

SSPNet

Social Signal Processing Network

<http://www.sspnet.eu/>

Funded under the 7th FP (Seventh Framework Programme)

Theme ICT-2007.2.2

[Network of Excellence]

D2.3: Second draft of the agenda on SSP research

Due date: 31/01/2011

Submission date: 31/01/2011

Project start date: 01/02/2009


Duration: 60 months

WP Manager: Maja Pantic

Revision: 1

Author(s): Maja Pantic, Roderick Cowie, Francesca DErrico, Dirk Heylen, Marc Mehu, Catherine Pelachaud, Isabella Poggi, Marc Schroeder, and Alessandro Vinciarelli

Project funded by the European Commission in the 7th Framework Programme (2009-2013)		
Dissemination Level		
PU	Public	No
RE	Restricted to a group specified by the consortium (includes Commission Services)	Yes
CO	Confidential, only for members of the consortium (includes Commission Services)	No



D2.3: Second draft of the agenda on SSP research

Abstract:

The exploration of how we react to the world and interact with it and each other remains one of the greatest scientific challenges. Latest research trends in cognitive sciences argue that our common view of intelligence is too narrow, ignoring a crucial range of abilities that matter immensely for how people do in life. This range of abilities is called social intelligence and includes the ability to express and recognise social signals produced during social interactions like agreement, politeness, empathy, friendliness, conflict, etc., coupled with the ability to manage them in order to get along well with others while winning their cooperation. Social Signal Processing (SSP) is the new research domain that aims at understanding and modelling social interactions (human-science goals), and at providing computers with similar abilities in human-computer interaction scenarios (technological goals). SSP is in its infancy, and the journey towards artificial social intelligence and socially-aware computing is still long. This research agenda is a twofold, a discussion about how the field is understood by people who are currently active in it and a discussion about issues that the researchers in this formative field face.

Contents

1	Introduction	1
2	Social Signals: Terminology, Definition, and Cognitive Modelling	3
2.1	Terminology	3
2.2	Working definition of Social Signals	7
3	Machine Analysis of Social Signals	10
4	Machine Synthesis of Social Signals	13
5	Summary and Additional Issues	16
6	Conclusion	18
7	References	18

1 Introduction

The exploration of how human beings react to the world and interact with it and each other remains one of the greatest scientific challenges. Perceiving, learning, and adapting to the world are commonly labelled as intelligent behaviour. But what does it mean being intelligent? Is IQ a good measure of human intelligence and the best predictor of somebody's success in life? There is now a growing research in cognitive sciences, which argues that our common view of intelligence is too narrow, ignoring a crucial range of abilities that matter immensely for how people do in life. This range of abilities is called social intelligence (Thorndike, 1920; Ambady & Rosenthal, 1992; Albrecht, 2005; Buchanan, 2007) and includes the ability to express and recognise social signals like turn taking, agreement, politeness, empathy, friendliness, conflict, etc., coupled with the ability to manage them in order to get along well with others while winning their cooperation. There is no common definition for the concept of social signal (as explained in section 2), and the definition that we adopt in this document refers to social signals as to signals produced during social interactions, that either play a part in the information and adjustment of relations and interactions between agents (human and artificial), or provide information about the agents. Social signals are manifested through a multiplicity of non-verbal behavioural cues including facial expressions, body postures and gestures, vocal outbursts like laughter, etc. (Fig. 1), that can be automatically analysed by technologies of signal processing (as discussed in section 3), or automatically generated by technologies of signal synthesis (as talked about in section 4).

When it comes to computers, however, they are socially ignorant (Pentland, 2005). Current computing devices do not account for the fact that human-human communication is always socially situated and that discussions are not just facts but part of a larger social interplay. However, not all computers will need social intelligence and none will need all of the related skills humans have. The current-state-of-the-art categorical computing works well and will always work well for context-independent tasks like making plane reservations and buying and selling stocks. However, this kind of computing is utterly inappropriate for virtual reality applications as well as for interacting with each of the (possibly hundreds) computer systems diffused throughout future smart environments (predicted as the future of computing by several visionaries such as Mark Weiser, 1991) and aimed at improving the quality of life by anticipating the users needs. Computer systems and devices capable of sensing agreement, inattention, or dispute, and capable of adapting and responding in real-time to these social signals in a polite, non-intrusive, or persuasive manner, are likely to be perceived as more natural, efficacious, and trustworthy. For example, in education, pupils social signals inform the teacher of the need to adjust the instructional message. Successful human teachers acknowledge this and work with it; digital conversational embodied agents must begin to do the same by employing tools that can accurately sense and interpret social signals and social context of the pupil, learn successful context-dependent social behaviour, and use a proper socially-adept presentation language (e.g., see Pelachaud et al, 2002) to drive the animation of the agent. The research area of machine analysis and employment of human social signals to build more natural, flexible computing technology goes by the general name of Socially-Aware Computing as introduced by Pentland (2004, 2005).

Although the importance of social signals in everyday life situations is evident, and in spite of recent advances in machine analysis and synthesis of relevant behavioural cues like gaze exchange, blinks, smiles, head nods, crossed arms, laughter, expressive prosody, and similar (Pantic et al., 2007, 2008; Zeng et al., 2009; Niewiadomski 2008; Schröder, 2009), the research



Fig. 1: Manifestations of social signals include a variety of non-verbal behavioural cues including facial expressions, body postures/ gestures, vocal outbursts like

efforts in machine analysis and synthesis of human social signals like attention, empathy, politeness, irting, (dis)agreement, etc., are still tentative and pioneering efforts. Nonetheless, the importance of studying social interactions and developing automated systems of social signals analysis from audiovisual recordings is indisputable. It will result in valuable multimodal tools that could revolutionise basic research in cognitive and social sciences by raising the quality and shortening the time to conduct research that is now lengthy, laborious, and often imprecise. The first results in the field attest that social interactions and behaviours, although complex and rooted in the deepest aspects of human psychology, can be analyzed automatically with the help of computers (for extensive overview of the past research in the field of automatic analysis of social signals, see Vinciarelli et al. 2009). In fact, the pioneering contributions in SSP (Pentland, 2004; Eagle & Pentland, 2006; Hung et al., 2007) have shown that social signals, typically described as so elusive and subtle that only trained psychologists can recognize them (Gladwell, 2005), are actually evident and detectable enough to be captured through sensors like microphones and cameras, and interpreted through analysis techniques like machine learning and statistics. At the same time, and as outlined above, tools for social signal synthesis in human-computer interactions (HCI) form a large step ahead in realising naturalistic, socially-aware computing and interfaces, built for humans, based on models of human behaviour. For example, combining synthetic speech with laughter influences the perception of social bonds (Trouvain & Schröder, 2004). Similarly, facial expressions influence a human user's evaluation of an Embodied Conversational Agent (Ruttkey & Pelachaud, 2004). Contingency of signals has a key role in creating rapport between human user and virtual agent (Gratch et al., 2007). Politeness cues (Wang et al., 2005) and empathic expressions (Niewiadomski et al., 2008) are perceived as more appropriate in interactive scenarios.

Social Signal Processing (SSP) (Pentland, 2005, 2007; Vinciarelli et al. 2008, 2009) is the

new research and technological domain that aims at providing computers with the ability to sense and understand human social signals. SSP is in its initial phase and the first step is to define the field and discuss issues facing the researchers in the field. This article attempts to achieve this. In section 2, an overview of the relevant terminology defined by the related human-science fields is provided. Next, in the absence of a uniquely accepted definition of social signals, a working definition of social signals is introduced. In section 3 and section 4, challenging issues facing researchers in automatic social signal analysis and synthesis are summarised. Section 5 summarizes the key goals of the SSP research overall, lists a number of issues that are of importance for the field but are still debated, and discusses the relevant ethical issues. Section 6 concludes the paper.

