

Design and validation of ECA gestures to improve dialogue system robustness

**Beatriz López, Álvaro Hernández, David Díaz,
Rubén Fernández, Luis Hernández**

GAPS, Signal, Systems and Radiocommunications
Department

Universidad Politécnica de Madrid
Ciudad Universitaria s/n, 28040 Madrid, Spain

alvaro@gaps.ssr.upm.es

Doroteo Torre

ATVS, Escuela Politécnica Superior
Universidad Autónoma de Madrid
Ciudad Universitaria de Cantoblanco,
28049 Madrid, Spain

Doroteo.torre@uam.es

Abstract

In this paper we present validation tests that we have carried out on gestures that we have designed for an embodied conversational agent (ECAs), to assess their soundness with a view to applying said gestures in a forthcoming experiment to explore the possibilities ECAs can offer to overcome typical robustness problems in spoken language dialogue systems (SLDSs). The paper is divided into two parts: First we carry out a literature review to acquire a sense of the extent to which ECAs can help overcome user frustration during human-machine interaction. Then we associate tentative, yet specific, ECA gestural behaviour with each of the main dialogue stages, with special emphasis on problem situations. In the second part we describe the tests we have carried out to validate our ECA's gestural repertoire. The results obtained show that users generally understand and naturally accept the gestures, to a reasonable degree. This encourages us to proceed with the next stage of research: evaluating the gestural strategy in real dialogue situations with the aim of learning about how to favour a more efficient and pleasant dialogue flow for the users.

1 Introduction

Spoken language dialogue systems and embodied conversational agents are being introduced in a rapidly increasing number of Human-Computer Interaction (HCI) applications. The technologies involved in SLDSs (speech recognition, dialogue design, etc.) are mature enough to allow the creation of trustworthy applications. However, robustness problems still arise in concrete limited dialogue systems because there are many error sources that may cause the system to perform poorly. A common example is that users tend to repeat their previous utterance with some frustration when error recovery mechanisms come into play, which does not help the recognition process, and as a result using the system seems slow and unnatural (Boyce, 1999).

At the same time, embodied conversational agents (ECAs) are gaining prominence in HCI systems, since they make for more user-friendly applications while increasing communication effectiveness. There are many studies on the effects – from psychological to efficiency in goal achievement – ECAs have on users of a variety of applications, see Bickmore et al. (2004) and Brave et al. (2005), but still very few (Bell and Gustafson, 2003) on the impact of ECAs in directed dialogue situations where robustness is a problem.

Our research explores the potential of ECAs to assist in, or resolve, certain difficult dialogue situations that have been identified by various leading authors in the field (Cassell and Thorisson, 1999; Cassell and Stone, 1999), as well as a few we our-

selves suggest. After identifying the problematic situations of the dialogue we suggest a gestural strategy for the ECA to respond to such problem situations. Then we propose an experimental framework, for forthcoming tests, to study the potential benefits of adding nonverbal communication in complex dialogue situations. In the study we present here we focus on preliminary validation of our gestural repertoire through user tests. We conclude by presenting our results and suggesting the direction our research will take from this point.

2 How ECA technology can improve interaction with SLDSs

There are many nonverbal elements of communication in everyday life that are important because they convey a considerable amount of information and qualify the spoken message, sometimes even to the extent that what is meant is actually the opposite of what is said (Krauss et al., 1996). ECAs offer the possibility to combine several communication modes such as speech and gestures, making it possible, in theory, to create interfaces with which human-machine interaction is much more natural and comfortable. In fact, they are already being employed to improve interaction (Massaro et al., 2000).

These are some common situations with SLDSs in which an ECA could have a positive effect:

Efficient turn management: The body language and expressiveness of agents are important not only to reinforce the spoken message, but also to regulate the flow of the dialogue, as Cassell points out (in Bickmore et al., 2004).

Improving error recovery: The process of recognition error recovery usually leads to a certain degree of user frustration (see Oviatt and VanGent, 1996). Indeed, it is common, once an error occurs, to enter into an error spiral in which the system is trying to recover, the user gets ever more frustrated, and this frustration interferes in the recognition process making the situation worse (Oviatt et al., 1998). ECAs may help reduce frustration, and by doing so make error recovery more effective (Hone, 2005).

