

A Robot's Experience of its User: Theory

Johan F. Hoorn (JF.Hoorn@fsw.vu.nl)

Center for Advanced Media Research Amsterdam

VU University

Buitenveldertselaan 3, 1082 VA Amsterdam, Netherlands

Abstract

To make a robot that emotionally is more human-like, it should be capable of balancing simultaneously occurring tendencies of positive and negative affect. This paper is about the modeling of robot attitudes towards the user, using a fuzzy trade-off algorithm. Empirical research showed that in humans, involvement and distance are not mutually exclusive and that a range of factors contribute (e.g., ethics, affordances). A hypercube of fuzzy sets deals with combining these factors into an overall value for the involvement-distance trade-off, which subsequently can form the input for affective response selection (Bosse et al. 2008).

Keywords: Robotics; experience design; affect; fuzzy sets; involvement-distance trade-off.

Introduction

There is quite a number of attempts to represent emotions and attitudes in robots. The early examples, such as the schizophrenic paranoid patient-simulator PARRY (Cerf 1973; Colby 1981) and Eliza, had no recognizable model of affect and neither did the story-line through affect analyzer BORIS (Dyer 1983), admirable work though it was. Later approaches implemented emotion generation in software agents (e.g., Bates, 1994; Velasquez, 1997) or attempted to model neurological details of the emotion process (Thayer and Lane, 2000). Recently, Marsella and Gratch (2003) simulated different strategies of coping with emotions such as 'positive reinterpretation', 'acceptance,' and 'denial.' Bosse, Pontier, and Treur (2007) modeled an existing theory of affect regulation (Gross, 2001) with dynamical systems representations.

However, the present contribution does not focus on the emotions per se but rather on the general involving and distancing trends from which positive and negative emotions can transpire (Hoorn and Konijn, 2003). That is, humans can at the same time find another person attractive but morally repulsive, useful to achieve certain goals but distasteful in his or her manners (e.g., Konijn and Hoorn, 2005). In other words, many specific positive and negative tendencies contribute to or can challenge the more wide-ranging concepts of involvement (e.g., becoming friends) and distance (e.g., staying aloof). A formal representation, therefore, should allow for the integration of small-scale judgments (e.g., morally bad, aesthetically pleasing, useful but rude) into two competing overall judgments of being involved as well as keeping a distance. One complication in such an attempt is that one single characteristic of a person may contribute to both involvement and distance all at the same time (e.g., eyes so beautiful and cruel). Thus, if we

regard 'eyes' as one of the features of a person,¹ that feature does not only participate in the set of beautiful things but partly also in the set of ethically bad things. From a conventional set theoretical viewpoint, features cannot be 'divided' between more sets. Fuzzy set theory, however, provides us with a way out.

In the remainder of the paper I will explain how users perceive of their robots to take this as a starting point to model the robot's experience of its user. The body of empirical research reported next is theoretically rooted in media psychology and computer-human interaction. Then I discuss the way fuzzy sets may be helpful to translate the empirical results obtained with humans in interaction with embodied agents into a computational model that the robot can use to interact with humans.

Interactive PEFiC

In earlier studies, we showed empirically that involvement-distance trade-offs are part of the user's experiences with fictional and virtual characters (e.g., Konijn and Hoorn 2005; Van Vugt, Hoorn et al. 2006; Van Vugt, Konijn et al. 2006; 2007). Such trade-offs lead to mixed emotions, which are natural to human assessment of others ("Love his looks, hate his behavior") (cf. conflicts between goals and principles in Elliot 1993). We repeatedly established the effect that involvement and distance were not mutually exclusive experiences but that they occurred in parallel. They affected each other but also functioned in a relatively independent way. The involvement-distance trade-off was the result of evaluating the features of an embodied agent on several dimensions, which we compiled into our Interactive model on Perceiving and Experiencing Fictional Characters (I-PEFiC, Figure 1) (Van Vugt, Hoorn et al. 2006).

The I-PEFiC model assumes that encountering an embodied agent is two-sided. For one thing the agent is a virtual person that you can engage with (the engagement process), for the other it is a tool you can interact with to achieve certain task-related goals (the interaction process). These processes consist of three stages: encoding, comparison, and response.

During the *encoding* of a character (Figure 1), the user appraises the agent's features for their level of being ethical (good-bad) (e.g., trustworthiness of an advisor), aesthetic (beautiful-ugly), and realistic (realistic-unrealistic) (we used the term "epistemics"). During the encoding, moreover, the user appraises in how far the agent system affords to be used

¹ In this paper, features are not meant as attributes of something that can take on a range of values but rather as the characteristics of a person or a character.

as a computer tool (aid-obstacle). Here also do opposites not exclude one another because features of the agent system (e.g., a search engine) can be perceived as partly handy (Ctrl F) and at the same time as partly unhandy (Edit > Search > Clippit dialog).