2 Social Signals: Terminology, Definition, and Cognitive Modelling

In order to anchor their discipline in the rich conceptual background developed in the behavioural sciences, SSP researchers are faced with the difficult task of defining a theoretical framework within which they will research the phenomena social signals they are eager to automatically detect, interpret, and synthesize. The major issue here is the diversity of conceptual ideas proposed about social signals and behaviour. Disciplines that dealt with the study of human psychological phenomena (mental states and behaviour) developed a myriad of ideas, definitions, and methods for the study of the same subject, human communication. In itself, this may be seen as a strength more than a weakness because having multiple approaches increases the potential for a good understanding of the complexities of human behaviour. However, this diversity may become a problem when people of different traditions come to work together on interdisciplinary research topic (such as SSP). The increasing specialization that characterizes most scientific disciplines can be a barrier to inter-disciplinarity, for it can hinder communication between scholars. For communication to be successful, one has to use terms that will be understood by various researchers working in the field. In an attempt to achieve this, we describe here the different approaches adopted by the human sciences to study social signals. Section 2.1 presents a (non-exhaustive) glossary of concepts generated by different fields (ethology, social psychology, linguistics, semiotics,...) to study communication. We present commonalities and differences between these approaches, so that scholars who are not familiar with the different disciplines can have a clearer idea of what the different positions are. The goal of this exercise is not to create a common definition for the concept of social signal because it would deny the specificities of each field and would constitute, for some, a loss of conceptual clarity. As SSP is a multi-disciplinary venture, our aim is to avoid the creation of a monolithic view that is unlikely to be adopted by the scientific community at large, or that may block the development of new ideas or research projects. Instead, our goal is to expose the diversity and be aware of it. However, as a definition of the studied phenomena is needed, and in the absence of a uniquely accepted definition of social signals, Section 2.2 introduces a working definition of social signals.

2.1 Terminology

Table 1 presents a (non-exhaustive) glossary of SSP-relevant concepts generated by different fields (ethology, social psychology, linguistics, semiotics, etc.) to study communication.

Table 1: Definitions of important concepts for social signal processing.

Ethology	
signal	An act or structure that affects the behaviour (or internal state) of another organism, which evolved because of that effect, and which is effective because the receiver's response has also evolved (Maynard Smith & Harper, 2003). A signal may (Maynard Smith & Harper, 2003) or may not (Rendall, Owren, & Ryan, 2009) convey reliable information.
cue	A feature of the world, animate or inanimate, that can be used by individuals as a guide to future action (Hasson, 1994).
display	Behaviour pattern that has been modified in the course of evolution to convey information (Beer, 1977). Displays are usually constituted of several components, like cues and signals. handicap A signal whose reliability is ensured because its cost is greater than those required to efficiently convey the information (Zahavi, 1975). The signal may be costly to produce, or have costly consequences (Vehrencamp, 2000).
index	A signal whose intensity is causally related with the information that is being signalled and that cannot be faked (Maynard Smith & Harper, 1995). Indices are equivalent to performance based signals (Enquist, 1985).
minimal-cost signal	A signal whose reliability does not depend on its cost (different from a handicap) and which can be made by most members of a population (different from an index) (Maynard Smith & Harper, 2003).
icon	A signal which form is similar to its meaning.
symbol	A signal whose form is unrelated to its meaning, e.g. conventional signal (Guilford & Dawkins, 1995).
Psychology	
cue	stimulus which serves as a sign or signal of something else, the connection having previously been learned (Woodworth, 1961).
indicator	No clear definition for nonverbal indicator, seems to be used in a loose fashion to reflect a connection between nonverbal behaviour and some underlying dimension.
signal	No precise definition of signal in social psychology, though some authors seem to imply that signals are intentionally communicative (Ekman, 1997). The category seems to include all nonverbal behaviours or morphological structure that convey information to a receiver (Poggi & D'Errico, 2010).
social signal	Communicative or informative signal that, either directly or indirectly, conveys information about social actions, social interactions, social emotions, social attitudes and social relationships (Poggi & D'Errico, 2010).
sign	Refers to an act that is informative but that was not necessarily produced to communicate information (Ekman, 1997).
emblem	Nonverbal act which has a direct verbal translation that is well-known by all members of a group, class, or culture (Efron, 1941; Ekman & Friesen, 1969).

illustrator	Movement directly tied to speech that illustrates what is said verbally (Efron, 1941; Ekman & Friesen, 1969).
regulator	Act that maintains and regulates the conversation between two or more individuals (Ekman & Friesen, 1969).
manipulator	Act that represents adaptive efforts to satisfy bodily needs, actions, to manage emotions, to develop interpersonal contacts, or to learn instrumental activities (see also adaptor in Ekman & Friesen, 1969).
emotional expression	Nonverbal act that is specific to a particular emotion (Tomkins, 1961; Ekman, 1992), or to an underlying emotional dimension (Scherer, 1992).
distal cues	Externalization of stable traits or transient states, can be motor expression or physical appearance (Brunswik, 1956; Scherer, 1978)
proximal percept	Mental representation resulting from the perceptual process of distal cues (Brunswik, 1956; Scherer, 1978)

Processes involved in the production and perception of social signals

code	Principle of correspondence between the act and its meaning (Ekman & Friesen, 1969). The code can be intrinsic, extrinsic, and iconic.
encoding	The process, taking place in the signaller, of relating the distal cue and its meaning. Transfer of information in one domain (e.g. thoughts, stances) to another domain (muscular contraction, blood concentration, ...).
decoding	The process taking place in the perceiver of relating the proximal percept to a semantic category or some other form of representation.

Linguistics and semiotics

turn taking	The order in which the participants in a conversation speak one after the other. The fulfillment or violation of turn-taking rules in a conversation provides cues about its cooperative or competitive structure (Sacks et al., 1974; Duncan, 1972).
backchannel	Feedback and comments provided by listeners during face-to-face conversation, through short verbalisations and nonverbal signals, showing how they are engaged in the speakers dialogue (HUMAINE glossary).

Miscellaneous

context	All the cues present in the physical and social environment of a perceiver as well as perceiver's characteristics that surrounds the signal.
information (mathematics)	Any physical property of the world that reduces uncertainty in the individual that perceives it (Shannon & Weaver, 1949).
meaning	The meaning of something is what it expresses or represents (Cambridge Advanced learner's dictionary)
ground truth	A term, with origins in cartography and aerial imaging, used to describe data that can be taken as definitive, and against which systems can be measured. Its application to emotion is controversial, since it is highly debatable whether emotions as they normally occur are things about which we can have definitive knowledge (HUMAINE glossary).

The main stream of research in animal and human communication acknowledges that signals convey information and/or meaning to a receiver. Although this could be considered as a commonality between all approaches, a critical analysis of research findings and theoretical developments in ethology suggest that the principles that are applied to the study of human language should not necessarily apply to the study of animal signals (Rendall, Owren, & Ryan, 2009) or to some aspects of human nonverbal behaviour (Owren & Bachorowski, 2003). In other words, the strong semantic component of human language may not necessarily be shared by other channels of communication. For this reason, the ethological definition of signals includes the possibility that they do not carry information.

Different disciplines adopt different ways of defining a signal. For example, ethologists define signals by their properties or nature, and by their function of influencing a perceiver's behaviour or internal state (e.g. Maynard Smith & Harper, 2003). On the other hand, Poggi & D'Errico (2010) define a social signal by its content (signals that have a social content, like a social attitude, a social emotion, etc.). In social psychology, scholars tend to use the term indicator, sign, signal, and display interchangeably (e.g. Burgoon & LePoire, 1999; Keltner, 1995), without specifying what they mean by the terms indicator, sign, signal, or display. We assume that this lack of specificity implies that these authors endorse a general dictionary definition of the word, their goal being to study the eliciting circumstances and information content of particular behaviour patterns rather than to develop different concepts for nonverbal communication. Other authors, however, decided to use different terms than signals to describe specific categories of nonverbal behaviour (Ekman & Friesen, 1967).

The variety in definitions may be partly explained by the use of different methodologies and the different empirical questions that have driven research activities in different fields. For example, ethology has always focused on the adaptive significance of behavioural patterns for the organisms displaying them and the selective pressures responsible for the evolution of signals (Eibl-Eibesfeldt, 1989; Hinde, 1975). Psychological science, however, has always placed a greater interest in discovering the significance or meaning of a particular behaviour in the mind of perceivers (for a critic, see Owren & Bachorowski, 2003). Finally, linguistic has been mostly preoccupied by the role of signals in the regulation of discourse and social interactions among members of conversational groups (Duncan, 1972; Hadar, 1985). The diversity in research methods and theoretical interests led scholars to use different terms to describe the same thing, or the same term to describe different ideas. By no means should this signify that one approach has more authority than the other, or that a research question is more relevant than another. The only drawback is that this state of affair may create confusion in scholars who are interested in social signals but are not familiar with the human and behavioural sciences. We hope that the overview provided here is helpful in that direction.

