Dual Ecologies of Robot as Communication Media: Thoughts on Coordinating Orientations and Projectability

Hideaki Kuzuoka*, Keiichi Yamazaki, Akiko Yamazaki***, Jun'ichi Kosaka*, Yasuko Suga**, Christian Heath******

*Institute of Engineering Mechanics and Systems
University of Tsukuba,
1-1-1 Tennoudai, Tsukuba, Ibaraki, Japan
{Kuzuoka, kosaka}@esys.tsukuba.ac.jp

**Department of Liberal Arts
Saitama University
255 Shimo-ohkubo, Sakura-ku, Saitama, Japan
{yamakei, ysuga}@post.saitama-u.ac.jp

***The Department of Media Architecture
Future University Hakodate
116-2 Kamedanakano-cho, Hakodate, Hokkaido, Japan
akikoy@fun.ac.jp

****WIT Research Group
King's College London
Franklin-Wilkins Building, London SE1 9NN UK
Christian.Heath@kcl.ac.uk

Abstract

The aim of our study is to investigate systems for supporting remote instruction via a mobile robot. In the real world, instructions are typically given through words and body orientations such as head movements, which make it possible to project others' actions. Projectability is an important resource in organizing multiple actions among multiple participants in co-ordination with one another. It can likewise be said that in the case of robot-human collaboration, it is necessary to design a robot's head so that a local participant can project the robot's (and remote person's) actions. GestureMan is a robot that is designed to support such projectability properties. It is argued that a remote controlled mobile robot, designed as a communication medium, makes relevant dual ecologies: ecology at a remote (robot operator's) site and at a local participant's (robot's) site. In order to design a robot as a viable communication medium, it is essential to consider how these ecologies can be mediated and supported.

Categories & Subject Descriptors: H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces – computer-supported cooperative work; H.4.1 [**Information Systems Applications**]: Office Automation – groupware; H.4.3 [Information Systems Application]: Communications Applications – computer conferencing, teleconferencing and videoconferencing

General Terms: Human Factor

Keywords: human-robot interaction, robot-mediated communication

INTRODUCTION

In recent years, the question of supporting remote collaboration in the real environment has been one of the major topics in CSCW [16]. Our answer to this problem is to develop robots that can serve as a mediator of communication between geographically distributed environments [12, 13, 22] (see also [17] for related work by Paulos and Canny). In particular, we are aiming at supporting robot-mediated remote instruction for performing physical tasks in the real world, such as fixing machinery and medical treatment.

Our latest development is a mobile robot named GestureMan. We conducted several remote instruction experiments with the robot. However, analyses of the results of the experiments revealed some major problems with the current system. One such problem is its difficulty to support the function of "projectability" in communication. By "projectability" we are broadly referring to the capacity of participants to predict, anticipate, or prefigure the unfolding of action [1, 21]. When introducing a remote-controlled robot, there are two related sites: a robot's operator's site (henceforth called the "remote" site) and the participants in the same site as the robot (henceforth the "local" site). We are interested in investigating how bodily orientation and projectability can be coordinated via the robot.

The problem of projectability and the remote-controlled robot raises another problem: the issue of what can be termed as "dual ecologies", i.e., the ecology of the remote site, and the ecology of the local site. This is a longstanding issue in the CSCW community. For example, Gaver [3] discussed this topic in relation to the "affordance" of communication media, Heath et al. [8] discussed it in relation to communicative asymmetry in media space, and Luff et al. [14] described it as "fractured ecologies". In this paper, we discuss these problems by specifically focus on the issue of projectability when tele-communication is mediated by a robot.

In this paper, in discussing the problem of dual ecologies, we first examine how orientations and projectability are organized in the real world, based on the observation of visitors' behavior in a museum. Second, we introduce a novel mobile robot named the GestureMan, and describe an experiment using the robot. Finally, based on our experiments, we discuss the characteristics and issues raised by the dual ecologies in which the robot is placed.