Figure 1: Interactive PEFiC

Our work showed that none of the dimensions we discerned in Figure 1 were superfluous (i.e. Konijn and Hoorn 2005; Konijn and Bushman 2007; Van Vugt, Hoorn et al. 2006) and that all were used to appraise the agent’s features.

In the *comparison* phase, the features are judged for similarity (similar-dissimilar) (e.g., “I am not as beautiful as she is”), relevance to user goals (relevant-irrelevant) (“A conversational agent is important for not feeling lonely”), and valence to goals (positive-negative outcome expectancies) (“I don’t expect the Office Assistant to provide me with any real help”). In this sense, whether the character has features that afford achieving goals (e.g., a search function) or that obstruct achieving goals (e.g., misplaced proactive behavior) not only affects use intentions but also the overall levels of involvement and distance. That is, the interaction and engagement process sometimes cross over.

The appraisals in the encode phase, mediated or moderated by the factors in the comparison phase determine the responses, that is, the levels of involvement with and distance towards the embodied agent. Moreover, the intention to use the agent as a tool may lead to actual use and together with involvement and distance this determines the overall satisfaction of the user with his or her agent system. Note that the curve in Figure 1 demarcates the line between covert behavior (what goes on in the user’s head) and overt behavior (what the user does). The overt behavior that I-PEFiC focused on was Use, which feeds back into the level of satisfaction with an agent but certainly other actions are operative in this domain as well.

From Covert to Overt: Decision Appraisals

To interact with users emotionally, robots should not only make ‘use’ of their users to achieve certain robot goals but

also be capable of selecting a situation that suits a socially desirable state. If in a meeting, the robot feels uncomfortable, it should be capable of walking away, change the subject, or turn to another conversation partner (cf. Gross, 2001). In our interpretation of Gratch (2000), appraisal of plans in a model of (emotion-based) decision-making is the outcome of a subconscious evaluation of certain options that someone has to choose from. People use these appraisals to decide which option ‘feels’ best.




Figure 2: Response phase of I-PEFiC extended with Decision Appraisals for overt actions such as Situation Selection

With regard to ‘Use’ (Figure 2), the Decision Appraisal is based on Intentions to Use (or not), which is a reasonable indicator of actual use (Davis, 1989; Venkatesh et al., 2003). With respect to ‘Situation Selection,’ the level of Satisfaction could be decisive. If Satisfaction stays below threshold, for example, because Distance remains too high, it is time for the robot to take action and turn to another conversation partner. In this paper and in its formal simulation (Bosse et al., 2008), we focus on the Decision Appraisal of Situation Selection.

Perceiving and Experiencing Interactive Users

If robots are to show some character, they should deem certain features of the user as good and/or bad, beautiful and/or ugly, aid and/or obstacle, etc. For the criteria the robot uses to judge the good or bad intent of a user, for example, one could look at the ethical reasoner THUNDER (Reeves 1991), Colby’s (1981) detection of malevolence, certain deontic logics systems, or even Asimov’s Laws. The criteria a robot uses to estimate the aesthetics of a user’s features (beautiful and/or ugly) could follow the norms for virtual plastic surgery (cf. Rabi and Aarabi 2006). The robot, then, should evaluate each feature in the user database on its relevance for the robot’s goals and concerns (cf. Frijda and Swagerman 1987; Elliot 1993). The robot’s concerns are implemented by the designer and could incorporate, for instance, the objective to recharge before

energy is down (“Hungry!”), a requirement for timely maintenance and overhaul (“Get me a doctor”), and the avoidance of memory overload (“Headache”). A user who wants to repair his robot may be seen as of good intent (“Good user”) but his clumsiness may flaw the software (“Bad performance”). Even one single feature may give rise to mixed emotions. If processing speed is a robot’s concern, it is happy to receive 10 extra megabytes of RAM. If this also means that time and again the user opens all available multimedia software all at the same time, the virtual memory runs low again. Then the initial joy about the new working memory (“Now I am a more reliable robot”) may be accompanied and compensated by disappointment (negative outcome expectancies) (“It’s of no use anyways”). Thus, a feature (“User gives me extra RAM”) may be involving for the robot (“User nice”) but at the same time distancing (“Extra RAM means extra programs on the system tray”).

This brings us to two challenges in the modeling of robot attitudes to the user. The robot may have certain features of the user participate in multiple sets (partly good, partly bad). That means that these sets are partially dependent or put differently, only relatively independent. Involvement and distance are estimated in contrast to one another (Konijn and Hoorn 2005; Van Vugt, Hoorn et al. 2006), which means that the robot should calculate a ratio between the feature sets that contribute to involvement and those sets that contribute to distance. However, there are serious theoretical and computational problems in classic accounts of set theory when one tries to relate a number of feature sets in this way (Hoorn and Konijn 2003).