Although we can see that research domains mostly differ in the detailed elaborations they made with regards to the nature of signals, their function, and their informative value, a few features and principles used to describe signals are shared among the different fields. First, some acts are considered functionally or intentionally communicative (e.g. signals, emblems, communicative signals); whereas others are simply considered as informative (cues, signs, informative signals), suggesting that information can be derived from them although they have not evolved, or are not intended, for communication¹. Most theories also recognize the

¹Note that within ethology, there's a division between those who consider signals as information carrier (e.g. Bradbury & Vehrencamp, 2000; Zahavi, 1975) and those who see signals as mostly influential (Dawkins & Krebs, 1984, Rendall et al., 2009).

existence of signals which meaning follows social conventions: symbols, conventional signals, and emblems. The iconic act also seems to meet agreement in the different fields, as it is defined by everyone as an act which meaning is defined by its form. Finally, the importance of multimodality is also recognized by all fields of research (Allwood, 2001; Bänziger & Scherer, 2007; Partan, 1999). Commonalities of this sort make collaborations between disciplines possible and create bridges that are necessary for inter-disciplinary research.

2.2 Working definition of Social Signals

As a definition of the studied phenomena is needed, and in the absence of a uniquely accepted definition of social signals, we provide here a working definition of what “social signals” are.

Social signal : Let us first define what a “signal” is. A signal is a perceivable stimulus PS a behaviour, a morphological trait, a chemical trace produced by an Emitter E. The Emitter E can be an individual or a group of people, a virtual character, an animal, or a machine. The signal is received by some Receiver R, who may interpret the signal and draw some information I from it (the signals meaning), whether E really intended to convey I or not. Taking this into account, we may define a social signal as follows.

A social signal is a signal that provides information about social facts, i.e., about social interactions, social emotions, social attitudes, or social relations.

We can further distinguish between informative and communicative signals. A communicative signal is a signal that the Emitter produces in order to convey a particular meaning (see the Speech Acts perspective: Grice, 1957; Cohen & Levesque, 1990; Poggi, 2007), while an informative signal is a signal from which the Receiver draws some meaning even if the Emitter did not intend to convey it (see the Semiotic perspective: Peirce, 1931). Let us explain these notions by means of an example.

Suppose that during a lunch break there is a group of children talking in a circle, where one of them is slightly outside of the circle. A prediction can be made that the child outside of the circle is at risk of being bullied or being dropped out from the group. The spatial positioning of children is a social signal that conveys information about the social relation between the child in question and the other children, without any of the children being aware that they convey this information. This signal is not a communicative signal, but an informative signal.

Furthermore, a distinction can be made between direct and indirect signals. Since social signals are produced (and understood) in context, information coming from the context may combine with the literal meaning of the signals (for a study on the literal meaning of behavioural signals, see Poggi, 2007) to introduce, through inferential processes, further indirect meanings of the displayed signals that differ from context to context. Let us explain this by means of an example.

Suppose that two people, A and B, sit together and both appear to be sad. This is not a social signal, just the fact that both people express the same emotion. However, if by showing sadness A wants to tighten her bond with B, then her display of sadness is an information signal representing an indirect social signal of her bond to B.



Fig. 2: “Social Facts”: social interactions, social emotions (e.g., compassion and empathy), social attitudes and social evaluation (e.g., approval and disapproval), and social

Taking these notions into account, we can redefine the definition of social signals as follows.

A Social signal is a communicative or informative signal that, either directly or indirectly, provides information about social facts, that is, about social interactions, social emotions, social attitudes, or social relations.

Hence, we define social signals as communicative and informative signals that concern social facts, namely, social interaction, social emotions, social attitudes and social relations. However, there is no strict definition of these notions. In what follows, we propose tentative definitions of these notions.

Social interactions : Social interaction is a specific event in which an agent A performs some social actions directed at another agent that is actually or virtually present. Social interactions may be mediated by communicative and informative signals. Typical communicative signals in social interactions are backchannel signals such as head nods, which inform the recipient that her interaction partner is following and understanding her (Heylen, 2005; Fig. 2).

Social emotions : A clear distinction can be made between individual and social emotions. The latter can be defined as an emotion that an Agent A feels toward and Agent B. Happiness and sadness are typical examples of individual emotions – we can be happy or sad on our own; our feelings are not directed to any other person. On the other hand, admiration, envy, and

compassion are typical examples of social emotions – we have these feelings toward another person. Signals revealing individual emotions of a person and those communicating social emotions both include facial expressions, vocal intonations and outbursts, body gestures and postures, etc. However, if a behavioural cue like a frown is displayed as a consequence of an individual emotion, then this cue is a behavioural signal but not a social signal. It is a social signal only if it displayed in order to communicate a social emotion. In addition, a signal of empathy (e.g., patting a companion on the shoulder to convey that we share his sadness, Fig. 2) is a social signal. A typical signal associated with empathy is mimicry. However, mimicry is not always unconscious, which is typical for sincere empathy, but can be deliberately displayed in order to gain acceptance or approval. In the latter case, mimicry does not convey empathy. Studying the role and the effects of both deliberate and unconscious mimicry is a challenge facing the researchers in the field.

Social evaluation : Social evaluation of a person relates to assessing whether and how much the characteristics of this person comply with our standards of beauty, intelligence, strength, justice, altruism, etc. We judge other people because based on our evaluation we decide whether to engage in a social interaction with them, what types of social actions to perform, and what relations to establish with them. Typical signals shown in social evaluation are approval and disapproval, at least when it comes to the evaluator (e.g., Fig 2). As far as the evaluated person is concerned, typical signals involve those conveying desired characteristics such as pride, self-confidence, mental strength, etc., which include raised chin, erected posture, easy and relaxed movements, etc.

Social attitudes : The notion of attitude has been widely investigated in Social Psychology. Social attitude can be defined as the tendency of a person to behave in a certain way toward another person or a group of people. Social attitudes include cognitive elements like beliefs, evaluations, opinions, and social emotions. All these elements determine (and are determined by) preferences and intentions (Fishbein & Ajzen, 1975).

Agreement and disagreement can be seen as being related to social attitudes. If two persons agree then this means that they have similar opinions, which usually entails an alliance, a commitment to cooperation, and a mutually positive attitude. In contrast, if two persons disagree, this typically implies conflict, non-cooperation, and mutually negative attitude. Typical signals of agreement and disagreement are head nods and head shakes, smile, lip wipe, crossed arms, wagging a hand, etc. (Bousmalis et al., 2009, 2010).

Persuasion is also closely linked to social attitudes; it is a kind of social influence aimed at changing other peoples attitudes towards a certain issue, by changing their opinions and evaluations about the target issue, and gaining agreement for the view he or she defends. Typical signals used in persuasion are persuasive words, gestures, gaze patterns, postures, as well as appropriate self-presentation aimed at eliciting the desired social evaluations.

Social relations : A social relation is a relation between two (or more) persons in which these persons have common or related goals, that is, in which the pursuit, achievement, or thwarting of a goal of one of these persons determines or is determined in some way by the pursuit, achievement, or thwarting of a goal of the other involved person (Kelley et al, 1983). Hence, not every relation is a social relation. Two persons sitting next to each other in a bus have a physical proximity relation, but this is not a social relation, although one can rise from

it (Festinger et al., 1950; Byrne, 1971). We can have many different kinds of social relations with other people: dependency, competition, cooperation, love, exploitation, etc.

Exchange Theory (Homans, 1961; Kelley & Thibaut, 1978) has attempted to describe all relations including love and friendship in terms of costs and benefits. According to this theory, a person stays in a relation until it is a satisfying relation. The factors influencing this satisfaction are: rewards (material and symbolic rewards computed in terms of costs and benefits), evaluation of possible alternatives (that affects commitment), and investment (of time, effort and resources). Several critics have challenged this view as being too close to classical utilitarianism, which does not account for the difference between material and symbolic rewards and rules out altruism (Hinde, 1979). Different typologies of relations have been proposed in terms of criteria like public vs. private, cooperation vs. competition, presence vs. absence of sexual relations, social-emotional support oriented vs. task oriented (e.g., Berscheid & Reiss, 1997). However, defining the notion of social relation and drawing a typology of social relations, such that they are conceptually sound while being useful for analysis and understanding of social signals, is yet to be attained. Also, assessing how social relations, social attitudes, social emotion, and social interaction overall, affect subsequent social relations is another challenge facing the researchers in the field.