ORIENTATIONS AND PROJECTABILITY IN THE REAL ENVIRONMENT

We have conducted an ethnographical study in a science museum to observe how people interact in the real environment when they experience an exhibition. We videotaped visitors' behavior and analyzed how orientations and projectability are organized in the real world. As a result, we made the following observations.

- Orientations toward an object are shown in multiple ways by the human body, such as the direction and movements of parts of the body: the head, the eyes, the torso, the hands and fingers, the legs, etc.
- Orientations of such parts of the body as well as words enable one to project/anticipate others' actions.
- Projectability is an important resource for coordinating multiple actions among multiple participants.

This kind of organization of bodily orientations is often discussed in ethnomethodological conversation analysis and interaction analysis [4, 5, 6, 11, 18, 19]. For example, Schegloff [18] shows that an utterance like 'Can I ask a question?' suggests not only a simple question but also an 'action projection' for participants to prepare for a question. Also Goodwin [4] analyzes how people give instructions at excavations in archeology classes. He points out that not only words but head (gaze) movements are also an important factor when giving instructions. Furthermore, Streek [20] clearly discusses how gesture relates with projectability. In this section, we go a step further to observe how bodily orientation is related to projectability.

Take an example which shows how multiple bodily orientations and projectability organize interaction among participants. Figure 1 shows visitors' experience of an exhibition called Elekitel which discharges sparks when a visitor rotates the handle. In this transcript, M1 stands for the father, B1 stands for the boy and G1 stands for the girl. Descriptions of bodily actions are supplied in parentheses. There are two lines for each utterance: the first line is the original Japanese talk and the second line is an English translation.

In line 2, the father (M1) looks to the left as he finishes talking (Ah so? 'Oh really?'), and steps backwards. Without speaking, the boy (B1) looks toward the same direction as M1. Then, the girl (G1) releases the handle and promptly follows M1. Finally, all three walk away from this exhibition at the same time.

In this occasion, the father's "preparatory action", i.e. head movement away from the exhibit, organizes the participation framework. The head movement is a resource for the two children to project what the father will focus on next. Therefore, the children can look toward the same direction concurrently as the father, which leads to their simultaneous action to leave the exhibit.

```
1  B1 : (nods and points at the ceiling)
        Kon'na akaruku cha yoku mie nai yo.
     It's hard to see 'cuz it's so bright in here.
```

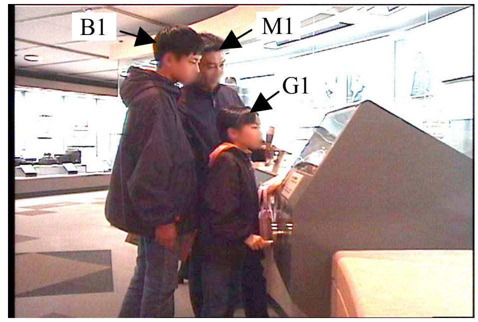


```
2  M1: A so? (looks to the left as he finishes
       talking, and steps backwards)
     Oh really?
```

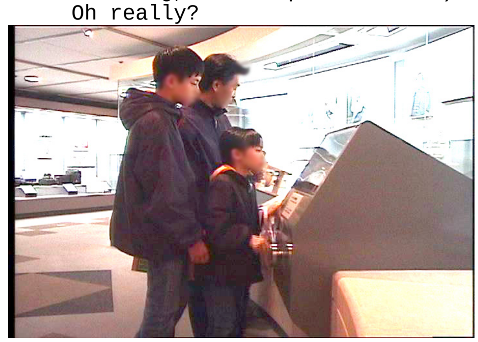

```
3  B1: (looks toward the same direction as M1)
4  G1: (releases the handle and promptly follows
       M1)
5  M1: (walking away) Otousan ni wa mieru
        Daddy can see it.
```

Figure 1. Projectability of head orientation in visitor's behavior.

From the above example, it can be seen that one's head movement projects what to see and what to do next. The head movement expresses one's subsequent focal point. This in turn permits co-participants to coordinate their focal point by moving their gaze toward the same direction. By sharing the same focal point (directionality), the next action becomes projectable for the co-participants.

