

Cognitive-model-based interpretation of emotions in a multi-modal dialog system

Michael Streit, Anton Batliner, Thomas Portele

Angaben zur Veröffentlichung / Publication details:

Streit, Michael, Anton Batliner, and Thomas Portele. 2004. "Cognitive-model-based interpretation of emotions in a multi-modal dialog system." In *Affective dialogue systems: Tutorial and Research Workshop, ADS 2004, Kloster Irsee, Germany, June 14-16, 2004*, edited by Elisabeth André, Laila Dybkjær, Wolfgang Minker, and Paul Heisterkamp, 65–76. Berlin: Springer. https://doi.org/10.1007/978-3-540-24842-2_7.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



Cognitive-Model-Based Interpretation of Emotions in a Multi-modal Dialog System

Michael Streit¹, Anton Batliner², and Thomas Portele³

¹ DFKI, Saarbrücken, Germany
`streit@dfki.de`,

² Chair for Pattern Recognition, University of Erlangen, Germany
`batliner@informatik.uni-erlangen.de`,

³ Philips Research Laboratories GmbH, Aachen, Germany
`Thomas.Portele@philips.com`

Abstract. The paper presents the cognitive-model-based approach of abductive interpretation of emotions that it is used in the multi-modal dialog system SmartKom¹. The approach is based on the OCC model of emotions, that explains emotions by matches or mismatches of the attitudes of an agent with the state of affairs in the relevant situation. It is explained how eliciting conditions, i.e. abstract schemata for the explanation of emotions, can be instantiated with general or abstract concepts for attitudes and actions, and further enhanced with conditions and operators for generating reactions, which allow for abductive inference of explanations of emotional states and determination of reactions. During this process concepts that are initially abstract are made concrete. Emotions may work as a self-contained dialog move. They show a complex relation to explicit communication. Additionally we present our approach of analyzing indicators of emotions and user state, that come from different sources.

1 Introduction

For a long period, the concept of rational agents that exchange rational arguments, was the predominant paradigm for research on dialog systems. In the last decade the scientific community became aware of the fact that emotions, moods and other attitudes play an important role in natural communication. While there are considerable advancements in generating affective artificial agents that display believable emotions in appropriate situations (cf. [1]), the recognition and interpretation of human emotions in dialog systems is still in its infancy. The term *emotion* normally aims at pronounced, clear forms of human states marked by strong feelings such as, e.g., anger, fear, sadness, joy, etc. – the so called “full-blown, big” n (n typically ranging between 4 and some twenty) –

¹ This work was funded by the German Federal Ministry of Education, Science, Research and Technology (*BMBF*) in the SmartKom project under Grant 01 IL 905. The responsibility for the contents of this study lies with the authors.

emotions. At a close look, however, almost nothing is that clear-cut: the underlying (bodily and cognitive) processes are not yet fully understood, emotions do often occur in mixed, not in pure forms, their marking can be overtly suppressed due to social constraints and rules [2,3,4,5], and there is no full agreement as for a catalogue of emotions, and of pivotal characteristics, telling emotions apart from other states as attitudes, mood, etc.

Research concerned with the generation of an affective and believable behavior of artificial agents is often based on the so-called **OCC model of emotions** [6] that explains emotion by cognitive processes relating the user's goals, standards, likes and dislikes to the actions of other agents and the state of the world that results from these actions.

Though mixing or suppressing emotions is a problem for the recognition of emotions as well as for the fine-tuning of the artificial generation of emotional behavior, the OCC model provides a systematic account for relating a certain situation to emotional states that fit to this situation. The logical structure of the situation that causes a certain emotion is not affected by the question of how intensive an emotion is or if it is displayed at all.