Correct understanding of the state of the dialogue: Sometimes the user doesn't know whether or not things are going normally (Oviatt, 1994). This sometimes leads the dialogue to error states that could be avoided. The expressive capacity of

ECAs could be used to reflect with greater clarity the state the system takes the dialogue to be in.

3 Suggesting ECA behaviour for each dialogue situation

A variety of studies have been carried out on behavioural strategies for embodied conversational agents (Poggi, 2001; Cassell et al., 2000; Cassell et al., 2001; Chovil, 1992; Kendon, 1990), which deal with behaviour in hypothetical situations and in terms of the informational goals of each particular interaction (be it human-human or human-machine). We direct our attention to the overall dialogue systems dynamics, focussing specifically on typical robustness problems and how to favour smooth sailing through the different stages of the dialogue. We draw from existing research undertaken to try to understand the effects different gestures displayed by ECAs have on people, and we apply this knowledge to a real dialogue system. In Table 1 we show the basic set of gestures we are using as a starting point. They are based mainly on descriptions in Bickmore (et al., 2004) and Cassell (et al., 2000), and on recommendations in Cassell and Thorisson (1999), Cassell (et al., 2001), Chovil (1992), Kendon (1990) and San-Segundo (et al., 2001), to which we have added a few suggestions of our own.

| Dialogue stage | ECA behaviour (movements, gestures and other cues) |
|------------------------------------|--|
| Initiation (welcoming the user) | 1. Welcome message: look at the camera, smile, wave hand 2. Explanation of the task: zoom in 3. Zoom out, lights dim |
| Give turn | Look directly at the user, raise eyebrows. Camera zooms out. Lights dim. |
| Take turn | Look directly at the user, raise hands into gesture space. Camera zooms in. Light gets brighter. |
| Wait | Slight leaning back, one arm crossed and the other touching the cheek shift of body weight |
| Help | Beat gesture with the hands. Change of posture |
| Error recovery with correction | Lean towards the camera, beat gesture |
| Confirmation (high confidence) | Nod, smile, eyes fully open |
| Confirmation (low confidence) | Slight leaning of the head to one side, stop smiling, mildly squint |

Table 1: Gesture repertoire for the main dialogue stages

3.1 Initiation

The inclusion of an ECA at this stage “humanises” the system (Oviatt and Adams, 2000). This is a problem, first because once a user has such high expectations the system can only end up disappointing her, and secondly because the user will tend to use more natural (and thus complex) communication, which the system is unable to handle, and the experience will ultimately be frustrating.

On the other hand, especially in the case of new users, contact with a dialoguing animated character may have the effect that the user’s level of attention to the actual information that is being given is reduced (Schaumburg, 2001; Catrambone, 2002). Thus the goal is to present a human-like interface that is, at the same time, less striking and thus less distracting at first contact, and one that clearly “sets the rules” of the interaction and makes sure that the user keeps it framed within the capability of the system.

We have designed a welcome gesture for our ECA based on the recommendations in Kendon (1990), to test whether or not it fosters a sense of ease in the user and helps her concentrate on the task at hand. Playing with the zoom, the size and the position of the ECA on the screen may also prove to be useful to frame the communication better (see Table 1).

3.2 Turn Management

Turn management involves two basic actions: taking turn and giving turn. Again, in Table 1 we show the corresponding ECA gestures we will start testing with. Note that apart from the ECA gestures, we also play with zoom and light intensity: when it’s the ECA’s turn to speak the camera zooms-in slightly and the light becomes brighter, and when it’s the user’s turn the camera zooms out and the lights dim. The idea is that, hopefully, the user will associate each camera shot and level of light intensity with each of the turn modes, and so know when she is expected to speak.

