Pareto-Optimal Offensive Player Positioning in Simulated Soccer

Vadim Kyrylov[1] and Serguei Razykov[2]

[1] Rogers State University; Claremore, OK 74017 USA
vkyrylov@rsu.edu
[2] Simon Fraser University – Surrey; Surrey, British Columbia V3T 0A3 Canada

Abstract. The ability by the simulated soccer player to make rational decisions about moving without ball is a critical factor of success. Here we limit our scope to the offensive situation, i.e. when the ball is controlled by own team, and propose a systematic method for determining the optimal player position. Existing methods for accomplishing this task do not systematically balance risks and rewards, as they are not Pareto optimal by design. This may result in overlooking good opportunities. One more shortcoming of these methods is over simplifications in predicting the situation on the field, which may lead to performance loss. We propose two new ideas to address these issues. Experiments demonstrate that this results in a substantial increase in the team performance.

1 Introduction

The ability by the soccer player making rational decisions about where to move without the ball is the critical factor of success both in a real-life and simulated game. With the total of 22 soccer players in two teams, an average player must be indeed purposefully moving to some place on the field more than 90 per cent of the time. Thus one should be expecting about 90 per cent impact on the whole game from any improvement in player positioning.

In this paper, we propose the algorithm for determining a good position on the soccer field for the artificial player when the ball is beyond its control. We limit here the consideration to the *offensive* positioning, i.e. when the ball is possessed by own team. In the offensive situation, the major purpose of co-coordinated moving to some individual positions on the field by the players without the ball is creating opportunities for receiving passes and scoring the opponent goal. Therefore, player positioning is not a standalone task; rather, it is part of a coordination mechanism. In this study, we deliberately isolate positioning from ball passing because in our recent paper we have proposed a Pareto-optimal method for ball passing [1]. Now we want the decisions about positioning to be optimized in the similar sense.

During the last 10 years, RoboCup scholars have addressed the positioning problem in different ways [2-8]. In order to demonstrate the unique features of our approach, here we propose a generic two-level model that includes the existing methods for player positioning.

The upper level determines so-called *reference* position for the player where it should be moving unless it is involved in intercepting or kicking the ball. Player coordination on

U. Visser et al. (Eds.): RoboCup 2007, LNAI 5001, pp. 228–237, 2008.

this level is addressed only implicitly. The lower level is responsible for fine tuning this reference position by transforming it into the *target* position. The decision-making rules and/or optimality criteria used in such methods are normally reflecting the algorithm designer's vision of the soccer game. Alternatively, the decision making rules are derived by learning algorithms.

The first ever comprehensive study of the player positioning problem was presumably made in [2]. In this method, the higher-level reference position can be regarded as a fixed point in the field assigned to each player with respect to its role in the team formation. The lower control level allowed the player to deviate from this default position within some area in order to move to the calculated target position with respect to current situation. Although this method implemented in *CMUnited* proved to be a good start, later on it was criticized in [3] for the limited flexibility on the upper control level. A method based on a set of logical rules proposed in [3] has addressed these shortcomings in the *FCPortugal* who outperformed *CMUnited*.

The development that followed, did not demonstrate much improvement on the lower control level, though. The next very successful team, *UvA Trilearn* [4] that had outplayed *FCPortugal* in several competitions, implemented somewhat simpler player positioning method. In our classification, the early version of this method dated 2001-2002 is completely located on the higher control layer. In particular, this method does not take into account fine details such as the opportunity to receive a pass. This shortcoming was later addressed using so-called coordination graphs [5]. This lower-level model combines decision making about passing the ball and receiving the pass in an elegant way; implemented in *UvA Trilearn*, it helped to become the World RoboCup winner in 2003. However, we believe that this model could be further improved, as it was using heuristics rather than rigor multi-criteria optimization.

One more group of the improvements to player positioning is a set of methods based on learning algorithms [6, 7, 8]. Like the coordination graph, some of these methods do not treat positioning as a standalone player skill, which makes theoretical comparisons difficult. More difficulties arise while trying to elicit meaningful decision making rules and especially address the convergence issue of learned rules to optimal decision making algorithms based on explicitly formulated decision criteria. Thus we are leaving methods based on learning beyond the scope of this study.