Social relations can be established not only with a single person, but with a group. Within group relations, particular challenges concern the definition and description of mechanisms of power, dominance, and leverage (Castelfranchi, 1990; Lewis, 2002). This relates to: (i) the allocation, change, and enhancement of power relations (e.g., through alliance, influence, and reputation), (ii) the interaction between gender and power relations, and (iii) the nature of leadership and the role of charisma in it. Clearly, all these issues are context and culture dependent. Typical signals revealing social relations include the manner of greeting (saying hello first signals the wish for a positive social relation, saluting signals belonging to a specific group like the army), the manner of conversing (e.g., using the word professor signals submission), mirroring (signalling wish to have a positive social relation, or displaying typical groups behaviour), spatial positioning (e.g., making a circle around a certain person distinguishes that person as the leader, touching another person indicates either affective relation or dominance, e.g., Fig. 2), etc. For group relationships, the manner of dressing, cutting ones hair, and mirroring, are the typical signals revealing whether a person belongs to a specific group or not. The emblems on the cloths, how elaborate is a hair dress or a crown, and the spatial arrangement of the members of the group are the typical signals revealing the status and the rank (i.e., power relations) of different members of the group.

3 Machine Analysis of Social Signals

Nonverbal behaviours like social signals cannot be read like words in a book (Knapp & Hall, 1972; Richmond & McCroskey, 1995); they are not always unequivocally associated to a specific meaning (although in general they are; Poggi, 2007) and their appearance can depend on factors that have nothing to do with social behaviour. For example, some postures correspond to certain social attitudes, but sometimes they are simply comfortable (Scheen, 1964). Similarly, physical distances typically account for social distances, but sometimes they are simply the effect of physical constraints (Hall, 1959). Moreover, as mentioned above, the same signal can correspond to different social behaviour interpretations depending on context and culture (Triandis, 1994), although many advocate that social signals are nat-

ural rather than cultural (Segerstrale & Molnar, 1997). In other words, social signals are intrinsically ambiguous, high-level semantic events, which typically include interactions with the environment and causal relationships. An important distinction between the analysis of high-level semantic events and the analysis of low-level semantic events like the occurrence of an individual behavioural cue like the blink, is the degree to which the context, different modalities, and time, must be explicitly represented and manipulated, ranging from simple spatial reasoning to context-constrained reasoning about multimodal events shown in temporal intervals. However, despite a significant progress in automatic recognition of audiovisual behavioural cues underlying the manifestation of various social signals (e.g., see Fig. 3), most of the present approaches to machine analysis of human behaviour are neither multimodal, nor context-sensitive, nor suitable for handling longer time scales (Pantic et al. 2007, 2008, 2010; Zeng et al. 2009). In turn, most of the social signal recognition methods reported so far are single-modal, context-insensitive and unable to handle long-time recordings of the target phenomena (Vinciarelli et al. 2009). This includes methods for recognition of social relations like social roles and social dominance (exemplified by the works reported by Salamin et al. 2009, Raducanu & Gatica-Perez 2010, and Hung et al. 2010), methods for recognition of social attitudes like agreement/ disagreement and politeness exemplified by the work reported by (Vinciarelli 2009, Bousmalis et al., 2010, and Brunet et al. 2010), as well as other methods for social signal recognition (for a comprehensive survey of the efforts in the field, please see the survey papers by Vinciarelli et al. 2009, 2010). Hence, the focus of future research efforts in the field should be primarily on tackling the problem of context-constrained analysis of multimodal behavioural signals shown in temporal intervals. As suggested by Pantic et al. (2007, 2008), this problem should be treated as one complex problem rather than a number of detached problems in human sensing, context sensing, and human behaviour understanding. More specifically, there are a number of scientific and technical challenges that we consider essential for advancing the state of the art in machine analysis of human behaviour like social signals.

Modalities: Which behavioural channels such as the face, the body and the tone of the voice, are minimally needed for realisation of robust and accurate human behaviour analysis? Does this hold independently of the target communicative intention (e.g., social interactions/ emotions/ relations) to be recognised? No comprehensive study on the topic is available yet. What we know for sure, however, is that integration of multiple modalities (at least facial and vocal) produces superior results in human behaviour analysis when compared to single-modal approaches. Numerous studies have theoretically and empirically demonstrated this (e.g., see the literature overview by Russell et al, 2003, for such studies in psychology, and the literature overview by Zeng et al., 2009, for such studies in automatic analysis of human behaviour). It is therefore not surprising that some of the most successful works in SSP so far use features extracted from multiple modalities (for an extensive overview of the past works, see Vinciarelli et al. 2009). However, other issues listed above are yet to be investigated. Also, note that some studies in the field indicate that the relative contributions of different modalities and the related behavioural cues to judgment of displayed behaviour depend on the targeted behavioural category and the context in which the behaviour occurs (Russell et al, 2003).



Fig. 3: Behavioural cues typical of disagreement (clockwise from top left): forefinger raise, forefinger wag, hand wag, and hands scissor (Bousmalis et al., 2009). The cues can be recognize with state-of-the-art human-action-recognition techniques like that proposed by Oikonomopoulos et al (2010).

Fusion: How to model temporal multimodal fusion which will take into account temporal correlations within and between different modalities? What is the optimal level of integrating these different streams? Does this depend upon the time scale at which the fusion is achieved? What is the optimal function for the integration? More specifically, most of the present audiovisual and multimodal systems in the field perform decision-level data fusion (i.e., classifier fusion) in which the input coming from each modality is modelled independently and these single-modal recognition results are combined at the end. Since humans display audio and visual expressions in a complementary and redundant manner, the assumption of conditional independence between audio and visual data streams in decision-level fusion is incorrect and results in the loss of information of mutual correlation between the two modalities. To address this problem, a number of model-level fusion methods were proposed that make use of the correlation between audio and visual data streams, and relax the requirement of synchronization of these streams (Zeng et al., 2009). However, how to model multimodal fusion on multiple time scales and how to model temporal correlations within and between different modalities is yet to be explored.

Fusion and context: Do context-dependent fusion of modalities and discordance handling, which are typical for fusion of sensory neurons in humans, pertain in machine context sensing? Note that context-dependent fusion and discordance handling were never attempted within an automated system. Also note that while W4 (where, what, when, who) is dealing only with the apparent perceptual aspect of the context in which the observed human behaviour

is shown, human behaviour understanding is about W5+ (where, what, when, who, why, how), where the why and how are directly related to recognising communicative intention including social signals, affect, and cognitive states of the observed person. Hence, SSP is about W5+. However, since the problem of context-sensing is extremely difficult to solve, especially for a general case (i.e., general-purpose W4 technology does not exist yet; Pantic et al., 2008, 2009), answering the why and how questions in a W4-context-sensitive manner when analysing human behaviour is virtually unexplored area of research. Having said that, it is not surprising that context-dependent fusion is truly a blue-sky research topic.

Technical aspects: Most methods for human sensing, context sensing, and human behaviour understanding work only in (often highly) constrained environments. Noise, fast movements, changes in illumination, etc., cause them to fail. Also, many of the methods in the field do not perform fast enough to support interactivity. Researchers usually choose for more sophisticated processing rather than for real-time processing. The aim of future efforts in the field should be the realisation of more robust, real-time systems, if they are to be deployed in anticipatory interfaces and social-computing technology defused throughout smart environments of the future.

4 Machine Synthesis of Social Signals

Automatic synthesis of social signals targets a human observer's or listener's perception of socially relevant information. While it may be true that much of social behaviour goes unnoticed (Gladwell, 2005), it appears that social signals still have an effect in terms of unconscious perception (Hyman, 1998) without being able to say exactly why, we either consider a person trustworthy, competent, polite, etc., or not. In automatic behaviour synthesis, the aim is thus to create this perception by timely generating suitable signals and behaviours in a synthetic voice, facial expressions and gestures of an Embodied Conversational Agent (ECA). Several efforts aimed at automatic generation of social signals have been reported recently. Those works include generation of social actions (e.g., turn taking generation by Bonaiuto & Thorisson 2008, and backchanneling generation by Kopp et al. 2007) and generation of social attitudes (e.g., politeness generation by de Jong et al. 2008 and Ochs et al., 2010, Fig. 4). For a comprehensive overview of works on social signal generation on virtual agents, see Vinciarelli et al. 2010.

