

How to Determine the Utility of Emotions

Matthias Scheutz

Artificial Intelligence and Robotics Laboratory
Department of Computer Science and Engineering
University of Notre Dame, Notre Dame, IN 46556, USA
mscheutz@cse.nd.edu

Abstract

In this paper, we describe a new methodology for determining the utility of emotions. After briefly reviewing the status quo of emotional agents in AI, we describe the methodology and demonstrate it by showing the utility of “anger” for biologically plausible foraging agents in an evolutionary setting.

Background on Emotions and AI

Evidence from psychology (Frijda 1986; Izard 1991; Scherer, Schorr, & Johnstone 2001), neuroscience (Damasio 1994; LeDoux & Fellous 1995; Panksepp 2000; Hamm, Schupp, & Weike 2003), and ethology (Lorenz & Leyhausen 1973; McFarland 1981) suggests that emotions play several crucial roles in biological organisms. Especially in humans, they seem to be deeply intertwined with cognitive processing (e.g., they can bias problem solving strategies in humans (Bless, Schwarz, & Wieland 1996; Schwarz) or help to evaluate a situation quickly (Kahneman, Wakker, & Sarin 1997; Damasio 1994; Clore, Gasper, & Conway 2001)). Finally, and most importantly, emotions are crucially involved in *social control* (Frijda 2000; Cosmides & Tooby 2000) ranging from signaling emotional states (e.g., pain) through facial expressions and gestures (Ekman 1993) to perceptions of emotional states that cause approval or disapproval of one’s own or another agents’ actions (relative to given norms), which can then trigger corrective responses (e.g., guilt). Yet, there is not even agreement among emotion researchers about how to construe *basic emotions* or whether the concept is coherent (Ortony & Turner 1990; Griffiths 1997).

The difficulties with emotion concepts are also reflected in AI, where different forms of emotions have been investigated to varying degrees ever since its beginnings (e.g., (Toda 1962; Simon 1967; Pfeifer & Nicholas 1982; Dyer 1987; Pfeifer 1988)). Over the recent years, various “believable synthetic characters and life-like animated agents” (e.g., (Bates 1994; Hayes-Roth 1995; Maes 1995; Lester & Stone 1997; Rizzo *et al.* 1997)), “emotional pedagogic agents” (e.g., (Gratch 2000; Shaw, Johnson, & Ganeshan 1999; Lester *et al.* 1997; Okonkwo & J. Vassileva 2001; Conati 2002)), “emotional virtual agents and robots” (e.g.,

(Bates, Loyall, & Reilly 1991; Velásquez 1999; Michaud & Audet 2001; Breazeal 2002; Arkin *et al.* 2003)), and “computational models of human emotion” (e.g., (Elliott 1992; Cănamero 1997; Wright 1997; Allen 2001; Marsella & Gratch 2002)) have been proposed.¹ Yet, there are divergent views among all these researchers about what it means to implement emotion in agents (e.g., (Ventura & Pinto-Ferreira 1999; Wehrle 1998; Picard 2001; Scheutz 2002a)).

Most of this work in AI has focused on what could be called *effect models* of emotion. Effect models implement only overt, observable effects of emotional behavior. They are intended to get the “input-output mapping” of a given behavioral description right. In the extreme case, such a mapping could be as simple as that employed in an animated shopping agent which displays a surprised face if the user attempts to delete an item from the shopping basket. Many architectures of so-called “believable agents” (e.g., (Hayes-Roth 1995; Scheutz & Römmer 2001; Rizzo *et al.* 1997; Loyall & Bates 1997) for simulated agents and (Shibata & Irie 1997; Breazeal 1998; Velásquez 1999; Michaud & Audet 2001; Murphy *et al.* forthcoming) for robots) are part of this group, where the primary goal is to induce the belief in the human observer that the agent is in a particular emotional state.