The main objective of this paper is to provide an improved solution for low-level decision making about player positioning based on the known tactical principles of the real-life soccer [9, 10]. Because the recent studies addressed the lower control level rather superficially, it looks like our closest prototype is described in [2]. That was indeed the first method that was using both control levels. On the lower level, it maximizes the following utility function:

$$U(P) = \sum_{i=1}^{n} dist(P,O_i) + \sum_{j=1}^{n-1} dist(P,T_j) - (dist(P,G))^2, \tag{1}$$

where

P – the desired position for the player in anticipation of a pass;
n – the number of agents on each team;
O_i– the current position of each opponent, $i = 1,\ldots, n$;
T_j – the current position of each teammate, $j = 1,\ldots, (n\text{-}1)$;
G – the position of the opponent goal;
dist(A,C) – distance between positions A and C.

One advantage of this utility function is robustness; the two sums suppress the sensor noise. It repulses the player without ball from other players and encourages advancing to the opponent goal while seeking its best position. This indeed is reflecting the soccer tactics.

We would also mention three shortcomings. First, from the tactical standpoint, criterion (1) does not encourage players to deliberately get open for receiving a pass; this can only happen by chance. Second, the contribution of the remotely located players is exaggerated; increasing distance from the closest opponent by say 5 meters has same effect as increasing distance by 5 meters for the opponent located on the other end of the field. Third, from the mathematical standpoint, the authors clearly indicate that this is a vector optimization problem; indeed, each addendum in (1) could stand for a criterion. However, the reduction to single-criteria optimization is questionable. Aggregating several criteria in a linear combination is indeed theoretically acceptable if all criteria functions and constraints are convex [11]. However, the first two addendums in (1) are multi-modal functions of P; hence they are non-convex. So this single-criterion optimization does not guarantee that all potentially acceptable player positions would be considered. This issue could be resolved by using the Pareto optimality principle.

Dynamics is one more difficulty with optimizing player position on low level. Reaching an optimal position takes some time, and the player must be persistent in doing so. In our experiments we have found that random factors in the simulated soccer tend to make the computation of the optimal target very unstable from cycle to cycle. In what follows, we discuss this issue.

Therefore, the unique contribution of this paper is in that we (1) resolve the time horizon issue; (2) propose a new set of decision making criteria based on real-life soccer tactics; and (3) implement a Pareto-optimal solution to the optimization problem with these criteria.

Section 2 addresses the time horizon issue. Section 3 explains how feasible alternative player positions could be generated. Section 4 introduces five optimality criteria and the algorithm for finding the Pareto-optimal solution. Section 5 provides experimental results and conclusions.

2 Predicting the Situation for Player Positioning

The issue is what the planning time horizon should be once the player decides to go to some *target* position. A straightforward approach is just setting some fixed time horizon τ, say 1 to 3 seconds, and trying to extrapolate the movements of the players and the ball for this time interval. However, our experiments have shown that this method does not work well. This is because of that, once the ball is kicked by some player, the rest players revise their previous intentions and change their directions of movement. Neglecting these abrupt changes if the next ball kick occurs before the prediction time expires would result in poor decisions. On the other hand, forecasting new movements by the players when the ball gets under close control is difficult.

While trying to resolve this issue, it makes sense to see how human players are acting in this situation [9, 10]. In the offensive positioning, the teammate with the ball may be in two different states: (a) chasing the freely rolling ball while staying the fastest player to it or (b) keeping the ball close to itself while being able to pass it at any time. In case (a) the human player without ball can easily predict the situation and estimate time remaining until the ball will be intercepted. So nothing abrupt will likely occur while the ball is rolling freely. Thus a human player predicts the situation and plans his actions accordingly until the ball is intercepted. This gives the idea of the time horizon that we should be using in our model. During this time the player without ball must concentrate on reaching good position before the teammate has gained close control of it. No communication is really necessary, and it does not make much sense substantially changing the originally optimized target position while the ball is rolling. Only minor corrections to it may be necessary as the world model is updated.

In case (b), however, the situation is hard to predict, as the player with the ball can kick it at almost any time in a wide range of directions. In this case human players, if they have chance to become pass receivers, are watching carefully the teammate with the ball in order to be prepared to intercept the ball. The time horizon can be communicated by the active player to the partners; so communication is highly important in this situation. With or without the communication, time horizon is very limited, anyway. So players in this case only can do short-time planning to adjust their positions; the major objective is implementing the team strategy by moving to what we call here the *reference* position determined by current situation. In many cases during the game player may be indeed too far away from this position; so the major concern is reaching it as fast as stamina permits.

