Modeling Affect Regulation and Induction

Eva Hudlicka
Psychometrix Associates, Inc.
1805 Azalea Drive, Blacksburg, VA 24060
hudlicka@acm.org

Abstract

The ability to regulate one's affective state and to induce an affective state in another person is a key component of social and emotional interaction. In this paper we describe an architecture design aimed at modeling these two closely related processes. We first define the high-level functional components of each process, and then use these to refine the specific knowledge necessary for their implementation. Based on this analysis we propose a cognitive architecture capable of modeling affect regulation in the self, and affect induction in the other. We illustrate the functioning of this architecture by way of an example from a peacekeeping training simulation. The architecture is currently being implemented in the context of this simulation training task.

Introduction

What do a caring mother, a beloved mentor, an effective manager, and a manipulative politician have in common? Whether intuitive and unconscious, or elaborately planned and carefully executed, each of these roles requires that the person be able to induce an appropriate affective state in another individual, possibly by consciously manipulating his or her own affective state. Two closely related phenomena play a role here: *affect regulation* in the self, and *affect induction* in the other.

The ability to identify and regulate one's own affective state, to correctly interpret others' affective states, and to appropriately express one's own to induce a desired state in another are the core are components of what has recently become known as emotional and social intelligence. Whether applied to self or to others, used consciously or unconsciously, to help or to control and manipulate, these processes are critical components of behavior management at both individual and interpersonal level.

The recent explosion of interest in affective phenomena and affective computing has resulted in much progress in a variety of areas, including affect recognition and expression (Canamero & Fredslund, 2000; Braezel & Fitzpatrick, 2000; Petrushin, 2000; Vyzas & Picard, 1998; Breese & Ball, 1998; Hudlicka, 2000b), models of cognitive appraisal processes (Elliot et al., 1999; Scherer, 1993; Andre et al., 1999), and models affective influence on cognition (Sloman, 2000; Hudlicka et al., 2000a) and behavior (Velasquez, 1998).

Not nearly as prevalent are efforts to model the two closely related phenomena that are central to emotional and social intelligence and effective social functioning: *affect regulation in the self* and *affect induction in others.*

This paper describes an on-going effort within this broad area of affect regulation and induction. Five primary objectives, both theoretical and applied, motivate our interest and approach:

- Elucidation of the mechanisms of affect regulation and induction through the development of computational models.
- Definition of specific functional components, and representational and inferencing requirements.
- Understanding how these processes are influenced and modified by individual and contextual variables.
- Increasing the realism of training simulations through explicit models of a) affect self-regulation, and b) affective interpersonal phenomena (e.g., team dynamics).
- Using these models as basis for investigation of the processes of "mindreading", which are central to social intelligence, and their disorders.

This work is being conducted in the context of a broader effort to develop a cognitive modeling testbed environment for modeling the influence of affective states and personality traits on cognitive processes: the MAMID testbed (Methodology for Analysis and Modeling of Individual Differences) (Hudlicka et al., 2000a; Hudlicka and Billingsley, 1999). The MAMID environment supports the modeling of multiple interacting agents, whose behavior is controlled by a highly parameterized cognitive architecture.

Our initial focus in the MAMID testbed development was on modeling individual differences in terms of variations in the cognitive architecture parameters. The current focus is enhancing the initial architecture by incorporating explicit models of affect regulation and affect induction.

The paper is organized as follows. *First*, we briefly illustrate the pervasiveness of affect regulation and induction and highlight some of the critical issues that must be addressed. *Second*, we define the functional components necessary to implement affect regulation and induction, in terms of specific processes and required knowledge. *Third*, we describe a cognitive architecture designed to model affect regulation and affect induction, and illustrate its functioning with an example. *Finally*, we conclude with a summary, conclusions, and a brief outline of future research directions.

Affect Regulation and Induction "In-the-Wild"

Consider the following vignettes. An infant cries and is immediately picked up and rocked by her mother until she calms down. A school teacher enters an unruly classroom, sternly demands silence, and glares at the pupils who are too slow to settle down. A diver stands still at the edge of the board and takes a few deep breaths as he prepares to jump. In spite of the variations in contexts, roles, and aims (to soothe, to control, to calm, to convince, to manipulate, to entertain), these vignettes share two factors in common: the actors attempt to induce a particular affective state, either in themselves or in others.