The main problem with effect models is that they are silent about the role of emotion in agent architectures. They may or may not actually implement emotional processes to achieve the desired overt behaviors. And if they do, the implemented states are often labeled with familiar terms, without specifying how the implemented states differ from those usually denoted with these terms (McDermott 1981; Scheutz 2002a). A state labeled “surprise”, for example, may have very little in common with the complex processes underlying notions of “surprise” in humans and various animals (i.e., the violation of a predicted outcome (Ortony, Clore, & Collins 1988; Macedo & Cardoso 2001)), if it is functionally defined to be triggered by loud noises (Velásquez 1997a; 1997b) (for such a state, “startle” would be the more appropriate label). Effect models are, therefore, inadequate for determining the utility of emotions in agent architectures.

¹This list is only a brief excerpt of the recent literature and by far not complete, see also (Trappl, Petta, & Payr 2001; Hatano, Okada, & Tanabe 2000; Pfeifer 1988).

Process models of emotion, on the other hand, are applicable as they are intended to model and simulate aspects of emotional processes (typically in humans) as they unfold (Peschl & Scheutz 2001), following predictions of psychological or neurological theories of emotion (Scherer 1993; Panksepp 1998; Ortony, Clore, & Collins 1988)). Process models are much more complex than effect models, given that they focus on the internal processes of an agent's control system, and are typically only implemented in simulated agents (e.g., (Wright 1997; Marsella & Gratch 2002; Cănamero 1997; Allen 2001; McCauley & Franklin 1998)).

The problem with current process models is threefold: for one, they do not provide or use a conceptual framework to characterize the implemented emotional states (i.e., what kind of state it is they implemented and what it takes in general to implement such a state), nor do they investigate variations of such states systematically. And most importantly, they typically do not evaluate emotional states with respect to their utility (neither by varying architectural parameters nor by comparing them to other implementations of the same task). Hence, the potential of the states, other than to be present in a particular model, remains unclear.

Evaluating the Utility of Emotional Architectures

We proposed a methodology that allows for an objective evaluation of the utility of emotions, which consists of a four step process: (1) emotion concepts are analyzed and defined in terms of architectural capacities of agent architectures (Sloman 2002). (2) Agent architectures with particular emotional states (as defined in (1)) are constructed for a given task, for which also a performance measure is defined. (3) "Experiments with agent architectures" (Pollack *et al.* 1994; Hanks, Pollack, & Cohen 1993)) are carried out with the so-defined emotional agents (either in simulations or on actual robots) and their performance is determined for a pre-determined set of architectural and environmental variations. The outcome then is a *performance space* that corresponds to the varied parameters. The last two steps are repeated with agents implementing non-emotional (or, in general, other) architectures. (4) All resulting performance spaces are then compared, in particular, with respect to the agents' *performance-cost tradeoffs*, i.e., their performance taken relative to the (computational) cost necessary to maintain and run the instantiated architecture. The last point is crucial as it may well be that emotional agents do not perform better than non-emotional ones on a given task in absolute terms, but that they do much better in relative terms, i.e., with fewer resources (which is usually believed to be the case by emotion researchers).

We have applied this methodology in various settings and tasks and found, for example, that *emotional action selection* can be very effective in the competition for resources in hostile multiagent environments (Scheutz 2000; Scheutz, Sloman, & Logan 2000; Scheutz under review). Emotional control mechanisms performed much better in a variety of foraging, survival, and object collection tasks in environments with little to no structure than agents with

much more sophisticated deliberative control systems (including A_e^* planning (Pearl 1982), plan executing methods with error feedback, and goal management mechanisms) if the "cost of deliberation" is taken into account (Scheutz & Logan 2001; Scheutz & Schermerhorn 2002; 2003). Furthermore, we found that emotional states like "fear" and "aggression" (Scheutz 2001)) are likely to evolve in a variety of competitive multiagent environments. Finally, in studies of the potential of *emotion expression and recognition for social control* we found that emotions can have a beneficial regulatory effect in social groups (Scheutz 2002b) and lead to superior conflict resolution strategies (Scheutz & Schermerhorn forthcoming).