Then, how can head movements be supported by a communicative robot? This issue is somewhat related to the work on how a person and a robot achieve joint attention, "i.e., they come to refer to the same real world situation [9]."

Most of such studies deal with the explicit interaction between a person and a robot, such as finger pointing. However, we are interested in supporting projectability by a robot's preparative actions before it takes explicit actions to achieve joint attention. More generally speaking, we are interested in how a robot's bodily actions can project area to be referred to, behavior of a robot, and time frame that such area and behavior lasts. In the next section, we briefly describe our newly designed robot named GestureMan, followed by a discussion of the problems of the system with respect to projectability.

GESTUREMAN

We have developed a series of robots to mediate telecommunication [12, 13, 22], and Figure 2 is our latest development named GestureMan. In order to support projectability properties, we considered following requirements in designing the robot.

- ***Supporting Mutual Observation***: Participants should be able to mutually observe each other's activities and orientations.
- ***Supporting body parts orientations***: Orientations of robot body parts should be clearly expressed so that local participant can easily recognize where a remote person is orienting to.

For a remote person to monitor local participant's orientation, it is quite effective to equip a robot with a camera which has a wide field of view. As shown in figure 2, GestureMan has a three-camera unit on its head. Since the lenses of the individual cameras have a 60-degree horizontal field of view, the remote participant will have a total horizontal field of view of about 180 degrees.

The remote person controls the robot by means of three LCD monitors facing him/her, making it possible to observe images that are captured by the three-camera unit (Fig. 3). GestureMan is also equipped with a laser pointer which is operated by a remote person and it can indicate locations of interest in a local space.

We have conducted several experiments using the system. Figure 4 is an example where the robot was used as a guide in a science museum. In this experiment, a remote person gave explanations about some of the exhibits to the visitors. Through these experiments, we found that the wide field-of-view monitor system was effective because:

- As seen in figure 4, local participants often stood sideways in relation to the robot, and therefore they are mostly seen in either the right or left monitor. Thus the wide field of view is important for a remote person to observe the local participants' orientations.
- A remote person needs to find various objects in the local participants' site in order to provide explanations for them. As can be easily imagined, it becomes much easier for a remote person to find objects in a local participants' site.

Figure 2. GestureMan

Figure 3: Three LCD monitors for a remote person.

Figure 4. Remote guide experiment at the science museum.

In order to support local participants to monitor a remote person's orientations, the robot should have resources to express such orientations. For this purpose, a helmet is mounted over the three-camera unit so that it looks like a head (Fig. 2). Also two ears and a visor are attached to the helmet. In this way the head gives the local participants clues as to where the middle camera is oriented to. Panning and tilting motions of the head is made possible by two motors at the neck. In this way, the robot can show both bodily orientation and head orientation independently. Indeed, during the experiment in the science museum, we found many instances in which visitors were able to project which exhibit the remote person was about to explain.

From these observations, we could confirm that the GestureMan has the ability to support local participant to project a remote person's orientation as long as its head is properly controlled.

However, there emerged another issue: a remote person's orientations cannot be accurately communicated to the GestureMan due to the difference of ecologies between the remote site and local participants' site. In the next section, we discuss the "dual" ecologies surrounding the robot as a communication medium.

MEDIATING DUAL ECOLOGIES

The advantages of the three monitors for a remote person occasionally led to inconsistencies in orientation between the remote person and the local participants. In other words, a remote person continuously seeks which object to give an explanation for next, or confirms the direction in which to proceed, without moving the head of the robot but simply by looking at the three monitors. At this point, the remote person's head is turned either toward the right or the left monitor. As seen in the previous section, a face-to-face environment offers local participants opportunities to project another person's next action by capturing these actions. However, with the GestureMan, this type of action is not always reflected in the robot's movement, and therefore a local participant cannot necessarily predict the remote person's actions. Of course, the remote person produces some actions within his/her own ecology, which may help the local participant to project the remote person's next movement but this is not mirrored in the local participant's ecology. This inconsistency between the ecologies of the remote person and the local participant results in the inability of the local participant to project the remote person's actions.

