Multi-Agent Explanation Strategies in Real-Time Domains

Kumiko Tanaka-Ishii

University of Tokyo,
7-3-1 Hongo Bunkyo-ku
Tokyo 113-8656
Japan
kumiko@ipl.t.u-tokyo.ac.jp

Ian Frank

Electrotechnical Laboratory
1-1-4 Umezono, Tsukuba
Ibaraki 305-0085
Japan
ianf@etl.go.jp

Abstract

We examine the benefits of using multiple agents to produce explanations. In particular, we identify the *ability to construct prior plans* as a key issue constraining the effectiveness of a single-agent approach. We describe an implemented system that uses multiple agents to tackle a problem for which prior planning is particularly impractical: real-time soccer commentary. Our commentary system demonstrates a number of the advantages of decomposing an explanation task among several agents. Most notably, it shows how individual agents can benefit from following different discourse strategies. Further, it illustrates that discourse issues such as controlling interruption, abbreviation, and maintaining consistency can also be decomposed: rather than considering them at the single level of one linear explanation they can also be tackled separately within each individual agent. We evaluate our system's output, and show that it closely compares to the speaking patterns of a human commentary team.

1 Introduction

This paper deals with the issue of high-level *vs* low-level explanation strategies. How should an explanation find a balance between describing the overall, high-level properties of the discourse subject, and the low-level, procedural details? In particular, we look at the difficulties presented by domains that change in real-time. For such domains, the balance between reacting to the domain events as they occur and maintaining the overall, high-level consistency is critical.

We argue that it is beneficial to decompose the overall explanation task so that it is carried out by

more than one agent. This allows a single agent to deal with the tracking of the low-level developments in the domain, leaving the others to concentrate on the high-level picture. The task of each individual agent is simplified, since they only have to maintain consistency for a single discourse strategy. Further, discourse issues such as controlling interruption, abbreviation, and maintaining consistency can also be decomposed: rather than considering them at the single level of one linear explanation they can be tackled separately within each individual agent and then also at the level of inter-agent cooperation.

We look at real-world examples of explanation tasks that are carried out by multiple agents, and also give a more detailed protocol analysis of one of these examples: World Cup soccer commentary by TV announcers. We then describe an actual implementation of an explanation system that produces multi-agent commentary in real-time for a game of simulated soccer. In this system, each of the agents selects their discourse content on the basis of importance scores attached to events in the domain. The interaction between the agents is controlled to maximise the importance score of the uttered comments.

Although our work focuses on real-time domains such as soccer, our discussion in §2 puts our contribution in a wider context and identifies a number of the general benefits of using multiple agents for explanation tasks. We chose the game of soccer for our research primarily because it is a multi-agent game in which various events happen simultaneously on the field. Thus, it is an excellent domain to study real-time content selection among many heterogeneous facts. A second reason for choosing soccer is that detailed, high-quality logs of simulated soccer games are available on a real-time basis from Soccer Server, the official soccer simulation system for the 'RoboCup' Robotic Soccer World Cup initiative (Kitano et al., 1997).

Figure 1: Common explanation tasks categorised according to the ease of planning them in advance

2 Explanation Strategies

In this paper, we use the term *explanation* in its broadest possible sense, covering the entire spectrum from planned lectures to commentating on sports events. Any such explanation task is affected by many considerations, including the level of knowledge assumed of the listeners and the available explanation time. However, the issue we mainly concentrate on here has not previously received significant attention: the benefits of splitting an explanation task between multiple agents.

2.1 Explanations and Multi-Agency

The general task of producing explanations with multiple agents has not been studied in depth in the literature. Even for the ‘naturally’ multi-agent task of soccer commentary, the systems described in the recent AI Magazine special issue on RoboCup (André et al., 2000) are all single-agent. However, one general issue that has been studied at the level of single agents is the trade-off between low-level and high-level explanations. For example, in tutoring systems (Cawsey, 1991) has described a system that handles real-time interactions with a user by separating the control of the content planning and dialogue planning.

We believe that the key issue constraining the use of high-level and low-level explanations in a discourse is the *ability to construct prior plans*. For example, researchers in the field of discourse analysis, (e.g., (Sinclair and Coulthard, 1975)) have found that relatively formal types of dialogues follow a regular hierarchical structure. When it is possible to find these kinds of *a priori* plans for a discourse to follow, approaches such as those cited above for tutoring are very effective. However, if prior plans are hard to specify, a single agent may simply find it becomes overloaded. Typically there will be two conflicting goals: deal with and explain each individual (unplanned) domain event as it occurs, or build up and explain a more abstract picture that conveys the overall nature of the explanation topic.

