

# Gaze-X: Adaptive, Affective, Multimodal Interface for Single-User Office Scenarios

Ludo Maat<sup>1</sup> and Maja Pantic<sup>2,3</sup>

<sup>1</sup> EEMCS, Delft University of Technology, Delft, The Netherlands

<sup>2</sup> Computing Dept., Imperial Collge London, London, UK

<sup>3</sup> EEMCS, University of Twente, Enschede, The Netherlands  
ludomaat@zonnet.nl, m.pantic@imperial.ac.uk

**Abstract.** This paper describes an intelligent system that we developed to support affective multimodal human-computer interaction (AMM-HCI) where the user's actions and emotions are modeled and then used to adapt the interaction and support the user in his or her activity. The proposed system, which we named Gaze-X, is based on sensing and interpretation of the human part of the computer's context, known as W5+ (who, where, what, when, why, how). It integrates a number of natural human communicative modalities including speech, eye gaze direction, face and facial expression, and a number of standard HCI modalities like keystrokes, mouse movements, and active software identification, which, in turn, are fed into processes that provide decision making and adapt the HCI to support the user in his or her activity according to his or her preferences. A usability study conducted in an office scenario with a number of users indicates that Gaze-X is perceived as effective, easy to use, useful, and affectively qualitative.

**Keywords:** Human Sensing, Context Sensing, Multimodal Interfaces, Affective Computing, Anthropocentric Interface Design.

## 1 Why Multimodal, Context-Sensitive User Interface Design?

We have entered an era of pervasive computing. Computers and the Internet have become so embedded in the daily fabric of our lives that we can no longer live without them [15]. We use them to work, study, communicate, shop, and entertain ourselves. With the ever-increasing diffusion of computers into society, human-computer interaction (HCI) is becoming increasingly essential to our daily lives.

Predicting the future of HCI is a difficult task, but one important source of help is the accumulated information about the preferences and limitations of humans interacting with computers. Principles can be drawn upon, which may explain why some interfaces survive and others become extinct. Rigid designs that assume that users will be explicit and fully attentive while interacting with the computer, that do not protect against errors, provide help at all times except at the right moment, and all in all make users frustrated, are likely to become quickly extinct due to their poor usability [28]. On the other hand, designs that include adequate attention to individual differences among users, support (natural) multimodal and context-sensitive

interaction, expend on designs for reliability and safety, provide access to the elderly and handicapped, and properly adapt to the user level of knowledge, skills, attention, preferences, moods, and intentions, are the kind of HCI designs that are likely to become the trend in computing technology [4], [41], [48], [39], [38], [51]. Although this list may not be complete, it points out important issues that are rather insufficiently addressed by the current initiatives [43].

### **1.1 The Evolution of Human-Computer Interfaces**

Around 1980, at the dawn of the personal computer age, many chaotic and rigid user interfaces were produced that turned the users into frustrated victims of machines they could not control. Typical examples of useless interfaces at that time could display a five-minute video without a stop button and generate choice sequences that could not be reversed or cancelled. As high-resolution displays and fast chips emerged, video and audio processing as well as animations flourished (particularly for video games), giving rise to a new generation of user interfaces, in which direct manipulation became the dominant form of interacting and WYSIWYG (what you see is what you get) became a guiding principle. The aim was: (i) to make operations visible, incremental, rapidly manageable by means of a keyboard, and reversible, as well as (ii) to prevent user errors by effective designs. During the late 80s and early 90s, direct-manipulation interfaces were enhanced with embedded menus in text and graphics, mice, and various joysticks as the devices of choice.

As remarked by Preece and Schneiderman [41] and Pentland [37], the mid 90s can be viewed as the dawn of pervasive computing that shed a new light on the future of computing and gave rise to novel requirements that useful user interfaces should fulfill. The growing availability of World Wide Web access with embedded menus providing links across the world led to an unusually rapid growth of Web servers and applications. The necessity of delivering new products in an ever-decreasing time frame affected, consequently, the quality of the issued products and interfaces. This Internet hype also blurred the essence of some paradigms, such as software agents since an increasing number of vaguely related applications needed legitimacy and sought it under the umbrella of the “agents”. Though it was unfortunate at one hand and accompanied by numerous shoddy Web-oriented applications, on the other hand this Internet hype initiated rapidly accelerating progress in facilitating accessibility, speed, and reduction of error and failure rates. Moreover, it changed our view on computing and commerce [50]. Above all, it forecasted the type of working environments and information-communication spaces we are about to use in our everyday activities. It clearly indicated that in the future, with the aid of computers, we will carry out our daily tasks, communicate, and entertain ourselves in cyberspaces across distance, cultures and time. Of course, the specifics of such cyberspaces including virtual cyber worlds and of the related interfaces, which should facilitate easy and natural communication within those spaces and with the variety of embedded computing devices, are far from settled. In this LNAI Special Volume on AI for Human Computing many relevant issues are discussed and debated but all agree that before this new generation of pervasive computing can be widely deployed, the users should experience it as being human-centered and universally usable.

The designers of older technologies such as postal services, telephones, and television, have reached the ultimate goal of having products and services that are universally usable, but developers of computing technology cannot claim the same. Schneiderman [48], [49], reports an average of 5.1 hours per week wasted by the users while trying to use computers. Consequently, despite visible progress in accessibility, increase of speed, and reduction of the error and failure rates, the primary experience of many computer users is dissatisfaction or even frustration. Common problems include incompatibility (of file formats, applications' versions, screen sizes, etc.) and low speed (e.g., due to varying network bandwidths and processor speeds). Although these issues are not of the least importance, the crucial problem, which is primarily responsible for the users' anxiety and dissatisfaction, is incomprehensibility of many currently available software packages and Internet services [47]. As remarked by several researchers (e.g., [48], [49], [47]), a large majority of HCI systems available today assume users' proficiency in computing. However, for a casual user, these systems are often cumbersome, lacking the adaptability necessary to accommodate users with various levels of computing skill and experience. Furthermore, virtually all "classic" HCI systems tend to confine the user to a less natural, single-modal means of interaction (e.g., a mouse movement, pressing of a key, speech input, or hand motion). For example, to manipulate a virtual object with a typical HCI system, the user is usually required to select the object by employing mouse motion, then point with the mouse at a control panel to change the object's color. On the other hand, in a more natural setup, the user would point at the object with his finger and say: "Make it red". Integration of more than one natural modality into an interface would potentially overcome the current limitations of HCI systems: it would ease the need for specialized training and ease the information- and command-flow bottleneck between the user and the computer. Besides, recent data shows that a multimodal HCI can be an effective means for reducing uncertainty of single-modally sensed data (such as speech or hand motion), thereby improving robustness [43]. Although the incorporation of all features of human-human interaction (i.e., an intricate interplay of thoughts, language, and non-verbal communicative displays) into human-computer interaction may be very complex and difficult to achieve, equipping HCI systems with a multimodal setup so that they can approach naturalness, flexibility and robustness of human-human communication will give them the potential to:

- transcend the traditional, cumbersome and rigid mouse/keyboard interaction, and
- yield a more effective and efficient information- and command-flow between the user and the computer system and, by that,
- approach universal usability.