But can we resolve the problem of projection by simply having a single monitor for the remote person to permit him/her to move the robot's head? In order to examine the effect of changing the width of the field of view in communication mediated by a robot, we conducted the following experiment which revealed how a remote person and a local participant actually produce actions induced by the inconsistencies in the dual ecologies.

Study Design

In this experiment, a remote person asked a local participant to touch specified objects in the local participant's environment. Overviews of a remote person's environment and local participant's environment are shown in figure 5. Three LCD monitors were placed in front of the remote person and images from the GestureMan's three-camera unit were displayed on them. A remote person controlled the GestureMan's head using a joystick.

In the local site, the local participant and the GestureMan stood face-to-face. A few tables were placed between the local participant and the GestureMan, and sixty objects such as dolls, cables, cameras, etc. were scattered over the tops of the tables. A local participant was asked to stand immediately in front of the GestureMan except when he/she was asked to touch a particular object.

Each pair (of remote person and local participant) was asked to perform the session under two different conditions. In condition 1 (single-monitor condition), only the center monitor was turned on and the right and left monitors were turned off. In condition 2 (three-monitor condition), a remote person could use three monitors to observe the local participant's site. A local participant was told that the robot was controlled by a remote person but he/she did not know the difference of the remote person's user interface between two conditions.

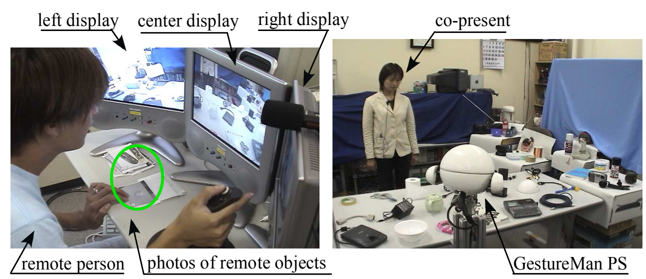

(a) Remote person's site (b) Local participant's site

Figure 5: Overview of the preliminary experiment.

At the beginning of each session, the remote person had photos of thirty out of the sixty objects in the local participant's environment. In order to eliminate the effect of order of presentation, the objects were chosen randomly in each session. During the sessions, the remote person looked at a photo of an object, controlled the GestureMan's head to find the object, and instructed the local participant to touch the object. During this process, both the remote person and local participant were allowed to talk freely using microphones. Because we wanted to investigate only the effect of the robot's head movement, we did not use the laser pointer for specifying objects.

Eight pairs of University of Tsukuba students served as remote persons and local participants. Each pair was asked to perform the task in both conditions. In order to minimize the effect of carry-over, four pairs performed the condition 1 session first, and the other four pairs performed the condition 2 session first.

Results

We measured the performance of the various sessions in terms of preparation time and task completion time. We defined the preparation time as the time for a remote person to find an object in the LCD monitor. The time was measured from the moment that a remote person started looking at a photo of an object until he/she started to verbally describe the object to be touched. During this time span, the remote person was asked to look for the object on the LCD monitor by controlling the robot's head. Task completion time was measured from the moment that a remote person started looking at the photo until a local participant touched the object. We defined the instruction time as the time from the moment that a remote person started to verbally describe the object until the moment that a local participant touched the object. Thus the task completion time is the sum of the preparation time and instruction time.

Although Wilcoxon signed rank test did not show the statistical significance, figure 6 shows the tendency that average preparation time was shorter when the remote person used three monitors.

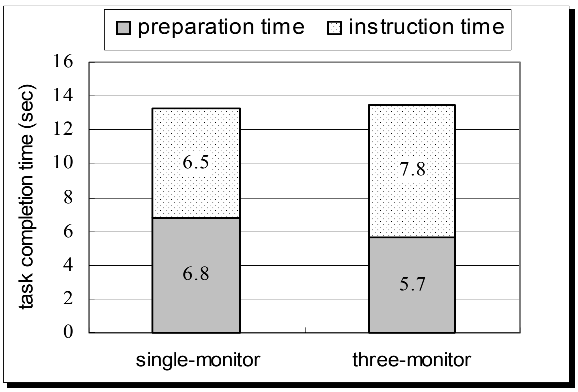

Figure 6. Average task completion time.