For research, which is concerned with the detection of problematic situations in communication by analyzing the user's behavior, **not only emotions** are relevant. This is the case, independently from the question whether the catalogue of emotions is completely defined or not. For instance, if the user is hesitant, she may need help, or if she is tired system probably should recommend the user to stop some activity that needs high attention. We use the term "(emotional) user states" to encompass all non-neutral, somehow marked behavior of the user within a human-machine-communication. From this point of view, user states as bored, stressed, irritated, tired, and so on, can and have to be addressed as well, irrespective of whether they belong to the one or the other psychological or physiological category.

In contrast, the psychological or physiological category of a state is relevant for its interpretation. The spirit of the approach, namely to consider what type of conditions elicit the affective state of the agent, extends to some non-emotional states, but not to all. For instance, explaining the state of being hesitant by a lack of information may lead to a helpful reaction, while identifying the cause of the state of tiredness (in the literal meaning) is - if possible at all - of limited use in a dialog system. Rather, the system has to consider the possible consequences of the user's state.

Some important **conditions** have to be met, however, if one wants to deal with user states in an automatic system:

- It must be possible to classify the state correctly up to a satisfying extent;
- thus there has to be a sufficiently large training sample, and
- the respective user state can be processed within the whole system, not only at the classification stage.

The first condition means that we should start with user states that are clearly marked. This rules out states such as "slightly irritated" - even if they might occur quite often and have a strong impact on the felicity of communication. We

are thus left with those pure emotions like anger or joy which do, alas, not occur often in real or in Wizard-of-Oz human-machine-communications [4,3,5,7].

Overview: The focus of the paper is to present the cognitive-model-based approach of abductive interpretation of emotions as it is used in the SmartKom system. We do not elaborate on the recognition methods that are utilized in SmartKom. We refer to [8] for a presentation of the recognition of emotions and user states by prosodic analysis and to [9] for a presentation of the recognition from facial expression.

We start with a brief description of the architecture of emotion analysis in SmartKom in the second section. In the third section we introduce the type of interaction that we want to realize in the system. In the fourth section we present our approach of calculating evidence for certain user states by combining indicators from different sources. The remaining sections are dedicated to the interpretation of emotions and user states and the generation of reactions to these states.

2 Emotion Processing in the SmartKom System

SmartKom (www.smartkom.org) is a multi-modal dialog system that provides access to multiple applications [10]. In addition to input-modalities that are used for intentional communication, the system accounts for the emotional state of the user as it is displayed by facial expression or by prosody. The processing of emotions and user states consists of three stages:

- At the first stage the emotional state of the user is recognized from facial expression and prosody.
- At the second stage indications of problematic situations and the emotional state of the user are collected from several sources and collectively evaluated. The component also analyzes the dialog with respect to the style of interaction and the task and paradigm knowledge of the user (cf. [11]).
- The interpretation of emotions and user states, and the generation of reactions to these states build the third stage. It is realized by the so-called dynamic help, a component that is dedicated to manage sub-dialogs to provide presentation specification and intention analysis in problematic situations that are not handled by the standard dialog component of SmartKom (cf. [12]).

3 The Use Cases

To demonstrate the added value of user state classification and its subsequent processing in the SmartKom system, we designed so called **Use Cases**. The first use case is intended to show how a merely emotional reaction, without explicit communication, can work as a self-contained dialog move. In this case, joy or anger are interpreted as positive or negative feedback. In the second use case

emotion works as a semantic operation that turns a positive feedback into a negative one, which is considered as a form of sarcasm.

In both use cases, the system suspects that the emotional reaction may be caused by a like or a dislike concerning the properties of the presented objects. If reasonable candidates of such likes or dislikes can be identified that are not already known by the system, it starts a preference update dialog.

If the system knows positive or negative preferences, it first presents objects that contain a preferred feature; objects that show a disliked feature will be shown last.²

user: *What's on on TV tomorrow?*

system: shows talk show at the top of the display, in the middle popular music, and crime at the bottom.

user: *And what's in the evening, in the First Program?*

system: shows a science fiction movie.