Another challenge in fashioning universally usable HCI systems is to make them context sensitive. The key idea is to account for individual differences of the users and for the overall situation in which the user acts. For computing technology applications, context can be defined as any information that can be used to explain the situation that is relevant to the interaction between users and the application [10]. For

a single-user scenario, the six questions that summarize the key aspects of the computer's context with respect to its human user are as follows:

- *Who?* (Who the user is?)
- *Where?* (Where the user is? For single-user desktop-computer scenarios this context question is superfluous as it is known where the user is – either in front of the computer or not.)
- *What?* (What is the current task of the user?)
- *How?* (How the information is passed on? Which interactive signals / actions have been used?)
- *When?* (What is the timing of displayed interactive signals with respect to changes in the computer environment?)
- *Why?* (What may be the user's reasons to display the observed cues? Except of the user's current task, the issues to be considered include whether the user is alone and what is his or her affective state.)

Based upon the user's identity and the knowledge about his or her environment and current task, it would be possible to retrieve information about the importance of that task in the given environment, the user's skills in performing that task, and the user's overall preferences. Together with the sensed user's affective state, this information could be employed to define the following.

*What form should the instruction manual have?* For example, for a novel user a lucid tutorial could be provided, to an average user constructive help files could be offered, for an expert user compact notes could prove to be sufficient. Yet the kind of help files provided to the user should not be rigid; it should be determined based upon the overall user's preferences and his/her current affective state. Namely, users might prefer compact help files even if they are novices and especially if they are in a hurry.

*When should the user be interrupted?* This is of particular importance for the systems that are designed such that they rely on the user's feedback at all times or offer assistance each time a certain task is to be performed or particular conditions are encountered. The user's current affective state and the importance of the current task given the context of use could be exploited to time the interrupts conveniently. For example, if the user is hurriedly writing e-mails, interrupting him to correct a syntax mistake can be postponed till the moment he tries to send the e-mail.

*When, in which part, and in what way should the system be adjusted?* The sensed user's affective state could be exploited to time the adjustment of the system, the information about the user's current task might form the target of the adjustment, and the complete information about the sensed context could be used to determine the adjustment properly. For instance, suppose that the user always browses through a particular application in the same way in order to come to a specific window and displays irritation each time the system starts the pertinent application by displaying its very first window. In that case, a proper adjustment might be to mark the window where the user commonly stops browsing and to start the pertinent application with that specific window (browsing through the preceding windows of the application does not have to be apparent to the user). Yet, suppose that the user always becomes frustrated if a certain person enters the office. In that case, no adjustment should be

made since the user's affective state is not caused by HCI but by an external (and for HCI) irrelevant event.

HCI design breakthroughs are necessary if the computing technology is to achieve the ultimate goal of being universally usable [43]. However, many research problems initially thought to be intractable have been proven manageable (e.g., sensing and detecting human interactive cues [33], affect-sensitive interpretation of those [34], [36], and context sensing in general [29], [33]). This initiated numerous efforts to design and develop multimodal and context-sensitive interfaces that are flexible and compatible with users' abilities and relevant work practices [30], [31].

## 1.2 The State of the Art

Several extensive survey and position papers have been published on vision-based [40], [52], multimodal [30], [31], [19], affective [24], [17], and context-sensitive interfaces [11], [29], [33]. Virtually all of these articles agree that approaching the naturalness of human-human interaction plays a central role in the future of HCI and that this objective can be approached by designing adaptive HCI systems that are affective, context-aware, and multimodal. However, many of these articles mention that the main application of this new technology is quickly changing from user interface models in which one user is sitting in front of the computer, to something else like ambient interface models in which multimodal multi-party interactions are observed and interpreted. We argue here that this statement may be very misleading. When the main application domain in a certain field has changed from *A* to *B*, this may imply (and it usually does imply) that problems in *A* have been researched, that they have been solved, and that the research has moved on to tackle other problems. In turn, the statement in question may imply that the realization of adaptive interfaces based on affective multimodal interaction models (AMM-HCI) can be considered a solved problem for single-user office scenarios. However, an extensive research of the large body of the related literature did not confirm this.

Only few works aimed at adaptive, affective, multimodal interfaces for single-user office scenarios have been reported up to date. The majority of past work in the field relate to multimodal, non-affective interaction with the user [30], [31]. Integration of multiple natural-interaction modalities such as speech, lip reading, hand gestures, eye-tracking, and writing into human-computer interfaces has long been viewed as a means for increasing the naturalness and, in turn, ease of use [3], [53], [54]. The research in this part of the field is still very active and the majority of the current work focuses on speak-and-point multimodal interaction (where pointing is conducted by either pen or hand gesture) and aims either at support of crisis management [27], [46], or at development of personal widgets [30].

A rather large body of research can also be found in the field of human affect sensing [34], [36], [33]. Most of these works are single modal, based either on facial or on vocal affect analysis. Recently, few works have been also proposed that combine two modalities into a single system for human affect analysis. The majority of these efforts aim mainly at combining facial and vocal expressions for affect recognition (e.g., [5], [55]), although some tentative attempts on combining facial expressions and body postures have been reported as well [14], [20]. In the same way that these methods do not tackle the problem of how the sensed user's affect can be

incorporated into the HCI, there is a large body of research that focuses on the role of human affect in HCI while assuming that user's affective states have been already sensed by the computer [7], [44]. Efforts that integrate these two detached research streams, and represent the related work to the one presented in this paper, are rare. Several works have been reported on developing adaptive interfaces that are based on sensing the user's affective states. These methods are usually single modal, based either on facial affect recognition (e.g., [2]) or on physiological affective reaction recognition (e.g., [18], [42]). To the best of our knowledge, the only exceptions from this rule are the single-user AMM-HCI systems proposed by Lisetti & Nasoz [23] and by Duric et al. [11]. The former combines facial expression and physiological signals to recognize the user's emotion (fear, anger, sadness, frustration, and neutral) and then to adapt an animated interface agent to mirror the user's emotion. The latter is a much more elaborate system, aimed at real-world office scenarios. It combines lower arm movements, gaze direction, eyes and mouth shapes, as well as the kinematics of mouse movements to encode the user's affective and cognitive states (confusion, fatigue, stress, lapses of attention, and misunderstanding of procedures). It then applies a model of embodied cognition, which can be seen as a detailed mapping between the user's affective states and the types of interface adaptations that the system supports, to adapt the interface in a reactive or a proactive manner to the user's affective feedback. The main drawback of this system is that it is not user-profiled, while different users may have different preferences for both the input modes per type of action and the type of interface adaptation. Another drawback is the employed model of embodied cognition, which is rigid and clumsy as it stores all possible combinations of inputs and outputs that are difficult to unlearn and reformat. Finally, as the system was not tested and evaluated, conclusions about its robustness, flexibility, adaptability, and overall usability cannot be drawn.