As for task completion time, however, there was little difference between the two conditions (figure 6). In other words, local participants could find objects faster in the single-monitor condition compared to the three-monitor condition. These assumptions, although they were not supported by the statistical analysis, suggested that if we could combine the advantages of both the single-monitor condition and the three-monitor condition, we could expect the task completion time to be as short as 12.2sec (=5.7+6.5). Thus, by observing the video of the experiment in detail, we were able to gain a better understand of the characteristics of interaction between a remote person and a local participant.

From the video observation, in the case of the single-monitor condition, since a remote person needs to look at an object in the center monitor, the robot's head becomes naturally oriented toward particular objects. By monitoring the robot head's orientation, a local participant was often seen to orient his/her head toward the appropriate object before a remote person provided a full substantive description of the location of the object. (Fig. 7). We use the term "correct anticipatory reaction" to describe the case in which the local participant was able to orient correctly toward the indicated object before the remote person gave a concrete verbal depiction of the location. We counted the number of similar correct anticipatory reactions and compared the two conditions (Fig. 8). As the figure shows, there were more correct anticipatory reactions observed in the single-monitor condition. Wilcoxons signed-rank test showed a significant difference between these two conditions ($p<0.01$).

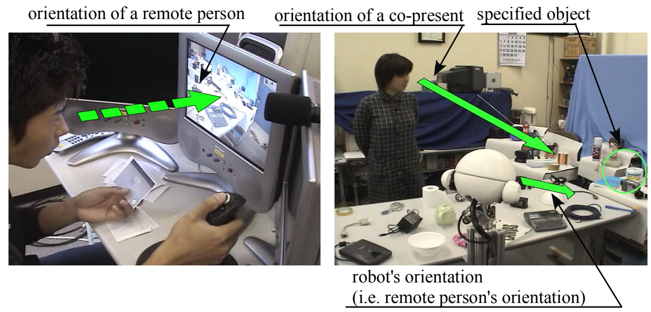

(left) remote site (right) local participant's site

Figure 7. An example of a correct anticipatory reaction.

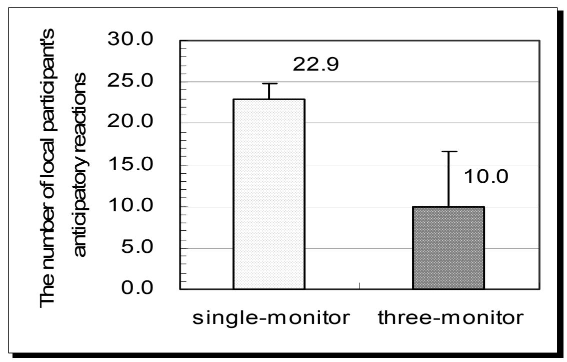

Figure 8. The average number of anticipatory reactions.

In the case of the three-monitor condition, the robot's head moved much less than in the single-monitor condition. Figure 9 shows the average duration that the robot's head was manipulated during one session. Wilcoxons singed-rank test showed that the robot's head was manipulated longer in the single-monitor condition than in the three-monitor condition ($p<0.01$). Thus the local participants could properly anticipate the objects to be selected.

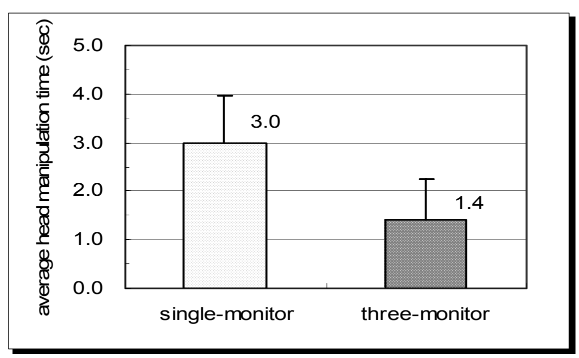

Figure 9. Average time the robot head was controlled per session.

