Automatic Soccer Commentary and RoboCup

Hitoshi Matsubara Ian Frank Kumiko Tanaka-Ishii
Itsuki Noda Hideyuki Nakashima Kôiti Hasida

Complex Games Lab, Language Learning Lab
Electrotechnical Laboratory (ETL)
Umezono 1-1-4, Tsukuba
Ibaraki, Japan 305

Abstract. This paper suggests that automated soccer commentary has a key role to play within the overall RoboCup initiative. Firstly, we identify soccer commentary as allowing and requiring investigation of a wide variety of research topics, many of which could not be addressed by the simple development of teams for the RoboCup leagues themselves. Secondly, we highlight a key task of soccer commentary: the expert analysis of a game. We suggest that this expert analysis task has the potential to make a significant impact on RoboCup challenges such as learning, teamwork, and opponent modeling. We illustrate our arguments by discussing the progress on soccer commentary systems to date, in particular reviewing our own system, MIKE.

1 Introduction

In real-life, commentary adds so much to the coverage of football that no TV company would contemplate screening a game without it. Indeed, sometimes the commentary itself is just as memorable as the scenes it describes (for example, the well-known "They think it's all over... It is now!" as Geoff Hurst scored in the dying seconds of extra time in the 1966 World Cup final, and the less well-known but equally noteworthy "17 minutes gone and already no goals!" [Jones 96]). At the first RoboCup tournament in Japan, however, the quantity of matches made it impractical to recruit human volunteers to describe all the games. Since RoboCup also does not (yet) generate crowd noise comparable to that of real soccer, it was therefore common to hear the comment that some games seemed 'flat'. It was to fill this gap, and to provide extra atmosphere, interest and context for RoboCup games that we originally started to develop our automatic commentary system, MIKE (Multi-agent Interactions Knowledgeably Explained). This system, which produces text or spoken commentary from the output of the Soccer Server [Noda *et al* 98], was first used to commentate public games at the Japan Open in April 1998, and will be used again in Paris this year to provide commentary for the second RoboCup contest.

We believe, however, that automated soccer commentary has a far greater role to play within the overall RoboCup initiative than simply adding atmosphere. In this paper we clarify this role by suggesting that soccer commentary allows for the investigation of a number of research topics that are outside the scope of the original RoboCup challenges. Also, within the framework of RoboCup itself, we identify how a key task of soccer commentary — the expert analysis of games — has the potential to make a significant impact on existing RoboCup challenges such as learning, teamwork, and opponent modeling. We demonstrate these points by examining the work to date on automatic commentary systems. We focus on our own system MIKE, but also review systems produced by others. We hope that our summary of expert analysis techniques, in particular, will be of general interest to a wide variety of researchers involved in RoboCup. We ourselves identify a number of specific tasks for which such techniques should be indispensable.

2 Why Soccer Commentary?

One of the strengths of RoboCup is its ability to attract researchers from many different domains. For example, [Kitano *et al* 97a] discusses how a large number of technologies need to be integrated to produce successful RoboCup teams. In the same vein, we see soccer commentary as providing challenges of its own. Let us illustrate this by examining the research topics involved in generating convincing game descriptions:

– **Natural language generation**. The generation of text or speech is the most obvious aspect of soccer commentary. The time-critical nature of the domain forces real-time decisions about what to utter, when the current commentary would be better interrupted by a new description, and how to maintain overall coherency.
– **Understanding of multi-agent systems**. To describe soccer it is necessary to understand how the multiple players interact. This understanding can be pursued on numerous levels, such as the recognition of low-level tactics or of high-level strategy, the following of a focus point (such as the ball), the analysis of the territories established by the players, or by a general analysis of the nature of teamwork. Note that the attribution of player *intentions* is also an issue here.
– **Machine vision**. For the robotic leagues (as for real-life soccer commentary), vision and image understanding is critical.
– **Presentation techniques**. Commentary is not necessarily restricted to natural language, but may also include the generation of replays to highlight interesting events, the incorporation of the visual display of statistics or graphs, and the superimposition of graphics on the screen to highlight interesting (or even upcoming) events.
– **Re-use of stored knowledge and experiences**. Many of the comments made by real-life announcers add context to a game by reviewing the past performances of a team and its players, or (in tournaments or league promotion/demotion situations) analyzing the consequences of the possible game results for each team. Also, past experience is often brought to bear to identify situations (goals, comebacks, sendings-off) similar to the current one,

and may even be used to make predictions ("Shearer scored from this position last weekend...").