One could argue that these types of affective manipulations constitute the essence of interpersonal interaction. While cultural differences exist in choice of affects to regulate, acceptable forms of affective expression, acceptable means of affect regulation, and appropriate contexts for affect inductions, the processes themselves, in multiple forms, are pervasive (Averill, 1994).

Affect regulation and induction have been written about broadly and extensively, across time, cultures and disciplines (e.g., Aristotle & the Stoics have addressed affect regulation in the self; Roman orators provided guidelines regarding the induction of affect in others). More recently, affect regulation has been subject of inquiry by a variety of emotion researchers. Psychologists have analyzed the process to more precisely define what exactly is being regulated in affect regulation (e.g., Levenson provides an analysis of the broad phenomenon of emotional control in terms of its constituent factors (e.g., direction, locus, onset, form, etc.) (Levenson, 1994). Neuroscientists have attempted to determine which specific channels of expressions can be controlled for which affects and to what extent (LeDoux, 1994). Clinicians have studied the role of affective regulation in the etiology, maintenance, and treatment of emotional disorders (e.g., Teasdale and Barnard, 1993). Affect induction, in particular induction by contagion, has been studied extensively (Hatfield et al., 1994). And of course there is the rich tradition of affect induction in the arts, and affect manipulation in a variety of more applied disciplines and practices, such as marketing and advertising.

Some critical questions and issues that must be addressed in modeling these processes are listed below.

- What are the limits of affective control and regulation in self and other?
- To what extent are affective regulation and induction conscious?
- What is the role of metacognition in affect regulation and induction?
- How are personality traits and contextual variables related to affect regulation and induction?
- Can affective profiles be constructed, for the self and other, that identify the best affective state required to induce particular cognitive / perceptual schemas?

Functional Components of Affect Regulation and Induction

The first step in designing a model of affect regulation and induction is to identify the structure of each process in terms of the distinct functional components, and the knowledge and inferencing requirements. Figure 1 provides a high-level "flow-chart" of possible process structure for affect induction (other-oriented). The flow chart for affect regulation in the self is a subset–with "Is own affective state appropriate?" and "Induce appropriate state in self" omitted. Below we define the role of several of these components within the overall process, and describe the necessary knowledge requirements.

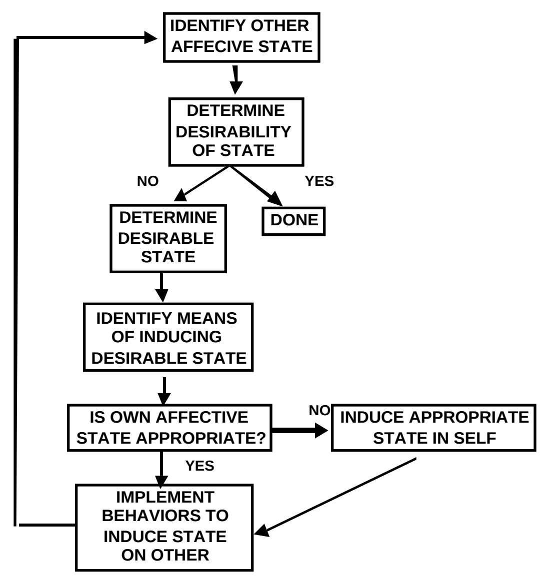

Figure 1: Affect Induction – "Other"

Recognition / Awareness of Affect

Role: The ability to recognize a particular affective state is the basis for further determination whether that affect is desirable or undesirable in a particular context, and represents the first step in both affect regulation and affect induction (first block of the process diagram in figure 1. In the self, this recognition is the basis for deciding whether to maintain or change own affective state (e.g., ability to recognize when one has entered an overly aggressive, passive, or anxious state). In the other, the recognition of a particular affective state is the basis for deciding whether to attempt to induce an alternative affect (e.g., ability to recognize that a teammate has become too anxious or too disengaged).

Knowledge / Inferencing required: In the self, data necessary for self-affect appraisal include self cues (e.g., awareness of mental and physiological correlates of the particular affective state), environment cues, and baseline data. In the other, data necessary for affect appraisal include baseline and current data for observable behavior (task, self, and interpersonal behaviors), and environment-context cues (e.g., task difficulty)). This type of information can be captured within a variety of representational formalisms, which need to be able to support flexible, differentially weighted integration of the multiple influences determining the final affective state, as well as reasoning under uncertainty (e.g., Bayesian belief nets, fuzzy logic, case-based reasoning, and rules capable of explicit representation of uncertainty).