In the case of three-monitor condition, when an object is displayed on the right or left monitor, the robot's head oriented toward quite a different direction from the remote person. In such a case, at the beginning of a session, a local participant occasionally looked toward a direction which is the opposite of the direction of the correct object. For example in Figure 10, the remote person found the object in the left monitor. Because the object was already visible on the left monitor, the remote person did not go out of his way to move the robot's head. However, when he located the object on the left monitor, he made the robot's head nod slightly, most likely in order to vertically align the camera to improve the image. It should be noted that at this point, the robot's head was turned horizontally slightly toward its right side (i.e. corresponding to the local participant's left hand side). As a consequence, when the remote person started to describe the object, the local participant turned her head toward her left, mistakingly assuming that the robot was indicating an object in that direction by the nod. The incidence of this kind of false anticipatory reaction is significantly greater in the three-monitor condition than in the single-monitor condition (Wilcoxons signed-rank test, $p<0.01$) (Fig. 11).

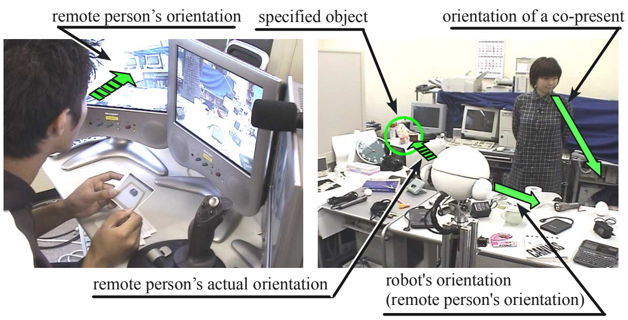

Figure 10. An example of a false anticipatory reaction.

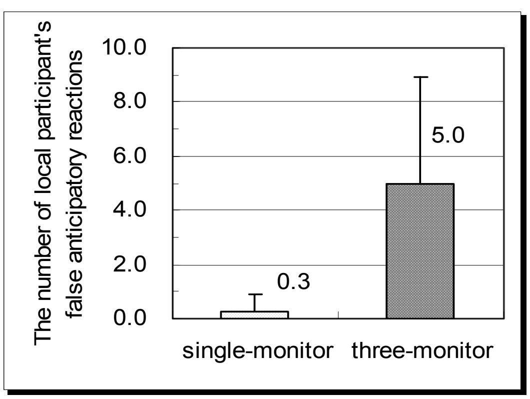

Figure 11. The average number of false anticipatory reactions.

In a questionnaire which was subsequently administered to the subjects, seven remote persons responded that the three-monitor condition was easier for giving instructions. On the other hand, seven local participants responded that the single-monitor condition was easier for finding specified objects.

From these results, it is clear that remote persons preferred the three-monitor condition because it both facilitates and expedites the process of finding objects. On the contrary, local participants preferred the single-monitor condition because the robot's head orientation enables them to project objects to be specified. Differences in the two ecologies generated the inconsistencies between the two conditions.

DICUSSIONS AND CONCLUSIONS

In this paper we have shown that preparatory actions are important resources for projection. In robot-mediated communication, by implementing resources for displaying orientations, some of the preparatory actions can be observed by local participants and the task of projection can be supported.

However, there are some cases in which preparatory actions in a remote site are not reflected in the robot's movements. For example, head movement is a significant resource for projection in real world communication. With our system, we have shown that the robot's head movement is also a significant resource for projection. However, because the remote person's head movement is oriented toward the monitors as an user interface, such orientations are not necessarily reflected in the robot's movement. To make the matter worse, since a local participant tends to "read or make sense of the actions of "the robot" through the ways in which they interweave conduct with particular features of the immediate environment [14]", a local person occasionally misunderstand a remote person's orientation. Inconsistency between a remote person's head orientation and the robot's head orientation leads to false anticipatory reactions. We have demonstrated this problem by comparing two conditions of monitors.

This fact led us to realize the existence of two ecologies. The remote site ecology which a remote person prefers (i.e. three-monitor condition) produces the local site ecology which a local participant does not prefer, and vice versa.