- **Rules of communication**. Depending on what the audience is assumed to know the appropriate commentary will change. For example, the English expression "Back to square one" comes from the early days of radio football commentary when listeners followed the game using a template of a football pitch divided into numbered regions: the square numbered one was the center spot. The flexibility to commentate for different audiences allows investigation of issues such as Grice's cooperative principles of conversation [Grice 75] (*e.g.*, the maxim of quantity, that your contribution should be no more and no less informative than is required). Rules of communication become especially significant if a commentary *team* is being modeled, for example with an announcer following the ball-by-ball action and an expert providing higher-level analysis.
- **Incorporation of emotion**. Soccer commentary offers plentiful opportunities for studying the expression of emotion in language. Meaning (and emotion) can be expressed in speech through prosody (the phrasing and the tones used for speech), and even via changes in vocabulary. Further, the effects of the natural bias of a commentator (in favor of teams from one country, of certain styles of play, or in support of the underdogs) can also be studied.
- **Focusing the complex**. Humans watching a game of soccer can view it as a single process. However, at any one time, the individual players involved will give rise to many possible focus points, both for the game itself and for a commentary. This interaction between a medium-sized number of adaptive agents meets the criteria of a complex, adaptive system described, for example, by [Casti 97]. As yet, there are no good mathematical theories for understanding the overall behavior of such systems.

Note that a significant proportion of the topics listed here are not addressed by simply developing teams for the RoboCup leagues themselves. Thus, we view the challenge of automated soccer commentary as extending the scope of the RoboCup goals. In fact, we additionally believe that research on commentary systems also has the potential to make a significant impact on the original goals of the RoboCup project. For instance, the 1997 RoboCup Synthetic Agent Challenge [Kitano *et al* 97b] presents three specific challenges for RoboCup research:

- **Learning**. The learning of individual agents and teams.
- **Teamwork**. Multi-agent team planning and plan execution in service of teamwork.
- **Opponent Modeling**. On-line tracking of opponents' aims, on-line strategy recognition, and off-line review.

We make the observation that the automatic production of *expert analysis* of soccer play can directly facilitate each of these challenges. This is illustrated in Figure 1. The expert analysis module in this figure is an automatic system

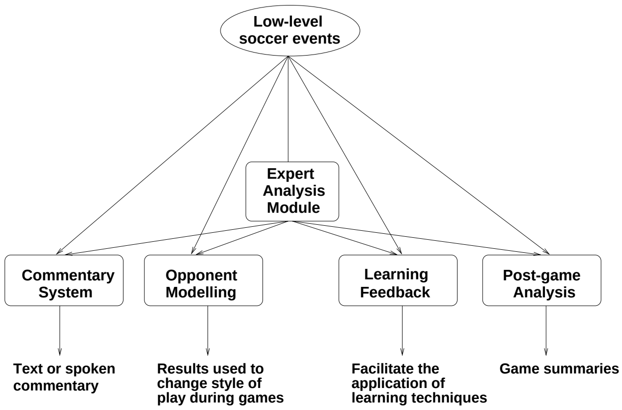

Fig. 1. Possible uses of an expert analysis module

that analyzes low-level events in a soccer game to produce an assessment of each team's playing style, tactics, strengths and weaknesses. We envisage the results of such analysis then being used for tasks such as:

- Informing a commentary system. This is our main focus in the current paper. Convincing commentary is probably not achievable without an expert analysis module.
- Opponent Modeling. This addresses one of the main RoboCup challenges — how a team should modify their play to deal with a particular opponent. Clearly, analyzing and understanding the opponents' play is a critical step in this process.
- Learning feedback. Another of the RoboCup challenges is learning. The 'credit-assignment problem' poses a serious barrier to automated learning techniques, however, since the number of goals scored by a team is only a rough indicator of how well it actually plays. An effective expert analysis module will make the use of learning techniques more efficient.
- Match analysis. An expert analysis module will also enable further interesting applications, such as post-game (or half-time) analysis. The ability to select highlights from a game and to demonstrate the strengths and weaknesses of each team (suggesting how a game was won and lost, and allowing the easy creation of game digests) will add to the interest in RoboCup competitions.

This list is only a broad characterization of the possible uses for expert analysis. However, we believe it demonstrates that an expert analysis module — in addition to being an integral part of a commentary system — directly enables work on the RoboCup learning, opponent modeling, and teamwork challenges described above. In our own work on the MIKE system, we have addressed some of the problems involved in producing such a module. In the following sections, we describe the analysis techniques we have developed, and also outline how their output is combined together within the framework of a coherent commentary.

3 MIKE — A Commentator System For Soccer Server

Although commentary systems can be envisaged for any of the RoboCup leagues, the MIKE system concentrates on just the simulation tournament. This is partly because, with current technology, teams in the simulator league produce more 'soccer-like' play than those in the robot leagues, and also partly because of the level of detail provided by the simulator.