In the following, we will briefly demonstrate this methodology.

Architectural Requirements and Mechanisms for Emotional Control

We start with a brief characterization of emotional states and show the difference between simple and complex versions in terms of architectural requirements and mechanisms.

Simple *emotions* are caused by some disparity between an agent's desire state and the state of the environment, and are themselves causes for actions that are intended to change the state of the environment so as to make it agree with the agents' desires (Sloman, Chrisley, & Scheutz forthcoming).² A simple "anger state", for example, is caused by the perception of a potentially threatening environmental condition (e.g., the approach of another agent) and causes the agent to change its behavioral dispositions so as to deal with the threat (e.g., to fight). It can be implemented by a controller integrating the frequency of perceptions of the threat over a given time interval (Scheutz 2001; under review), e.g., using the differential equation $\partial Output / \partial t = Output \cdot (G_{sensor} \cdot S_e - G_{discount})$, where *Output* is the output of the controller, G_{sensor} is the gain for the sensor input and $G_{discount}$ is the discount value for the past output.

More complex emotional states (such as "worrying about whether a grant proposal can be completed in time") can be caused by a combination of perceptions and processes internal to the agent (e.g., results of complex deliberations about the utility of trying to achieve a particular goal compared to alternatives).

Such emotions may include any of the following components and possibly more (based on the analysis in (Beaudoin & Sloman 1993; Ortony, Clore, & Collins 1988; Wehrle & Scherer 2001) and others):

1. an elicitor (e.g., *the grant proposal*)
2. an eliciting condition (e.g., *the possibility of (1) not being completed by the deadline*)
3. criteria for the evaluation of (2) based on various factors such as beliefs, goals, norms, standards, tastes, attitudes, etc. (e.g., *completing (1) by the deadline is crucial to research career*)
4. an evaluation of (2) in terms of (3) (e.g., *(2) is undesirable*)

²Often, emotions are themselves the states that the agent does or does not desire.

5. possible causes for (2) (e.g., *deadline approaching rapidly, work progressing too slowly, etc.*)
6. a hedonic attitude towards (2) (e.g., *displeasure*)
7. a measure of the urgency to act on (1) given (2) (e.g., *urgent*)
8. a set of strategy to cope with (2) (e.g., *cancel meetings, focus attention on (1), etc.*)
9. a set of motivations to be instantiated based on (8) (e.g., *being able to continue one's research, being able to fund students, etc.*)
10. a set of emotions to be instantiated based on (4) through (8) (e.g., *distress*)
11. the selected motivation (if any) based on (4) through (8) (e.g., *being able to continue research*)

Consequently, complex representational and processing mechanisms (e.g., frames (Minsky 1975) or scripts (Schank & Abelson 1977) combined with pattern matching and rule instantiation mechanisms) are required for architectures to be able to support complex emotions.

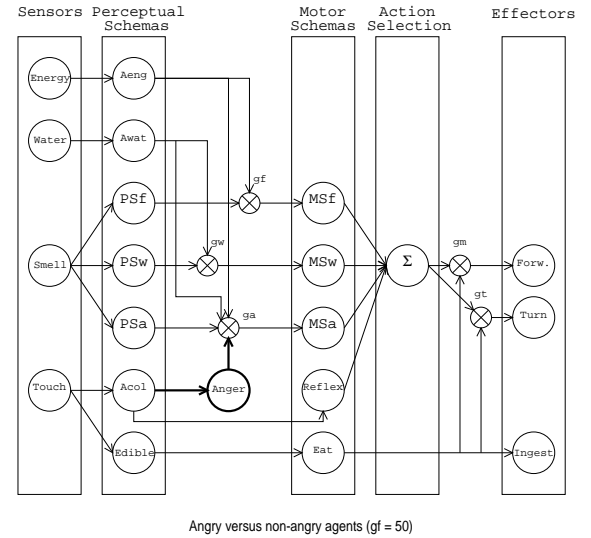
The Utility of Anger in Biological Settings