Of course some persons may adapt to such inconsistencies through practice [2]. However, most of the preparatory actions are conducted unintentionally. Therefore, it may be difficult to make a remote person intentionally control some kind of user interface to express such unintentional actions. Therefore, a user interface that can detect such unintentional preparatory actions should be developed.

To sum up, in the case of a robot-mediated communication system, designers of a system always have to consider;

- the existence of dual ecologies,
- what kind of user interface provides preferable ecology for a local person and a remote participant,
- and the ways to resolve inconsistencies between two ecologies.

In case of supporting projectability by head orientation, the remote person's side needs an environment (user interface) that naturally renders the head movement to be conspicuous enough to be picked up by the system. The local co-participant's side needs an indicator that is controlled in response to the remote person's head movement so that the local participant can read the remote person's intentions correctly.

Some may think this result is not surprising. However, it is equally surprising that most existing remote controlled robot systems cannot cope with this problem. Still, many user interfaces for robot control repeat a similar approach to ours.

One of the possible solutions is to have a remote person wear a head-mounted display (HMD) with a wide field of view. Also a 3D motion tracker should be attached to an HMD in order to control the robot's head in response to the remote person's head motions. However, an HMD with such a wide field of vision tends to be massive and very heavy, and therefore would be cumbersome to wear on one's head.

Therefore, we are currently developing a system as shown in Figure 12. A camera is mounted on the robot's body, not on its head. A remote person sits in front of a three-monitor system and wears a small 3D motion tracker on his/her head. Since a three-monitor system has a wide spread, when the remote person looks around, his/her head naturally moves to the right and left. This head movement can be traced by the 3D motion tracker. The robot's head is controlled and moved around in accordance with the remote person's head movements. It must be noted that the robot's head moves so that the local participant may see where the remote person is orienting toward, and that the direction of the camera does not change unless the remote person changes the robot's body orientation using a joystick.

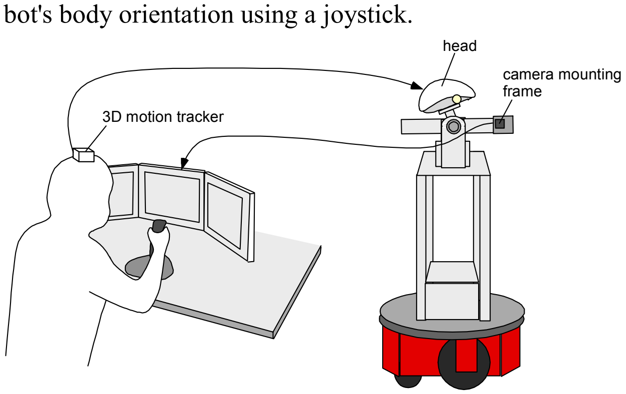

Figure 12. A new GestureMan system that considers dual ecologies.

The Mutually-Immersive Mobile Telepresence [10] is taking a similar approach but the remote person's real head image is shown on LCD monitors mounted on the robot's head. Although this advanced approach is very interesting, it is known that a remote person's head orientation that is seen on a flat monitor "is interpreted much the same way by most people, independent of position -- a type of Mona Lisa effect in which a person passing in front of the famous face feels that those eyes are following them[15]." Since local participants tend to stand sideways of a robot (Fig. 4), the Mona Lisa effect can be problematic. We are expecting that the physical head of our new robot will be able to alleviate this problem and effectively support projectability.

We think we can support projectability not only by head movement but by movements of other parts of the body, such as torso orientation and arm movement. For this purpose, we are considering consistencies of two ecologies. By supporting projectability, we think we can realize more efficient tele-communication systems.

ACKNOWLEDGEMENTS

The research was supported by Scalable Virtual Reality project of the Telecommunications Advancement Organization of Japan, Oki Electric Industry Co. Ltd., Venture Business Laboratory (VBL) of Ministry of Education, Culture, Sports Science and Technology, The Telecommunications Advancement Foundation, and International Communications Foundation. We extend our thanks to Science Museum, Tokyo and Communications Museum (Tei-Park) for supporting our experiments.