Thus, for any changing domain in which it is hard to plan the overall discourse, it can be beneficial to divide the explanation task between mul-

tiple agents. Especially for real-time domains, the primary benefit of decomposing the explanation task in this way is that it allows each agent to use a different discourse strategy to explain different aspects of the domain (typically, high-level or low-level). However, we can see from Figure 1 that even some activities that are highly planned are sometimes carried out by multiple agents. For example, business presentations are often carried out by a team of people, each of which is an expert in some particular area. Clearly, there are other benefits that come from decomposing the explanation task between more than one agent. We can give a partial list of these here:

- Agents may start with different abilities. For example, in a panel session, one panellist may be an expert on Etruscan vases, while another may be an expert on Byzantine art.
- It can take time to observe high-level patterns in a domain, and to explain them coherently. Having a dedicated agent for commenting on the low-level changes increases the chance that higher-level agents have a chance to carry out analysis.
- A team of agents can converse together. In particular, they can make explanations to *each other* instead of directly explaining things to the listeners. This can be a more comfortable psychological position for the listener to accept new information.
- The simple label of “expert” adds weight to the words of a speaker, as shown convincingly by the research of (Reeves and Nass, 1996). The use of multiple agents actually gives a chance to *describe* individual agents as “experts” on specific topics.
- Even a single agent speaking in isolation could describe itself as an expert on various topics. However, (Reeves and Nass, 1996) also show that self-praise has far less impact than the same praise from another source. Rather than describing themselves as experts, a multi-agent framework allows agents to describe the *other* agents as experts.

To illustrate the different roles that can be taken by multiple agents in an explanation task, we carried out a simple protocol analysis of an example from the far left of our scale of Figure 1: soccer commentary.

2.2 Soccer Protocol Analysis

We analysed the video of the NHK coverage of the first half 1998 World Cup final. This commentary was carried out by a team of two people who we call the ‘announcer’ and the ‘expert’. The figures in Table 1 demonstrate that there are clear differences between the roles assumed by this commentary team. Although both use some background knowledge to fill out their portions of the commentary, the announcer mostly comments on low-level events, whilst the expert mostly gives higher-level, state-based information. Further, we can see that the announcer asked questions of the expert with a high frequency.

Overall, there is a clear indication that one agent follows the low-level events and that the other follows the high-level nature of the game. Accordingly, their discourse strategies are also different: the announcer tends to speak in shorter phrases, whereas the expert produces longer analyses of any given subject. The commentary team collaborates so that the consistency between high-level, low-level, and background comments is balanced within the content spoken by each individual, and also within the overall commentary.

2.3 A First Implementation

As a first step towards a multi-agent explanation system based on the above observations, the following sections describe how we implemented a commentary system for a game of simulated soccer. Our experience with this system reflected the discussion above in that we found it was very difficult to consistently manage all the possible discourse topics within a single-agent framework. When changing to a multi-agent system, however, we found that a small number of simple rules for inter-agent interaction produced a far more manageable system. We also found that the system was behaviourally very similar to the protocol of Table 1.

3 An Architecture For Multi-Agent Soccer Commentary

Figure 2 shows the basic architecture of our soccer commentator system. As we mentioned in the Introduction, this system is designed to produce live commentary for games played on RoboCup’s Soccer Server. Since the Soccer Server was orig-

inally designed as a testbed for multi-agent systems (Noda et al., 1998), we call our commentator MIKE (“Multi-agent Interactions Knowledgeably Explained”). Typically, MIKE is used to add atmosphere to games played in the RoboCup tournaments, so we assume that the people listening to MIKE can also see the game being described.

Figure 2: MIKE — a multi-agent commentator

The Soccer Server provides a real-time game log of a very high quality, sending information on the positions of the players and the ball to a monitoring program every 100msec. Specifically, this information consists of 1) player location and orientation, 2) ball location, and 3) game score and play modes (throw ins, goal kicks, etc).

