1} g(y;\sigma)\,dy = G(y_2;\sigma) - G(y_1;\sigma) \tag{3}$$

where $G(y;\sigma)$ is 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 2(b)). This case is more involved than the previous one, because the ball can travel different distances before it reaches the goal. Since different travelled distances imply different deviations according to (1), the ball distribution along the goal line is no longer Gaussian. This makes an exact calculation of the total probability difficult.[4]

The key observation however, is that we want to compute probability masses and for equal masses, the particular shape of the distribution that produces these masses is irrelevant. This observation directly motivates our solution to the problem: instead of computing the distribution of the ball along the goal line analytically and then integrating to find its probability mass between the two goalposts, we compute the probability mass from the identity

$$\mathrm{P}\{\text{goal}\} = 1 - \mathrm{P}\{\text{not goal}\} \tag{4}$$

where P{not goal} denotes the probability that the ball will miss the goal, going out from the left or the right goalpost. This probability mass is easier to compute, to a good approximation, from the tails of the Gaussian distributions corresponding to the two goalposts.

This is shown in Figure 2(b): when the ball reaches the left goal post, it has effectively travelled distance l_l and its corresponding distribution perpendicular to the shooting line is

[4] A detailed analysis would involve bounding the corresponding diffusion process appropriately. Computing statistics for such a process is, however, a formidable task.

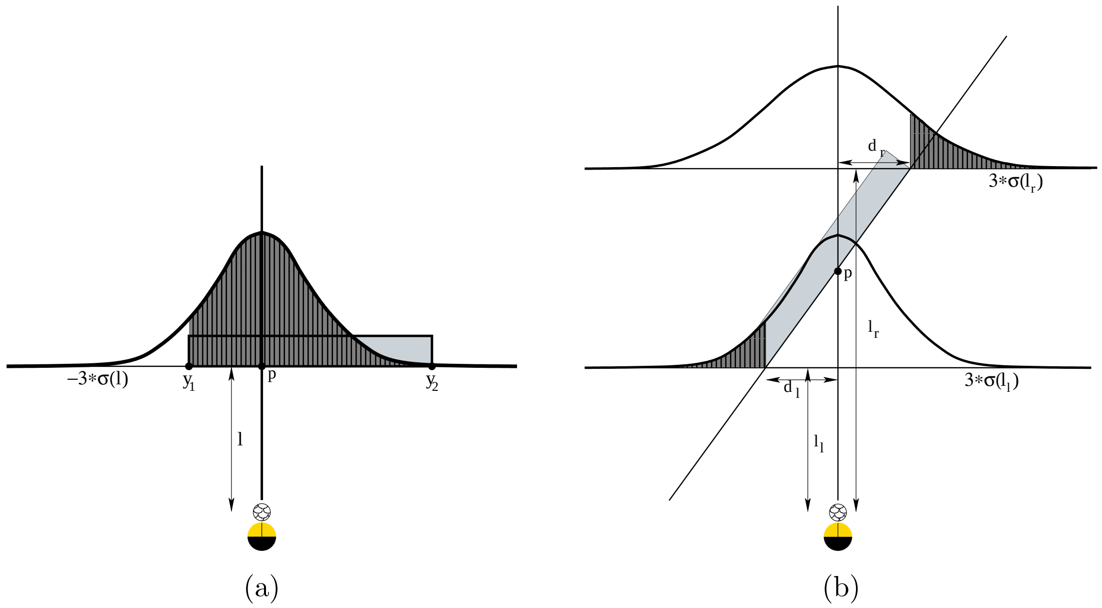

Figure 2: Two situations of shooting to the goal (light gray) together with the associated probability distributions. (a) Shooting perpendicularly. (b) Shooting at an angle.

Gaussian with deviation $\sigma(l_l)$ from (1). The probability that the ball will go out from the left goalpost is approximately[5] equal to the shaded area on the left. Thus

$$\mathrm{P}\{\text{out from left}\} \approx \int_{-\infty}^{-d_l} g(y; \sigma(l_l))\, dy \tag{5}$$

where the integration runs up to $-d_l$, the (negative) shortest distance from the left goalpost to the shooting line.