However, how to generate suitable behavioural signals is by no means clear, mainly due to the following two reasons. Firstly, too little is known about the types of socially relevant information conveyed in everyday human-to-human interactions, as well as about the signals and behaviours that humans naturally use to convey them. A first step in this direction would be to acknowledge the complexity of the phenomena, as has been done for emotion-related communication (Douglas-Cowie et al., 2005). Then, different contexts and effects could be studied based on suitable data, and the findings could be described in terms of explicit markup language (Martin et al., 2005) or in terms of statistical, data-driven models (Hofer et al., 2008). Secondly, it is not self-evident that synthetic agents should behave in the same way as humans do, or that they should exhibit faithful copy of human social behaviours. On the contrary, evidence from the cartoon industry (Bates, 1994) suggests that, in order to be believable, cartoon characters need to show strongly exaggerated behaviour. This suggests further that a trade-off between the degree of naturalness and the type of (exaggerated)

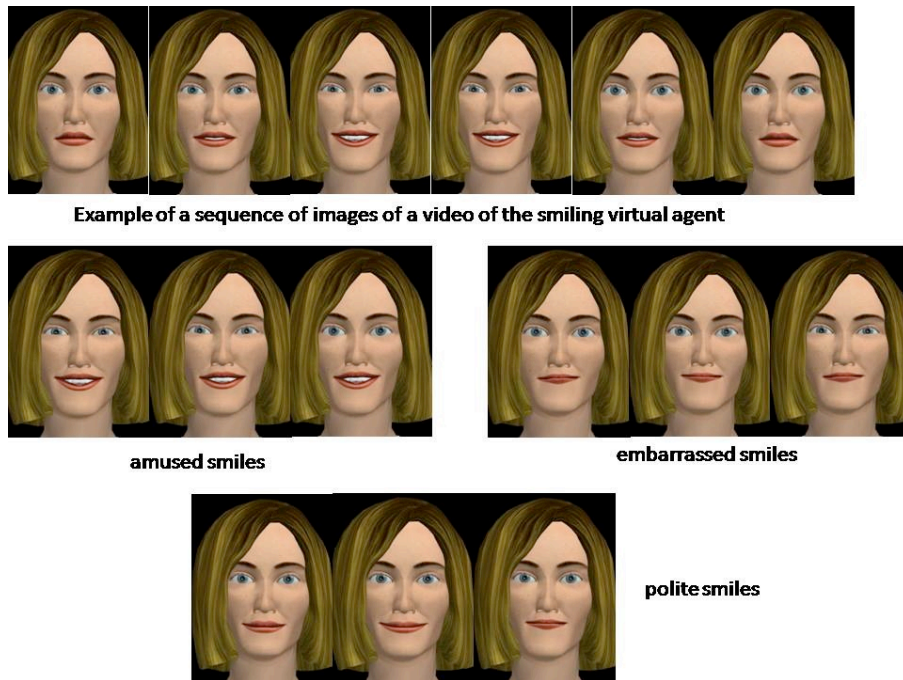


Fig. 4: A variety of smiles of a virtual agent (Ochs et al., 2010)

gestural and vocal expression may be necessary for modelling a believable ECAs behaviour.

In addition, a number of aspects of social signals are particularly relevant and challenging when it comes to synthesis of human-like behaviour.

Continuity: Unlike traditional dialogue systems, in which verbal and non-verbal behaviour is exhibited only when the system has the turn, socially-aware systems need to be continuous in terms of non-verbal behaviour to be exhibited. In any socially relevant situation, social signals are continuously displayed, and lack of such displays in an automatic conversational system is interpreted as social ignorance (Vinciarelli et al., 2009). The omission of social signals, typical for today's technology, is a social signal in itself, indicating the lack of social competence. Yet, continuous synthesis of socially-appropriate social signals is yet to be attempted.

Complex relations between social signals form and meaning: As explained above, relationships between social signals and their meaning are intrinsically complex. Firstly, the meaning of various signals is often not additive: when signals with meanings x and y are shown at the same time, the meaning of this complex signal may not be derivable from x and y alone. In addition, context plays a crucial role for the choice and interpretation of social signals. For example, environmental aspects such as the level of visibility and noise influence the choice of signals to be shown. On the other hand, societal aspects such as the formality of the situation and previously established roles and relations of the persons involved, and individual aspects such as the personality and affective state influence not only the choice of signals to be shown but the interpretation of the observed signals as well. Hence, context-

sensitive synthesis of human behaviour is needed but it still represents an entirely blue-sky research topic.

Timing: Social signals are not only characterised by the verbal and nonverbal cues by means of which they are displayed but also by their timing, that is, when they were displayed in relation to the signals displayed by other communicators involved in the interaction. Thus, social signals of an ECA need to be produced in anticipation, synchrony, or response to the actions of the human user with whom the character engages in the social interaction. This requires complex feedback loops between action and perception in real-time systems. This is another entirely unexplored, yet highly relevant, research topic.

Consistency: In general, it appears that human users are very critical when it comes to the consistency of a virtual character (Isbister & Nass, 2000). This relates to the challenge of multimodal synchronisation, that is, to timing between facial expression, gesture, and voice conveying a coherent and appropriate message. Research on this aspect is still ongoing there is no consensus on whether multimodal cues need to be fully synchronised, whether the redundancy of information coming from multiple cues is required, or whether it is also possible for one modality to compensate for the lack of expressiveness in other modalities (e.g., de Gelder & Vroomen, 2000). Consistency may also play a role in Mori's notion of an uncanny valley (Mori, 1970) - a robot that looks like a human but does not behave like one is perceived as unfamiliar and strange. Similarly, behaviour that may be consistent with a photo-realistic character may not be perceived as natural for a cartoon-like character, and vice versa.

Technical aspects: While it will take decades to fully understand and be able to synthesise various combinations of social signals that are appropriate for different contexts and different ECAs, we expect that it will soon be possible to model some limited but relevant phenomena. One example could be a model of politeness taking into account various modalities that, for a given ECA in a given context, contribute individually and jointly to the perception of a polite or rude behaviour (e.g., see the work by de Jong et al. 2008). There is an obvious relevance for applications: just like their human models, service robots/ ECAs should exhibit polite behaviour, whereas rescue robots should be able to insist on security-related requests. However, even when it is clear what signals and behaviours to generate, a practical challenge remains: current technology still lacks flexible models of expressivity and it usually does not operate in real-time. Expressive synthetic speech, for example, is a research topic that despite two decades of active research is still somewhat in its infancy (Schröder, 2009). Existing approaches are either capable of domain-specific natural-sounding vocal expressivity for a small number of possible expressions, or they achieve more flexible control over expressivity but of lower quality. Similarly, fully naturalistic movements of virtual agents can be attained when human movements recorded using motion capture technology are played back (Moeslund et al., 2006), but movements generated based on behaviour markup language (Kopp, 2006) tend to look less natural (Foster, 2007). These problems are not specific to synthesis of social signals, and they do not form insurmountable obstacles to research; however, they slow down the research, by making it substantially more time-consuming to create high-quality examples of the targeted expressions. Given the above-mentioned importance of timing, the lack of real-time systems impedes the realisation of timely appropriate social behaviours. Even a slight

delay in the analysis and synthesis of signals hinders dynamic adaptation and synchrony that are crucial in social interaction. Furthermore, the technological limitations pose serious difficulties for exploitation of research results in end-user applications, where fast adaptation to new domains is an important requirement. Therefore, enhancing the existing technology remains an important challenge facing the researchers in the field, independently of whether the aim is to develop socially-adapt ECAs or robots with no need of social awareness.

5 Summary and Additional Issues

Based on the enumeration of goals and challenges facing the researchers in the SSP domain as discussed in the previous chapters, the goals of the SSP research overall can be summarized under three headings: technological goals, human science goals, and practical impact goals.

Technological Goals:

- To develop systems capable of detecting and interpreting behavioural patterns that carry information about human social activity (analysis).
- To develop systems capable of synthesising behavioural patterns that carry socially significant information to humans (synthesis).
- To develop systems capable of spotting patterns of the user's behaviour that carry socially significant information to synthesise appropriate behaviours in an interaction with the user (system responsiveness).
- To develop sophisticated tools for instrumenting human science research.