Thus we come to the idea of the variable time horizon during that player behavior should be staying persistent. While the ball is rolling freely, its movement follows the laws of physics and is easily predictable. All players (if they are rational) are acting in the ways that are also rather easy to predict if their decision making model is known. Predicting teammate movements thus is possible with high precision, Figuring out the upper-level positioning algorithm of the opponent team is possible using opponent modeling implemented in the online coach. This problem could be simplified by concentrating only on the situations when the ball is rolling freely.

Therefore, for these situations we are using the prediction horizon τ that is equal to the time remaining until the ball will be intercepted by a teammate. (If the fastest to the ball is the opponent player, the situation changes to defensive; this lies outside the scope of this study.) With this approach, once the ball has been kicked, the player estimates the interception point and time τ. This is the time remaining to plan its actions and move to the best possible position on the field before the ball will be kicked in the new direction. Because τ is known, the player selects only such optimal target position that could be indeed reached during this time. While the ball is rolling, the player is persistently moving to this position.

The model for predicting the situation comprises three components: the ball, the friendly and the opponent player movements. The ball movement can be predicted with high precision, as the underlying physics is simple. The movements by teammates can be also predicted with precision, as their control algorithms and to some extent their perceived states of the world are available to the player in question. The fastest teammate to the ball will be intercepting it by moving with the maximal speed;

so its position can be predicted even more precisely. The rest teammates will be moving towards the best positions determined with yet another precisely known algorithm which could be used for predicting their positions. However, in our method we do not use such a sophisticated approach; in regards of each teammate without the ball, we assume that it is just moving with a constant velocity. This velocity is estimated by the player using own world model. Of the opponent players, the fastest one will be presumably also chasing the ball by following the trajectory that can be predicted fairly precisely. For the rest opponents possible options include assuming same positioning algorithm as for the teammates.

If the ball is kickable by the teammate, we would suggest that τ should be set to a small constant value $\tau_{\min}$ which should be determined experimentally. So while the ball is closely controlled by the teammate, this would allow the player to continue adjusting its position within this limited time horizon.

3 Identifying the Feasible Options

Decision making is always a choice from a set of alternatives. In the discussed problem, the player first generates a set of feasible options and evaluates those using different criteria. Then the multi-criteria algorithm is applied to find the single best option by balancing the anticipated risks and rewards.

Once the ball interception point and the remaining time τ have been determined by the player without ball, it generates a set of alternative positions in the vicinity of the reference position. Because the decision space (xy-plane) is continuous, it contains infinite number of such alternatives. With highly nonlinear and non-convex decision criteria, searching such space systematically would be hardly possible. Therefore, we use a discrete approximation, with the alternative positions forming on the xy-plane a grid about the default position. To reduce computations, we would like to keep the number of points in this grid minimal. The grid step determines the total number of the alternatives to be searched. The rule of thumb is setting the step equal to the radius of the player reach for kicking the ball. Increasing it might result in lost opportunities. Using a smaller step makes sense only if we have enough computational time to further fine tune the balance of different optimality criteria (see more about it in the next section).

Of these alternative positions, the player is only interested in those that could be reached in time τ. This allows eliminating part of the grid that is lying beyond the player reach. One more constraint that helps eliminating poor alternatives is the maximal distance from the reference position. The alternatives lying outside the field or creating risk of offside are also eliminated.

Figure 1 displays the field when the ball is kicked by red player #5. Arrows show the predicted positions of all objects at the moment when the ball is intercepted by red player #7. Highlighted are the alternative positions for red player #8. The area of responsibility is filled with small gray points; the reference position being the center of this area. The bigger blue points show the reachable positions of which this player must choose the best.