The situation that the ball will go out from the right post is analogous. The only difference is that the ball will have to travel a larger distance, thus its deviation will be larger, and the corresponding Gaussian will be flatter. The respective probability is approximately equal to the shaded area on the right. Thus

$$\mathrm{P}\{\text{out from right}\} \approx 1 - \int_{-\infty}^{d_r} g(y; \sigma(l_r))\, dy \tag{6}$$

where the integration now runs up to d_r, the shortest distance from the right goalpost to the shooting line, and the corresponding Gaussian has deviation $\sigma(l_r)$ which is computed for travelled distance l_r from (1).

Concluding, the probability that the ball ends up inside the goal becomes

$$\begin{aligned}\mathrm{P}\{\text{goal}\} &= 1 - \mathrm{P}\{\text{not goal}\} \\ &= 1 - \mathrm{P}\{\text{out from left}\} - \mathrm{P}\{\text{out from right}\} \end{aligned} \tag{7}$$

which can be computed directly using (5) and (6).

2.2 Subproblem 2: Probability of Passing the Goalkeeper

The second problem 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. Clearly,

[5]There is a probability (albeit small) that the ball will end up to the right of the left goalpost, after having travelled an 'illegal' trajectory outside the field.

this problem depends heavily on the opponent goalkeeper and, unless a provably optimal goalkeeper behavior has been implemented[6], the experiments have to be based on existing goalkeeper implementations. In our experiments we used the goalkeeper of the *Robocup-2000* winner, *FC-Portugal2000*, since it appeared to be one of the best available goalkeepers.

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 player and the goalkeeper positions (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.

To this end, 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 player. A data set was formed by recording 10.000 situations, together with a boolean indicating whether the goalkeeper had intercepted the ball or not. Through experiments, the relevant features for classification turned out to be[7] (i) the absolute angle a between the goalkeeper and the shooting point in the goal as seen by the player, and (ii) the distance b between the player and the goalkeeper. These two values form a two-dimensional feature vector on which the classification has been based.

The recorded data set is shown in Figure 3(a) where we note that there is an almost linear discriminant function between the two classes. We determined this discriminant via regression on the class indicator boolean variable, a procedure which is known to give the optimal *Fisher's linear discriminant*. For details we refer the reader to [4, Ch. 3.2]. This function is characterized by the equation

$$u = (a - 26.1) * 0.043 + (b - 9.0) * 0.09 - 0.2 \tag{8}$$

for distance values b between 3 and 15. This equation can be interpreted as follows: for a new angle-distance pair (a, b), the sample mean $(26.1, 9.0)$ is first subtracted from it, and then the inner product (projection) with the vector $(0.043, 0.09)$ is carried out; this vector is perpendicular to the discriminant boundary. The offset 0.2 shifts the boundary appropriately. The pairs for which (8) equals zero form the discriminant boundary between the two classes. This is plotted by a dotted line in Figure 3(a).

Projecting all (a_i, b_i) pairs perpendicularly to the discriminant line via (8), we get a set of one-dimensional points u_i that describe, to a good approximation, the two classes. The histogram class distributions of these points are plotted in Figure 3(b): the upper one corresponds to the situations where the goalkeeper did not succeed in intercepting the ball and the lower one to situations where the goalkeeper did intercept the ball. Instead of trying to model these two distributions parametrically, we note that the relevant range for classification is only where the two histograms *overlap*, i.e., the interval $[-0.5..0.5]$. It is easy to see that the non-interception posterior probability will be zero for approximately $u < 0.5$, will be one for $u > 0.5$ and will increase smoothly from zero to one in the interval in between.

Thus, we fit two univariate Gaussian functions on these two classes in the overlapping region, as shown in Figure 3(c). Having a Gaussian model for the class-conditional density function $p(u|C)$ for a class C, we can easily compute the posterior probability for this class using the Bayes rule[8]

$$\mathrm{P}(C|u) = \frac{p(u|C)p(C)}{p(u|C)p(C) + p(u|\bar{C})p(\bar{C})} \tag{9}$$

which is a sigmoid-like function. Since this is a simple two-class problem, $\bar{C}$ refers to the 'other' class, while the prior probability $p(C)$ of a class C is computed by the proportion of points u_i in

[6] We are currently working in this direction.