This section illustrates how a simple form of “anger” can be implemented in a biologically plausible way in a schema-based agent architecture and used to influence the agent’s *action selection*. Furthermore, it demonstrates the systematic exploration of the performance space defined by this architectures by systematically varying architectural parameters. Comparing the performance space of “angry agents”, for simplicity sake, with those of “non-angry” agents in a survival task, the (relative) benefits of emotional control can be determined. The left part of Figure 1 shows a schema-based architecture (Arkin 1989; Arbib 1992) for a foraging agent that needs to find food and water in a hostile multi-agent environment in order to survive. The small crossed circles indicate gains of schemas that are taken as architectural parameters: the degree to which an agent is attracted to food (gf), to water (gw), and to other agents (ga).

The bold-face circle labeled “Anger” represents a schema that is only present in the “angry agents” (non-angry agents do not have it nor the associated links). It is connected to an “alarm schema” (Acol), which is triggered if an agent touches another agents and implements a simple emotional control circuit as described above by virtue of influencing the gain of a motor schema that changes the agent’s propensity to fight other agents: the higher the output of the controller, the more likely the agent will fight (for details see (Scheutz under review)).

The right part of Figure 1 shows the performance space for both agent kinds using “average number of survivors after 10000 cycles” (averaged over 40 simulation runs) as the performance measure. As can be seen from the graph, angry agents reach a global maximum at $ga = 10$ and $gw = 30$ (which is statistically marginally significant: t-test, $p < 0.09$ for $\alpha = 0.05$). Consequently, in the kinds of environments studied, being (capable of being) “angry” does prove useful for survival.

The experimentation and evaluation method demonstrated here with a simple example can be straightforwardly applied to more complex agents, tasks and environments (as



average number of survivors




Figure 1: Left: a schema-based emotional architecture for simulated agents. Right: a performance comparison of emotional and non-emotional agents along two architectural dimensions.

we have done in the past (Scheutz & Schermerhorn 2002; 2003)). It is also worth mentioning that the emotional subsystem of the proposed architecture was based on what emotion researchers presume to be the functional organization of the emotional “fear/anger system” in many animals, given that animals are typically taken to exhibit either a “fight” or a “flee” behavior (e.g., (Berkowitz 2003)). Since fight and flee behaviors are very directly linked with their emotional makeup (fear will lead to flee, anger to fight behavior), these two emotions are typically taken to be incompatible, i.e., they cannot be present at the same time, e.g., (Anderson, Deuser, & DeNeve 1995)). While the above architectures are even more restrictive in that an agent cannot have fear and anger at different times, only a few simple modifications need to be applied to the architecture to allow an agent to be capable of having both kinds of emotions (either an addition of a simple switch system that flips the sign on the gain in certain circumstances would allow the agent to have

fear sometimes and anger other times, or another emotional controller could be added as mentioned above allowing for fear and anger to be present at the same time, an architecture that would find support from recent results, e.g., (Berkowitz 2003)).

Conclusion

The methodology proposed in the paper will allow researchers to study the utility of emotions in a very general way that applies to biological organisms and artificial agents alike. It is built on a systematic way of defining emotional states in terms of capacities of agent architectures and exploring their utility for the control of agents in experiments with agent architectures by systematically varying architectural (and environmental) parameters for the given task. By comparing the resulting performance spaces of emotional and non-emotional agents (in particular, using performance-cost tradeoffs), we believe that it will be possible to answer important open questions about the utility of emotions for both biological and artificial agents in a great variety of tasks.