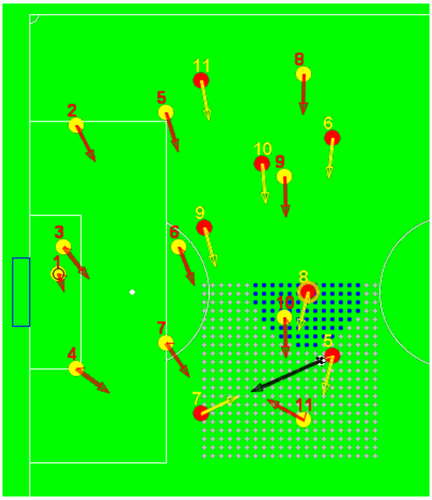

Fig. 1. The area of responsibility for red player #8 (gray dots) and reachable positions (blue dots) before the ball is intercepted by red #7

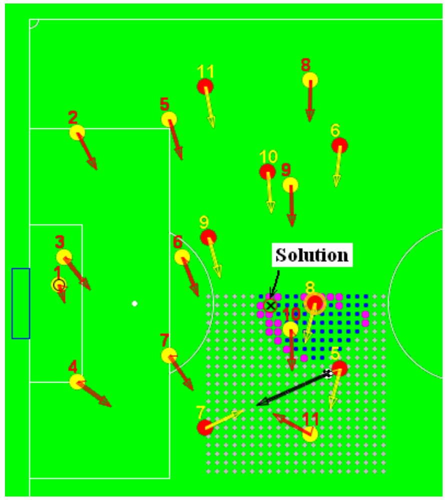

Fig. 2. The Pareto set for red player #8 (bigger dots) and the optimal solution

4 Criteria for Decision Making and the Optimization Algorithm

Each feasible alternative position has its pros and cons that an intelligent player is taking into account while choosing the best option. These decision criteria should be reflecting the soccer tactics; in particular they should be measuring anticipated rewards and risks. We propose slightly different criteria sets for attackers, midfielders, and defenders because their tactical roles differ indeed [9, 10].

For the attackers the criteria set is, as follows.

1. All players must maintain the formation thus implementing the team strategy. So the distance between the point in the feasible set and the reference position should be minimized.

2. All attackers must be open for a direct pass. Thus the angle between the direction to the ball interception point and the direction to the opponent located between the evaluated position and the interception point must be maximized.

3. All players must maintain open space. This means that the distance from the evaluated point to the closest opponent should be maximized.

4. The attackers must keep an open path to the opponent goal to create the scoring opportunity. So the distance from the line connecting the evaluated point and the goal center to the closest opponent (except the goalie) should be maximized. This criterion is only used in the vicinity of the opponent goal.

5. The player must keep as close as possible to the opponent offside line to be able to penetrate the defence. So, the player should minimize the x-coordinate distance between the point in the feasible set and the offside line (yet not crossing this line).

Note that each criterion appears to have equal tactical importance; this observation will be used while discussing the optimization procedure below.

Criteria for midfielders and defenders differ in that they do not contain criteria 4 and 5 that encourage the opponent defense penetration. Instead, these players should be creating opportunities for launching the attack. This is achieved by minimizing the opponent player presence between the evaluated position and the direction to the opponent side of the field.

These criteria are conflicting, as it is hardly possible to optimize all them simultaneously. This situation is well known in the literature on systems analysis and economics; a special paradigm called the *Pareto optimality principle* allows to eliminate wittingly inappropriate so-called dominated alternatives [11]. These are the points that could be outperformed by at least some other point in the feasible set by at least one criterion. So only the non-dominated alternatives making so-called Pareto set should be searched for the 'best' balance of all criteria. Balancing requires additional information about the relative importance of these criteria, or their weights. If the criteria functions and the feasible set are all convex, then the optimal point could be found by minimizing the weighed sum of the criteria (assuming that they all must be minimized) [11]. However, because in our case the criteria functions may have several extremes, there is no hope for such a simple solution.

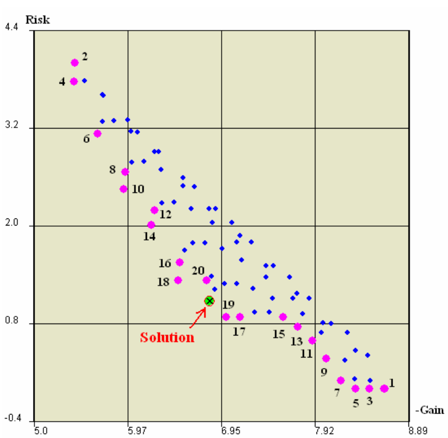

Fig. 3. The criteria space. Numbers at the points in the Pareto set show the elimination order.