[7] There are principled methods for automatic feature extraction, see [4].

[8] A more principled way would be to use logistic discrimination [4, Ch. 3.5] and fit directly a posterior sigmoid from the data with maximum likelihood. However, the low dimensionality of the problem allows for the proposed simpler solution.

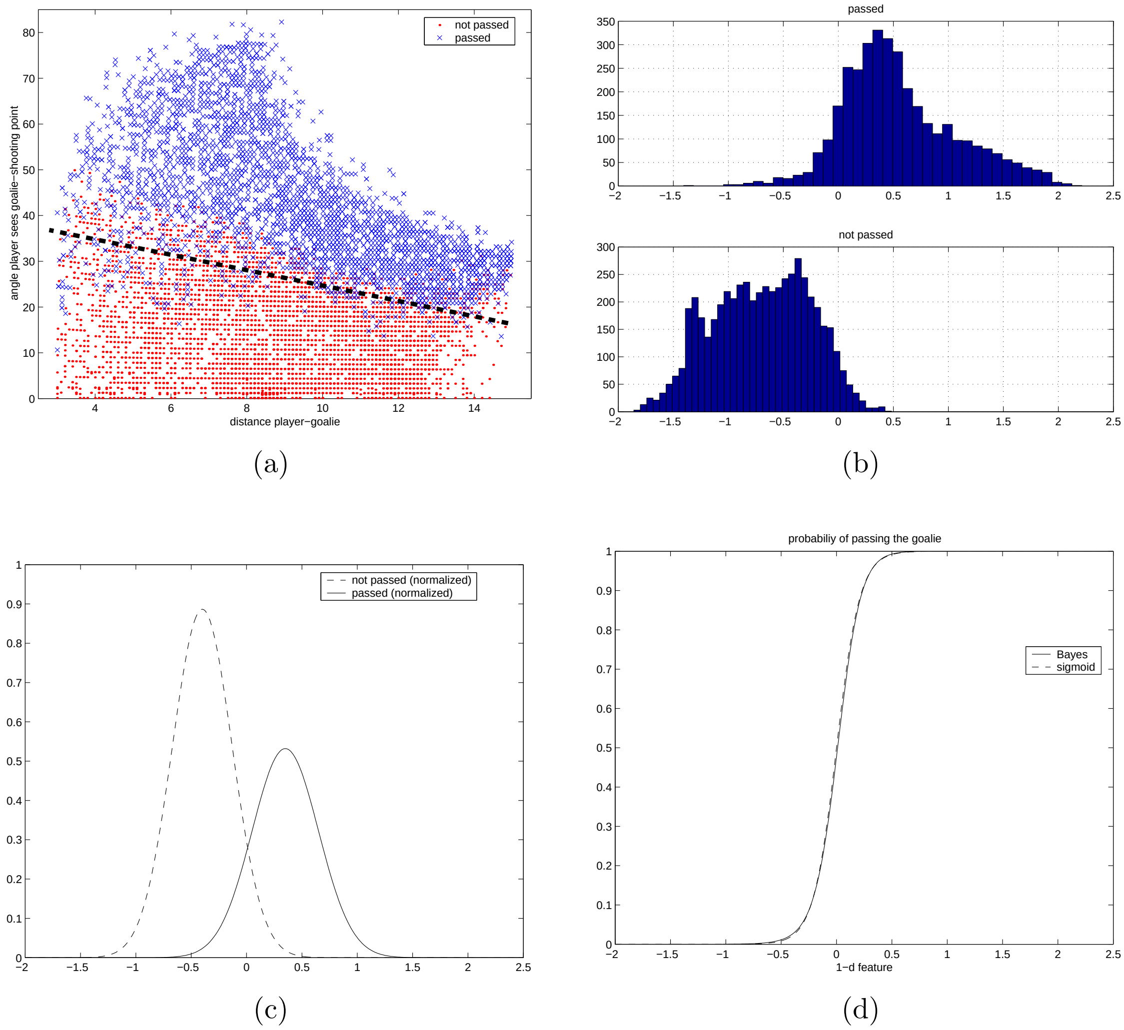

Figure 3: (a) Data set and discriminant function. (b) 1-d class histograms. (c) Gaussian approximations near discriminant. (d) Estimated posterior probability of non-interception.