The following are some typical examples of problem situations together with further considerations about ECA behaviour that could help avoid or recover from them:

- The user tries to interrupt at a point at which the barge-in feature is not active. If this happens the system does not process what the user has said, and when the system

finally returns to listening mode there is silence from both parts: the system expects input from the user, and the user expects an answer. Often both finally break the silence at the same time and the cycle begins again, or, if the system caught part of the user’s utterance, a recognition error will most likely occur and the system will fall into a recognition error recovery subdialogue that the user does not expect. To help avoid such faulty events the ECAs demeanour should indicate as clearly as possible that the user is not being listened to at that particular moment. Speaking while looking away, perhaps at some object, and absence of attention cues (such as nodding) are possible ways to show that the user is not expected to interrupt the ECA. Since our present dialogue system produces fairly short utterances for the ECA, we are somewhat limited as to the active strategies to build into the ECA’s behaviour. However, there are at least three cues the user could read to realise that the system didn’t listen to what she said. The first is the fact that the system carries on speaking, ignoring the user’s utterance. Second, at the end of the system’s turn the ECA will perform a specific give-turn gesture. And third, after giving the turn the ECA will remain still and silent for a few seconds before performing a waiting gesture (leaning back slightly with her arms crossed, shifting the body weight from one leg to another; see Table 1). In addition, if the user still remains silent after yet another brief waiting period the system will offer help. It will be interesting to see at which point users realise that the system didn’t register their utterance.

- A similar situation occurs if the Voice Activity Detector (VAD) fails and the system doesn’t capture the user’s entire utterance, or when the user simply doesn’t say anything when she is expected to (“no input”). Again, both system and user end up waiting for each other to say something. And again, the strategy we use is to have the ECA display a waiting posture.
- It can also happen that the user doesn’t speak but the VAD “thinks” she did, perhaps after detecting some background noise

(a “phantom input”). The dialogue system’s reaction to something the user didn’t say can cause surprise and confusion in the user. Here the visible reactions of an ECA might help the user understand what has happened and allow her to steer the dialogue back on track.

3.3 Recognition Confidence Scheme

Once the user utterance has been recognised, information confirmation strategies are commonly used in dialogue systems. Different strategies are taken depending on the level of confidence in the correctness of the user locution as captured by the speech recognition unit (San-Segundo et al., 2001). Our scheme is as follows:

- *High confidence*: if recognition confidence is high enough to safely assume that no error has occurred, the dialogue strategy is made more fluent, with no confirmations being sought by the system.
- *Intermediate confidence*: the result is regarded as uncertain and the system tries implicit confirmation (by including the uncertain piece of information in a question about something else.) This, combined with a mixed initiative approach, allows the user to correct the system if an error did occur.
- *Low confidence*: in this case recognition has probably failed. When this happens the dialogue switches to a more guided strategy, with explicit confirmation of the collected information and no mixed initiative. The user’s reply may confirm that the system understood correctly, in which case the dialogue continues to flow normally, or, on the other hand, it may show that there was a recognition error. In this case an error recovery mechanism begins.

In addition to the dialogue strategies, ECAs could also be used to reflect in their manner the level of confidence that the system has understood the user, in accordance with the confirmation dialogue strategies. While the user speaks, our ECA will, if the recognition confidence level is high, nod her head (Cassell et al., 2000), smile and have her eyes fully open to give the user feedback that everything is going well and the system is understanding. If, on the other hand, confidence is low,

in order to make it clearer to the user that there might be some problem with recognition and that extra care should be taken, an option might be for the ECA to gesture in such a way as to show that she isn’t quite sure she’s understood but is making an effort to. We have attempted to create this effect by having the ECA lean her head slightly to one side, stop smiling and mildly squint. Our goal, once again, is to find out whether these cues do indeed help users realise what the situation is. This is especially important if it helps to avoid the well-known problem of falling into error spirals when a recognition error occurs in a spoken dialogue system (Bulyko et al., 2005). In the case of intermediate recognition confidence followed by a mixed initiative strategy involving implicit confirmation, specific gestures could also be envisaged. We have chosen not to include specific gestures for these situations in our first trials, however, so as not to obscure our observations for the high and low confidence cases. A neutral stance for the intermediate confidence level should be a useful reference against which to compare the other two cases.