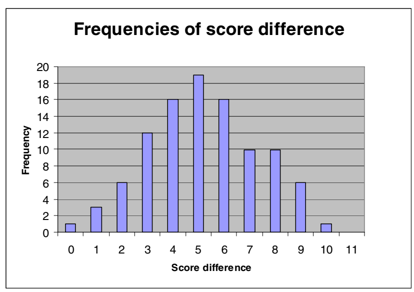

Fig. 4. A histogram of the score difference in 100 games

The way out has been proposed in our recent work [1], where a method for searching the balanced optimal point in the finite Pareto set was presented. This method is based on the sequential elimination of the poorest alternatives using just one criterion at a time. With N alternatives in the Pareto set, it requires N-1 steps. The criterion for the elimination on each step is selected randomly with the probability proportional to the weight of this criterion. Hence more important criteria are being applied more frequently. The sole remaining option after N-1 steps is the result of this optimization. This method works for any non-convex and even disconnected Pareto set. Its computational complexity is $O(N^2)$.

In this application, we have further simplified the decision making procedure by assuming that all criteria have equal importance. Thus instead of randomly selecting the criteria on each step of elimination, our procedure is looping through the criteria in the deterministic order.

If the total number of the alternatives is too small, this would result in only near-optimal decision. Better balancing of the conflicting criteria is possible with increased *N*. So we propose to estimate the available computational time in current simulation cycle and select larger *N* if time permits. This optimization algorithm is scalable indeed. It is also robust, because even with small *N* the decisions returned by it are still fairly good.

Although we actually have five optimality criteria, for the purpose of illustration we have aggregated them all in just two: the *Risk*, which is combination of criteria 2 and 3, and *Gain* which aggregates criteria 1, 4, and 5. The signs of the individual criteria in these aggregates were chosen so that both Risk and -Gain must be minimized.

Figures 2 and 3 illustrate the configuration of the Pareto set in the decision and criteria space, respectively. Of the total of 21 points in the Pareto set 20 are eliminated as shown in Figure 3; the remaining point is the sought solution. Note that the Pareto frontier is non-convex.

The optimal point is reachable and is located at less than the maximal distance of the reference position. It is lying on the way towards the opponent goal and far away from the predicted positions of the two threatening opponents, yellow #10 and #6. This point is open for receiving the pass from the anticipated interception point. This is indeed a well-balanced solution to the positioning problem for the red player #8. With non-aggregated five criteria we can only expect even better decisions.

5 Experimental Results and Conclusion

We have conducted experiments with the purpose to estimate the sole contribution of the proposed method for the lower-level optimized player positioning compared with only strategic, higher-level positioning.

Measuring the player performance using existing RoboCup teams is difficult because new features always require careful fine tuning with the existing ones. For this reason, we decided to compare two very basic simulated soccer teams. The only difference was that the experimental team had player positioning on two levels and the control team just on one level. Players in both teams had rather good direct ball passing and goal scoring skills and no dribbling or holding the ball at all. Thus any player, once gaining the control of the ball, was forced to immediately pass it to some teammate. In this setting, the ball was rolling freely more than 95 per cent of the time, thus providing ideal conditions for evaluating the proposed method.

To further isolate the effects of imperfect sensors, we decided to use *Tao of Soccer*, the simplified soccer simulator with complete information about the world; it is available as the open source project [12]. Using the RoboCup simulator would require prohibitively long running time to sort out the effects of improved player positioning among many ambiguous factors.

The higher-level player positioning was implemented similar to used in *UvA Trilearn* [4]; this method proved to be reasonably good indeed. Assuming that both goals are lying on x-coordinate axis, the coordinates of the reference position for *i*-th player are calculated as follows:

$$x_i = w*xhome_i + (1-w)*xball + \Delta x_i,$$
$$y_i = w*xhome_i + (1-w)*yball, \quad (2)$$

where w is the weight ($0<w<1$), ($xhome_i$, $yhome_i$) and ($xball$, $yball$) are the fixed home and the current ball positions respectively, Δx_i is the fixed individual adjustment of x-coordinate whose sign differs for the offensive and defensive situations and the player role.

Because players in the control team were moving to the reference positions without any fine tuning, ball passing opportunities were occurring as a matter of chance. In the experimental team, rather, players were creating these opportunities on purpose.