References

- Allen, S. 2001. *Concern Processing in Autonomous Agents*. Ph.D. Dissertation, School of Computer Science, The University of Birmingham.
- Anderson, C.; Deuser, W.; and DeNeve, K. 1995. Hot temperatures, hostile affect, hostile cognition, and arousal: Tests of a general model of affective aggression. *Personality and Social Psychology Bulletin* 21:434–448.
- Arbib, M. 1992. In Shapiro, S., ed., *The Handbook of Brain Theory and Neural Networks*. MIT Press. chapter Schema Theory, 830–833.
- Arkin, R.; Fujita, M.; Takagi, T.; and Hasegawa, R. 2003. An ethological and emotional basis for human-robot interaction. *Robotics and Autonomous Systems* 42:3–4.
- Arkin, R. C. 1989. Motor schema-based mobile robot navigation. *International Journal of Robotic Research* 8(4):92–112.
- Bates, J.; Loyall, A. B.; and Reilly, W. S. 1991. Broad agents. In *AAAI spring symposium on integrated intelligent architectures*. American Association for Artificial Intelligence. (Repr. in SIGART BULLETIN, 2(4), Aug. 1991, pp. 38–40).
- Bates, J. 1994. The role of emotion in believable agents. *Communications of the ACM* 37(7):122–125.
- Beaudoin, L., and Sloman, A. 1993. A study of motive processing and attention. In Sloman, A.; Hogg, D.; Humphreys, G.; Partridge, D.; and Ramsay, A., eds., *Prospects for Artificial Intelligence*. Amsterdam: IOS Press. 229–238.
- Berkowitz, L. 2003. Affect, aggression, and antisocial behavior. In (Davidson, Scherer, & Goldsmith 2003). 804–823.
- Bless, H.; Schwarz, N.; and Wieland, R. 1996. Mood and the impact of category membership and individualizing information. *European Journal of Social Psychology* 26:935–959.
- Breazeal, C. 1998. Regulating human-robot interaction using ‘emotions’, ‘drives’, and facial expressions. In *Proceedings of Autonomous Agents 98*, 14–21.
- Breazeal, C. L. 2002. *Designing Sociable Robots*. MIT Press.
- Cānamero, D. 1997. Modeling motivations and emotions as a basis for intelligent behavior. In Johnson, L., ed., *Proceedings of the First International Symposium on Autonomous Agents (Agents’97)*, 148–155. New York, NY: ACM.
- Clore, G.; Gasper, K.; and Conway, H. 2001. Affect as information.
- Conati, C. 2002. Probabilistic assessment of user’s emotions in educational games. *Journal of Applied Artificial Intelligence, special issue on “Merging Cognition and Affect in HCI”*.
- Cosmides, L., and Tooby, J. 2000. Evolutionary psychology and the emotions. In Lewis, M., and Haviland-Jones, J. M., eds., *Handbook of Emotions*. NY: Guilford, 2nd edition. 91–115.
- Damasio, A. R. 1994. *Descartes’ Error: Emotion, Reason, and the Human Brain*. New York, NY: Gosset/Putnam Press.
- Davidson, R. J.; Scherer, K. R.; and Goldsmith, H. H., eds. 2003. *Handbook of Affective Sciences*. New York: Oxford University Press.
- Dyer, M. G. 1987. Emotions and their computations: Three computer models. *Cognition and Emotion* 1(3):323–347.
- Ekman, P. 1993. Facial expression and emotion. *American Psychologist* 48(4):384–392.
- Elliott, C. 1992. *The Affective Reasoner: A process model of emotions in a multi-agent system*. Ph.D. Dissertation, Institute for the Learning Sciences, Northwestern University.
- Frijda, N. H. 1986. *The emotions. Studies in Emotion and Social Interaction*. Cambridge: Cambridge University Press.
- Frijda, N. H. 2000. The psychologists’ point of view. In (Lewis & Haviland-Jones 2000). 59–74.
- Gratch, J. 2000. Emile: Marshalling passions in training and education. In *4th International Conference on Autonomous Agents*, 325–332.
- Griffiths, P. 1997. *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: Chicago University Press.
- Hamm, A. O.; Schupp, H. T.; and Weike, A. I. 2003. Motivational organization of emotions: Autonomic changes, cortical responses, and reflex modulation. In (Davidson, Scherer, & Goldsmith 2003). 187–211.
- Hanks, S.; Pollack, M. E.; and Cohen, P. 1993. Benchmarks, testbeds, controlled experimentation, and the design of agent architectures. *AI Magazine* 14(4):17–42. <http://www.cs.pitt.edu/pollack/distrib/testbeds.ps>.