the data set belonging to C (this is reflected through the height of the corresponding Gaussians in Figure 3(c)). In Figure 3(d) we plotted the posterior probability for the non-intercepting class as given from the Bayes rule above, together with the sigmoid approximation

$$\mathrm{P}(\text{pass goalkeeper } |u) = \frac{1}{1 + \exp(-9.5u)} \tag{10}$$

which allows for an easy implementation.

2.3 Determining the Best Scoring Point

Having computed the probability that the ball will end up in the goal (7) and the probability that the goalkeeper will not intercept it (10), the assumption of independence gives the total probability as the product of these two values. In order to determine the best shooting point in the goal, we discretize the goal interval $[-7.01..7.01]$ and compute the total probability that the ball will end up in each discretized bin. This total probability is a bell-shaped function, representing the probability that the ball will enter the goal, with a valley around the position of the goalkeeper (see Figure 4). The global maximum of this curve determines the best shooting point. The curve will have only two local maxima, corresponding to the left and the right starting

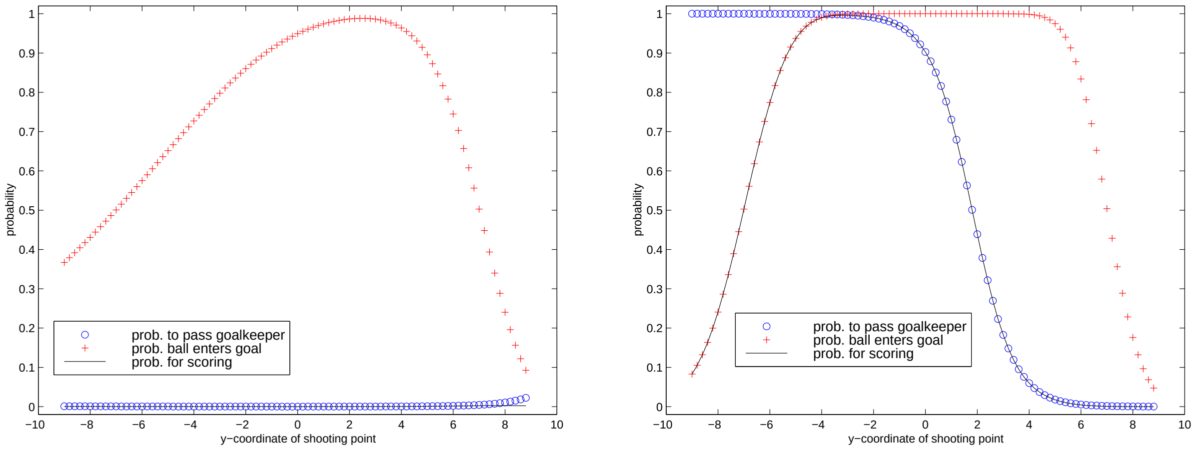

(a) (b)

Figure 4: Scoring probability curves. (a) Player is shooting from the right while the goalkeeper covers his goal well. (b) Player is shooting from straight in front of the goal while the goalkeeper is slightly off to the right.

point of the valley, which can be located with a simple hill-climbing algorithm. The maximum of these two is selected as the best shooting point. In practice, this point is selected only if the respective single probability that the ball will enter the goal is larger than a specified threshold.[9] This ensures that, independent of the particular goalkeeper behavior, the probability that the ball will end up inside the goal is high enough.