3.1 MIKE's Input and Output

Games in RoboCup's simulator league are conducted using the Soccer Server [Noda *et al* 98]. In this simulation, the soccer field and all objects on it are 2-dimensional so that, unlike in real football, the ball cannot be kicked in the air and the players cannot make use of skills such as heading or volleying. The Soccer Server provides a real-time game log of a very high quality, sending detailed information on the positions of the players and the ball to a monitoring program every 100ms. Specifically, this information consists of: player location and orientation (for all players), ball location, and the game score and play modes (such as throw-ins, goal kicks, *etc*).

From the Soccer Server input, MIKE creates a commentary that can consist of any combination of the possible repertoire of remarks shown in Figure 2. Currently, this output is produced with the simple mechanism of template matching, converting the system's internal language into appropriate expressions in either

- **Explanation of complex events**. Formation changes, position change, and advanced plays.
- **Evaluation of team play**. Average formations, formations at a certain moment, player locations, indication of active or problematic players, winning passwork patterns, wasteful movements.
- **Suggestions for improving play**. Loose defense areas, better locations for inactive players, and 'should-have' comments about failed passes.
- **Predictions**. Prediction of passes and shots at goal. Also, prediction of game result by comparing team performance metrics against statistics compiled from a database of played matches.
- **Set pieces**. Goal kicks, throw ins, kick offs, corner kicks, and free kicks.
- **Passwork**. Basic tracking of the ball-by-ball play.

Fig. 2. MIKE's commentary repertoire

English, Japanese or French. To reduce repetition, this matching process is non-deterministic, and several templates are available for each decision. An example of MIKE's English language commentary might be "Interception by the Yellow-Team,... Yellow10 shoots!... Red4,... Yellow11's shot!... The Yellow-Team's 7th goal!! 7 to 0! Another goal by Yellow11!" MIKE also uses off-the-shelf text-to-speech packages to produce spoken commentary.

3.2 Summary of MIKE's Overall Structure

MIKE's design is described in detail in [Tanaka-Ishii *et al* 98b]. Here, we therefore give just a brief overview by presenting Figure 3. In this figure, the rectangles represent data and the ovals represent processes, which run concurrently, carrying out the following tasks:

- **Communicator**. Receives log data from the Soccer Server every 100ms and writes it into shared memory.
- **Soccer Analyzer Modules**. There are six Soccer Analyzers, of which three analyze basic events (shown in the figure as the 'Basic', 'Techniques', and 'Shoot' processes), and three carry out more high-level analysis (shown as the 'Bigram', 'Voronoi', and 'Statistic' processes). These six processes analyze the information posted to the shared memory by the Communicator, communicate with each other via the shared memory, and also post *propositions* to the proposition pool. Propositions are MIKE's internal represen-

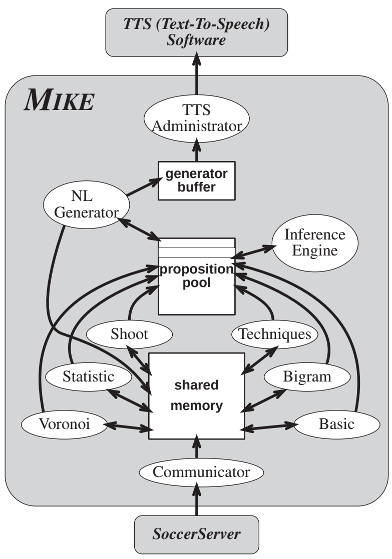

Fig. 3. MIKE's structure

tation of commentary fragments, and consist of a tag accompanied by some attributes. For example, a kick by player No.5 is represented as `(Kick 5)`, where `Kick` is the tag and `5` is the attribute. In total, MIKE has about 80 different kinds of propositions. Every proposition posted to the pool is also given an *importance* value by the Analyzer that generated it. The importance of a proposition decreases with time until it reaches zero, when it is deleted from the proposition pool without having been uttered.
- **Inference Engine**. A collection of over 50 forward chaining inference rules that identify relations between propositions in the pool. Successful firing of a rule may add new propositions (logical consequences), change the importance values of existing propositions (logical subsumption or identification of state change), or identify relations between two or more propositions (second-order relations).
- **Natural Language Generator**. Selects the proposition from the proposition pool that best fits the current state of the game (considering both the situation on the field, the importance values of the available propositions, and the commentary currently being made). Translates the proposition into natural language, using pattern-matching.
- **Text-To-Speech Administrator**. Synchronizes the Natural Language Generator's output with a text-to-speech software program.