This information is placed in MIKE’s shared memory, where it is processed by a number of ‘Soccer Analyser’ modules that analyse higher-level features of a game. These features include statistics on player positions, and also ‘bigrams’ of ball play chains represented as first order Markov chains. The Voronoi analyser uses Voronoi diagrams to assess game features such as defensive areas. Note that we do not consider the Soccer Analysers to be ‘agents’; they are simply processes that manipulate the information in the shared memory. The only true ‘agents’ in the system are the Announcer and the Analyser, which communicate both with each other and with the audience.

All information in MIKE’s shared memory is represented in the form of commentary fragments that we call *propositions*. Each proposition consists of a tag and some attributes. For example, a pass from player No.5 to No.11 is represented as (Pass 5 11), where *Pass* is the tag, and the

Commentary Feature	Announcer	Expert	Note
Background comment (<i>e.g.</i> , on stadium, or team backgrounds)	7%	20%	(predefined plan)
Event-based comment	82%	3%	(low-level)
State-based comment	11%	77%	(high-level)
Average length of comment	1.3sec	3.8sec	(consistency)
Asks a question to the other	30	0	(new explanation mode)
Interrupts the other	5	0	(priority of roles)
Announcer describes expert as expert	0	n/a	(adds weight to expert)

Table 1: Protocol analysis of announcer and expert utterances in professional TV coverage of soccer

	Local	Global
E v e n t	Kick Pass Dribble ShootPredict	ChangeForm SideChange
S t a t e	Mark PlayerPassSuccessRate ProblematicPlayer PlayerActive	TeamPassSuccessRate AveragePassDistance Score Time

Table 2: Examples of MIKE’s proposition tags

numbers 5 and 11 are the attributes. MIKE uses around 80 different tags categorised in two ways: as being local or global and as being state-based or event-based. Table 2 shows some examples of categorised proposition tags.

The operation of the Announcer and the Analyst agents is described in detail in the following section. Basically, they select propositions from the shared memory (based on their ‘importance scores’) and process them with inference rules to produce higher-level chains of explanations. The discourse control techniques of interruption, repetition, abbreviation, and silence are used to control both the dialogue strategies of each individual agent and also the interaction between them.

To produce variety in the commentary, each possible proposition is associated with several possible commentary templates (output can be in English or Japanese). Figure 3 shows the overall repertoire of MIKE’s comments. The actual spoken commentary is realised with off-the-shelf text-to-speech software (Fujitsu’s Japanese Synthesiser for Japanese, and DecTalk for English).

4 Multi-Agent NL Generation

In this section, we describe how MIKE uses importance scores, real-time inferencing, and discourse control strategies to implement — and control the interaction between — agents with differing expla-

- **Explanation of complex events.** Formation and position changes, advanced plays.
- **Evaluation of team plays.** Average formations, formations at a certain moment, players’ locations, indication of active or problematic players, winning passwork patterns, wasteful movements.
- **Suggestions for improving play.** Loose defence areas, better locations for inactive players.
- **Predictions.** Passes, game results, shots.
- **Set pieces.** Goal kicks, throw ins, kick offs, corner kicks, free kicks.
- **Passwork.** Tracking of basic passing play.

Figure 3: MIKE’s repertoire of statements

nation strategies.

To form a single coherent commentary with multiple agents we extended the single-agent framework of (Tanaka-Ishii et al., 1998). The basic principle of this framework is that given a set of scores that capture the information transmitted by making any utterance, the most effective dialogue is the one that maximises the total score of all the propositions that are verbalised. We therefore created two agents with different strategies for content scheduling. One agent acts as an announcer, following the low-level events on the field. This agent’s strategy is biased to allow frequent topic change and although it uses inference rules to look for connections between propositions in the shared memory, it only uses short chains of inference. On the other hand, the second agent acts as an ‘expert analyst’, and is predominantly state based. The expert’s strategy is biased to have more consistency, and to apply longer chains of inference rules than the announcer.

4.1 Importance Scores

In MIKE, importance scores are designed to capture the amount of information that any given

proposition will transmit to an audience. They are not fixed values, but are computed from scratch at every game step (100msec). The importance score of each proposition depends on three factors: 1) the elapsed time since the proposition was generated, 2) for event-based propositions, a comparison of the place associated with the proposition and the current location of the ball, and 3) the frequency that the proposition has already been stated. To keep the number of comments in the shared memory to a manageable number they are simply limited in number, with the oldest entries being removed as new propositions are added.