We believe that the main reason for the lack of research on single-user AMM-HCI is a twofold. First, the highly-probable misconception that the problem in question has been solved could cause the lack of interest by researchers and research sponsors. Second, it seems that the vast majority of researchers treat the problem of adaptive, affective, multimodal interfaces as a set of detached problems in human affect sensing, human interactive signals processing, human-human interaction modeling for HCI, and computer-human interface design. In this paper, we treat this problem as one complex problem rather than a set of detached problems and we propose a single-user AMM-HCI system similar to that proposed by Duric et al. [11]. A difference between the two systems is that ours uses a dynamic case base, an incrementally self-organizing event-content-addressable memory that reflects user preferences and allows easy reformatting each time the user wishes so. In addition, our system, which we call Gaze-X, is based on sensing and interpretation of the human part of the computer's context, known as W5+ (who, where, what, when, why, how) and, in turn, is user- and context-profiled.

The state of the art in context-aware applications ensues from two streams of research: the one on context sensing [29], [33], which focuses on sensor-signal processing (audio, visual, tactile), and the other on context modelling [10], which focuses on specifying procedures and requirements for all pieces of context information that will be followed by a context-aware application. Gaze-X integrates those two detached poles of the research. It uses a face recognition system to answer

*who* the user is and to retract his or her profile (i.e., user-profiled case base). It employs an eye-tracking system and a speech recognizer in combination with event handling of standard HCI events like mouse movements, keystrokes, and active software identification to answer *what* is the current task of the user. In addition to these input modalities, a system that recognizes prototypic facial expressions of six basic emotions (anger, fear, happiness, surprise, sadness, and disgust) is used to answer the *how* context question. To answer the *when* context question, we simply keep a log of the time and the cost (in time) of HCI events associated with various input modalities. To answer the *why* context question, which is the most complex context question, we use case-based reasoning that enables evaluation of encountered events based upon the user preferences and the generalizations formed from prior input. Based upon the conducted evaluation, Gaze-X executes the most appropriate user-supportive action. The system was initially proposed by the authors in [25].


### 1.3 Organization of the Paper

The paper is organized as follows. Section 2 gives an overview of the Gaze-X architecture. Section 3 presents the system's input modalities. The utilized case-based reasoning is explained in detail in section 4. Adaptive and user-supportive actions of the interface are discussed in section 5. The Graphical User Interface (GUI) of the Gaze-X is presented in section 6. The usability study that we carried out is discussed in section 7. Section 8 concludes the paper.

## 2 System Architecture

The outline of the system is illustrated in Fig. 1. The main modules are the multimodal input module, the reasoning module, and the feedback module. The user of Gaze-X experiences an adaptive interface, which changes as a function of the currently sensed context. This function is represented by the cases of the utilized dynamic case base, having the system's and the user's state as the input (represented in terms of exhibited multimodal interactive actions and cues) and the adaptive and user-supportive changes in the interaction as the output. The fact that the changes in the interaction do not have to occur in each time instance (e.g., if the user's preference is to remain undisturbed while working with a certain application), triggering of the feedback module is optional and illustrated using a dashed line.

Gaze-X has been implemented as an agent-based system. The main reason for doing so is to support concepts of *concurrency* (which allows sub-systems to operate independently and yet at the same time), *modularity/scalability* (which allows easy upgrade through inclusion of additional sub-systems), *persistence* (which ensures robust performance by saving intermediate settings), and *mobility* (which enables transport of agents to another host). We used Fleeble Agent Framework [32] to develop the Gaze-X. Fleeble can be seen as a common programming interface defining the behavior of all the agents build with the framework. The main characteristics of Fleeble and Fleeble-based multi-agent systems (MAS) can be summarized as follows (for details see [32]).



**Fig. 1.** Overview of the Gaze-X system: the input modalities, modules, and agents

*Fleeble enables easy development of agents and MAS.* The framework can instantiate and configure an agent and then start it up in a separate thread. The agent is autonym, although it is running in the framework's processing space. An agent can also instruct the framework to start up another agent. The framework keeps track of all agents and their parent agents. So, a single agent can be created which starts up the appropriate agents for each MAS. This *kickoff* agent is then the parent agent of all agents that form the MAS in question.

*Fleeble supports simple processing of events coming from the outside world and other agents.* Fleeble offers a message distribution system for communication between agents that is based on a Publish/Subscribe system, which is centered on the concept of a *channel*. Channels are named entities that allow a single message to be delivered to any number of agents. An agent informs Fleeble that it is interested in events pertaining to a specific channel (i.e., it subscribes to that channel). An agent can ask Fleeble to deliver a message to a channel (i.e. it publishes to the channel in question). Fleeble creates a "handler" thread for each agent that has subscribed to the channel in question. All handler threads are started at the same time and deliver the message to the subscribed agents. Hence, e.g., the user's facial cues can be simultaneously processed by both the Identity Agent and the Emotion Agent (Fig. 1).

*Fleeble supports the concept of concurrency needed to allow agents to operate independently and yet at the same time.* This has been achieved by starting each agent in a separate thread, allowing it to access the delivery system described above at its own convenience.

*Fleeble supports data and state persistence.* Fleeble agents can instruct the framework to store values referenced by a key. The framework stores this (key, value) pair and allows access to it at any time, even when the execution of the framework has ceased in the meantime. State persistency allows the user to shut down a single agent (or MAS) and to restore it from the point where it was suspended later on, even



when the PC has been shut down in the meantime. This makes Fleeble-based MAS very robust.

*Fleeble supports the concepts of distribution and mobility.* It establishes socket-to-socket connections between frameworks residing on different computers and manages these connections, i.e., creates and closes them as needed. Connections can be used to transport agents to another host, allowing agents to physically move. The process of being moved to another host can be started either by the agent or by any parent agent.

### 3 Input Modalities

The front end of Gaze-X consists of the multimodal input module, which processes images of the user's face, gaze direction, speech, and actions done while interacting with the computer including the mouse movements and the keystrokes.

The choice of input modes was largely influenced by findings in cognitive science and HCI-related research on multisensory perception and multimodal coordination. These findings reveal that people have strong preference to interact multimodally when engaged in spatial location commands, whereas they tend to interact unimodally when engaged in actions that lack a spatial component. More specifically, it seems that people prefer speak-and-point interaction in visual-spatial domains like Graphical User Interfaces [15], whereas they prefer speech input alone for issuing commands for actions and for describing objects, events, and out-of-view objects [6]. Although researchers discarded the idea of using gaze fixations as a mouse replacement, especially for extended periods, gaze direction is exploited nowadays as a reliable index of the user's interest [25]. Finally, the face is our primary means to identify other members of the species, and to communicate and understand somebody's affective state and intentions on the basis of shown facial expressions [21]. Hence, integrating face, facial expression, speech, mouse-pointing, and gaze direction input modes into the Gaze-X, promised a suitably complementary modality combination that could support the intended, adaptive, context-sensitive (user-, task-, and affect-sensitive) interaction with the user.