3.4 Recognition Problems

We will consider those situations in which the system finds the user’s utterance incomprehensible (no-match situations) and those in which the system gets the user’s message wrong (recognitions errors). When a no-match occurs there are two ways in which an ECA can be useful. First, what the character should say must be carefully pondered to ensure that the user is aware that the system didn’t understand what she said and that the immediate objective is to solve this particular problem. This knowledge can make the user more patient with the system and tolerate better the unexpected lengthening of the interaction (Goldberg, 2003). Second, the ECAs manner should try to keep the user in a positive attitude. A common problem in no-match and error recovery situations is that the user becomes irritated or hyperarticulates in an attempt to make herself understood, which in fact increases the probability of yet another no-match or a recognition error. This we should obviously try to avoid.

The ECA behaviour strategy we will test in no-match situations is to have the character lean towards the camera and raise her eyebrows (the idea being to convey a sense of surprise coupled with friendly interest). We have based our gesture on

one given in (Fagerberg et al., 2003). If the user points out to the system that there has been a recognition error in a way that gives the correct information at the same time, then the ECA will confirm the corrected information with special emphasis in speech and gesture. For this purpose we have designed a beat gesture with both hands (see Table 1).

3.5 Help offers and request

It will be interesting to see whether the fact that help is offered by an animated character (the ECA) is regarded by users to be more user-friendly than otherwise. If users feel more comfortable with the ECA, perhaps they will show greater initiative in requesting help from the system; and when it is offered by the system (when a problem situation occurs), the presence of a friendly ECA might help control user frustration. While the ECA is giving the requested information, she will perform a beat gesture with both hands for emphasis, and she will also change posture. The idea is to see whether this captures the interest of the user, makes her more confident and the experience more pleasant or, on the contrary, it distracts the user and makes help delivery less effective.

Figure 1 illustrates a dialogue sequence including the association between the different dialogue strategies and the ECA gesture sequences after a user's utterance.

4 Experimental set up

Gestures and nonverbal communication are culture-dependent. This is an important fact to take

into account because a single gesture might be interpreted in different ways depending on the user's culture (Kleinsmith et al., 2006). Therefore, a necessary step prior to the evaluation of the various hypotheses put forward in the previous section is to test the gestures we have implemented for our ECA, within the framework designed for our study. This implies validating the gestures for Spanish users, since we have based them on studies within the Anglo-Saxon culture.

4.1 Procedure

For the purpose of testing the gesture repertoire developed for our ECA we have conceived an evaluation environment that simulates a realistic mobile videotelephony application that allows users to remotely check the state (*e.g.*, on/off) of several household devices (lights, heating, etc.). Our dialogue system incorporates mixed initiative, error recovery subdialogues, context-dependent help and the production of guided or flexible dialogues according to the confidence levels of the speech recogniser. Our environment uses Nuance Communications' speech recognition technology (www.nuance.com). The ECA character has been designed by Haptik (www.haptik.com).

During the gesture validation tests users didn't interact directly with the dialogue system. We first asked the users to watch a system simulator (a video recording of a user interacting with the system), so that they could see the ECA performing the gestures in the context of a real dialogue.

After watching the simulation the users were asked to fill out a questionnaire. The questionnaire allowed users to view isolated clips of each