Human Science Goals:

- To develop theories regarding the use of social signals during human-human interactions that can inform artificial agent behaviour, and can inform human-computer interactions.
- To contribute to the human science literature by modifying current theories and proposing new theories informed by the computational research in SSP.
- To create databases suitable for the analysis of human-human interactions, and suitable for training synthesis systems.
- To develop representational systems that describe human social behaviour and cognition in ways that are appropriate to technological tasks (such as labelling databases).
- To develop methods of measuring & evaluating social interactions (human/ human and human/ machine).

Practical Impact Goals: Application of the research on SSP is not restricted to a narrowly predefined set of issues like the ones listed above. It aims to address practical problems in a range of areas. Natural application areas include artificial agents and companions, human-computer interfaces, ambient intelligence, assisted living, entertainment, education, social skills training, and multimedia indexing. Applications have the important advantage of linking

the effectiveness of detection/ synthesis of social signals to the reality. For example, one of the earliest applications was the prediction of the outcome in transactions recorded at a call centre, and the results show that the number of successful calls can be increased by around 20

The Key Challenge: Based on the discussion so far, it should be clear that the SSP research has a specific challenge arising from the nature of the research it requires a strong collaboration between human sciences and technology research. This challenge should not only be achievable, but should be considered paramount to the success of SSP research. Besides the challenges discussed in the previous sections, there are a number of issues with a significant bearing on the character of the field that are still a matter of debate. Although they have not been decisively resolved, the profile of technological activities in the field implies that it tilts towards a particular kind of balance. Key examples are the following.

- Should linguistic information be included? From a human science standpoint, language is the social signal par excellence, and should obviously be included. Technologically, there is an obvious motive to avoid it. To wit, findings in basic research like those reported by Furnas et al. (1987) and Ambady & Rosenthal (1992) indicate that linguistic messages are rather unreliable means to analyse human behaviour, and it is very difficult to anticipate a persons word choice and the associated intent in affective and socially-situated expressions. In addition, the association between linguistic content and behaviour (e.g. emotion) is language-dependent and generalising from one language to another is very difficult to achieve.
- Naturalness vs. artificiality: Research in some related areas (e.g., affective computing) has relied heavily on data from actors or laboratory tasks, because naturalistic data and the related ground truth is too difficult to acquire. In return, some critics imply that only research on totally natural data is of any value. The balance implicit in the SSP research is that naturalness is a matter of degree, especially when it comes to learning the behaviour-synthesis models. Simulation is acceptable and, probably, in some cases practically necessary, so long as the signs in question are actually being used in an appropriate kind of interaction. Although such acted data can be used to learn how to synthesise certain behaviours, deliberately displayed data should be avoided when it comes to training machine learning methods for automatic analysis of social signals. Increasing evidence suggests that deliberate or posed behaviour differs in appearance and timing from that which occurs in daily life (Cohn & Schmidt, 2004; Valstar et al. 2006, 2007; Pantic 2009). Approaches that have been trained on deliberate and often exaggerated behaviours may fail to generalize to the complexity of expressive behaviour found in real-world settings.
- What are the appropriate validity criteria? Research in computer science, especially in computer vision and pattern recognition, insists that data should be associated with a clear ground truth. In SSP that leads to very difficult demands asking, for instance, what a person really felt or intended in a particular situation. A common alternative is to require high inter-rater agreement. That, too, is problematic, because it is a feature of some social signals that different people read them in different ways. The balance implicit in SSP is that the appropriate test depends on the actual application.

An additional challenging issue that has not been discussed so far relates to the fact that SSP deals with issues that are ethically sensitive. As a result, SSP has a range of ethical obligations. Many are standard, but some are not. Obligations that are shared with many other fields include: avoiding distress, deception and other undesirable effects on participants in studies, maintaining the confidentiality and where appropriate anonymity of participants involved in the research, avoiding the development of systems that could reasonably be regarded as intrusive, and limiting opportunities for abuse of the systems that they develop (e.g., through licensing arrangements). Particular obligations arise from the combination of complexity and sensitivity that is associated with social signals. The general requirement is sensitivity to the ways that social communication can affect people. Applying that to specific cases depends on intellectual awareness of individual issues (personality, age, etc.), of cultural issues (norms, specific signs, etc.), and of general expectations (what is disturbing, humiliating, etc.). Communicating about the area to non-experts raises particular issues. People are prone to systematic misunderstanding of SSP-type systems, so that they rely on them when they ought not to, fear them when they have no need to, and so on. Obligations relevant to offsetting are honesty (i.e., ensuring that what is said about a system is true), modesty (i.e., taking pains to ensure that its limitations as well as its achievements are understood), and public education (i.e., trying to equip people with the background knowledge to grasp what a particular system might or might not be able to do).

6 Conclusion

Social Signal Processing (Pentland, 2005, 2007; Vinciarelli et al. 2008, 2009) is the new research and technological domain that aims at providing computers with the ability to sense and understand human social signals. SSP is in its initial phase and the first step is to define the field and discuss issues facing the researchers in the field, which we attempted to achieve in this article.

Despite being in its initial phase, SSP has already attracted the attention of the technological community: the MIT Technology Review magazine identifies reality mining (one of the main applications of SSP so far), as one of the ten technologies likely to change the world (Greene, 2008), while management experts expect SSP to change organization studies like the microscope has changed medicine a few centuries ago (Buchanan, 2007). What is more important is that the first results in the field attest that social interactions and behaviours, although complex and rooted in the deepest aspects of human psychology, can be analyzed and synthesised automatically with the help of computers (Vinciarelli et al. 2009, 2010). However, although fundamental, these are only the first steps, and the journey towards artificial social intelligence and socially-aware computing is still long.

7 References

1. K. Albrecht. *Social Intelligence: The new science of success*. John Wiley & Sons Ltd, 2005.
2. J. Allwood. *Cooperative Multimodal Communication*, volume 2155. Berlin, Heidelberg: Springer, 2001.
3. N. Ambady and R. Rosenthal. Thin slices of expressive behavior as predictors of interpersonal consequences: a meta-analysis. *Psychological Bulletin*, 111(2): 256-274, 1992.

4. T. Bnziger and K.R. Scherer. Using actor portrayals to systematically study multimodal emotion expression: the GEMEP corpus. *Lecture Notes in Computer Science* 4738: 476-487, 2007.
5. J. Bates. The role of emotion in believable agents. *Communications of the ACM*, 37: 122-125, 1994.
6. C.G. Beer. What is a display? *American Zoologist*, 17(1): 155-165, 1977.
7. E. Berscheid and H.T. Reiss. Attraction and close relationships. In *Handbook of social psychology*, D. Gilbert, S. Fiske and G. Lindzey, eds., pp. 193-281. New York: McGraw Hill, 1997.
8. T. Bickmore and R. Picard. Establishing and Maintaining Long-Term Human-Computer Relationships. *ACM Trans. Computer Human Interaction*, 59(1): 21-30, 2005.
9. J. Bonaiuto and K.R. Thorisson. Towards a neurocognitive model of real time turn taking in face-to-face dialogue. *Embodied Communication in Humans And Machines*, G.K.I. Wachsmuth, M. Lenzen, Eds., Oxford University Press, 2008, pp. 451-484.
10. K. Bousmalis, M. Mehu and M. Pantic. Spotting agreement and disagreement: A Survey of nonverbal audiovisual cues and tools. *International Conference on Affective Computing and Intelligent Interfaces Workshops (ACIIW09)*, 2009.
11. K. Bousmalis, M. Mehu and M. Pantic. Spotting agreement and disagreement based on nonverbal audiovisual cues: A Survey. *IEEE Trans. Affective Computing*, submitted for publication.
12. P.M. Brunet, M. Charfuelan, R. Cowie, M. Schroeder, H. Donnan and E. Douglas-Cowie. Detecting politeness and efficiency in a cooperative social interaction. *International Conference on Spoken Language Processing (Interspeech)*, 2010, pp. 2542-2545.
13. E. Brunswik, E. Perception and the representative design of psychological experiments. Berkeley: University of California Press, 1956.
14. M. Buchanan. The science of subtle signals. *Strategy+Business*, 48: 68-77, 2007.
15. J.K. Burgoon and B.A.L. Poire. Nonverbal cues and interpersonal judgments: Participant and observer perceptions of intimacy, dominance, composure, and formality. *Communication Monographs*, 66(2): 105-124, 1999.
16. D. Byrne. *The Attraction Paradigm*, New York: Academic Press, 1971.
17. C. Castelfranchi. Social Power: a missed point in DAI, MA and HCI. In Y. Demazeau and J.P. Mueller (Eds.), *Decentralized AI*, pp. 49-62. North-Holland: Elsevier, 1990.
18. T.L. Chartrand and J.A. Bargh. The chameleon effect: the perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, 76(6): 893-910, 1999.
19. P.R. Cohen and H.J. Levesque. Performatives in a rationally based speech act theory. *Annual Meeting of the Association for Computational Linguistics*, 1990, pp. 79-88.
20. J. F. Cohn and K. L. Schmidt. The timing of facial motion in posed and spontaneous smiles. *J. Wavelets, Multi-resolution and Information Processing*, 2(2): 121-132, 2004. pages 97-102, 2004.
21. J.R. Curhan and A. Pentland. Thin slices of negotiation: predicting outcomes from conversational dynamics within the first 5 minutes. *Journal of Applied Psychology*, 92(3): 802-811, 2007.
22. M. de Jong, M. Theune and D. Hofs. Politeness and alignment in dialogues with a virtual guide. *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, 2008, pp. 207-214.
23. E. Douglas-Cowie, L. Devillers, J.C. Martin, R. Cowie, S. Savvidou, S. Abrilian and C. Cox. Multimodal databases of everyday emotion: Facing up to complexity. *International*