To process the images of the user's face acquired by a standard web-cam, we use a commercially available the Face Reader system for face and facial affect recognition produced by Vicar Vision.<sup>1</sup> This system operates as follows. It detects candidate face regions in the input scene by comparing image regions to a number of prototype faces. These prototypes are representative of a large database of human faces. An Active Appearance Model (AAM) [8] is then fitted to the detected face region (Fig. 2). The variations in the shape and texture of the AAM, caused by fitting the AAM to the detected face, serve as a unique identifier of the individual (identifiers for different users are stored in a database). Once the user is identified, AAM fitting is employed to detect the affective state of the user. The variations in the shape and texture of the AAM, caused by fitting the AAM to the user's face in the current frame, are fed to a neural network trained to recognize the prototypic facial expressions of six basic emotions (surprise, fear, anger, happiness, sadness, and disgust) [22], defined in classic psychological studies on human emotions [21]. The

---

<sup>1</sup> Vicar Vision BV, 2004. <http://www.vicarvision.nl>



**Fig. 2.** Person identification of the Face Reader system [22]. From left to right: input image, automatic placement of facial landmarks, initial positioning of the AAM, best-fitting face.

output layer of the neural network consists of 7 nodes (one for each emotion and one for the neutral state), each of which outputs a continuous value. This enables detection of blends of emotions (e.g., surprise and happiness in an expression of delight) and low-intensity emotions (e.g., a frown is recognized as a low intensity of anger).

To detect the user's gaze direction, we employ a commercial system for remote eye tracking produced by SMI GmbH.<sup>2</sup> The utilized iView X remote eye-tracking system consists of two main components, the infrared pan tilt camera and the software used for the calibration and eye tracking. To determine the direction of the gaze, the system employs the so-called red-eye effect, i.e., the difference in reflection between the cornea and the pupil. This direction determines the point on the screen, which is the focus of the user's attention (Fig. 3). As a cursor can be displayed at this point, eye-tracking can be employed to free the user from the constraints of the mouse and, in combination with spoken commands, to allow a hands-free interaction with the computer.

A variety of speech-recognition systems are now available either as commercial products or as open source. Although some novel approaches to speech recognition has been proposed recently [9], most of the existing methods utilize acoustic information of speech contained in frame-by-frame spectral envelopes which are statistically classified by Hidden Markov Models. Gaze-X utilize Sphinx 4<sup>3</sup> speech recognizer, which is a Java-implemented speech recognizer that recognizes a predefined set of words (vocabulary) based on acoustic features and a HMM architecture.

To monitor standard HCI events including keystrokes, currently active software and currently visited web-site, we utilize the Best Free Keylogger software.<sup>4</sup> It monitors website visiting/blocking, e-mail visiting, keystrokes, and application activity, and writes these data into log files. We also monitor and log the locations of the mouse cursor. In the case that the user's preference is to use eye-tracker as an alternative for the mouse, the system does not log the mouse movements.

<sup>2</sup> SensoMotoric Instruments GmbH, 2002. <http://www.smi.de>

<sup>3</sup> Sphinx-4, 2004. <http://cmusphinx.sourceforge.net/sphinx4>

<sup>4</sup> Best Free Keylogger, 2006. <http://sourceforge.net/projects/bfk/>

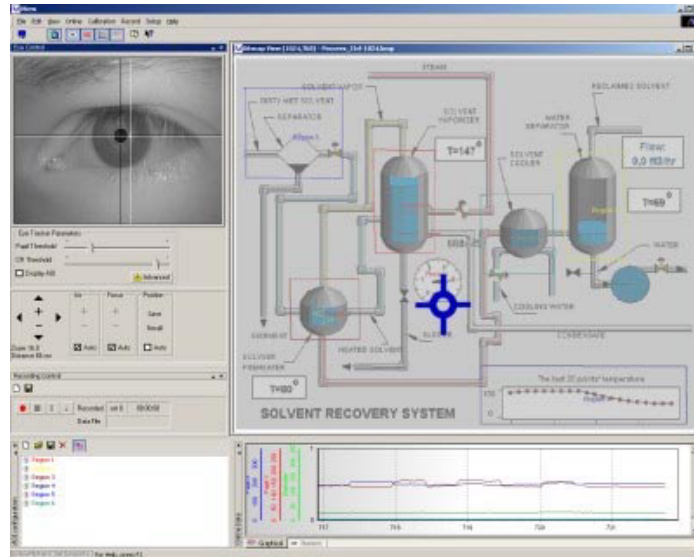


Fig. 3. Determining the user's gaze direction (focus of attention) by iView X eye-tracker

The user's identity, his or her displayed affective state, current gaze direction, and uttered words, delimit the current user's state. The HCI events including the mouse movements, keystrokes, and the currently active software delimit the current system's state. These two states form further the input to the reasoning module.

#### 4 Case-Based Reasoning

Since the Gaze-X can have different users, each of which can be using different applications in his or her daily work with the PC, while showing emotions and a variety of interactive patterns using the standard and the natural interactive modalities that the Gaze-X supports, the mapping of the system's multimodal input onto a large number of adaptive and user-supportive changes in interface in a user-profiled manner is an extremely complex problem. To tackle this problem, one can apply either eager or lazy learning methods. Eager learning methods such as neural networks extract as much information as possible from training data and construct a general approximation of the target function. Lazy learning methods such as case-based reasoning store the presented data and generalizing beyond these data is postponed until an explicit request is made. When a query instance is encountered, similar related instances are retrieved from the memory and used to classify the new instance. Hence, lazy methods have the option of selecting a different local approximation of the target function for each presented query instance [1]. Eager methods using the same hypothesis space are more restricted since they must choose their approximation of the target function before query instances are observed. In turn, lazy methods are usually more appropriate for complex and incomplete problem domains than eager methods, which replace the training data with abstractions obtained by generalization

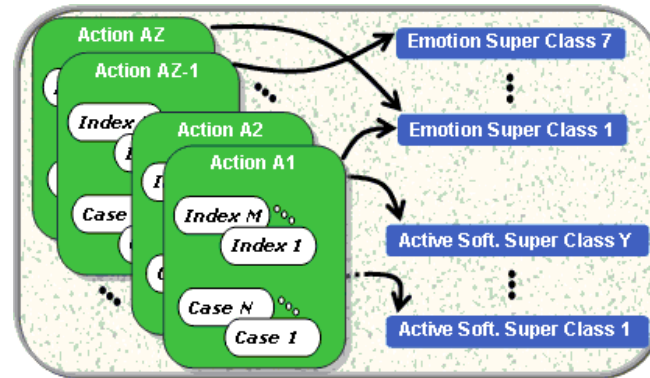
and which, in turn, require excessive amount of training data. Hence, we chose to implement the inference engine of the Gaze-X as the case-based reasoning on the content of a dynamic memory. The memory is dynamic in the sense that, besides generating user-supportive feedback by analogy to that provided to the user in similar situations “experienced” in the past, it is able to unlearn feedback actions that the user liked once but now tends to dislike and to learn new feedback actions according to the instructions of the user, thereby increasing its expertise in user-profiled, user-supportive, adaptive user-computer interaction.