- Hatano, G.; Okada, N.; and Tanabe, H., eds. 2000. *Affective Minds*. Amsterdam: Elsevier.
- Hayes-Roth, B. 1995. Agents on stage: Advancing the state of the art of AI. In *Proc 14th Int. Joint Conference on AI*, 967–971.
- Izard, C. E. 1991. *The Psychology of Emotions*. New York: Plenum Press.
- Kahneman, D.; Wakker, P.; and Sarin, R. 1997. Back to bentham? explorations of experienced utility. *Quarterly Journal of Economics* 112:375–405.
- LeDoux, J., and Fellous, J. 1995. Emotion and computational neuroscience. In Arbib, M. A., ed., *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press. 356–360.
- Lester, J. C., and Stone, B. A. 1997. Increasing believability in animated pedagogical agents. In Johnson, W. L., and Hayes-Roth, B., eds., *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, 16–21. Marina del Rey, CA, USA: ACM Press.
- Lester, J.; Converse, S.; Kahler, S.; Barlow, T.; Stone, B.; and Bhogal, R. 1997. The persona effect: Affective impact of animated pedagogical agents. In *Conference on Human Factors in Computing Systems CHI '97*, 359–366.
- Lewis, M., and Haviland-Jones, J. M., eds. 2000. *Handbook of Emotions*. New York: The Guilford Press, 2nd edition.
- Lorenz, K., and Leyhausen, P. 1973. *Motivation and Animal Behavior: An Ethological View*. New York: Van Nostrand Co.
- Loyall, B., and Bates, J. 1997. Personality-rich believable agents that use language. In *Proceedings of the 1st International Conference on Autonomous Agents (Agents'97)*. ACM Press.
- Macedo, L., and Cardoso, A. 2001. Modeling artificial forms of surprise in an artificial agent. In *23rd Annual Conference of the Cognitive Science Society*.
- Maes, P. 1995. Artificial life meets entertainment: Lifelike autonomous agents. *CACM* 38(11):108–114.
- Marsella, S., and Gratch, J. 2002. Modeling the influence of emotion on belief for virtual training simulations. In *Proceedings of the 11th Conference on Computer-Generated Forces and Behavior Representation*.
- McCauley, T. L., and Franklin, S. 1998. An architecture for emotion. In *AAAI Fall Symposium entitled Emotional and Intelligent: The Tangled Knot of Cognition*, 122–127.
- McDermott, D. 1981. Artificial intelligence meets natural stupidity. In Haugeland, J., ed., *Mind Design*. Cambridge, MA: MIT Press. 143–160.
- McFarland, D. 1981. *The Oxford Companion to Animal Behavior*. Oxford: Oxford University Press.
- Michaud, F., and Audet, J. 2001. Using motives and artificial emotion for long-term activity of an autonomous robot. In *5th Autonomous Agents Conference*, 188–189. Montreal, Quebec: ACM Press.