3 Results

We have implemented this scoring policy in our team *UvA Trilearn 2001* [5] as follows. When the agent has control of the ball, the first test in the decision procedure is to check whether the total scoring probability is higher than a specified threshold.[10] When this is the case, the agent tries to score. Otherwise he tries different alternative options, like passing or dribbling, which are performed when the predicted success rate is high enough. When all alternatives fail and the agent is at a close distance to the goal, he shoots to the best scoring point anyhow.

A simple example is depicted in Figure 4. The horizontal axis represents the y-coordinate on the goal line, where the left and right post are located at y-coordinates -7.01 and 7.01, respectively. In the left figure, the agent is shooting the ball from the right side of the field, while the opponent goalkeeper is covering his goal well. The total probability value is almost zero for all shooting points. The agent thus decides *not* to shoot but to pass to a teammate that is standing free in front of the goal. When this teammate receives the ball, the opponent goalkeeper still stands slightly off to the right. The scoring probability curves for this agent are shown in the right figure. The left slope of the total scoring probability (the solid line) is bounded by the probability that the ball enters the goal, while the right slope is bounded by the probability that the goalkeeper intercepts the ball. For the point around -3.8 the total scoring probability is almost one and the agent decides to shoot there and scores.

We have participated at *RoboCup-2001* in Seattle and reached fourth place in this competition. We have gathered statistics concerning the percentage of successful scoring attempts during the second group stage and knock-out stage of *RoboCup-2001*. Table 1 shows these per-

[9] 65% in our current implementation
[10] 90% in our current implementation.

centages for the top four teams in the competition.[11] This shows that the success rate for *UvA Trilearn* is higher than for the other teams.

Team	**Percentage**	**Team**	**Percentage**
Tsinghuaeolus	80.0% (56 of 70)	FC Portugal	77.41% (72 of 93)
Brainstormers	58.97% (23 of 39)	UvA Trilearn	80.95% (34 of 42)

Table 1: Percentage of scoring attempts that resulted in a goal.

4 Conclusions and Future Work

We have described a methodology that allows a simulated soccer agent to determine the scoring probability when he shoots the ball to a specific point in the goal. The single probability that the ball enters the goal depends on various server parameters that control the movement noise of the ball, the shooting power, the size of the goal, etc. Our approach is general because it can 'learn' this probability even when these parameters change, e.g. in a future server version.

However, the probability of passing the goalkeeper depends on the particular goalkeeper and different opponent goalkeepers exhibit different behaviors. In our current implementation we based the probability of passing the goalkeeper on the goalkeeper of *FCPortugal 2000*. Since this is a good goalkeeper, our method is useful against other goalkeepers as well. Nevertheless, 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 during the game. It is therefore important to extract the most relevant features and parametrize an opponent goalkeeper's intercepting behavior in a compact manner that permits on-line learning. This is one line of research we are currently pursuing, through the use of statistics collected by the coach.

References

[1] Ehsan Foroughi, Fredrik Heintz, Spiros Kapetanakis, Kostas Kostiadis, Johan Kummeneje, Itsuki Noda, Oliver Obst, Patrick Riley, Timo Steffens, USTC9811 Group. *RoboCup Soccer Server User Manual; Soccer Server version 7.06 and later.*

[2] Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, Eiichi Osawa. *RoboCup: The Robot World Cup Initiative.* In Proceedings of the IJCAI-95 Workshop on Entertainment and AI/Alife, 1995.

[3] A. Papoulis. *Probability, Random Variables and Stochastic Processes.* 3rd ed., McGraw-Hill, 1991.

[4] Ripley, B. D. *Pattern Recognition and Neural Networks.* Cambridge University Press, Cambridge, U.K., 1996.

[5] The *UvA Trilearn 2001* website at *http://www.science.uva.nl/˜jellekok/robocup.*

[11] These statistics have been generated by *RoboBase*; see `http://www.cs.man.ac.uk/~searj6/`.

IAS reports

This report is in the series of IAS technical reports. The series editor is Stephan ten Hagen (stephanh@science.uva.nl). Within this series the following titles appeared:

See: http://www.science.uva.nl/research/ias/tr/

You may order copies of the IAS technical reports from the corresponding author or the series editor. Most of the reports can also be found on the web pages of the IAS group (see the inside front page).