The utilized dynamic memory of experiences is based on Schank’s theory of functional organization of human memory of experiences [45]. According to this theory, for a certain event to remind one spontaneously of another, both events must be represented within the same dynamic chunking memory structure, which organizes the experienced events according to their thematic similarities. Both events must be indexed further by a similar explanatory theme that has sufficient salience in the person’s experience to have merited such indexing in the past. Indexing, in fact, defines the scheme for retrieval of events from the memory. The best indexing is the one that will return events most relevant for the event just encountered.

In the case of Gaze-X memory of experiences, each event is one or more micro-events, each of which is an interactive cue (a part of the system’s multimodal input) displayed by the user while interacting with the computer. Micro-events that trigger a specific user-supportive action are grouped within the same dynamic memory chunk. The indexes associated with each chunk comprise individual micro-events and their combinations that are most characteristic for the user-supportive action in question. Finally, micro-events of each dynamic memory chunk are hierarchically ordered according to their typicality: the larger the number of times the user was satisfied when the related user-supportive action was executed as the given micro-event occurred, the higher the hierarchical position of that micro-event within the given chunk. Certain user-supportive action can be preferred by the user for both different software applications that he or she is usually using and different affective states that he or she is displaying. Hence, to optimize the search through the case base, affective states and currently active software are not treated as other micro-events but are used as containers of various memory chunks. More specifically, memory chunks representing user-supportive actions that are triggered when the user is displaying a certain emotion are grouped in a super class representing that emotion. Hence, each chunk may contain a pointer to an emotion-identifying super class. Similarly, each chunk may contain a pointer to an active-software-identifying super class. A schematic representation of Gaze-X case base organization is given in Fig. 4.

To decide which user-supportive action is to be executed (if any) given an input set of interactive cues displayed by the user, the following steps are taken:

- Search the dynamic memory for similar cases based on the input set of observed interactive cues, retrieve them, and trigger the user-supportive action suggested by the retrieved cases.
- If the user is satisfied with the executed action, store the case in the dynamic memory and increase its typicality. If the user is not satisfied with the executed action, adapt the dynamic memory by decreasing the typicality of the just executed action.



**Fig. 4.** Schematic organization of the case base utilized by the Gaze-X. Each user has his or her own personal case base (user profile).

The simplest form of retrieval is to apply the first nearest neighbor algorithm, that is, to match all cases of the case base and return a single best match. This method is usually too slow. A pre-selection of cases is therefore usually made based on the indexing structure of the utilized case base. Our retrieval algorithm employs a pre-selection of cases that is based on the clustered organization of the case base (super classes and memory chunks), the indexing structure of the memory, and the hierarchical organization of cases within the memory chunks according to their typicality.

Gaze-X can run in two modes, an unsupervised and a supervised mode. In the unsupervised mode, the affective state of the user is used to decide on his or her satisfaction with the executed action. If a happy or a neutral expression is displayed, Gaze-X assumes that the user is satisfied. Otherwise, if the user emotes negatively, Gaze-X assumes that he or she is dissatisfied. In the supervised mode, the user explicitly confirms that an action of his preference has been executed. If the user is not satisfied with the executed action, the dynamic memory is adapted. In the unsupervised mode, the typicality of the relevant case is decreased. In the supervised mode, the typicality of the relevant case is decreased and the user may provide further feedback on the action of his/her preference that should be executed instead.

## 5 Interaction Adaptation

The final processing step of Gaze-X is to adapt the user-computer interaction based on current needs and preferences of the user and according to adaptive and user-supportive changes suggested by the system's dynamic memory of experiences. General types of interface adaptations supported by Gaze-X include the following.

*Help provision* – Examples include the following. When the open-file-dialogue is open for a long time, help can be provided by highlighting the file names that were opened in the past in combination with the currently open files. Alternatively, desktop search application can be started. If the user selected a column in a table, and (s)he is

scrutinizing the menu bar for a long time, help can be given by highlighting table-related menu options. Alternatively, help-menu option can be highlighted.

*Addition/removal of automation of tasks* – Examples include automatic opening of all windows of an application that the user opens each time he or she starts up the application in question, automatic error correction, automatic blockage of websites that are similar to already blocked sites, and removal of such an automation if the user disapproves of it.

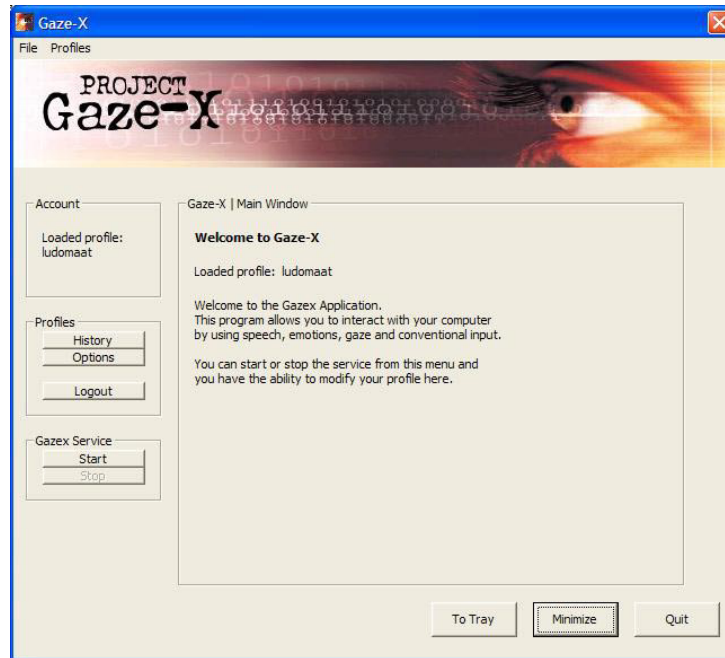
*Changing information presentation* – Automatic selection of most relevant features/options to be displayed given the user's current task, automatic font size increase/decrease according to the user's preferences, usage of eye tracking and speech as an alternative to mouse movements, automatic sound play, etc, are typical examples of this type of interface adaptation.

Gaze-X carries out interface changes in a rather conservative way. More specifically, when operating in the unsupervised mode, it executes adaptive and user-supportive actions one at the time and in a rather slow pace. The underlying philosophy is not to make an ever-changing interface a source of user's frustration in itself. In order to allow the user some time to get accustomed to the idea of a self-adaptive interface and to initialize the system using a (small) set of interface changes that the user considers helpful, Gaze-X is initially set to operate in the supervised mode. As soon as the user considers the system profiled enough, he or she can set the system to operate in the unsupervised mode.

## 6 GUI of the System

Gaze-X has a simple, easy to understand GUI. The goal was to develop a direct-manipulation GUI in which WYSIWYG (what you see is what you get) would be the guiding principle. A simple self-explanatory GUI was developed that is easy to understand and use. As can be seen in Fig. 5, the main window of Gaze-X GUI visualizes who the current user is (i.e., whose user-profiled case base is currently used), allows loading, creation, and adaptation of the profiles, and enables initiation of adaptive interaction mode supported by the system. Gaze-X automatically loads the user's profile for a known user based on the output of the Face Reader system as explained in section 3. However, an option to load a profile manually, using a valid username and password, also exists. Besides, since Gaze-X is developed as Fleeble-based MAS, the agents that constitute the system including the reasoning agent (representing the profile of the user) are mobile and can be moved to another host. Hence, users' profiles can be transmitted as needed to any computer where Gaze-X is installed.

