Towards an optimal scoring policy for simulated soccer agents

Jelle Kok Remco de Boer Nikos Vlassis

Computer Science Institute, University of Amsterdam
Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
{jellekok,remdboer,vlassis}@science.uva.nl

Abstract. This paper describes the scoring policy used by the agents of our simulation robot soccer 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. The ball motion can be regarded as a geometrically constrained continuous-time Markov process. Our main contribution is an approximate method for learning relevant statistics of such a process.

1 Introduction

RoboCup 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. 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 this paper we describe the implementation of a scoring policy that was used by the agents of our *UvA Trilearn 2001* team in the *RoboCup-2001* simulation soccer world championship.

An interesting aspect of the *soccer server* simulator is that, although the motion noise added to the ball is known, an analytical solution of the corresponding diffusion process (position of the ball in each time step) is difficult for two reasons: (1) the noise added by the server is by construction non-white, and (2) the process is geometrically constrained (the ball must end up inside the goal). We propose an approximate solution to the problem of learning the statistics of such a geometrically-constrained continuous-time Markov process, which we believe can also be useful in other applications.

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 problem 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.

2.1 Subproblem 1: Probability that the Ball Enters the Goal

When the ball is shot to a point somewhere inside the goal, it can miss the goal due to motion noise (introduced by the server). We are interested in the probability that the ball will end up *somewhere* inside the goal when shot at a *specific* point. To this end 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 ball velocity in each simulation cycle.[1] 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.

Treating the ball motion as a continuous-time Markov process, computing exact statistics for each time step would require the solution of a corresponding Fokker-Planck equation [2], further complicated by the fact that the motion noise is non-white. Moreover, the solution will not be generic but depends on the current values of the server parameters. To avoid both problems, we propose to *learn* the statistics of the ball motion directly from experiments, and to compute the required probabilities from these statistics.

We estimated the *cumulative* noise added to the ball perpendicular to the shooting direction as a function of the travelled distance d along this direction. This function was learned by repeating an experiment in which a player was placed at various distances 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 where the ball entered the goal. We empirically found that, to a good approximation, the standard deviation σ of the ball perpendicular to the shooting direction was given by a monotone increasing function

$$\sigma(d) = -1.88 * \ln(1 - d/45) \qquad (1)$$

with $\ln(\cdot)$ the natural logarithm. Moreover, the Central Limit Theorem [2] indicates that the ball distribution along the goal line will be approximately Gaussian with zero mean and standard deviation $\sigma(d)$ from (1), as shown in Fig. 1(a). The scoring probability is then given by the area of the Gaussian density between the two goalposts.

When the ball is shot at an angle to the goal, the ball can travel different distances (implying different deviations) before it reaches the goal line, causing the distribution along the goal line to be non-Gaussian. The key observation is that we want to compute probability *masses* and for equal masses the particular shape of the distribution that gives rise to these masses is irrelevant. Therefore, instead of computing the distribution of the ball along the goal line analytically (by solving the constrained diffusion process equations) and then integrating to find its probability mass between the two goalposts, we compute the probability mass from the identity

$$\mathrm{P}\{\text{goal}\} = 1 - \mathrm{P}\{\text{not goal}\} = 1 - \mathrm{P}\{\text{out from left}\} - \mathrm{P}\{\text{out from right}\} \qquad (2)$$

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 from the tails of the Gaussian distributions corresponding to the two goalposts.

[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 $= 0.05 * ||(v_x^t, v_y^t)||$.

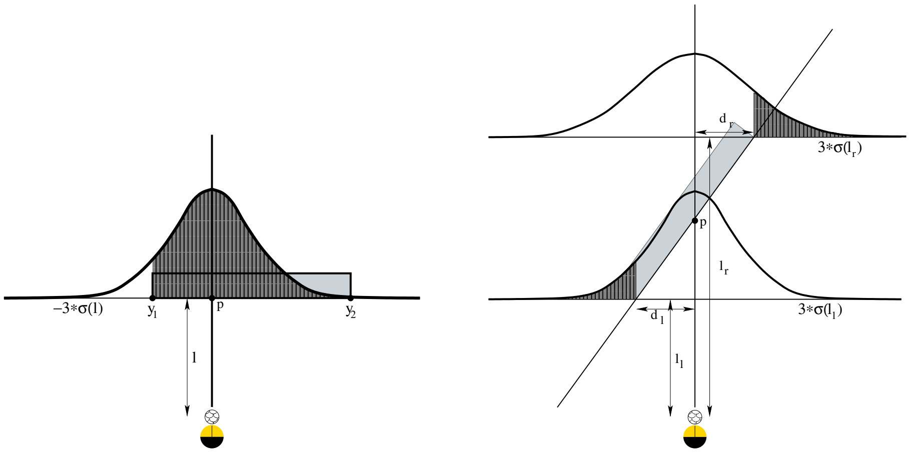

(a) (b)

Figure 1: Two situations of shooting to the goal (light gray) together with the associated probability distributions. (a) Shooting perpendicularly. (b) Shooting at an angle.