- Minsky, M. L. 1975. A framework for representing knowledge. In Winston, P. H., ed., *The psychology of computer vision*. New York: McGraw-Hill. 211–280.
- Murphy, R. R.; Lisetti, C.; Tardif, R.; Irish, L.; and Gage, A. forthcoming. Emotion-based control of cooperating heterogeneous mobile robots. *IEEE Transactions on Robotics and Automation*.
- Okonkwo, C., and J. Vassileva. 2001. Affective pedagogical agents and user persuasion. In *Proceedings of the 9th International Conference on Human-Computer Interaction*, 397–401.
- Ortony, A., and Turner, T. 1990. What's basic about basic emotions? *Psychological Review* 97:315–331.
- Ortony, A.; Clore, G.; and Collins, A. 1988. *The Cognitive Structure of the Emotions*. New York: Cambridge University Press.
- Panksepp, J. 1998. *Affective neuroscience-The Foundations of Human and Animal Emotions*. Oxford: Oxford University Press.
- Panksepp, J. 2000. Emotions as natural kinds within the mammalian brain. In (Lewis & Haviland-Jones 2000). 137–156.
- Pearl, J. 1982. A_{ϵ}^* —an algorithm using search effort estimates. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 4, 392–399.
- Peschl, M., and Scheutz, M. 2001. Explicating the epistemological role of simulation in the development of theories of cognition. In *Proceedings of the Seventh International Colloquium on Cognitive Science*, 274–280. Institute for Logic, Cognition, Language and Information, The University of the Basque Country.
- Pfeifer, R., and Nicholas, D. W. 1982. Toward computational models of emotion. In *Proc. of the 5th ECAI*, 269–271.
- Pfeifer, R. 1988. Artificial intelligence models of emotion. In Hamilton, V.; Bower, G. H.; and Frijda, N. H., eds., *Cognitive Perspectives on Emotion and Motivation*, volume 44 of *Series D: Behavioural and Social Sciences*. Netherlands: Kluwer Academic Publishers. 287–320.
- Picard, R. 2001. What does it mean for a computer to “have” emotions? In Trapp, R.; Petta, P.; and Payr, S., eds., *Emotions in Humans and Artifacts*. MIT Press.
- Pollack, M. E.; Joslin, D.; Nunes, A.; Ur, S.; and Ephrati, E. 1994. Experimental investigation of an agent commitment strategy. Technical Report 94-31, University of Pittsburgh. <http://www.cs.pitt.edu/pollack/distrib/tw.ps>.
- Rizzo, P.; Veloso, M.; Miceli, M.; and Cesta, A. 1997. Personality-driven social behaviors in believable agents. In *Proceedings of the AAAI Fall Symposium on Socially Intelligent Agents*, 109–114.
- Schank, R., and Abelson, R. R. 1977. *Scripts, Plans, Goals, and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Scherer, K.; Schorr, A.; and Johnstone, T. 2001. *Appraisal theories of emotions: Theories, methods, research*. New York: Oxford University Press.