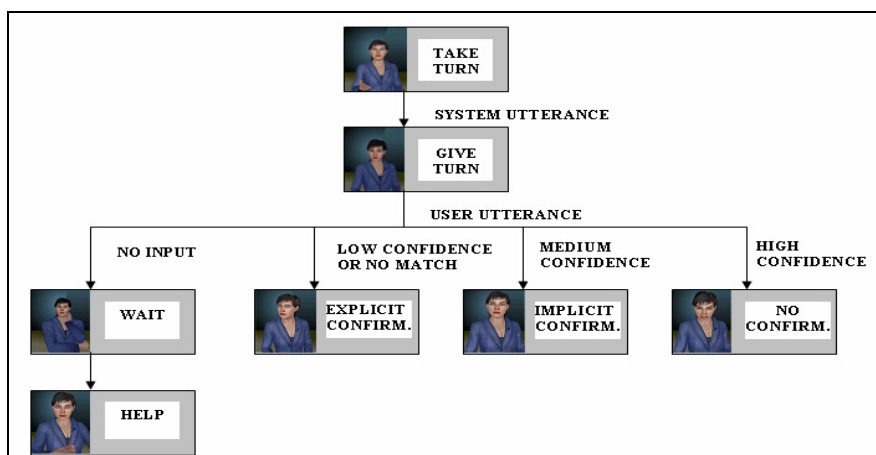


Figure 1: Dialogue strategies and related gesture sequence

of the dialogue gestures (the eight that had appeared in the video). To each gesture clip were associated questions basically covering the following three aspects:

- *Gesture interpretation*: Users are asked to interpret each gesture, choosing one from among several given options (the same options for all gestures). The aim is to see whether the meaning and intention of each gesture are clear. In addition users told us whether they thought they had seen the gesture in the previous dialogue sample.
- *Gesture design*: Do users think the gesture is well made and does it look natural? To answer this question we asked users to rate the quality, expressiveness and clarity of the ECAs gesture (on a 9-point Likert scale).
- *User expectations*: Users rated how useful they thought each gesture was (on a 9-point Likert scale). The idea is to juxtapose the utility function of the gestures in the users' mental model to our own when we designed them, and evaluate the similarity. In addition we collected suggestions as to how the users thought the gestures could be improved.

4.2 Results

We recruited 17 test users (most of them students between 20 and 25 years of age) for our trial. The results concerning the three previously mentioned aspects are shown in Table 2. In the case of the *gesture interpretation*, we present the percentage of the users who interpreted each gesture "correctly" (*i.e.*, as we had intended when we designed them). Depending on this percentage we label each gesture as "Good", "Average", or "Bad". For each of the parameters for *gesture design* and *user expectations* we give the mean and the standard deviation of the Likert scale scores. We label the average scores as "Low" (Likert score between 1 and 3), Medium (4-6) or "High" (7-9).

We now discuss the results separately for each of the dimensions:

Regarding user expectations, the values for each gesture are High except for two of them, valued as Medium. These two gestures are the welcome gesture and the gesture for offering help. In the case of the welcome gesture, users probably believe the

beginning of the dialogue is already well enough defined when the ECA starts to speak. If so, users might see an element of redundancy in the welcome gesture, lowering its perceived utility in the dialogue process. On the other hand, the help gesture utility might be valued lower than the rest because many users didn't seem to understand its purpose (the clarity of the Help gesture was the least valued of all, $\mu=5.117$). Nevertheless, the general user impressions of the utility of the evaluated gesture repertoire fairly high.

In relation to gesture design, we can see that, overall, the marks for quality and expressiveness are high. This implies our gesture design is, on the whole, adequate. Regarding the clarity of the gestures, three of them are valued as Medium. These are the gestures expressing Give Turn, Error Recovery and Help offer. This could be related to the prevailing opinion among users that there are a few confusing gestures, although they are better understood in the context of the application, when you listen to what the ECA says.

Only half of the gestures were properly interpreted by the users. Those that weren't (Give Turn, Take Turn, Error Recovery and the Help gesture) are, we realize, the subtlest in the repertoire, so we asked ourselves if there could be relation between a bad interpretation of the gesture and the whether that user didn't remember seeing the gesture in the dialogue. In Figure 2 we show the number of users who claimed they hadn't seen the ECA gestures during the dialogue sample. The coloured bars represent the overall accuracy in the interpretation of the gesture. We may observe that the gestures that a larger number of users hadn't seen in the dialogue, and therefore, hadn't an image of in proper context, tended also to be considered more unclear.