Identification of the Most Appropriate Affective State for a Given Task Context (Self and Other)

Role: Sufficient understanding of one's own and others' affective profile is necessary to identify the most appropriate affective state to accomplish goals in a particular task / social context (second and third blocks of the process diagram in figure 1). Such understanding is the basis for determining the most appropriate behaviors, and the affective state triggering these behaviors (e.g., recognize that a deadline-induced heightened anxiety helps project completion).

Knowledge / Inferencing Required: A broad range of knowledge is required to support this function This includes (for both self and other): the commonly experienced affective states; possible transitions among them; external / internal triggers of each state; goals, biases, expectations, beliefs, and behaviors associated with each state. This knowledge is captured within a structure we have termed Affective / Motivational Profile (see figure 2), which is structured as a state transition diagram. The structure of this profile is loosely modeled on a structure termed role-relationship model, used in clinical psychology to characterize the individual's multiple social roles and associated affective states (Horowitz et al., 1991).

On the inferencing side, an ability to match current task, interpersonal, and intrapersonal context and objectives with specific behaviors and goals contained in the profile is necessary. This then enables the selection of the most desirable affective state to trigger the required behaviors. Given the generally noisy data, representation and reasoning under uncertainty are critical, again, Belief nets, fuzzy logic and fuzzy case-based reasoning are good candidates.

Figure 2: Affective / Motivational Profile

Architecture Design

The requirements outlined above led to a cognitive architecture design aimed at modeling affect regulation in the self and affect induction in the other. A high-level outline of the architecture is shown in figure 3. The detailed architecture design is described elsewhere (Hudlicka and Billingsley, 1999; Hudlicka, 2001).

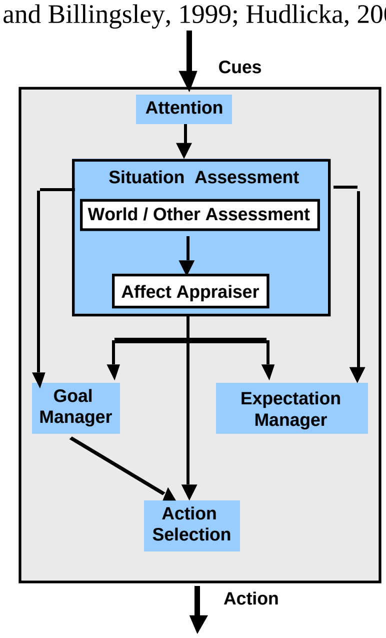

Figure 3: Cognitive Architecture

A key feature of this design is that no explicit modules are necessary which are dedicated exclusively to affective regulation and induction, other than the affect appraiser component of the situation assessment module. The affective regulation and induction are integrated with the non-affective task processing and behaviors (see detailed description below). This design then allows the modeling of a variety of phenomena where affective processing either enhances or interferes with task processing, particularly in multi-agent team settings. For example, the commonly observed performance reduction associated with extreme states of anxiety can be modeled not as a *reduction* in available resources, as it is often conceptualized, but rather as a *redirection* of attention and processing to the self.

Overview: The architecture implements a staged, sequential processing model of cognition (see description below). Its initial design was motivated by the desire to model individual differences, including affective states and personality traits, in terms of architecture parameter manipulations. The augmented design described here provides further representational and inferencing enhancements to support more extensive modeling of additional affective phenomena, including regulation of affect in the self, induction of affective states in the other, and shifts among behaviors oriented towards the task (world), other agents, or the self.

The enhancements include the following: a distinct *Affect Appraisal* module, containing a parameterized model of affect appraisal process, embedded within the *Situation Assessment* module; explicit representation of goals and expectations as critical factors influencing affect appraisal; distinct *Goal Manager* module and *Expectation Manager* module; and *Affective / Motivational Profiles* for self and other agents within the simulation (refer to figure 2).

Integrated World, Other, and Self Processing: This design uses the same processing structure to process diverse categories of stimuli; those relevant to the task (e.g., state of the terrain, unit configuration within the terrain); those relevant to others (e.g., superiors' and subordinates' affective states); and those relevant to self (e.g., own affective state, status of self goals). The switching among these stimuli, and thus observable shifts in behavioral focus from task to others to self, occur as a result of the combined effects of affective state and incoming stimuli, along with current goals and expectations, all of which influence the attentional filtering process and situation assessment, which then determine the final goal and behavior selection.