- Scherer, K. R. 1993. Neuroscience projections to current debates in emotion psychology. *Cognition and Emotion* 7(1):1–41.
- Scheutz, M., and Logan, B. 2001. Affective versus deliberative agent control. In Colton, S., ed., *Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing*, 1–10. York: Society for the Study of Artificial Intelligence and the Simulation of Behaviour.
- Scheutz, M., and Römmer, B. 2001. Autonomous avatars? from users to agents and back. In de Antonio, A.; Aylett, R.; and Ballin, D., eds., *Intelligent Virtual Agents, Third International Workshop, IVA 2001, Madrid, Spain, September 10-11, 2001, Proceedings*, volume 2190 of *Lecture Notes in Computer Science*, 61–71. Springer.
- Scheutz, M., and Schermerhorn, P. 2002. Steps towards a theory of possible trajectories from reactive to deliberative control systems. In Standish, R., ed., *Proceedings of the 8th Conference of Artificial Life*. MIT Press.
- Scheutz, M., and Schermerhorn, P. 2003. Many is more but not too many: Dimensions of cooperation of agents with and without predictive capabilities. In *Proceedings of IEEE/WIC IAT-2003*. IEEE Computer Society Press.
- Scheutz, M., and Schermerhorn, P. forthcoming. The role of signaling action tendencies in conflict resolution. *Journal of Artificial Societies and Social Simulation*.
- Scheutz, M.; Sloman, A.; and Logan, B. 2000. Emotional states and realistic agent behaviour. In Geril, P., ed., *Proceedings of GameOn 2000, Imperial College London*, 81–88. Delft: Society for Computer Simulation.
- Scheutz, M. 2000. Surviving in a hostile multiagent environment: How simple affective states can aid in the competition for resources. In Hamilton, H. J., ed., *Advances in Artificial Intelligence, 13th Biennial Conference of the Canadian Society for Computational Studies of Intelligence, AI 2000, Montréal, Québec, Canada, May 14-17, 2000, Proceedings*, volume 1822 of *Lecture Notes in Computer Science*, 389–399. Springer.
- Scheutz, M. 2001. The evolution of simple affective states in multi-agent environments. In Cañamero, D., ed., *Proceedings of AAAI Fall Symposium*, 123–128. Falmouth, MA: AAAI Press.
- Scheutz, M. 2002a. Agents with or without emotions? In Weber, R., ed., *Proceedings of the 15th International FLAIRS Conference*, 89–94. AAAI Press.
- Scheutz, M. 2002b. The evolution of affective states and social control. In Hemelrijk, C. K., ed., *Proceedings of International Workshop on Self-Organisation and Evolution of Social Behaviour*.
- Scheutz, M. under review. Schema-based architectural approach towards implementing affective states in autonomous agents.
- Schwarz, N. Feelings as information: Informational and motivational functions of affective states.
- Shaw, E.; Johnson, W. L.; and Ganesan, R. 1999. Pedagogical agents on the web. In *Third International Conference on Autonomous Agents*, 283–290.
- Shibata, T., and Irie, R. 1997. Artificial emotional creature for human-robot interaction: A new direction for intelligent systems. In *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. paper number 47, in CD-ROM Procs.
- Simon, H. A. 1967. Motivational and emotional controls of cognition. Reprinted in *Models of Thought*, Yale University Press, 29–38, 1979.
- Sloman, A.; Chrisley, R.; and Scheutz, M. forthcoming. The architectural basis of affective states and processes. In Fellous, J., and Arbib, M., eds., *Who needs emotions? The Brain Meets the Machine*. New York: Oxford University Press.
- Sloman, A. 2002. Architecture-based conceptions of mind. In *Proceedings 11th International Congress of Logic, Methodology and Philosophy of Science*, 397–421. Dordrecht: Kluwer. (Synthese Library Series).
- Toda, M. 1962. The design of the fungus eater: A model of human behavior in an unsophisticated environment. *Behavioral Science* 7:164–183.
- Trapp, R.; Petta, P.; and Payr, S., eds. 2001. MIT Press.
- Velásquez, J. 1997a. Cathexis: A computational model of emotions. In *Proceedings of the First International Conference on Autonomous Agents*, 518–519. Marina del Rey, CA: ACM Press.
- Velásquez, J. 1997b. Modeling emotions and other motivations in synthetic agents. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 10–15. Menlo Park, CA: AAAI.
- Velásquez, J. 1999. When robots weep: Emotional memories and decision-making. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 70–75. Menlo Park: AAAI, CA.
- Ventura, R., and Pinto-Ferreira, C. 1999. Emotion-based agents: Three approaches to implementation. In Velásquez, J., ed., *Third International Conference on Autonomous Agents, Workshop on Emotion-Based Agent Architectures*, 121–129.
- Wehrle, T., and Scherer, K. 2001. chapter Towards computational modeling of appraisal theories, 350–365.
- Wehrle, T. 1998. Motivations behind modeling emotional agents: Whose emotion does your robot have? In Nummaoka, C.; Cañamero, D.; and Petta, P., eds., *Grounding Emotions in Adaptive Systems, SAB'98 (5th International Conference of the Society for Adaptive Behavior)*. Zurich: MIT Press.
- Wright, I. 1997. *Emotional agents*. Ph.D. Dissertation, School of Computer Science, The University of Birmingham.