The primary research goals in developing MIKE have been to investigate the automatic analysis of multi-agent systems, and to study the task of natural language generation in a fast-moving, real-time domain. For a detailed discussion of the natural language issues involved (such as interruption control), readers are referred to [Tanaka-Ishii *et al* 98a]. In this paper, we will concentrate on the central theme identified in §2: the generation of expert soccer analysis.

4 Expert Analysis of Soccer Within MIKE

In MIKE, the workings of the multi-agent system that is soccer are tracked and interpreted by six Soccer Analyzers. Three of these Analyzers are very low-level, 'event-based' modules that identify individual incidents on the field. These are the 'Basic' module, which identifies simple events such as ordinary kicks, the 'Techniques' module, which identifies passes, dribbles, interceptions, one-two passes, and through passes (all defined in terms of patterns of successive kicks), and the 'Shoot' module, which details shots and goals. The remaining three Analyzers are more high-level 'state-based' modules that keep track of and interpret accumulated events on the field of play. It is these that provide the expert analysis of a game.

4.1 Simple Statistics

Even basic statistics can be very useful for interpreting a game and assessing the performance of teams. MIKE's Statistics Analyzer is responsible for maintaining the following figures:

– Each player's average location and its variance. For example, Figure 4 shows the average positions of the players in the first half of the RoboCup'97 World Cup Final held in Nagoya, Japan (players numbered from 1 to 11 play left to right, and those numbered 12 to 22 play right to left). Such analysis is used by MIKE to make simple guesses about the roles each player is taking.

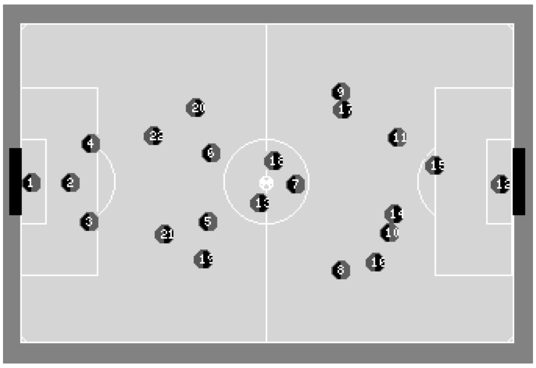

Fig. 4. Average positions of players in the RoboCup'97 final

– The average position of *all* of a team's players taken as a whole (and its variance), expressed as a distance (positive or negative) past the centerline. This figure indicates whether a team favors defense or attack. (Note that the average position of a team may also vary towards the left or right wing.)
– Average position of ball. Again, expressed as a distance past the centerline.
– The average duration of each team's possession spells, and each team's overall percentage share of the possession.
– Number of free kicks and corners taken by each team.
– Average pass distance (for each team).
– Average length of ball-play chains (for each team).

Much information on passing (such as the number of successful passes, the number of successful steals, or the length of pass sequences) is absent from this list. This is because, for MIKE, such information is the responsibility of the Bigrams Analyzer.

4.2 Analysis Based on Ball-play Chains (Bigrams)

In MIKE, ball-play chains are regarded as first-order Markov processes. Matrices representing these processes are automatically maintained by MIKE and used to describe the activity of each player. These 24×24 matrices (22 players and 2 goals) are referred to as *pass bigrams*, and record the *numbers* of each possible ball transition. As an example, consider Table 1, which shows the pass distribution and reception bigrams produced by MIKE for part of a real game (again, the first half of the RoboCup'97 World Cup Final). Only those transitions occurring more than once, plus all goals, are included in a bigram. In our examples, the left column shows the champion team, Humboldt, whose players are numbered from 1 to 11, and the right column shows Andhill, whose players are numbered 12 to 22. The letter 'G' denotes the scoring of a goal. Transitions involving giving the ball to an opponent are marked with a '+' sign.

MIKE can use these bigrams to interpret why Humboldt outperform Andhill. Of course, Humboldt scored seven goals, whereas Andhill scored none, but more than this, MIKE can assess the difference in the passing abilities of the teams by comparing the number of transitions marked with a '+' sign. Also, playmakers can be identified. For example, player No.7 (who can be seen to be a midfielder from analysis of Figure 4), successfully passes to players No.11 and No.10 (forwards), and even scores a goal himself. The main goalscorer, Player No.11, scored four times even though Figure 4 suggests he could have been marked by players No.14 and No.15. On the other hand, Andhill did not score any goals despite