When installing Gaze-X, a directory needs to be found where the main system and the supporting systems and sensors (Fleeble, the FaceReader face and facial expression detector, the iView-X remote eye tracker, a web cam, and a microphone) can be installed with proper access rights. Also a specific version of Java Runtime Environment (version 1.5 or later) needs to be properly installed before Gaze-X can run. To make this process as easy as possible for the user, a setup wizard has been implemented. It executes automatically many of the required steps to set up Gaze-X and leads the user through the rest of the required steps to ensure that the required hardware and software are properly setup.



**Fig. 5.** The main window of the Graphical User Interface of the Gaze-X

Any new way of thinking about programming takes some getting used to, and Gaze-X is no exception to this axiom. To aid users in working out questions that they may have, Gaze-X provides a tutorial. It shows the basic functionalities and usage of Gaze-X in a step by step demonstration. The tutorial can be started by means of ‘To Try’ button illustrated in Fig. 5.

## 7 Usability Study

To make a preliminary assessment of effectiveness, usability, usefulness, affective quality, and ethical issues relevant to Gaze-X, we conducted a small evaluation study with the help of six participants. The participants were 18 to 61 years old, 33% female, 50% Caucasian, 33% Asian, 17% African, 17% expert, 50% intermediate, and 33% novel computer users. We asked them to install Gaze-X first and then to use it as demonstrated by the tutorial integrated into the system. For each session with another user, we used either a Linux or a Windows machine from which the Gaze-X was removed. We did not require the users to use the Gaze-X for a certain period of time, each of them was engaged in exploring Gaze-X for as long as he or she wanted. We also did not require from the participants to work with specific software applications. Note, however, that the software installed on the machines we used for the experiments had an Internet browser, an e-mail handler, a text editor, Adobe Reader, and a number of multimedia handlers such as music and movie players.

We used a custom-made questionnaire to elicit users' attitudes towards the Gaze-X-supported interaction with the computer. The questionnaire includes questions soliciting users' attitudes toward:

- the effectiveness of the HCI design (i.e., whether the interaction with computer is more natural than it is the case with standard HCI designs and whether it is robust enough, [43]),
- the usability of Gaze-X (i.e., whether technological variety and user diversity are supported and whether gaps in the user's computer knowledge play an important role, [48]),
- the usefulness of the system (i.e., whether the utility of system's functionalities and the utility of interface adaptation supported by the system are obvious to the users, whether they will choose to use the system if it was publicly available),
- the ethical issues related to the HCI design (i.e., do users feel uncomfortable under the scrutiny of machines that monitor their affective states, work- and interaction patterns, will users adapt soon to emote just to make the computer do something, [24]),
- the affective quality of the HCI design (i.e., whether the GUI of the Gaze-X is aesthetically qualitative in terms of orderly, clear, and creative design, [56]).

The utilized questionnaire also invites participants' suggestions on how to improve the system in any of the aspects mentioned above. It employs a 5-point Likert scale ranging from strongly disagree (1), via neutral (3), to strongly agree (5). 'I do not know' is also a possible answer. The main points of the obtained survey results are listed in Table 1.

**Table 1.** Users' satisfaction with effectiveness, usability, usefulness, affective quality, and ethical issues relevant to Gaze-X. The percentages in the table show the percentage of *agree* and *strongly agree* answers. EU stands for expert users (total: 17%) and NU stands for intermediate and novel users (total: 83%). × indicates an 'I do not know' answer.

Survey question	EU	NU
The interaction is more natural than in standard HCI	0%	100%
The interaction is robust enough	0%	66%
Gaze-X supports technological variety	×	100%
Gaze-X supports users of different age, skills, culture...	100%	100%
Gaze-X is easy to use even if users lack IT knowledge	100%	100%
Face identification is a useful functionality	100%	100%
Monitoring affective states is a useful functionality	0%	66%
Having mobile user profiles is useful	100%	100%
Having multimodal interaction is useful	0%	66%
Adaptive and user-supportive interface is useful	0%	100%
Gaze-X makes the interaction with computer easier	0%	100%
I do not mind to be observed by the camera	100%	100%
I like the aesthetics of the Gaze-X GUI design	100%	100%
I would use Gaze-X if it was publicly available	0%	66%



All participants seem to agree on the usability, affective quality, and ethical issues relevant to Gaze-X. The perceived usability is directly related to the following properties of Gaze-X: it runs on various platforms like Linux, Windows, and Mac Os X (as a result of being Java-implemented), it accommodates users of different age, gender, culture, computing skills and knowledge, and it bridges the gap between what the user knows about the system and HCI in general and what he or she needs to know (it provides a setup wizard and a tutorial that shows how to use Gaze-X in a step by step demonstration). The perceived affective quality of Gaze-X is directly tied to aesthetics qualities of the Gaze-X GUI: (i) it has an orderly and clear design in accordance to the rules advocated by usability experts, and (ii) it reveals designers' creativity, originality, and the ability to break design conventions manifested by, for example, multimodal design, affective design, tutorial demonstration, etc. These findings are consistent with research finding of Zhang and Li [56], who suggested that the perceived affective quality of a software product is directly tied to aesthetics qualities of that product and who argued that affectively qualitative products have significantly larger chance to be widely accepted technology. The perceived ethical issues relevant to Gaze-X were somewhat surprising as it seems that the users have no problem with being continuously observed by a web cam. Standard concern that the user's behavioral and affective patterns could be used to mind-read and manipulate him or her was not mentioned a single time. However, all participants in our study stressed the importance of privacy and asked about measures taken to prevent hacking and intrusion. Ultimately, if users have control about whether, when, and with whom they will share their private information such as the observations of their behavioral patterns while interacting with the PC (as stored in their personal profile), fears from big-brother-is-watching-you scenarios vanish.

However, as can be seen from Table 1, the question that remains is whether or not an individual user values the capability of Gaze-X to be aware of his or her behavioral and affective patterns. Our study suggests that there might be a very large difference between the expert and non expert computer users when it comes to this question. While all users agreed that having a face-identification-based access to the system is a very useful functionality, expert users found other functionalities of Gaze-X less appealing and in some instances irritating (e.g., popup widows used in supervised operating mode to ask the feedback about the user's preferences). On the other hand, less experienced computer users perceived Gaze-X as very useful since it provides support when needed, in a way needed and preferred, making the interaction with the computer more efficient and easier (e.g., less time was spent on searching various functionalities and undoing erroneous actions). In turn, it seems that Gaze-X is very suitable for novel and intermediate computer users but is much less so for experienced users. As suggested by the expert computer user who participated in our study, much more sophisticated user support should be provided to an experienced user than it is currently the case. For example, support should be provided only when a new application is installed and used for the first couple of times, no support should be provided for long installed software applications. Note, however, that only one expert computer user has participated in the present usability study. A much more elaborate survey must be conducted with experienced users if some firm conclusions are to be made about the ways to make Gaze-X useful and appealing to experienced computer users.