Architecture Modules: The augmented MAMID architecture thus consists of the following six modules: *Attention*, filtering the incoming cues and selecting a subset for further processing; *Situation Assessment* integrating low-level cues into high-level situations (task, other, and self related) and containing an embedded *Affect Appraiser* which integrates various factors that influence the affect appraisal process; both static (traits, individual history) and dynamic (current affective state, current situation, goals, expectations); *Goal Manager*, integrating current situation assessment (including assessment of own and others' affective states), existing goals, and expectations to derive the next most critical goal, *Expectation Manager,* combining situation and affective state information to derive the next expectation, and *Action Selection,* selecting the most suitable action for achieving the current goal within the current situation context and using the available behavioral repertoire.

Knowledge Representation: The key representational structures within the architecture are *cues*, *situations*, *expectations*, *goals*, and *behaviors*, each representing relevant information about the external task / world, others, and the self. Cues are highly abstracted perceptual features representing status of relevant entities (e.g., location of Unit x, affective state of Agent B, incoming status report from Agent B, heart rate of self, etc.). Situations, expectations, and goals consist of {entity | attribute | value} 3-tuples, that capture the relevant world, other, and self information. Situations represent actual state of the world, other, or self (e.g., hostile crowd at location x,y; subordinate S's morale is low; self affective state is calm). Goals represent desired state of the world, other, or self. Expectations represent the agent's expected developments within each of these categories.

As discussed above, a key knowledge structure for implementing affect regulation and induction is a detailed model of both self and other agents' *Affective / Motivational Profile* (figure 2). This profile provides detailed information necessary to determine, for both self and other agents, the most desirable affect for a particular situational context, and the best means of inducing a different affect.

Inferencing: MAMID uses a number of inferencing mechanisms. The *Attention Module* uses procedural knowledge representation and inferencing to filter out less relevant cues by prioritizing the incoming cues according to current goals and expectations (priming); threat level (anxiety-induced attentional biases); and familiarity, ease of recognition, and task-relevance (reflection of skill level and training). The *Situation Assessment / Affect Appraiser* module uses Bayesian belief nets as the primary means of integrating a variety of uncertain evidence to derive the final high-level assessment of the current world situation, and self or other affect appraisal. The *Affect Appraiser* performs affect appraisal both for the self, and for other agents, using the same inferencing mechanism (BN's) but different sources and types of data. Belief nets are well-suited for combination of uncertain evidence and have been used by several researchers to implement affect assessments (e.g., Breese and Ball, 1998; Conati, 2001).

For *self-appraisal*, the module aims to provide an integrated model of theories that postulate primacy of cognition over emotion on the one hand (the cognitive appraisal theories (e.g., Ortony et al., 1988; Lazarus, 1991)), and those that postulate primacy of emotion (e.g., Zajonc, 1985).

For appraisal of other agents' affective state, a variety of factors are combined, including observed behavioral manifestations (e.g., help requests, frequency of communication, sudden cessation of expected activities), knowledge of the other agent (e.g., their trait profile, skill level, history), and awareness of the overall world situation as it impacts the other agent (e.g., knowing they are under attack). This appraisal process implements what we refer to as knowledge-based appraisal and builds upon an affect assessment procedure implemented in a related project: the Affect and Belief Adaptive Interface System (Hudlicka, 2000b).

The *Goal Manager*, *Expectation Manager*, and *Action Selection* modules use a combination of belief net and rule based inferencing to derive the next most critical goal and expectation, and to select the most appropriate action to satisfy the current goal.

The functioning of this architecture is illustrated below by way of an example from the current MAMID implementation: peacekeeping training simulation.

Example

The demonstration scenario simulates the behavior of multiple interacting agents conducting a peacekeeping operation that involves interaction among multiple agents: friendly and adversary military unit commanders, and neutral and hostile group leaders. Ample opportunities exist for both affective self-regulation and affect induction (e.g., a commander might become too anxious and avoid engagement; too aggressive and take too many risks; subordinate might become too anxious and need to be calmed down; superior might become too passive and hinder mission progress, etc.). The example below illustrates the process of affect regulation in the self.