4.2 Real Time Inference

MIKE's commentary propositions are the results of large amounts of real-time data processing, but are typically low-level. A commentary based solely on these propositions would be rather detailed and disconnected. Thus, to analyse the play more deeply, MIKE gives the commentary agents access to a set of forward-chaining rules that describe the possible relationships between the propositions. In total, there are 145 of these rules, divided into the two classes of logical consequences and second order relations. We give a representative example from each class here:

- **Logical consequence:**
 (PassSuccessRate *player* *percentage*)
 (PassPattern *player* Goal) $\rightarrow$
 (active *player*)
- **Second order relation:**
 (PassSuccessRate *player* *percentage*)
 (PlayerOnVoronoiLine *player*) $\rightarrow$
 (Reason @1 @2)

The basic premise of the announcer's dialogue strategy is to follow the play by repeatedly choosing the proposition with the highest importance score. Before stating this proposition, however, the announcer checks any applicable inference rules in a top down manner, in an attempt to produce higher-level commentary fragments and background related information. In contrast to this, the expert agent has a library of *themes* (e.g., pass statistics, formation, stamina) between which it chooses based on the propositions selected by the announcer so far. It then uses inference rules to try to construct a series of high-level inferences related to the theme. The expert applies rules until it succeeds in constructing a single coherent piece of structured commentary. When it is the agent's turn to speak it can then send this commentary to the TTS software.

4.3 Discourse Control Strategies

Consider a passage of commentary where the announcer is speaking and a proposition with a much larger importance score than the one being uttered appears in the shared memory. If this occurs, the total importance score may become larger if the announcer immediately interrupts the current utterance and switches to the new one. As an example, the left of Figure 4 shows (solid line) the change of the importance score with time when an interruption takes place (the dotted line represents the importance score without interruption). The left part of the solid line is lower than the dotted, because we assume that the first utterance conveys less of its importance score when it is not completely uttered. However, the right part of the solid line is higher than the dotted line, because the importance of the second utterance will be lower by the time it is uttered without interrupting the commentary. Note that after selecting a proposition to be uttered, its importance score is assumed to decrease with time (as indicated in the figure, the decrease is computed dynamically and will be different for each proposition, and often not even linear). The decision of whether or not to interrupt is based on a comparison of the area between the solid or dotted lines and the horizontal axis.

Similarly, it may happen that when the two most important propositions in shared memory

Figure 4: Change of importance score on interruption and abbreviation

Figure 5: Increase in importance scores caused by emphasis of repeating a proposition

are of similar importance, the amount of communicated information can best be maximised by quickly uttering the most important proposition and then moving on to the second before it loses importance due to some development of the game situation. This is illustrated in the second graph of Figure 4. Here, the left hand side of the solid line is lower than that of the dotted because an abbreviated utterance (which might not be grammatically correct, or whose context might not be fully given) transmits less information than a more complete utterance. But since the second proposition can be uttered before losing its importance score, the right hand part of the solid line is higher than that of the dotted. As before, the benefits or otherwise of this modification should be decided by comparing the two areas made by the solid and the dotted line with the horizontal axis.

We originally designed these techniques just to improve the ability of the announcer agent to follow the play. With two commentary agents, however, both interruption and abbreviation can be adapted to control inter-agent switching. In MIKE, the default operation is for the announcer to keep talking while the ball is in the final third of the field, or while there are important propositions to utter. When the announcer has nothing to say, the expert agent can speak or both agents can remain silent. If the expert agent chooses to speak, it may happen that an important event on the field makes the announcer wants to speak again. We model both interruption and abbreviation as multi-agent versions of the graphs of Figure 4: the agent speaking the first utterance is the expert and the agent speaking the second is the announcer.

We use two further discourse control techniques in MIKE: repetition and silence. Repetition is depicted in Figure 5. Sometimes it can happen that the remaining un-uttered propositions in the shared memory have much smaller scores than

any of those that have already been selected. In this case, we allow individual agents to repeat things that they have previously said. Also, we allow them to repeat things that the *other* agent has said, to increase the effectiveness of the dialogue between them. Finally, we also model silence by adding bonuses to the importance scores of the propositions uttered by the commentators. Specifically, we add a bonus to the scores of propositions uttered directly before a period where both commentators are silent (the longer that a commentary continues uninterrupted, the higher the silence bonus). This models the benefit of giving listeners time to actually digest the commentary. Also, a period of silence contributes a bonus to the importance scores of the immediately following propositions. This models the increased emphasis of pausing before an important statement.