Finally, all participants remarked that the robustness of the system can be improved. Two issues are of importance here. The first is the sensitivity of the Face Reader system, used for face and facial expression analysis, to changes in lighting conditions. The second is the sensitivity of the iView-X remote eye tracker to changes in the user's position. More specifically, the user is expected to remain in front of the computer and is not allowed to shift his or her position in any direction for more than 30 cm. Otherwise, the iView-X system should be recalibrated before it can be used again. These findings indicate that in the future more robust systems for facial expression analysis and eye tracking should be considered for inclusion in the Gaze-X.

From other remarks mentioned by the participants in the present study, arguably the most important one relates to the choice of the affective states to be tagged by Gaze-X. Most of participants said that they may experience confusion, frustration, understanding, tiredness, and satisfaction while interacting with the computer. However, the currently employed Face Reader system recognizes only facial expressions of six basic emotions including disgust, fear, and sadness, for which the participants in our study said that they are not likely to be experienced in a HCI setting like office scenarios. This indicates that in the future we should employ either automatic analyzers of attitudinal and non-basic affective states like attentiveness [12] and fatigue [13], or systems for user-profiled interpretation of facial expressions [35].

## 8 Conclusions

In this paper we proposed one of the first systems for adaptive, affective, multimodal human-computer interaction in standard office scenarios. Our system, Gaze-X, is based on sensing and interpretation of the human part of the computer's context, known as W5+ (who, where, what, when, why, how) and, in turn, is user- and context-profiled. The user of Gaze-X experiences an adaptive interface, which changes as a function of the currently sensed context. This function is represented by the cases of the utilized dynamic case base, having the system's and the user's state as the input (represented in terms of exhibited multimodal interactive actions and cues) and the adaptive and user-supportive changes in the interaction as the output. A usability study conducted in an office scenario with the help of six users indicates that Gaze-X is perceived as effective, easy to use, and useful by novice and less experienced users and as usable and affectively qualitative by all participants in the present study. In turn, Gaze-X seems to be very suitable for novel and intermediate computer users but is much less so for experienced users. As only one experienced user has participated in the present usability study, a much more elaborate survey must be conducted with experienced users if some firm conclusions are to be made about the ways to make Gaze-X useful and appealing to this group of users. Ultimately, as the majority of current software products are still designed for experienced frequent users and designing for a broad audience of unskilled users is still seen as a far greater challenge [48], we are very glad and proud that Gaze-X was perceived as useful and easy to use by novice and less experienced users, who in general still experience computing technology as too difficult to use. Except of a more elaborate usability study with a large number of various users, the focus of our future research will also be aimed at enabling the system to tag attitudinal and non-basic affective states such as confusion, stress, fatigue and satisfaction.

## Acknowledgments

The authors would like to thank Marten den Uyl of Vicar Vision BV for providing the Face Reader system. The work has been conducted at Delft University of Technology. The work of M. Pantic has been supported by the Netherlands Organization for Scientific Research (NWO) Grant EW-639.021.202. The cooperation with Vicar Vision BV has been carried out in the scope of the Dutch BSIK-MultimediaN-N2 project on Interaction.

## References

1. Bartsch-Sporl, B., Lez, M., Hubner, A.: Case-based reasoning – survey and future directions. *Lecture Notes in Artificial Intelligence*, Vol. 1570 (1999) 67-89
2. Bianchi-Berthouze, N., Lisetti, C.L.: Modeling multimodal expression of user's affective subjective experience. *User Modeling and User-Adapted Interaction*, Vol. 12, No. 1 (2002) 49-84
3. Bolt, R.A.: Put-that-there: Voice and gesture at the graphics interface. *Computer Graphics*, Vol. 14, No. 3 (1980) 262-270 (*Proc. ACM SIGGRAPH'80*)
4. Browne, D., Norman, M., Riches, D.: Why Build Adaptive Interfaces? In: Browne, D., Totterdell, P., Norman, M. (eds.): *Adaptive User Interfaces*. Academic Press, London, UK (1990) 15-57
5. Caridakis, G., Malatesta, L., Kessous, L., Amir, N., Raouzaoui, A., Karpouzis, K.: Modeling Naturalistic Affective States via Facial and Vocal Expressions Recognition. *Proc. Int'l Conf. Multimodal Interfaces* (2006) 146-154
6. Cohen, P., Oviatt, S.L.: The role of voice input for human-machine communication. *Proc. National Academy of Sciences*, Vol. 92 (1995) 9921-9927
7. Conati, C.: Probabilistic assessment of user's emotions in educational games. *Applied Artificial Intelligence*, Vol. 16, No. 7-8 (2002) 555-575
8. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 23, No. 6 (2001) 681-685
9. Deng, B.L., Huang, X.: Challenges in adopting speech recognition. *Communications of the ACM*, Vol. 47, No. 1 (2004) 69-75
10. Dey, A.K., Abowd, G.D., Salber, D.: A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *J. Human-Computer Interaction*, Vol. 16, No. 2-4 (2001) 97-166
11. Duric, Z., Gray, W.D., Heishman, R., Li, F., Rosenfeld, A., Schoelles, M.J., Schunn, C., Wechsler, H.: Integrating perceptual and cognitive modeling for adaptive and intelligent human-computer interaction. *Proceedings of the IEEE*, Vol. 90, No. 7 (2002) 1272-1289
12. El Kaliouby, R., Robinson, P.: Real-Time Inference of Complex Mental States from Facial Expressions and Head Gestures. *Proc. Int'l Conf. Computer Vision & Pattern Recognition*, Vol. 3 (2004) 154-
13. Gu, H., Ji, Q.: An automated face reader for fatigue detection. *Proc. Int'l Conf. Face & Gesture Recognition* (2006) 111-116
14. Gunes, H., Piccardi, M.: Affect Recognition from Face and Body: Early Fusion vs. Late Fusion. *Proc. Int'l Conf. Systems, Man and Cybernetics* (2005) 3437- 3443
15. Hauptmann, A.G.: Speech and gestures for graphic image manipulation. *Proc. ACM Int'l Conf. Human Factors in Computing Systems* (1989) 241-245