- Conference on Spoken Language Processing (Interspeech), 2005, pp. 813-816.
24. S. Duncan, S. Some signals and rules for taking speaking turns in conversations. *Journal of Personality and Social Psychology*, 23(2): 283-292, 1972.
 25. N. Eagle and A. Pentland. Reality mining: sensing complex social signals. *Journal of Personal and Ubiquitous Computing*, 10(4): 255-268, 2006.
 26. D. Efron. *Gesture and environment*. New York: King's Crown, 1941.
 27. I. Eibl-Eibesfeldt, I. Human ethology. New York: Aldine De Gruyter, 1989.
 28. P. Ekman, P. Are there basic emotions? *Psychological Review*, 99(3): 550-553, 1992.
 29. P. Ekman, P. Should we call it expression or communication? *Innovations in Social Science Research*, 10(4): 333-344, 1997.
 30. P. Ekman and W.V. Friesen. The repertoire of nonverbal behavior: Categories, origins, usage, and coding. *Semiotica*, 1(1): 49-98, 1969.
 31. M. Enquist. Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, 33(4): 1152-1161, 1985.
 32. L. Festinger, S. Schachter and K. Back. *Social Pressures in Informal Groups; a Study of Human Factors in Housing*. Palo Alto, California: Stanford University Press, 1950.
 33. M. Fishbein and I. Ajzen. *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley, 1975.
 34. M.E. Foster. Comparing Rule-based and Data-driven Selection of Facial Displays. *Workshop on Embodied Language Processing*, 2007, pp. 1-8.
 35. B.D. Gelder and J. Vroomen. The perception of emotions by ear and by eye. *Cognition & Emotion*, 14: 289-311, 2000.
 36. M. Gladwell. *Blink: The Power of Thinking without Thinking*. Little Brown & Company, 2005.
 37. D. Goleman. *Social intelligence*. Hutchinson, 2006.
 38. J.E. Grahe and F.J. Bernieri. The importance of nonverbal cues in judging rapport. *Journal of Nonverbal Behavior*, 23(4): 253-269, 1999.
 39. J. Gratch, N. Wang, J. Gerten, E. Fast and R. Duffy. Creating Rapport with Virtual Agents. *International Conference on Intelligent Virtual Agents*, 2007, pp. 197-210.
 40. K. Greene. 10 emerging technologies 2008. *MIT Technology Review*, February 2008.
 41. S.W. Gregory, K. Dagan and S. Webster. Evaluating the relation of vocal accommodation in conversation partners fundamental frequencies to perceptions of communication quality. *Journal of Nonverbal Behavior*, 21(1): 23-43, 1997.
 42. H.P. Grice. Meaning. *The Philosophical Review*, 66: 377-388, 1957.
 43. T. Guilford and M.S. Dawkins. What are conventional signals? *Animal Behaviour*, 49: 1689-1695, 1995.
 44. U. Hadar, T. Steiner and F.C. Rose. Head movement during listening turns in conversation. *Journal of Nonverbal Behavior*, 9(4): 214-228, 1985.
 45. E.T. Hall. *The silent language*. Doubleday, 1959.
 46. O. Hasson. Cheating signals. *Journal of Theoretical Biology*, 167: 223-238, 1994.
 47. D. Heylen. Challenges Ahead - Head movements and other social acts in conversations. *International Conference on Intelligent Virtual Agents*, 2005, pp. 25-36.
 48. R. Hinde. The concept of function. In G. Baerends and A. Manning (eds) *Function and evolution in behaviour*, (pp 3-15). Oxford: Clarendon Press, 1975.
 49. G.C. Homans. *Social behavior: its elementary forms*. Harcourt Brace, Orlando, 1961.
 50. H. Hung, D. Jayagopi, C. Yeo, G. Friedland, S. Ba, J.M. Odobez, K. Ramchandran, N. Mirghafori and D. Gatica-Perez. Using audio and video features to classify the most domi-

- nant person in a group meeting. *International Conference Multimedia*, 2007, pp. 835-838.
51. H. Hung, Y. Huang, G. Friedland and D. Gatica-Perez. Estimating Dominance In Multi-Party Meetings Using Speaker Diarization. *IEEE Transactions on Audio, Speech, and Language Processing*, 2010.
 52. S.E. Hyman. A new image for fear and emotion. *Nature*, 393: 417-418, Jun. 1998.
 53. K. Isbister and C. Nass. Consistency of personality in interactive characters: verbal cues, non-verbal cues, and user characteristics. *International Journal of Human-Computers Studies*, 53: 251-267, 2000.
 54. H.H. Kelley and J. Thibaut. *Interpersonal relations: A theory of interdependence*. New York: Wiley, 1978.
 55. D. Keltner. Signs of appeasement: Evidence for the distinct displays of embarrassment, amusement, and shame. *Journal of Personality and Social Psychology*, 68(3): 441-454, 1995.
 56. M. Kimura and I. Daibo. Interactional synchrony in conversations about emotional episodes: A measurement by the between-participants pseudosynchrony experimental paradigm. *Journal of Nonverbal Behavior*, 30(3): 115-126, 2006.
 57. M.L. Knapp and J.A. Hall. *Nonverbal Communication in Human Interaction*. Harcourt Brace College Publishers, 1972.
 58. S. Kopp, T. Stocksmeier and D. Gibbon. Incremental multimodal feedback for conversational agents. *International Conference on Intelligent Virtual Agents*, 2007, pp. 139-146.
 59. J.L. Lakin, V.E. Jeeris, C.M. Cheng and T.L. Chartrand. The Chameleon effect as social glue: Evidence for the evolutionary significance of nonconscious mimicry. *Journal of Nonverbal Behavior*, 27(3): 145-162, 2003.
 60. R.L. Lewis. Beyond Dominance: the importance of Leverage. *The Quarterly Review of Biology*, 77(2): 149-164, 2002.
 61. J. Martin, S. Abrilian, L. Devillers, M. Lamolle, M. Mancini and C. Pelachaud. Levels of Representation in the Annotation of Emotion for the Specification of Expressivity in ECAs. *International Conference on Intelligent Virtual Agents*, 2005, pp. 405-417.
 62. J. Maynard-Smith and D.G. Harper. *Animal signals: models and terminology*. *Journal of Theoretical Biology*, 177: 305-311, 1995.
 63. J. Maynard-Smith and D.G. Harper. *Animal signals*. Oxford University Press, Oxford, 2003.
 64. I. McCowan, D. Gatica-Perez, S. Bengio, G. Lathoud, M. Barnard and D. Zhang. Automatic analysis of multimodal group actions in meetings. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 27(3): 305-317, 2005.
 65. T.B. Moeslund, A. Hilton and V. Krger. A survey of advances in vision-based human motion capture and analysis. *Computer Vision and Image Understanding*, 104: 90-126, 2006.
 66. M. Mori. The Uncanny Valley. *Energy*, 7: 33-35, 1970.
 67. D. Morris. *People watching*. Vintage, 2007.
 68. R. Niewiadomski, M. Ochs and C. Pelachaud. Expressions of empathy in ECAs. *International Conference on Intelligent Virtual Agents*, 2008, pp. 37-44.
 69. M. Ochs, R. Niewiadomski and C. Pelachaud. How a virtual agent should smile? Morphological and dynamic characteristics of virtual agents smiles. *International Conference on Intelligent Virtual Agents (IVA10)*, 2010, pp. 427-440.
 70. A. Oikonomopoulos, I. Patras and M. Pantic. Discriminative space-time voting for joint recognition and localization of actions. *International ACM Conference on Multimedia, Workshops (ACM-MM-W10)*, 2010.
 71. M.J. Owren and J.A. Bachorowski. Reconsidering the evolution of nonlinguistic commu-