4.3.1 Communication Templates

To improve the smoothness of the transfer of the commentary between the two agents we devised a small number of simple *communication templates*. The phrases contained in these templates are spoken by the agents whenever MIKE realises that the commentary is switching between them. For the purposes of keeping the two agents distinct, the expert agent is referred to by the announcer as E-MIKE (“Expert Mike”). To pass the commentary to the Expert, the Announcer can use a number of phrases such as “E-MIKE, over to you”, “Any impressions, E-MIKE?”, or just “E-MIKE?”. The announcer can also pass over control by simply stopping speaking. If the commentary switches from Announcer to Expert with a question, the Expert will start with “Yes...” or “Well...”.

The communication templates for passing the commentary in the other direction (Expert to Announcer) are shown in Table 3. To help listeners distinguish the dialogue between the Announcer and Expert better, we also use a female voice for one agent and a male voice for the other.

5 Evaluation

MIKE is robust enough for us to have used it to produce live commentary at RoboCup events, and to be distributed on the Internet (it has been downloaded by groups in Australia and Hungary and used for public demonstrations). A short example of MIKE’s output is shown in Figure 6.

To evaluate MIKE more rigorously we carried out two questionnaire-based evaluations, and also a log comparison with the data produced from the real-world soccer commentary in §2. For the first of the questionnaire evaluations, we used as sub-

Question	Scale	Results
Is the game better with or without commentary?	(5=with, 1=without)	4.97
Was the commentary easy to understand?	(5=easy, 1=hard)	3.44
Were the commentary contents accurate?	(5=correct, 1=incorrect)	3.25
Was the commentary informative?	(5=yes, 1=no)	3.53
Did you get tired of the commentary?	(5=no, 1=quickly)	3.97

Table 4: Average responses of 20 subjects to first questionnaire evaluation of (two-agent) MIKE

Question	Scale	1-agent	2-agent	Diff
Is the game better with or without...?	(5=with, 1=without)	4.45	4.45	0%
Was the commentary easy to understand?	(5=easy, 1=hard)	2.95	3.25	+10%
Were the commentary contents accurate?	(5=correct, 1=incorrect)	2.65	2.95	+11%
Was the commentary informative?	(5=yes, 1=no)	3.15	3.35	+6%
Did you get tired of the commentary?	(5=no, 1=quickly)	2.35	3.35	+43%

Table 5: Difference in response with ten subjects when viewing 1-agent and 2-agent versions of MIKE

Announcer interrupts expert
Sorry, E-MIKE.
Have to stop you there E-MIKE.
Oh!...
But look at this!
Announcer speaks when expert stops
Thanks.
That's very true.
Thanks E-MIKE.
Maybe that will change as the game goes on.
OK...

Table 3: Phrases used by announcer when interrupting the expert, or when speaking after the expert agent has simply stopped (no interruption)

jects twenty of the attendees of a recent RoboCup Spring camp. All these subjects were familiar with the RoboCup domain and the Soccer Server environment. We showed them an entire half of a RoboCup game commentated by MIKE and colated their responses to the questions shown in Table 4. These results largely show that the listeners found the commentary to be useful and to contain enough information to maintain their attention. We also included some open-ended questions on the questionnaire to elicit suggestions for features that should be strengthened or incorporated in future versions of MIKE. The most frequent responses here were requests for more background information on previous games played by the teams (possible in RoboCup, but to date we have only done this thoroughly for individ-

Announcer: yellow 9,in the middle of the field,yellow team (<i>a set play happened here</i>). Any impressions, E-MIKE?
Analyser: Well, here are statistics concerning possessions, left team has slightly smaller value of possession, it is 43 percent versus 56. right team has rather higher value of territorial advantage, Overall, right team is ahead there. (<i>Score is currently 0-0. E-MIKEjudges that red team is doing better</i>).
Announcer: Really. dribble , yellow 3, on the left, great long pass made to yellow 1, for red 6, red 2's pass success rate is 100 percent. E-MIKE?
Analyser: Looking at the dribbles and steals, red team was a little less successful in dribbling, red team has a lower value of dribble average length, left is 21 meters whereas right is 11, right team has a few less players making zero passes, yellow team has made slightly less stealing,
Announcer: wow (<i>interruption because red 11 made a shot</i>), red 11, goal, red 11, Goal ! It was red 10, And a pass for red 11 ! The score is 0 1!