Affect Regulation: Reducing Anxiety Level in Self In this situation the friendly commander encounters a hostile crowd blocking the movement of his unit. The commander is not experienced in this type of operations, and expects further casualties, both of which contribute to his already heightened level of anxiety. His anxious state focuses his attention on the threatening stimuli (e.g., hostile crowd) and increasingly on his own affective state and its manifestations (e.g., rapid heart rate). (This is in contrast to a non-anxious commander, who might focus instead on his own combat superiority, forthcoming help, and high skill level.) The output of the *Attention* module thus consists primarily of self-related cues. These cues are passed on to the *Situation Assessment / Affect Appraiser* module. Here they are combined into integrated assessments of the situation, which, again, emphasize the commander's own affective state, which at this point has reached the high anxiety level. The processing continues in the *Goal Manager* module, which receives the affect assessment (high anxiety), and identifies the desired affect for this context, as specified in the *Affective / Motivational Profile* for that agent (low anxiety). The *Goal Manager* compares the current state (high anxiety) with the desired state (low anxiety) and marks low anxiety as the current goal. Note that the goal is a self goal and defines the desired affective state. (This is in contrast to a non-anxious commander, whose goal might be to calm the crowd, ask for help, withdraw etc. task related goals.) This information is then passed on to the *Action Selection* module, which identifies the best means of accomplishing this goal by identifying the specific strategies for lowering the anxiety level, as specified in the agent's *Affective / Motivational* profile. For this particular commander and context, this involves extensive communication behaviors to his superiors (asking for help), peers (reporting situation), and subordinates (requesting extensive status reports).

Summary, Conclusions and Future Work

We described a research effort aimed at producing models of affect regulation and affect induction. Based on analysis of existing theories and data regarding these processes, we first defined a high-level structure of each process, consisting of distinct functional components. We then analyzed these components and identified specific knowledge and inferencing requirements necessary for their implementation. Finally, we described a candidate architecture for their implementation and illustrated its functionality with an example from a peacekeeping training simulation scenario.

An initial proof-of-concept MAMID cognitive architecture has been implemented (Hudlicka and Billingsley, 1999; Hudlicka et al., 2000a). We are currently in the process of augmenting this architecture with the enhancements described above (e.g., Affect Appraiser, Goal Manager, and Expectation Manager), to enable the implementation and evaluation of the affective regulation and induction processes described in this paper.

Conclusions Clearly the proposition to sufficiently understand one self, let alone another individual, to enable affect regulation and induction is at best a daunting enterprise.

We do not suggest that the model described in this paper is an adequate, generic representation of these processes. Rather we consider this a first step in investigating these complex processes, and a means of generating empirically testable hypotheses regarding the specific mechanisms and structures that mediate them. In addition to serving this basic research function, these models also contribute to an increased realism of autonomous agents, by enhancing their ability to respond in more socially-aware manner within their "interpersonal" environments.

Requirements / Constraints for Successful Implementation A number of requirements were identified as necessary for implementing an affect regulation and induction model. These include: Limiting the number, type, and resolution of affective states; Using highly individualized, multiple sources of data for affect appraisal; Constraining the environment in terms of possible situations, expectations, goals, and behavioral alternatives.

Although the current MAMID simulation is implemented within a military training simulation, we believe that the architecture, the affect appraisal processes, and the affective / motivation profile structures, apply to a variety of domains involving social-affective interactions and behaviors.

Future Work Once the full implementation of the design is completed, future work falls into three categories: *Evaluation* of the full implementation to determine: a) whether specific empirical data can be accommodated within the model, and b) whether more realistic interpersonal and self behaviors emerge as a result of augmenting the architecture with affect regulation and induction.; *Broadening the modeling approach* to focus on additional details of the regulation / induction mechanisms; Application areas such as models supporting the development of virtual reality therapeutic environments (e.g., group therapy for Asperger syndrome children).

Acknowledgements

The research described in this paper was supported in part by US Army Contract DASW01-98-M-1985. We would like to acknowledge continued support of the COTR, Dr. Bob Witmer. J. Billingsley and J. Broglio implemented the MAMID prototype, Mr. Ted Fichtl of The Compass Foundation was the primary subject matter expert.

References

Andre, E., Klesen, M., Gebhard, P., Allen, S. , Rist, T. (1999). Integrating Models of Personality and Emotions in Lifelike Characters. In *Proceedings of IWAI*. Italy.