16. Hoffman, D.L., Novak, T.P., Venkatesh, A.: Has the Internet become indispensable? *Communications of the ACM*, Vol. 47, No. 7 (2004) 37-42
17. Hudlicka, E.: To feel or not to feel: The role of affect in human-computer interaction. *Int'l J. Human-Computer Studies*, Vol. 59, No. 1-2 (2003) 1-32
18. Hudlicka, E., McNeese, M.D.: Assessment of user affective/belief states for interface adaptation. *User Modeling & User-Adapted Interaction*, Vol. 12, No. 1 (2002) 1-47
19. Jaimes, A., Sebe, N.: Multimodal human computer interaction: A survey. *Proc. IEEE ICCV Int'l Workshop on HCI in conjunction with Int'l Conf. Computer Vision (2005)*
20. Kapoor, A., Picard, R.W.: Multimodal affect recognition in learning environments. *Proc. ACM Int'l Conf. Multimedia (2005)* 677-682
21. Keltner, D., Ekman, P.: Facial expression of emotion. In: Lewis, M., Haviland-Jones, J.M. (eds.): *Handbook of Emotions*. The Guilford Press, New York (2000) 236-249
22. van Kuilenburg, H., Wiering, M., den Uyl, M.: A model-based method for automatic facial expression recognition. *Lecture Notes in Artificial Intelligence*, Vol. 3720 (2005) 194-205
23. Lisetti, C.L., Nasoz, F.: MAUI: A multimodal affective user interface. *Proc. Int'l Conf. Multimedia (2002)* 161-170
24. Lisetti, C.L., Schiano, D.J.: Automatic facial expression interpretation: Where human-computer interaction, AI and cognitive science intersect. *Pragmatics and Cognition*, Vol. 8, No. 1 (2000) 185-235
25. Maat, L., Pantic, M.: Gaze-X: Adaptive affective multimodal interface for single-user office scenarios. *Proc. Int'l Conf. Multimodal Interfaces (2006)* 171-178
26. Maglio, P.P., Barrett, R., Campbell, C.S., Selker, T.T.: SUITOR: An attentive information system. *Proc. Int'l Conf. Intelligent User Interfaces (2000)* 169-176
27. Marsic, I., Medl, A., Flanagan, J.: Natural communication with information systems. *Proceedings of the IEEE*, Vol. 88, No. 8 (2000) 1354-1366
28. Nielsen, J.: *Multimedia and hypertext: The Internet and beyond*. Academic Press, Cambridge, USA (1995)
29. Nock, H.J., Iyengar, G., Neti, C.: Multimodal processing by finding common cause. *Communications of the ACM*, Vol. 47, No. 1 (2004) 51-56
30. Oviatt, S.: User-centred modelling and evaluation of multimodal interfaces. *Proceedings of the IEEE*, Vol. 91, No. 9 (2003) 1457-1468
31. Oviatt, S., Darrell, T., Flickner, M.: Multimodal Interfaces that Flex, Adapt, and Persist. *Communications of the ACM*, Vol. 47, No. 1 (2004) 30-33
32. Pantic, M., Grootjans, R.J., Zwitserloot, R.: Teaching Ad-hoc Networks using a Simple Agent Framework. *Proc. Int'l Conf. Information Technology Based Higher Education and Training (2005)* 6-11
33. Pantic, M., Pentland, A., Nijholt, A., Huang, T.S.: Human computing and machine understanding of human behaviour: A survey. *Proc. Int'l Conf. Multimodal Interfaces (2006)* 239-248
34. Pantic, M., Rothkrantz, L.J.M.: Toward an affect-sensitive multimodal human-computer interaction. *Proceedings of the IEEE*, Vol. 91, No. 9 (2003) 1370-1390
35. Pantic, M., Rothkrantz, L.J.M.: Case-based reasoning for user-profiled recognition of emotions from face images. *Proc. Int'l Conf. Multimedia and Expo (2004)* 391-394
36. Pantic, M., Sebe, N., Cohn, J.F., Huang, T.S.: Affective multimodal human-computer interaction. *Proc. ACM Int'l Conf. Multimedia (2005)* 669-676
37. Pentland, A.: Looking at people: Sensing for ubiquitous and wearable computing. *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 22, No. 1 (2000) 107-119
38. Pentland, A.: Perceptual intelligence. *Communications of the ACM*, Vol. 43, No. 3 (2000) 35-44

39. Picard, R.W.: *Affective Computing*. The MIT Press, Cambridge, USA (1997)
40. Porta, M.: Vision-based user interfaces: methods and applications. *Int'l J. Human-Computer Studies*, Vol. 57, No. 1 (2002) 27-73
41. Preece, J., Schneiderman, B.: Survival of the fittest: Evolution of multimedia user interfaces. *ACM Computing Surveys*, Vol. 27, No. 4 (1995) 557-559
42. Prendinger, H., Ishizuka, M.: The empathic companion: A character-based interface that addresses users' affective states. *Applied Artificial Intelligence*, Vol. 19, No. 3-4 (2005) 267-285
43. Reeves, L.M., Lai, J., Larson, J.A., Oviatt, S., Balaji, T.S., Buisine, S., Collings, P., Cohen, P., Kraal, B., Martin, J.C., McTear, M., Raman, T.V., Stanney, K.M., Su, H., Wang, Q.Y.: Guidelines for multimodal user interface design. *Communications of the ACM*, Vol. 47, No. 1 (2004) 57-59
44. Ruttkay, Z., Pelachaud, C. (eds.): *From brows to trust: Evaluating embodied conversational agents*. Kluwer Academic Publishers, Norwell, USA (2004)
45. Schank, R.C.: *Memory based expert systems*. AFOSR.TR. 84-0814, Comp. Science Dept., Yale University (1984)
46. Sharma, R., Yeasin, M., Krahnstoever, N., Rauschert, I., Cai, G., Maceachren, A.M., Sengupta, K.: Speech-gesture driven multimodal interfaces for crisis management. *Proceedings of the IEEE*, Vol. 91, No. 9 (2003) 1327-1354
47. Sharp, H., Rogers, Y., Preece, J.: *Interaction Design*, 2<sup>nd</sup> edition. John Wiley & Sons, Chichester, UK (2007)
48. Shneiderman, B.: Universal usability. *Communications of the ACM*, Vol. 43, No. 5 (2000) 85-91
49. Shneiderman, B.: CUU: Bridging the Digital Divide with Universal Usability. *ACM Interactions*, Vol. 8, No. 2 (2001) 11-15
50. Shoham, Y.: What we talk about when we talk about software agents. *IEEE Intelligent Systems and Their Applications*, Vol. 14, No. 2 (1999) 28-31
51. Tennenhouse, D.: Proactive computing. *Communications of the ACM*, Vol. 43, No. 5 (2000) 43-50
52. Turk, M.: Computer vision in the interface. *Communications of the ACM*, Vol. 47, No. 1 (2004) 61-67
53. Vo, M.T., Waibel, A.: Multimodal human-computer interaction. *Proc. Int'l Symposium on Spoken Dialogue* (1993)
54. Waibel, A., Vo, M.T., Duchnowski, P., Manke, S.: Multimodal Interfaces. *Artificial Intelligence Review*, Vol. 10, No. 3-4 (1995) 299-319
55. Zeng, Z., Hu, Y., Roisman, G.I., Wen, Z., Fu, Y., Huang, T.S.: Audio-Visual Emotion Recognition in Adult Attachment Interview. *Proc. Int'l Conf. Multimodal Interfaces* (2006) 139-145
56. Zhang, P., Li, N.: The importance of affective quality. *Communications of the ACM*, Vol. 48, No. 9 (2005) 105-108