First constellation: emotion-only:

user: displays joy via facial gestures.

system: *Do you like science fiction? Shall I account for that in future presentations?*

Second constellation: emotionally marked verbal communication:

user: *That's really a beautiful program!* (She produces this sentence with an angry prosody. The positive feedback is analyzed as being sarcastic)

system: *You don't like science fiction? Shall I account for that in future presentations*

user: *Yes./No.*

system: *OK. I'll take care of that!*

(Suppose, the user's answer was yes: In the first constellation **science fiction** will be presented at the beginning of a presentation, in the second constellation at the end.)

user: *Please, again tomorrow's program!*

system: Shows **science fiction** at the beginning (at the end) of a presentation.

Instead of answering *no* or *yes* the user may also correct the supposed like or dislike, e.g. by saying *No, I like crime movies* or she may just ignore the question, by moving to a different topic. In such cases, the system will simply not re-arrange the order of presentation.

² It is possible that an object has liked and disliked attributes, e.g., there may be a movie with a preferred genre, in which a disliked actor plays.

4 Analysis of Indicators of Emotions and Problematic Situations

We introduced in SmartKom a component, the *interaction module*, that collects and evaluates indications of emotions, problematic situations and other aspects of the interaction. Indicators can have values between 0 and 1 and these values may change in time.

The interaction module provides a set of **models** each representing a certain user state or a feature of the interaction as output. Each model value is also in the range between 0 and 1. Several models support the recognition of emotions and try to detect problematic situations during a dialog. The indicator values are mapped to the models by means of a matrix multiplication.

Problematic situations and user state information are expressed by three models:

- One model describes the likelihood that the user is angry by combining scores from facial expression analysis, emotion extraction from prosody, and use of certain words.
- A second model combines confidence values from recognizers and similar scores from speech analysis, domain model, discourse history and intention recognition as well as differences in the distribution of these values among concurring hypotheses; this model is supposed to indicate problems in the analysis part of the system.
- A third model estimates the dialog progress. Here, the ratio of new information items to total information items (after completion by discourse analysis) is employed as one important indicator.

5 Cognitive-Model-Based Interpretation of Emotions

Our approach to the analysis of emotions is based on the OCC model of emotions developed by Ortony, Clore and Collins. Following the OCC model, emotions are characterized by their eliciting conditions. These conditions consist of a certain combination of

- the goals of the agent in this situation
- her attitudes to certain events (mainly likes and dislikes)
- the standards that she uses to (morally) judge an event
- the facts that hold in a certain situation
- the actions (of other agents) that caused these facts

For triggering an emotion, it is important how facts are related to the goals and the likes and dislikes of the user. Especially, it is interesting if they coincide or not. Standards are important for emotions as anger or gratitude that contain criticism or praise of another agent based on her actions. Eliciting conditions can be viewed as expressing the cause of an emotion by providing a cognitively comprehensible explanation of an emotion. The following eliciting condition for anger is taken from [13]:

```

anger(Agent1,Agent2,State,Sit) if
    holds(did(Agent2,Action),Sit),
    causes(Action,State,Sit0),
    wants(Agent1, non_State,Sit),
    blameworthy(Agent1,Action)) ,
    (Sit0 < Sit)

```

This condition means that the agent is angry, if she believes that another agent caused some state of affairs that contradicts her goals by performing an action that is not acceptable according to the user's standards (expressed by the *blameworthy* predicate). By the situation variables Sit, Sit0, one can express how the elements of the conditions are connected with respect to the sequence of situations that occur (subsequently we will omit situation variables).

Recognizing the intensity of emotions could provide additional valuable information, e.g., slight anger may occur at the beginning of a problem, while strong anger may indicate an enduring problem. But the recognition of the situation that caused the emotion and the generation of appropriate reaction is basically the same whether emotions are displayed slightly or strongly.