We may conclude that some gestures need to be evaluated in context. In any case, and in spite of the uncertainty we have found regarding the interpretation of certain gestures, we believe the positive evaluation by the users for the expressiveness and the quality of the gestures justifies us in validating our gestural repertoire for the next research stage where we will evaluate how well our ECA gestures function under real interaction conditions (taking into account objective data related to dialogue efficiency).

| | INTERPRETATION | DESIGN | | | EXPECTATIONS |
|---|-------------------------|-----------------------|-------------------------|-----------------------|-------------------------|
| | Good Interpretation (%) | Quality | Clarity | Expressiveness | Usefulness |
| G1 Wellcome | 88.23 Good | 7.117 (0.927) High | 7.588 (1.277) High | 6.764 (1.147) High | 5.647 (2.119) Medium |
| G2 Give Turn | 35.29 Average | 6.647 (1.057) High | 5.823 (1.333) Medium | 6.470 (1.007) High | 6.588 (1.543) High |
| G3 Take Turn | 23.53 Bad | 7.117 (1.166) High | 6.705 (1.447) High | 6.941 (1.444) High | 6.647 (1.271) High |
| G4 Wait | 82.35 Good | 7.058 (1.088) High | 7.176 (1.185) High | 7.176 (0.727) High | 6.588 (1.622) High |
| G5 Confirmation (Low confidence) | 76.47 Good | 8.294 (0.587) High | 8.058 (1.028) High | 8.058 (1.028) High | 7.941 (1.028) High |
| G6 Confirmation (High confidence) | 94.11 Good | 7.529 (1.124) High | 7.529 (1.124) High | 7.705(1.263) High | 7.588 (1.175) High |
| G7 Error Recovery | 41.17 Average | 6.941 (1.088) High | 5.588 (2.032) Medium | 6.529 (1.462) High | 6.058 (1.390) High |
| G8 Help | 35.29 Average | 6.823 (1.185) High | 5.117 (1.932) Medium | 6.058(1.560) High | 5.529 (1.771) Medium |

Table 2: Results of the gesture validation tests.




Figure 2: Interpretation vs. 'visibility' of the gestures.

5 Conclusions and future lines of work

In this article we have identified a range of problem situations that may arise in dialogue systems, and defined various strategies for using an ECA to improve user-machine interaction throughout the whole dialogue. We have developed an experimental set up for a user validation of ECA gestures in the dialogue system and have obtained quantitative results and user opinions to improve the design of the gestures. The results of this validation allow us to be in a position to begin testing our dialogue

system and evaluate our ECA gestures in the context of a real dialogue.

In future experiments we will attempt to go one step further and analyse how empathic emotions vs. self-oriented behaviour (see Brave et al., 2005) may affect the resolution of a variety of dialogue situations. To this end we plan to design ECA prototypes that incorporate specific emotions, hoping to learn how best to connect empathically with the user, and what effects this may have on dialogue dynamics and the overall user perception of the system.

References

- Linda Bell and Joakim Gustafson, 2003. *Child and Adult Speaker Adaptation during Error Resolution in a Publicly Available Spoken Dialogue System*. Proceedings of Eurospeech 03, Geneva, Schweiz.
- Timothy W. Bickmore, Justine Cassell, Jan van Kuppevelt, Laila Dybkjaer and Niels Ole Bernsen, 2004. *Natural, Intelligent and Effective Interaction with Multimodal Dialogue Systems*, chapter Social Dialogue with Embodied Conversational Agents. Kluwer Academic.
- Susan J. Boyce, 1999. *Spoken natural language dialogue systems: user interface issues for the future*. In Human Factors and Voice Interactive Systems. D. Gardner-Bonneau Ed. Norwell, Massachusetts, Kluwer Academic Publishers: 37-62.
- Scott Brave, Clifford Nass, Kevin Hutchinson, 2005. *Computers that care: investigating the effects of ori-*