Table 1. RoboCup'97 final: ball-play transitions represented as bigrams

Pass Distribution					
Humboldt			Andhill		
from	to	freq	from	to	freq
2	4	2	+12	10	4
3	5	2	+12	11	3
5	8	2	+13	5	4
+5	19	2	+13	6	4
6	7	3	+13	7	3
+6	18	2	13	14	2
7	8	2	+14	6	2
7	10	2	+14	7	2
7	11	3	+14	10	5
+7	18	2	+15	7	2
+7	20	2	+15	11	2
7	G	1	15	14	2
8	10	2	16	14	3
+8	12	2	+17	7	2
+8	14	3	19	13	3
8	G	1	+20	4	2
10	8	2	+21	5	3
+10	12	5			
+10	15	3			
+10	16	3			
10	G	1			
+11	12	3			
+11	15	3			
11	G	4			

Pass Reception					
Humboldt			Andhill		
to	from	freq	to	from	freq
4	2	2	+12	10	5
+4	20	2	+12	11	3
+5	13	4	+12	8	2
+5	21	3	13	19	3
5	3	2	14	13	2
+6	13	4	14	15	2
+6	14	2	14	16	3
+7	13	3	+14	8	3
+7	14	2	+15	10	3
+7	15	2	+15	11	3
+7	17	2	+16	10	3
7	6	3	+18	6	2
8	10	2	+18	7	2
8	5	2	+19	5	2
8	7	2	+20	7	2
+10	12	4			
+10	14	5			
10	7	2			
10	8	2			
+11	12	3			
+11	15	2			
p11	7	3			

the ball reaching player No.21, which Figure 4 identifies as a forward. Players No.12, 13, 14, 15, 20, and 21 often give the ball to the opponents, so they are either severely marked, or just very bad at passing.

To help interpret the information in a bigram, MIKE introduces the notion of a *winning passwork pattern*. This is defined as any chain of three players from the same team A, B, C, such that both the transition from A to B and the transition from B to C appear in the bigram, and C scores at least one goal. In the example of Table 1, Humboldt has six winning passwork patterns and Andhill have none, as they failed to score any goals. An example winning passwork pattern is from Player No.6 to No.7 to No.11. MIKE uses winning passwork patterns to predict passwork and shots on goal during a game.

4.3 Analysis Based on Voronoi Diagrams

To analyze the area of influence of each player, MIKE divides the field of play into *Voronoi regions*. These regions are technically defined in terms of a set of n points (or *sites*) in a plane. Given such a set of sites, the Voronoi region associated with site p is the locus of all points in the plane closer to p than to any of the other $n-1$ sites. A graphical rendering of the borders of the Voronoi regions for a given set of sites is known as the Voronoi diagram for that set. In MIKE, such diagrams are calculated for each team by using the team's players as sites. To do this efficiently, the technique of [Oishi & Sugihara 95] is used, which has a cost of $O(n)$ on average, and $O(n^2)$ at worst. This technique allows MIKE to calculate complete diagrams for each team every simulation cycle (100ms).

Using Voronoi diagrams, MIKE can determine the defensive areas covered by players and also assess overall positioning. As an example, Figure 5 shows the Voronoi diagram of a moment from the first half of the RoboCup'97 final. Humboldt's players are represented by a '+' sign, Andhill's players by a '$\diamondminus$', and the ball by a '$\boxminus$'. In this figure, the positions of Humboldt's No.7 and No.10 have also been marked. We have already seen how MIKE can use bigrams to identify Humboldt's No.7 as a playmaker. Now we see that at the instant depicted by the Voronoi diagram, No.7 is close to the meeting point of a number of the opponent's Voronoi regions, and is thus as far as possible from the neighboring Andhill players. Also, No.10 has an excellent chance to shoot at goal: there are no opponents near him and he has the ball. Further, right in front of the left goal, Humboldt has created a triangular Voronoi region. As the average shape of Voronoi regions is hexagonal, this area demonstrates Humboldt's tight defense.

Note that [Taki *et al* 96] have proposed a method of calculating players' areas of influence that takes into account each player's current speed and orientation. This model generates defensive borders in the form of higher-ordered curves. In MIKE, however, the speed and orientation of a player are taken into consideration via a simple yet very effective short-cut: before calculating the Voronoi regions for a team's players, the location of the sites is modified by displacing each player in the direction they are currently facing, for a distance proportional to their current speed.