5.1 Abductive Interpretation of Eliciting Conditions

The OCC model is mainly used for the generation of the behavior of an animated agent. In this case, one can deliberately define the agent's likes, dislikes and standards in advance. If we want to interpret emotions that are displayed by an agent, we have to find out, which combination of facts, attitudes and standards may have caused the emotion. Our approach is to achieve this by analyzing eliciting conditions in an abductive manner. Abduction as a form of practical inference is introduced by Peirce [14]. Abduction is often characterized as inference to the best explanation: Suppose, we observe some fact A, which is surprising for us. If we know the rule

$$B, C \rightarrow A$$

(i.e. A is true if B and C are true), then we may suspect that also B and C are true, because this would plausibly explain A. If we know that there is another rule

$$D \rightarrow A$$

then D is another candidate for explaining A. Hence we need a criterion to decide which explanation is better. The quality of an explanation depends on two factors: Do we know all relevant rules (i.e., explanations)? Do we possess criterions to choose from explanations?. With eliciting conditions we have the advantage to possess schemata that claim to characterize all possible explanations of an emotion.

5.2 Problems with Abductive Interpretation

Eliciting conditions are abstract schemata that cannot be used directly to infer possible causes of emotions. To perform abductive reasoning on eliciting conditions, we have to identify concepts that could be filled into the schemata.

Seemingly, we are in a problematic situation. The system has no information about the user's standards, likes and dislikes in advance. It can get information about her goals from the user's input. But, on the one hand, this information may be based on misunderstanding, and, on the other hand, the user may have goals which cannot be recognized from her utterances. Similar problems occurs with the actions of the system. Action that are based on misunderstanding are not relevant for the analysis of the user's emotion³.

5.3 Abstract Goals and Actions for Emotion Interpretation

To overcome the problems mentioned in the last paragraph, we introduce meta-goals concerning general principles of communication and abstract goals concerning user needs that (to some extent) depend on the application. For every meta-goal or abstract goal we introduce an abstract action that satisfies the goal.

For instance, to account for misunderstandings, we introduce *understanding* as an action on the meta-level and *to be understood* as a goal on the meta-level. To account for user preferences, we introduce the concept that *a presentation accounts for the user's preferences* as an abstract action of the system — let it be called *presentByPreferences* — and accordingly the possible abstract fact or user goal *isPresentedByPreferences*⁴. This goal is abstract and under-specified because we do not know the concrete preferences of the user. Further, the relevant types of preferences depend on the type of the application.

Reasonable goals (facts, actions, likes, standards) have to be identified by careful analysis of general principles of communication and the needs of the user with respect to the type of applications she is working with. This needs empirical validation, which could not be provided within the scope of the SmartKom project. Which set of concepts is chosen, depends also on practical decisions: which goals will the system support at all, will the system possibly recognize goals that it is not able to handle, will the system react on any recognized emotion in some way (e.g. by regretting as a default in case of anger), or will it only react to emotions to which it can provide a repair or other meaningful cooperative reaction? We demonstrate the approach by the example of **anger**.

General Concepts. We first look for actions or facts that may contradict the user's wishes, likes, dislikes, or standards on a general level. Important candidates for abstract actions that contradict the user's wishes are **misunderstanding**, **slow processing** and **requests with a negative or disliked outcome**. Accordingly, we stipulate abstract or general goals, e.g., the goal to be understood properly.

³ Although the type of the action that the system wrongly performs, may influence the intensity of the user's negative feelings

⁴ For convenience we often identify the name of the fact and the name of the goal to make this fact true.

Application Dependent Concepts: Problematic Results of Database Queries. According to our use cases we concentrate on requests with liked or disliked outcome as a source of negative or positive emotions. We identified four types of disliked results:

- the result is empty,
- the majority of retrieved objects show features that are not liked by the user,
- the objects are presented in a way that is contrary to the preferences of the user, e.g, by presenting disliked objects first, and
- the user query resulted in a recall, which is too large. The user may need help for further specification possibilities.