The team performance was measured by the game score difference. Figure 4 shows the histogram based on 100 games each 10 minutes long.

Only one game has ended in a tie; in all the rest 99 games the experimental team won. The mean and the standard deviation of the score difference are 5.20 and 2.14, respectively. By approximating with Gaussian distribution, we get 0.9925 probability of not loosing the game. The probability to have the score difference is greater than 1 is 0.975 and greater than 2 is 0.933. This gives the idea of the potential contribution of the low-level position optimization. With the smaller proportion of the time when the ball is rolling freely, this contribution will decrease. So teams favoring ball passing would likely benefit from our method more than teams that prefer dribbling.

The experimental results demonstrate that, by locally adjusting their positions using the proposed method, players substantially contribute to the simulated soccer team performance by scoring on the average about five extra goals than the opponent team that does not have this feature. This confirms that optimized player positioning in the simulated soccer is the critical factor of success. Although this method has been developed for simulated soccer, we did not rely much on the specifics of the simulation league. Therefore, we believe that the main ideas presented in this work could be reused with minor modifications in other RoboCup leagues.

Acknowledgements

The authors would like to thank Peter Stone and Marian Lekavy for their help with identifying the early studies of the player positioning methods.

References

1. Kyrylov, V.: Balancing Rewards, Risks, Costs, and Real-Time Constraints in the Ball Passing Algorithm for the Robotic Soccer. In: Lakemeyer, G., Sklar, E., Sorrenti, D.G., Takahashi, T. (eds.) RoboCup 2006: Robot Soccer World Cup X. LNCS (LNAI), vol. 4434. Springer, Heidelberg (2007)

2. Stone, P., Veloso, M., Riley, P.: The CMUnited-98 Champion Simulator Team. In: Asada, M., Kitano, H. (eds.) RoboCup 1998. LNCS (LNAI), vol. 1604, pp. 61–76. Springer, Heidelberg (1999)
3. Reis, L.P., Lau, N., Oliveira, E.C.: Situation Based Strategic Positioning for Coordinating a Team of Homogeneous Agents. In: Hannebauer, M., Wendler, J., Pagello, E. (eds.) ECAI-WS 2000. LNCS (LNAI), vol. 2103. Springer, Heidelberg (2001)
4. De Boer, R., Kok, J.: The Incremental Development of a Synthetic Multi-Agent System: The UvA Trilearn 2001 Robotic Soccer Simulation Team. Master's Thesis. University of Amsterdam (2002)
5. Kok, J., Spaan, M., Vlassis, N.: Multi-Robot Decision Making Using Coordination Graphs. In: Proceedings of the International Conference on Advanced Robotics (ICAR), Coimbra, Portugal, pp. 1124–1129 (June 2003)
6. Andou, T.: Refinement of Soccer Agents' Positions Using Reinforcement Learning. In: Kitano, H. (ed.) RoboCup 1997. LNCS, vol. 1395, pp. 373–388. Springer, Heidelberg (1998)
7. Nakashima, T., Udo, M., Ishibuchi, H.: Acquiring the positioning skill in a soccer game using a fuzzy Q-learning. In: Proceedings of IEEE International Symposium on Computational Intelligence in Robotics and Automation, July 16-20, 2003, vol. 3, pp. 1488–1491 (2003)
8. Kalyanakrishnan, S., Liu, Y., Stone, P.: Half Field Offense in RoboCup Soccer: A Multiagent Reinforcement Learning Case Study. In: Lakemeyer, G., Sklar, E., Sorrenti, D.G., Takahashi, T. (eds.) RoboCup 2006: Robot Soccer World Cup X. LNCS (LNAI), vol. 4434. Springer, Heidelberg (to appear, 2007)
9. Beim, G.: Principles of Modern Soccer. Houghton Mifflin Company, Boston (1977)
10. Vogelsinger, H.: The Challenge of Soccer. Allyn and Bacon, Inc., Boston (1973)
11. Miettinen, K.: Nonlinear Multiobjective Optimization. Kluwer Acaemic Publishers, Berlin (1998)
12. Zhan, Y.: Tao of Soccer: An Open Source project (2006), `https://sourceforge.net/projects/soccer/`
13. Fisher, R.A., Bennett, J.H.: Statistical Methods, Experimental Design, and Scientific Inference. Oxford University Press, Oxford (1990)