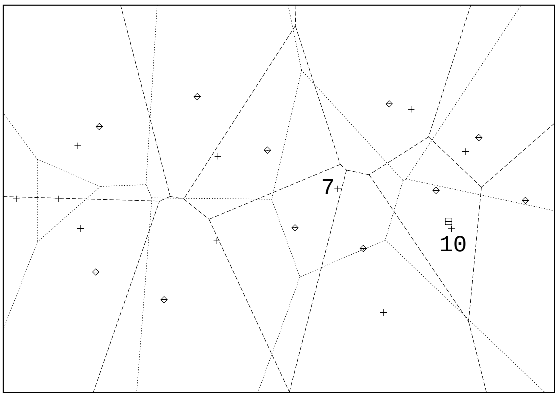

Fig. 5. Voronoi diagram (RoboCup'97 final)

4.4 Re-use of Stored Game Logs

To demonstrate the general applicability of MIKE's analyses, we processed the logs of the 26 games from the 1998 Japan Open tournament (the tournament conducted with the version of Soccer Server most similar to the one that will be used for the 1998 RoboCup in Paris). To identify the statistics most useful for highlighting differences between teams, we collected figures separately for the winners and losers of each half of each game (looking at halfs of games reduces the number of draws, and therefore also reduces the amount of data that is discarded). The results, shown in Table 2, demonstrate that, with the exception of the number of steals and the number of corners, many of the statistics have

Table 2. Average statistics, non-drawn games from the Japan Open 1998

	Winners	Losers
Goals	5.3	0.8
Shots	16.1	12.9
Average distance of a team's players past centerline (m)	-3.91	-10.84
Average number of players in ball-play chains	2.1	1.8
Total successful passes	16.45	10.36
Total number of dribbles	1.7	0.6
Total successful steals	17.6	22.9
Total winning passwork patterns	0.31	0
Number of corners taken by each team	0.23	0.13

good predictive value for identifying the stronger team. In the case of corners, the low values are explained by the nature of collisions in the simulator; rather than calculating deflections, the Soccer Server simply *stops* moving objects when their paths cross.

Base values such as those in Table 2 are employed by MIKE for predicting which team is likely to win a game. This is a simple example of how the re-use of stored knowledge, as mentioned in §2, can be applied to analyze games in real-time.

5 Related Systems

To date, there has been only limited research on automatically analyzing soccer. In terms of expert analysis techniques, the main drive for progress has been in commercial systems to assist professional coaches. Academic research on soccer commentary has also been carried out, however, focusing on the multimedia presentation of visual data and the incremental recognition of events.

5.1 SECOND LOOK

A comprehensive commercial package for analyzing football games is SECOND LOOK, a system that enables users to produce detailed statistical analyses. SECOND LOOK has been developed by Zvi Friedman, an international soccer coach, and Jon Kotas, a computer programmer, who founded a company in 1993 to market their product. On their Web pages [Sof] they claim to provide services for the American Major League Soccer (MLS), the European soccer finals in England, and to have provided the Women's Olympics Soccer Team with analysis of their opponents at the 1996 Olympics Games. A 1997 job advert found on the Internet also suggests that Manchester United may also now be making use of their system (`http://www.umist.ac.uk/INDEX_ONE/softsport.html`).

SECOND LOOK requires a user to input a game by sequentially recording the *position* on the pitch of each kick of the ball, and also the player in possession. This restriction to individual kicks facilitates the process of data entry, since the positions of the remaining players (who don't have the ball) do not have to be specified. The user-entered data is processed to automatically generate graphics, charts and detailed pitch diagrams that analyze the game. In the previous section we have already described the expert analysis capabilities of MIKE. Now we can extend the list of possible soccer analyses by summarizing the further game interpretations suggested by SECOND LOOK.

- **Pass bigrams**. Like MIKE, SECOND LOOK maintains pass bigrams, but adds the option of displaying these bigrams for individual players rather than for a whole team. This emphasis on individual players is repeated in many of SECOND LOOK's analyses. For example, the number of different players that each team member manages to pass to is compared. Also, each player's passing success rate and ball possession time is maintained separately (and compared to a database giving statistics collected over many matches for defenders, midfielders, attackers and goalkeepers).

– **Shots on goal**. How many shots does each player make? SECOND LOOK also analyses what proportion of a team's shots come from defenders, attackers, or midfielders, and where a team's shots come from on the field. In the latter case, a graphical display can be particularly revealing, indicating for example whether a team favors a particular side of the field, or close-range or long distance strikes (an example of this analysis is shown in Figure 6). Note that identifying a 'shot' is not a straightforward task for automated systems, since it involves reasoning about players intentions. In MIKE's bigrams, for instance, shots are not recorded unless they produce goals or go to the opposing keeper.
– **Impact passes**. SECOND LOOK introduces this useful concept, defined as all those passes made by a team in a series that results in a shot on goal. A graphic displaying all the chains of impact passes can reveal information such as 'all plays started from the right-hand side of the Mexican attack'. For real soccer, such pass chains are also known to often display directional changes. A shot or a goal is also counted as an impact pass.
– **Completed passes**. A completed pass is a pass between two players of the same team. Completed passes can be displayed on a pitch diagram for individual players or for an entire team. Analyzing these diagrams can lead to analysis such as which side of the pitch a team is favoring: 'the Mexican