We assume for the list of topics above, that the disliked or problematic results are not due to misunderstanding. Misunderstanding is taken as evoking its own class of constellations. If misunderstanding is involved, the result is not relevant for the analysis.

User specified goals and system initiated actions. As far as no misunderstanding is involved, the SmartKom system will usually simply follow the user's specification. If this works, no anger should arise with respect to the fact that the system tries to achieve this goal (but perhaps instead *joy*). In specific situations the system may initiate actions that are necessary from the point of view of the system, but may be disliked or even considered blameworthy by the user. For instance, the system may require a biometric verification, which the user dislikes. Such actions are relevant for explaining negative emotions, but are not considered in our implementation.

As an anonymous reviewer pointed out, that inappropriate or undesired reactions on emotions could also be a cause for anger (or for being bored). In fact this is a subcase disliked system initiated actions.

6 Analyzing and Handling of Pure Emotion

With the concepts introduced in the last section, we are able to build instantiations of eliciting conditions that allow to infer combinations of goals, facts, actions, likes and dislikes that possibly explain the user's emotion. We call instantiations of eliciting condition schemata **eliciting constellation**. To get criterions for selecting the relevant constellation, we augment constellations with conditions and organize these conditions internally as a decision tree⁵.

Further, the system has to determine reactions that are appropriate for

- resolving the situation that caused the negative emotion,
- avoiding negative emotions in future in similar situations, and
- promoting the occurrence of positive emotions in similar situations.

⁵ As mentioned in the conclusion we could perform testing only in limited way. Thus no training of the decision tree was possible

It is also desirable to include methods that provide abstract under-specified goals and actions with presumable values. Such values are not only used for determining concrete system reactions, they serve as a part of the constellation conditions.

According to our use cases, we have to consider database queries that retrieve disliked objects. The system offers as repair that it will regard the likes and dislikes of the user in its presentations.

A constellation for handling anger according to our use cases is given below (leaving out some minor details) in a Prolog style notation. It applies to browsing television program or cinema program. For these applications preferences are actually taken into account for the presentation. This rule are basically processed in the following manner: First the conditions are tested (internally the conditions are processed in a decision tree like order). Then the cause of the emotion, which is represented by the clauses above the conditions, is considered as a reasonable explanation, whereby the variables are filled by the result of the condition processing. Then the system action is performed.

```
anger(thisConstellation,user,system) if
  holds(did(system, non_presentByPreference(dislike(user,X))),
  causes(non_presentByPreference(dislike(user,X)),
    non_isPresentedByPreference(dislike(user,X))),
  wants(user, isPresentedByPreferences(dislike(user,X))),
  blameworthy(user, non_presentByPreferences(dislike(user,X))),
  conditions(thisConstellation,X),
  (proposed system action:) update(dislike(user,X)).
```

The constellation expresses, that there is a concrete reading of the goal, that there is a concrete reading of the goal *presentByPreferences* that may be a goal of the user, that this goal is not satisfied, and that ignoring the goal is against the standards of the user. The constellation contains facts and actions that are not concretely specified. For instance we do not know whether the presentation contains some possibly disliked feature, and we do not know which feature it is.

We test the salience of the constellation by establishing the following conditions. The predicate *presentationEntriesContainCommonFeature(X)* also delivers a concrete presumable instance of the user's dislike.

```
conditions(thisConstellation,X) if
  presentationEntriesContainCommonFeature(X),      (1)
  non_specified(user,X),                            (2)
  non_knows(system,like(user,X)),                   (3)
  non_knows(system,dislike(user,X)).                (4)
```

(1) verifies if the user perceives *too many* objects with the supposed disliked feature. (It also excludes the case that there is no result at all, which would support a different explanation for anger). It is important for the other tests, that the predicate delivers a hypothesis for the disliked feature. (2) excludes that the user is angry about the occurrence of features that she has specified in her

request (there is a possibility of misunderstanding). (3) excludes, that the user is angry about a feature, that she has already declared to like. (4) excludes that the system in fact tried to present the disliked feature appropriately but just did not find other objects.