- entation of emotion exhibited by an embodied computer agent. *Int. J. Human-Computer Studies*, Nr. 62, Issue 2, pp. 161-178.
- Ivan Bulyko, Katrin Kirchhoff, Mari Ostendorf, Julie Goldberg, 2005 *Error correction detection and response generation in a spoken dialogue system*. *Speech Communication* 45, 271-288.
- Justine Cassell, Kristinn R. Thorisson, 1999. *The power of a nod and a glance: envelope vs. emotional feedback in animated conversational agents*. *Applied Artificial Intelligence*, vol.13, pp.519-538.
- Justine Cassell and Matthew Stone, 1999. *Living Hand to Mouth: Psychological Theories about Speech and Gesture in Interactive Dialogue Systems*. Proceedings of the AAAI 1999 Fall Symposium on Psychological Models of Communication in Collaborative Systems, pp. 34-42. November 5-7, North Falmouth, MA, 1999.
- Justine Cassell, Timothy W. Bickmore, Hannes Vilhjálmsón and Hao Yan, 2000. *More than just a pretty face: affordances of embodiment*. In Proceedings of the 5th international Conference on intelligent User interfaces.
- Justine Cassell, Yukiko I. Nakano, Timothy W. Bickmore, Candace L. Sidner and Charles Rich, 2001. *Non-verbal cues for discourse structure*. In Proceedings of the 39th Annual Meeting on Association For Computational Linguistics.
- Richard Catrambone, 2002 *Anthropomorphic agents as a user interface paradigm: Experimental findings and a framework for research*. In: Proceedings of the 24th Annual Conference of the Cognitive Science Society (pp. 166-171), Fairfax, VA, August.
- Nicole Chovil, 1992. *Discourse-Oriented Facial Displays in Conversation*. *Research on Language and Social Interaction*, 25, 163-194.
- Petra Fagerberg, Anna Ståhl, Kristina Höök, 2003. *Designing Gestures for Affective Input: an Analysis of Shape, Effort and Valence*. In Proceedings of Mobile Ubiquitous and Multimedia, Norrköping, Sweden.
- Julie Goldberg, Mari Ostendorf, Katrin Kirchhoff, 2003. *The impact of response wording in error correction subdialogs*, In EHSD-2003, 101-106.
- Kate Hone, 2005. *Animated Agents to reduce user frustration*, in The 19th British HCI Group Annual Conference, Edinburgh, UK.
- Adam Kendon, 1990. *Conducting interaction: patterns of behaviour in focused encounters*, Cambridge University Press.
- Andrea Kleinsmith, P. Ravindra De Silva, Nadia Bianchi-Berthouze, 2006 *Cross-cultural differences in recognizing affect from body posture* *Interacting with computers* 10 1371-1389
- Robert M. Krauss, Yihsiu Chen and Purnima Chawla, 1996 *Nonverbal behavior and nonverbal communication: What do conversational hand gestures tell us?* In M. Zanna (Ed.), *Advances in experimental social psychology* (pp. 389 450).San Diego, CA: Academic Press.
- Dominic W. Massaro, Michael M. Cohen, Jonas Beskow and Ronald A. Cole, 2000.*Developing and evaluating conversational agents*. In *Embodied Conversational Agents* MIT Press, Cambridge, MA, 287-318.
- Sharon Oviatt. 1994. *Interface techniques for minimizing disfluent input to spoken language systems*. In Proc. CHI'94 (pp. 205-210) Boston, ACM Press, 1994
- Sharon Oviatt and Robert VanGent, 1996, *Error resolution during multimodal humancomputer interaction*. Proc. International Conference on Spoken Language Processing, 1 204-207.
- Sharon Oviatt, Margaret MacEachern, and Gina-Anne Levow, G.,1998. *Predicting hyperarticulate speech during human-computer error resolution*. *Speech Communication*, vol.24, 2, 1-23.
- Sharon Oviatt, and Bridget Adams, 2000. *Designing and evaluating conversational interfaces with animated characters*. *Embodied conversational agents*, MIT Press: 319-345.
- Isabella Poggi, 2001. *How to decide which gesture to make according to our goals and our contextual knowledge*. Paper presented at Gesture Workshop 2001 London 18th-20th April, 2001
- Ruben San-Segundo, Juan M. Montero, Javier Ferreiros, Ricardo Córdoba, Jose M. Pardo, 2001 *Designing Confirmation Mechanisms and Error Recover Techniques in a Railway Information System for Spanish*. SIGDIAL. Septiembre 1-2, Aalborg (Dinamarca).
- Heike Schaumburg, 2001. *Computers as tools or as social actors?the users' perspective on anthropomorphic agents*.*International Journal of Cooperative Information Systems*.10, 1, 2, 217-234.