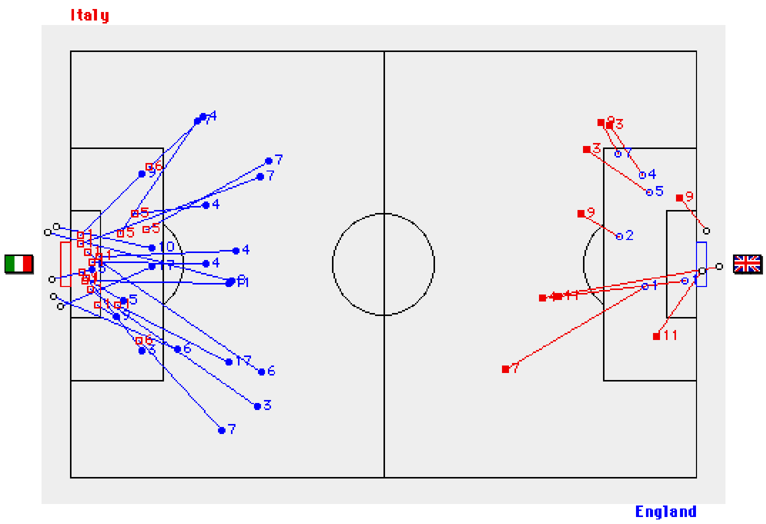

Fig. 6. Example screenshot from SECOND LOOK, showing all the shots on goal taken by Italy (left to right) and England (right to left) in their first World Cup qualifier in 1996. This image is reproduced by kind permission of SoftSport Inc.

team had a definite strategic shift to the left side of the Brazilian defense'. The number of completed passes within the opponent's third of the pitch (and in the penalty area) are also indicators of attacking strength. Further useful statistics to maintain are the percentage of each player's passes that are forward, and whether the passes received by attackers are long or short.
- **Turnover analysis**. SECOND LOOK monitors where the majority of the changes in possession occur.
- **Range**. SECOND LOOK can also produce a pitch diagram displaying every location that a player touched (kicked) the ball. This is very useful for examining how much of the pitch is covered by each player.
- **Corner kicks**. How many corner kicks did each player force? This is harder than simply counting the corner kicks taken by each team, since automatically identifying the player forcing the kick may not be trivial.

All the statistics produced by SECOND LOOK should be reproducible in the RoboCup environment. In fact, the level of information in the logs of the Soccer Server should provide even more opportunities for analysis (and also make the task of analysis still more challenging). As a first step, we plan to combine the MIKE and the SECOND LOOK expert analysis techniques described in this paper into a single program for automatically analyzing Soccer Server games. As we stated in §2, we believe that such a resource will be a powerful tool for researchers working on RoboCup.

5.2 SOCCER and ROCCO

In terms of academic research on automated soccer commentary, we are aware of only one other project: the SOCCER system described in [Andrè *et al* 88] (and more recently in [Andrè *et al* 94]). However, this system addresses a number of the research issues we identified in §2. For instance, the original version of SOCCER tackles machine vision and scene interpretation problems by using *geometrical scene descriptions* [Neumann 89] to interpret short sections of video recordings of real soccer games. It also employs multimedia presentation techniques to combine text, graphics and video in its presentations of these recordings. An updated version of SOCCER, called ROCCO, is being designed specifically for RoboCup, and is to be described in the proceedings of the 1997 RoboCup workshop (not yet in print at the time of writing this paper). Both ROCCO and SOCCER tackle the problem of incremental recognition of events, developing a model that allows the initiation of events to be recognized and commented upon before the event is actually completed.

For the task of simultaneous description, SOCCER has a control loop that selects events according to their *topicality*, where topicality is determined by the salience of an event and the time that has passed since its occurrence — a similar notion to the importance values used by MIKE. However, in SOCCER it seems that inferences are not used to identify higher-level, state-based explanations. MIKE not only does this, but also uses its importance values to guide its handling of natural language issues such as interruption, abbreviation and repetition.

Unlike MIKE, however, SOCCER can also carry out post-game analysis, for which the system has more complicated control structures, plus techniques for searching for events of certain types, modifying the viewpoint from which action is viewed, and enhancing an image to concentrate attention on important features, such as the superimposition of an arrow to indicate the forthcoming movements of ball or players. In terms of simple text, the system produces output such as 'Bommer, the midfield player, passes the ball... to Bosch the outside left... He is attacked by Müller, the outside right.'