Much gratitude is extended to Dr. Hiroko Tanaka, University of Essex for her kind assistance in revising the paper.

REFERENCES

1. Auer, P., Projection in interaction and projection in grammar, InLiSt ,N0. 33, Freiburg (2002).
2. Dourish, P., Adler, A., Bellotti, V., Henderson, A. Your place or mine? Learning from long-term use of audio-video communication. Kluwer Academic Publishers Norwell, MA, USA, 33-62 (1996).
3. Gaver, W., The Affordances of Media Spaces for Collaboration. Proc. CSCW'92, ACM Press, 17-24 (1992).
4. Goodwin, C., Professional vision. American Anthropologist 96, 606-33 (1994).
5. Goodwin, C., Transparent vision. In: Ochs, E., E.A. Schegloff, S.A. Thompson, eds., Interaction and Grammar. Cambridge, Cambridge University Press, 370-404 (1996).
6. Goodwin, C., Action and embodiment within situated human interaction. Journal of Pragmatics, 32, 1489-1522 (2000).
7. Heath, C, Luff, P., Media space and communicative asymmetries: preliminary observations of video mediated interaction, Human-Computer Interaction 7, 315-46 (1992).
8. Heath, C., Luff, P., Kuzuoka, H., Yamazaki, K., Creating coherent environments for collaboration, Proc. ECSCW2001, 119-128 (2001).

9. Imai, M., Ono, T., Ishiguro, H., Physical Relation and Expression: Joint Attention for Human-Robot Interaction, IEEE Trans. on Industrial Electronics, Vol. 50, No. 4, 636-643 (2002).

10. Jouppi, N., First steps towards mutually-immersive mobile telepresence, Proc. CSCW2002, ACM Press, 354-363 (2002).

11. Kendon, A., Conducting Interaction: Patterns of Behavior in Focused Interaction, Cambridge, Cambridge University Press (1990).

12. Kuzuoka, H., Kosuge, T., and Tanaka, M., GestureCam: A Video Communication System for Sympathetic Remote Collaboration, Proc. of CSCW'94, 35-43 (1994).

13. Kuzuoka, H., Oyama, S., Yamazaki, K., Suzuki, K., Mitsuishi, M., GestureMan: A Mobile Robot that Embodies a Remote Instructor's Actions, Proc. CSCW2000, ACM Press, 155-162 (2000).

14. Luff, P., Heath, C., Kuzuoka, H., Hindmarsh, J., Yamazaki, K., and Oyama, S., Fractured ecologies: creating environments for collaboration, Special Issue of the HCI Journal: 'Talking About Things: Mediated Conversation about Objects', 18, 1 & 2, 51-84 (2003).

15. Morikawa, O. and Maesako, T., HyperMirror: Toward Pleasant-to-use Video Mediated Communication System, CSCW98, 149-158 (1998).

16. Nardi, B., Schwarz, H., Kuchinsky, A., Leichner, R., Whittaker, S., and Sclabassi, R., Turning away from talking heads: the use of video-as-data in neurosurgery, Proc. CHI'93, ACM Press, 327 – 334 (1993).

17. Paulos, E. and Canny, J., ProP: Personal Roving Presence, Proc. CHI' 98, ACM Press, 296-303 (1998).

18. Schegloff, E. A., Preliminaries to preliminaries: "Can I ask you a question?"', Sociological Inquiry 50, 104-52 (1980).

19. Schegloff, E. A., 'Body torque', Social Research, 65, 535-86 (1998).

20. Streeck, J., On Projection. In: E. Goody, ed., Interaction and Social Intelligence, Cambridge, Cambridge University Press, 84-110 (1995).

21. Tanaka, H., Turn-Taking in Japanese Conversation: A Study in Grammar and Interaction. Amsterdam: John Benjamins (1999).

22. Yamazaki, K., Yamazaki, A., Kuzuoka, H., Oyama, S., Kato, H., Suzuki, H., and Miki, H., GestureLaser and Gesturelaser Car: development of an Embodied Space to Support Remote Instruction, Proc. of ECSCW'99, 239-258 (1999).