Figure 6: Example of MIKE's commentary from RoboCup'98 final

ual games), more conversation between the agents (we plan to improve this with more communication templates), and more emotion in the voices of the commentators (we have not yet tackled such surface-level NLG issues). We also asked what the ideal number of commentators for a game would be; almost all subjects replied 2, with just two replying 3 and one replying 1.

The above results are encouraging for MIKE, but to show that the use of multiple agents was actually one of the reasons for the favourable audience impression, we carried out a further test. We

Commentary feature	Announcer	Expert	Note
Background comment	16%	22%	(predefined plan)
Event-based comment	64%	0%	(low-level)
State-based comment	20%	78%	(high-level)
Average length of comment	1.1sec	2.9sec	(consistency)
Asks a question to the other	12.2	0	(new explanation mode)
Interrupts the other	8.6	0	(priority of roles)
Announcer describes expert as expert	0	n/a	(adds weight to expert)

Table 6: Breakdown of MIKE’s agent utterances over ten randomly selected RoboCup half-games

created a single-agent version of MIKE by switching off the male/female voices in the TTS software and disabling the communication templates. This single-agent commentator comments on almost exactly the same game content as the multi-agent version, but with a single voice. We recruited ten volunteers with no prior knowledge of RoboCup and showed them both the single-agent and multi-agent versions of MIKE commentating the same game as used in the previous experiment. We split the subjects into two groups so that one group watched the multi-agent version first, and the other watched the single-agent version first. Table 5 shows that the average questionnaire responses over the two groups were lower than with the subjects who were familiar with RoboCup, but that the multi-agent version was more highly evaluated than the single-agent version. Thus, even the superficially small modification of removing the agent dialogue has a measurable effect on the commentary.

Finally, we analysed MIKE’s commentary using the same criteria as our protocol analysis of human soccer commentary in §2.2. We selected ten half-games at random from the 1998 RoboCup and compiled statistics on MIKE’s output with an automatic script. The results of this analysis (Table 6) show a marked similarity to those of the human commentators. This initial result is a very encouraging sign for further work in this area.

6 Conclusions

We have argued for superiority of producing explanations with multiple, rather than single, agents. In particular, we identified the *difficulty of producing prior plans* as the key issue constraining the ability of a single agent to switch between high-level and low-level discourse strategies.

As a first step towards a multi-agent explanation system with solid theoretical underpinnings, we described the explanation strategies used by our live soccer commentary system, MIKE. We showed how a set of importance scores and infer-

ence rules can be used as the basis for agents with different discourse strategies, and how the discourse control techniques of interruption, abbreviation, repetition and silence can be used not just to moderate the discourse of an individual agent, but also the interaction between agents. We evaluated MIKE’s output through listener surveys, showing that it represents an advance over existing commentary programs, which are all single-agent. We also found that the discourse strategies of MIKE’s agents closely resembled those revealed by the protocol analysis of a team of real-life soccer commentators.

References

- E. André, K. Binsted, K. Tanaka-Ishii, S. Luke, G. Herzog, and T. Rist. 2000. Three RoboCup simulation league commentator systems. *AI Magazine*, 21(1):57–66, Spring.
- A. J. Cawsey. 1991. Generating intreractive explanations. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, pages 86–91.
- H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara. 1997. RoboCup: A challenge problem for AI. *AI Magazine*, pages 73–85, Spring.
- I. Noda, H. Matsubara, K. Hiraki, and I. Frank. 1998. Soccer Server: a tool for research on multi-agent systems. *Applied Artificial Intelligence*, 12(2–3):233–251.
- B. Reeves and C. Nass. 1996. *The Media Equation*. CSLI Publications.
- J. Sinclair and R. Coulthard. 1975. *Towards an Analysis of Discourse: The English Used by Teachers and Pupils*. Oxford University Press.
- K. Tanaka-Ishii, K. Hasida, and I. Noda. 1998. Reactive content selection in the generation of real-time soccer commentary. In *Proceedings of COLING-ACL’98*, pages 1282–1288, Montreal.