5.3 BYRNE

Another commentary-related research project is the talking head BYRNE system described elsewhere in this volume [Binsted 98]. BYRNE uses the domain of soccer to investigate the effects of personality and emotion on the 'believability' of a designed character — in this case, a soccer commentator. BYRNE assumes that the identification of events in the domain is carried out externally and that they have 'relevance scores' that decay with time. It then generates *emotional structures* that describe the character's reactions to these events. These emotional structures have seven different possible types: fear, anger, sadness happiness, disgust and surprise, and interest. The process for generating these structures takes into account a set of *static characteristics* of the commentator, thus allowing different personalities to be explored. The final system output is then realized via a text-to-speech synthesizer and a facial animation system. BYRNE is due to be demonstrated at the 1998 RoboCup; we anticipate that it will provide an excellent example of how the domain of soccer commentary facilitates the investigation of a diverse range of research issues.

6 Conclusions

We have examined the role of automatic soccer commentary within the RoboCup initiative. Initially, we identified soccer commentary as allowing and requiring investigation of an unusually wide variety of research topics, many of which could not be addressed by the simple development of teams for the RoboCup leagues themselves. Next, we highlighted the crucial role of expert analysis, not just for the development of a commentary system, but also for the tackling of RoboCup challenges such as learning, opponent modeling and teamwork. We reviewed our own soccer commentary system, MIKE, paying special attention to its techniques for expert analysis. We also looked at related systems, highlighting the research issues that have been tackled to date, and also any further expert analysis techniques. Far from simply providing extra atmosphere in live games, then, we hope we have demonstrated that soccer commentary has the potential to play an indispensable role within the overall RoboCup initiative.

References

[Andrè *et al* 88] E. Andrè, G. Herzog, and T. Rist. On the simultaneous interpretation of real world image sequences and their natural language description: The system SOCCER. In *Proc. of the 8th ECAI*, pages 449–454, Munich, 1988.

[Andrè *et al* 94] E. Andrè, G. Herzog, and T. Rist. Multimedia presentation of interpreted visual data. In *Proceedings of AAAI-94, Workshop on Integration of Natural Language and Vision Processing*, pages 74–82, Seattle, WA, 1994.

[Binsted 98] K. Binsted. Character design for soccer commentary. In *Proc. of the Second International Workshop on RoboCup*, pages 22–36, 1998.

[Casti 97] John L. Casti. *Would-be Worlds: how simulation is changing the frontiers of science.* John Wiley and Sons, Inc, 1997.

[Grice 75] H. P. Grice. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics: Speech Acts*, volume 3, pages 41–58. Academic Press, 1975.

[Jones 96] R. Jones. Quoted in Carlisle United fanzine, August 1996.

[Kitano *et al* 97a] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara. RoboCup: A challenge problem for AI. *AI Magazine*, pages 73–85, Spring 1997.

[Kitano *et al* 97b] H. Kitano, M. Tambe, P. Stone, M. Veloso, S. Coradeschi, E. Osawa, H. Matsubara, I. Noda, and M. Asada. The RoboCup synthetic agent challenge 97. In *Proceedings of IJCAI-97*, pages 24–29, Nagoya, Japan, 1997.

[Neumann 89] B. Neumann. Natural language description of time-varying scenes. In D.L. Waltz, editor, *Semantic Structures: Advances in Natural Language Processing*, pages 167–207. Lawrence Erlbaum, 1989. ISBN 0-89859-817-6.

[Noda *et al* 98] I. Noda, H. Matsubara, K. Hiraki, and I. Frank. Soccer Server: a tool for research on multi-agent systems. *Applied Atificial Intelligence*, 12(2–3):233–251, 1998.

[Oishi & Sugihara 95] Y. Oishi and K. Sugihara. Topology-oriented divide-and-conquer algorithm for Voronoi diagrams. *Graphical Models and Image Processing*, 57(4):303–314, July 1995.

[Sof] SECOND LOOK soccer analysis software. The makers, SoftSport Inc., can be found on the Web at `http://www.softsport.com/`.

[Taki *et al* 96] T. Taki, J. Hasegawa, and T. Furukawa. Development of motion analysis system for quantitative evaluation of teamwork in soccer games. In *Proc. of ICIP'96*, 1996.

[Tanaka-Ishii *et al* 98a] K. Tanaka-Ishii, K. Hasida, and I. Noda. Reactive content selection in the generation of real-time soccer commentary. In *Proceedings of COLING-ACL'98*, pages 1282–1288, Montreal, 1998.

[Tanaka-Ishii *et al* 98b] K. Tanaka-Ishii, I. Noda, I. Frank, H. Nakashima, K. Hasida, and H. Matsubara. MIKE: An automatic commentary system for soccer. In *Proceedings of ICMAS-98*, pages 285–292, 1998.