emotion. For emotions displayed by facial expressions, we prove if the emotion emerges in a certain time interval after the presentation was displayed. With prosodically displayed emotion we prove if the verbally expressed content was compatible with the explanation of the emotion. It turned out that it is not sufficient to test if there are already stored preferences. It should additionally be proved, if a user has not agreed with storing a preference. This has to be remembered, otherwise the system may propose the same preference repeatedly.

The action *update(dislike, user, X)*, which is attached to the constellation, initiates a sub-dialog that verifies if the user has the supposed dislike. It is not only a repair action, but takes part in the explanation process.

The conditions mentioned so far are not sufficient to **discriminate competing explanations**. Such competing explanations have to be modeled, even if no reaction is foreseen for these cases. We distinguished three main sources of anger: misunderstanding, slow processing, and requests with a negative or disliked outcome. Evidence for problems in the analysis part is detected by the interaction module (cf. section 4). Slow processing is a possible explanation for anger, if anger occurs during the analysis. Also the absolute duration of processing is a criterion. These dates are accessible via a module (the so called watchdog) that monitors the processing state of the system.

7 Emotions and Communicative Acts

Emotions that are signaled by facial expressions do not need to be accompanied by additional communication at all. Emotions expressed by voice are naturally related to some acoustic output. In the extreme, this output is only a container for the expressed emotion, but usually it contains a certain semantic content. The analysis of the relation between semantic content and underlying emotions is in its infancy, compared, e.g., with the relation between verbally communicated semantic content and pointing gestures. The latter is sufficiently known to build practical application. We distinguish in the following between communicative acts with semantic content, that are provided by speech and gestures, on the one hand, and emotions on the other hand.

The interpretation of pointing gestures and verbal utterances can be conceived as a fusion process, which unifies pieces of information. Semantic contradictions between pointing gestures and verbally provided information are indications for errors. The relation between emotions and communicative acts is much more complicated. We give a presumably non-exhaustive classification of types of interaction between displayed emotion and communicated semantic content.

Redundancy. Semantic content redundantly expresses a simultaneously displayed emotion as *that makes me angry* or *I'm glad about that* or semantic

content expresses an attitude that corresponds to the direction of the emotion (whether it is positive or negative) as *great*, *bad*.

Contribution to the Explanation of the Emotion. Semantic content expresses a concrete attitude (like or dislike) that is involved in triggering the emotion as *I don't like thrillers* or *great movies*, or semantic content addresses the facts and actions that caused the emotion as *you didn't understand me* or *that takes too much time* or simply by uttering *thrillers* accompanied by a positive or negative emotion.

The *thriller* example contributes the concrete feature, that may fill the abstract goal of being presented accordingly preferences. But this example does not necessarily express a like or dislike as *great movies*. With a negative emotion, the example may also belong to the topic *Semantic Content as Repair Action*.

Semantic Content as Repair Action. The semantic information is provided to repair the state of affairs, that has caused the emotional state of the user. The example *thriller* works also here: *thriller* could be a correction of a misunderstanding of genre. There is no direct relation between the content of the utterance and the displayed emotion.

This is very common and important in human-machine-dialog as well as in human human dialog: the dialog partner repeats or reformulates her request and concurrently displays a negative emotion. With overt anger, it could also be expected that the user cancels the interaction as a final form of repair.

Change of Semantic Content. The user displays a negative emotion and communicates verbally a positive attitude *marvelous*, *great movies*. The direction of the valenced attitude that is communicated verbally is changed by the direction of the displayed emotion. This is a simple form of sarcasm.