Averill, J.R. (1994). Emotions Unbecoming and Becoming. In *The Nature of Emotions* (P. Ekman, R.Davidson, eds.). Oxford: Oxford.

Braezel, C. and Fitzpatrick, P. (2000). That Certain Look: Social Amplification of Animate Vision. In *Proceedings of the AAAI FS Socially Intelligent Agents: The Human in the Loop*. TR FS-00-04. Menlo Park, CA: AAAI Press.

Breese, J. and Ball, G. (1998). Bayesian Networks for Modeling Emotional State and Personality. In *Proceedings of Emotional and Intelligent*. AAAI FS, TR FS-98-03. Menlo Park, CA: AAAI Press.

Canamero, L.D. and Fredslund, J. (2000). How Does It Feel? Emotional Interaction with Humanoid LEGO Robot. In *Proceedings of the AAAI FS Socially Intelligent Agents*. TR FS-00-04. Menlo Park, CA: AAAI Press.

Conati, C. (2001). Probabilistic Assessment of User's Emotions During the Interaction with Educational Games. In *Proceedings of Workshop on Attitudes, Personality and Emotions in User-Adapted Interaction*, Germany.

Elliot, C., Lester, J., & Rickel, J. (1999). Lifelike pedagogical agents and affective computing: An exploratory synthesis. In M. Wooldridge and M. Veloso, (Eds.). *AI Today*. Lecture Notes in AI. NY: Springer-Verlag.

Hatfield, E., Cacioppo, J.T., and Rapson, R.L. (1994). *Emotional Contagion.* Cambridge, UK: Cambridge University Press.

Horowitz, M.J., Merluzzi, M.E., Ghannam, J.H., Hartley, D. and Stinson, C.H. (1991). *Person Schemas and Maladaptive Interpersonal Patterns*. Chicago: The University of Chicago Press.

Hudlicka, E., Zacharias, G., & Psotka, J. (2000a). Increasing Realism of human agents by modeling individual differences. In *Proceedings of the AAAI FS Simulating Human Agents*. TR FS-00-03. Menlo Park, CA: AAAI Press.

Hudlicka, E. (2000b). *Increasing SIA Architecture Realism by Modeling and Adapting to Affect and Personality*. In *Proceedings of the AAAI Symposium Socially Intelligent Agents*. TR FS-00-04. Menlo Park, CA: AAAI Press.

Hudlicka, E. and Billingsley, J. (1999). Representing Behavior Moderators in Military Human Performance Models. In *Proceedings of the 8th Conference on Computer Generated Forces and Behavioral Representation*, Orlando, FL.

Lazarus, R. S. (1991). *Emotion and Adaptation*. Oxford: Oxford.

LeDoux, J. (1994). The Degree of Emotional Control Depends on the King of Personal System Involved. In *The Nature of Emotions*. (P. Ekman, R.Davidson, eds.). Oxford: Oxford.

Levenson, R.W. (1994). Emotion Control: Variations and Consequences. In *The Nature of Emotions*. (P. Ekman, R. Davidson, eds.). Oxford: Oxford.

Ortony, A., Clore, G.L., and Collins, A. (1988). *The Cognitive Structure of Emotions*. Cambridge: Cambridge.

Petrushin, V. (2000). Emotion Recognition Agents in Real World. In *Proceedings of the AAAI FS Socially Intelligent Agents*. TR FS-00-04. Menlo Park, CA: AAAI Press.

Sloman, A. (2000). Architectural Requirements for Human-like Agents Both Natural and Artificial. In *Human Cognition and Social Agent Technology*. NY: John Benjamin.

Teasdale, J.D. and Barnard, P.J. (1993). *Affect, Cognition an Change*. Hove: Erlbaum.

Velaszuez, J.D. (1998). Modeling Emotion-Based Decision-Making. In *Proceedings of Emotional and Intelligent*,. AAAI FS, TR FS-98-03. Menlo Park, CA: AAAI Press.

Vyzas, J.D. and Picard, R. W. (1998). Affective Pattern Classification. In *Proceedings of Emotional and Intelligent*.. AAAI FS, TR FS-98-03. Menlo Park, CA: AAAI Press.

Zajonc R.B. (1985) Emotion and Facial Efference: A theory reclaimed.*Science*, 228, 15-21.