This is shown in Figure 1(b). For example, when the ball reaches the left post, it has effectively travelled distance l_l and its distribution perpendicular to the shooting line is Gaussian with deviation $\sigma(l_l)$ from (1). The probability that the ball will go out from the left goalpost is approximately[2] equal to the shaded area on the left. Thus

$$\mathrm{P}\{\text{out from left}\} \approx \frac{1}{\sigma(l_l)\sqrt{2\pi}} \int_{-\infty}^{-d_l} \exp\left[-\frac{y^2}{2\sigma^2(l_l)}\right] dy \tag{3}$$

where d_l is the 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, which will make its deviation larger and the corresponding Gaussian flatter. Finally, using (2) we can determine the probability that the ball will enter the goal.

2.2 *Subproblem 2: 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. In our experiments we used the goalkeeper of *Robocup-2000* winner *FC Portugal*, since it appeared to be one of the best available goalkeepers. Due to lack of space we only briefly describe the proposed method, and refer for details to [1].

The main observation is that ball interception can be regarded as a two-class *classification* problem: given the player and goalkeeper position (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. Formalizing the problem in this way allows for direct application of a variety of methods from the field of statistical pattern recognition [3].

To collect a training set, 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 at 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

[2] We neglect the probability that the ball will end up to the right of the left goalpost after having travelled an 'illegal' trajectory outside the field.

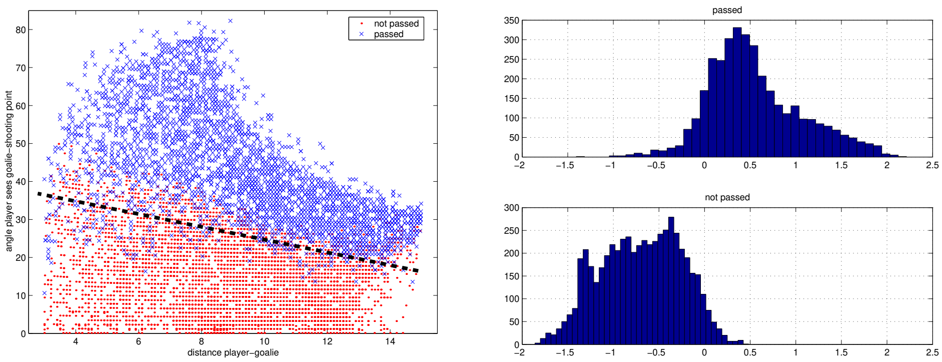

(a) (b)

Figure 2: (a) Data set and discriminant function. (b) 1-d class histograms.

intercepted the ball or not. The relevant features for classification turned out to be (i) the absolute angle between the goalkeeper and the shooting point in the goal as seen by the player, and (ii) the distance between the player and the goalkeeper. These two values formed a two-dimensional feature vector on which the classification was based.

The recorded data set is shown in Figure 2(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* [3, Ch. 3.2]. Projecting all data points perpendicularly to the discriminant line, 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 2(b). Then, we fit a univariate Gaussian function $p(u|C)$ on each class C in the overlapping region, from which we can compute the posterior probability for this class using the Bayes rule, which approximately gives the required posterior P(pass goalkeeper $|u$) as a simple sigmoid function [1].

2.3 Determining the Best Scoring Point

Having computed the probability that the ball will end up inside the goal (subproblem 1) and the probability that the goalkeeper will not intercept it (subproblem 2), the assumption of independence gives the total probability as the product of these two probabilities. 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. The curve will have only two local maxima, corresponding to the left and the right starting point of the valley, which can be located with a simple hill-climbing algorithm. The global maximizer of this function is selected as the best shooting point.

References

[1] J. Kok, R. de Boer, and N. Vlassis. Towards an optimal scoring policy for simulated soccer agents. Technical Report IAS-UVA-01-06, Computer Science Institute, University of Amsterdam, The Netherlands, Oct. 2001.

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

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