- nication: the case of laughter. *Journal of Nonverbal Behavior*, 27(3): 183-200, 2003.
72. M. Pantic. Machine analysis of facial behaviour: Naturalistic and dynamic behaviour. *Philosophical Transactions of Royal Society B*, 364 (12 December): 3505-3513, 2009.
73. M. Pantic, A. Pentland, A. Nijholt and T. Huang. Human computing and machine understanding of human behavior: A survey. *LNAI*, 4451: 47-71, 2007.
74. M. Pantic, A. Pentland, A. Nijholt and T. Huang. Human-centred intelligent human-computer interaction (HCI2): How far are we from attaining it? *Intl J. Autonomous and Adaptive Communications Systems*, 1(2): 168-187, 2008.
75. S.R. Partan and P. Marler. Communication goes multimodal. *Science*, 283(5406): 1272-1273, 1999.
76. C.C. Peirce. *Collected chapters*. Cambridge: Cambridge University Press, 1931-1935.
77. C. Pelachaud, V. Caroglio, B. De Carolis, F. de Rosi and I. Poggi. Embodied contextual agent in information delivering application. *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)*, 2002, pp. 758-765.
78. A. Pentland. *Social dynamics: Signals and behavior*. *International Conference Developmental Learning*, 2004.
79. A. Pentland. Socially aware computation and communication. *IEEE Computer*, 38(3): 33-40, 2005.
80. A. Pentland. *Social Signal Processing*. *IEEE Signal Processing Magazine*, 24(4): 108-111, 2007.
81. I. Poggi. *Mind, hands, face and body: Goal and belief view of multimodal communication*. Berlin: Weidler, 2007.
82. I. Poggi and F. D'Errico. *Cognitive Modelling of Human Social Signals*. *Social Signal Processing Workshop*, in conjunction with *International Conference on Multimedia*, 2010.
83. G. Psathas. *Conversation Analysis - The study of talk-in-interaction*. Sage Publications, 1995.
84. B. Raducanu and D. Gatica-Perez. Inferring competitive role patterns in reality TV show through nonverbal analysis. *Multimedia Tools and Applications J.*, 2010.
85. D. Rendall, M.J. Owren and M.J. Ryan. What do animal signals mean? *Animal Behaviour*, 78(2): 233-240, 2009.
86. V.P. Richmond and J.C. McCroskey. *Nonverbal Behaviors in interpersonal relations*. Allyn and Bacon, 1995.
87. J.A. Russell, J.A. Bachorowski, and J.M. Fernandez-Dols. Facial and vocal expressions of emotion. *Annual Reviews in Psychology*, 54(1): 329-349, 2003.
88. Z. Ruttkay and C. Pelachaud. *From brows to trust: Evaluating Embodied Conversational Agents*. Kluwer Academic Publishing, 2004.
89. H. Sacks, E.A. Schegloff and G. Jefferson. A simplest systematics for the organization of turn taking for conversation. *Language*, 50: 696-735, 1974.
90. H. Salamin, S. Favre, and A. Vinciarelli, Automatic role recognition in multiparty recordings: Using social affiliation networks for feature extraction. *IEEE Transactions on Multimedia*, 11(7): 1373-1380, 2009.
91. A.E. Scheen. The significance of posture in communication systems. *Psychiatry*, 27: 316-331, 1964.
92. K.R. Scherer. Personality inference from voice quality: The loud voice of extroversion. *European Journal of Social Psychology*, 8(4): 467-487, 1978.
93. K.R. Scherer. What does facial expression express? In K. T. Strongman (Ed.), *International review of studies of emotion*, volume 2, (pp. 139-165). Chichester, U.K.: Wiley, 1992.

94. K.R. Scherer, T. Bnziger and E. Roesch. *Blueprint for affective computing: A source-book*. Oxford, England: Oxford University Press, 2010.
95. M. Schröder. *Expressive Speech Synthesis: Past, Present, and Possible Futures*. In *Affective Information Processing*, J. Tao and T. Tan, eds., pp. 111-126. London: Springer, 2009.
96. M. Schröder and R. Cowie. *Developing a consistent view on emotion-oriented computing*. LNCS, 3869: 194-205, 2006.
97. U. Segerstrale and P. Molnar, Eds. *Nonverbal communication: where nature meets culture*. Lawrence Erlbaum Associates, 1997.
98. C.E. Shannon and W. Weaver. *The mathematical theory of information*. Urbana: University of Illinois Press, 1949.
99. A.F. Smeaton, P. Over, and W. Kraaij. *Evaluation campaigns and TRECVID*. Proc. ACM Intl Workshop on Multimedia Information Retrieval, pp. 321-330, 2006.
100. E. L. Thorndike. *Intelligence and its use*. Harpers Magazine, 140: 227-235, 1920.
101. S.S. Tomkins. *Consciousness, imagery and affect*, volume 1. New York: Springer, 1962.
102. H.C. Triandis. *Culture and social behavior*. McGraw-Hill, 1994.
103. J. Trouvain and M. Schröder. *How (Not) to Add Laughter to Synthetic Speech*. LNAI, 3068: 229-232, 2004.
104. M.F. Valstar, H. Gunes and M. Pantic. *How to distinguish posed from spontaneous smiles using geometric features*. International Conference Multimodal Interfaces, 2007, pp. 38-45.
105. M.F. Valstar, M. Pantic, Z. Ambadar and J.F. Cohn. *Spontaneous vs. posed facial behaviour: Automatic analysis of brow actions*. International Conference Multimodal Interfaces, 2006, pp. 162-170.
106. S.L. Verhencamp. *Handicap, index, and conventional signal elements of bird song*. In Y. Edpmark, T. Amundsen, & G. Rosenqvist, G. (eds), *Animal Signals: Signalling and Signal Design in Animal Communication*, (pp. 277-300). Trondheim, Norway: Tapir Academic Press, 2000.
107. A. Vinciarelli. *Capturing order in social interactions*. IEEE Signal Processing Magazine, 26(5): 133-137, 2009.
108. A. Vinciarelli, M. Pantic, H. Bourlard, and A. Pentland. *Social Signal Processing: State-of-the-art and future perspectives of an emerging domain*. International Conference Multimedia, 2008, pp. 1061-1070.
109. A. Vinciarelli, M. Pantic, and H. Bourlard. *Social Signal Processing: Survey of an emerging domain*. Image and Vision Computing Journal, 27(12): 1743-1759, 2009.
110. A. Vinciarelli, M. Pantic, D. Heylen, C. Pelachaud, I. Poggi, F. DErrico and M. Schroeder. *Bridging the Gap Between Social Animal and Unsocial Machine: A Survey of Social Signal Processing*. IEEE Transactions on Affective Computing, 2010, submitted.
111. N. Wang, W.L. Johnson, P. Rizzo, E. Shaw and R.E. Mayer. *Experimental Evaluation of Polite Interaction Tactics for Pedagogical Agents*. International Conference Intelligent User Interfaces, 2005, pp. 12-19.
112. M. Weiser. *The computer for the 21st century*. Scientific American, Special Issue on Communications, Computers, and Networks, 265(3): 95-104, 1991.
113. R.S. Woodworth. *Dynamics of behavior*. Holt New York, 1961.
114. G. Yule. *Pragmatics*. Oxford University Press, 1996.
115. A. Zahavi. *Mate selection: selection for a handicap*. Journal of Theoretical Biology, 53: 205-214, 1975.
116. Z. Zeng, M. Pantic, G.I. Roisman and T.H. Huang. *A survey of affect recognition methods: audio, visual and spontaneous expressions*. IEEE Trans. Pattern Analysis and Machine

Intelligence, 31(1): 39-58, 2009.