8 Results and Conclusion

A complete implementation of the whole processing chain was available at the end of the project. There was no opportunity for systematic tests, which require high effort. For instance, the recognition of facial expression needs careful preparation of the environment in respect to lighting conditions to work. Our limited testing shows, that, provided recognition is correct, the emotion interpretation generates the reactions that are requested by the use case specification.

We implemented successfully a cognitive-model based approach for analyzing emotions and other affective states of a user that participates in a multi-modal human-machine dialog. This is a success, but it will still take considerable effort to make it practically useful. The approach is based on an elaborated theory, which covers a broad range of phenomena. This is promising with respect to the extensibility of the approach. It is an important advantage of the approach that it generates conceivable explanations of emotions, that allow for well directed

system reactions. The approach is not restricted to handle classical emotions, but extends to other affective states. Also it is not restricted to states, that are displayed non-verbally. Affective verbal feedback, as *I like this*, can be explained along similar lines.

References

1. Picard, R., ed.: Affective Computing. The MIT Press (1997)
2. Cornelius, R.R.: Theoretical Approaches to Emotion. In Cowie, R., Douglas-Cowie, E., Schröder, M., eds.: Proc. ISCA Workshop on Speech and Emotion, Newcastle, Northern Ireland (2000) 3–10
3. Batliner, A., Fischer, K., Huber, R., Spilker, J., Nöth, E.: How to Find Trouble in Communication. *Speech Communication* **40** (2003) 117–143
4. Batliner, A., Zeissler, V., Frank, C., Adelhardt, J., Shi, R.P., Nöth, E.: We are not amused - but how do you know? User states in a multi-modal dialogue system. In: Proc. European Conf. on Speech Communication and Technology. Volume 1., Geneva, Switzerland (2003) 733–736
5. Batliner, A., Hacker, C., Steidl, S., Nöth, E., Haas, J.: User States, User Strategies, and System Performance: How to Match the One with the Other. In: Proc. of an ISCA tutorial and research workshop on error handling in spoken dialogue systems, Chateau d'Oex, ISCA (2003) 5–10
6. Ortony, A., Clore, G., Collins, A., eds.: The Cognitive Structure of Emotions. Cambridge University Press (1988)
7. Ang, J., Dhillon, R., Krupski, A., Shriberg, E., Stolcke, A.: Prosody-based automatic detection of annoyance and frustration in human-computer dialog. In: Proc. ICSLP 2002. (2002) 2037–2040
8. Adelhardt, J., Shi, R., Frank, C., Zeißler, V., Batliner, A., Nöth, E., Niemann, H.: Multimodal User State Recognition in a Modern Dialogue System. In: Proc. of the 26th German Conference on Artificial Intelligence, KI'03, Chateau d'Oex, Springer (2003) 591–605
9. Frank, C., Nöth, E.: Optimizing Eigenfaces by Face Masks for Facial Expression Recognition. In: Computer Analysis of Images and Patterns – CAIP 2003, Lecture Notes in Computer Science. Springer, Heidelberg (2003) 1–13
10. Wahlster, W.: SmartKom: Symmetric multimodality in an adaptive and reusable dialog shell. In: Proc. Human Computer Interaction Status Conference, Berlin (2003)
11. Portele, T.: Interaction Modeling in the SmartKom system. In: Proc. of an ISCA tutorial and research workshop on error handling in spoken dialogue systems, Chateau d'Oex, ISCA (2003)
12. Streit, M.: Context Dependend Error Handling by a three Layered Processing Model. In: Proc. of an ISCA tutorial and research workshop on error handling in spoken dialogue systems, Chateau d'Oex, ISCA (2003)
13. Prendinger, H., Descamps, S., Ishizuka, M.: Scripting affective Communication With Life-Like Characters in Web-Based Interaction Systems. *Speech Communication* **16** (2002) 519–553
14. Peirce, C.: Abduction and Induction. In: Philosophical Writings of Pierce. Springer, New York (1995) 150–156