Another point is that evaluative judgments are liable to the order of processing the factors or factor levels depicted in Figure 1. Put simply, preceding judgments affect following judgments. For instance, when we had an embodied agent judged for its looks, the better-looking agent raised higher expectations about its performance than the worse looking agent although the affordances of both agents were of equal quality (Van Vugt, Hoorn et al. 2006). In view of its aesthetics, the users probably regarded the affordances of the good-looking agent of better quality than of the ugly agent.

In our empirical work, order effects were also established for the involvement-distance trade-off while activating peer-group norms prior to individual norms (or vice versa). For instance, users felt little involvement with the Bonzi Buddy character when they judged him prior to the activation of the norms of their peers. When their peer group norms were activated first, however, individual involvement with Bonzi Buddy increased (Hoorn and Van Vugt 2006).

Thus, a robot should be made sensible to the context of its own judgment formation. If the user is of good intent, poor affordances are less destructive to sympathetic judgments than if the user willfully installs malicious software – although the damage may be the same.

In sum,

- Certain features participate in multiple sets (“So ugly that it is beautiful again”)
- Involvement and distance are co-occurring and partially related (it is a trade-off)
- The engagement and the interaction process affect each other
 - Judgments that are relative to one another (e.g., the level of involvement affects distance and vice versa) are misrepresented by the division of (sub) sets (Hoorn and Konijn 2003)
 - Prior evaluations affect subsequent evaluations. The order of processing information contributes to the outcome of a trade-off
 - The level of satisfaction determines whether the robot changes its situation or not (Decision Appraisal)

In view of this bulleted list, I will show how the utilization of fuzzy sets helps to calculate the involvement-distance trade-off that the robot ‘experiences’ with its user (Figure 1). This process is a multi-factorial hypercube in which the value of the outcome of the trade-off contributes to the overall level of Satisfaction of the robot with its user. Based on this level, a Decision Appraisal determines whether the robot continues the interaction with its user.

Fuzzy Involvement-Distance Trade-off

Robots should start their judgment formation by perceiving and generating features for the user and for the robot’s ‘self.’ The robot should activate a feature set X for the user and a set Y for robot’s self, including situational context and stored situations (e.g., Elliot 1993: Affective Reasoner; O’Rourke and Ortony 1994: AbMal; Rollenhagen and Dalkvist 1989: SIES). A user, then, is a database of semantically associated features that are registered online (e.g., click logs, Web cam) or that are inputted by human hand. Features are conceived of in the broadest sense; they may be descriptive (e.g., nose, hair, clothes), figurative (e.g., symbols, metaphors), relations (e.g., analogies, inferences), or behaviors.

Each feature not only carries a percentage of membership for each fuzzy sub set in which it participates (good, ugly, realistic, etc.); it also has a fuzzy membership function for its contribution to involvement (\tilde{I}) and to distance (\tilde{D}). This way, judgments of involvement and distance can concern a minor detail (e.g., one feature: “User, don’t pull the plug like that”) or can form complete generalizations across feature sets or factors (“Bad user”).

In earlier work we showed for human users that involvement and distance are two relatively independent measures. They operate in parallel but are yet related. To account for the trade-off between the judgments ‘involved’ and ‘at a distance,’ I propose to employ the compensatory_AND (Zimmermann 1994, p. 395) or the fuzzy_AND-operator γ (Werners 1988, p. 297; Zimmermann, 1994, p. 36, p. 396). The γ -operator accounts for the trade-off between conflicting options (e.g.,

‘involvement’ vs. ‘distance’), while compensation is allowed but never complete (Zimmermann 1994, p. 35, p. 286). This way, distance may be partly compensated by involvement and vice versa, quite like our empirical findings with humans assessing embodied agents (e.g., Van Vugt, Hoorn et al. 2006).

Hence, each feature u in the trade-off has a membership function μ in the involvement (\tilde{I}) and distance (\tilde{D}) stimulating fuzzy sets, allowing to map between the minimum and maximum degree of membership of these sets. In this case, the γ -operator indicates the “degree of nearness to the strict logical meaning of ‘and,’ and comprises different weights as well as different grades of compensation” (Zimmermann 1994, p. 398):

$$\mu_{and}(\mu_{\tilde{I}}(u), \mu_{\tilde{D}}(u)) = \gamma \cdot \min\{\mu_{\tilde{I}}(u), \mu_{\tilde{D}}(u)\} + ((1 - \gamma)(\mu_{\tilde{I}}(u) + \mu_{\tilde{D}}(u)) / n),$$

where $u \in U$, $\gamma \in [0,1]$, and n is the number of fuzzy sets for which the mean is calculated (here 2).

With Dollard and Miller (1950, chap. 22), Figure 3 assumes that the intercept of distance (\tilde{D}) is smaller than that of involvement (\tilde{I}) but that the slope of distance is steeper. The result of the trade-off decides whether and to what degree ‘involvement’ and ‘distance’ are in balance. According to threshold value, either ‘involvement’ or ‘distance’ may occur if

$$\mu_{\tilde{I}}(u) \text{ and } \mu_{\tilde{D}}(u)$$

are almost mutually exclusive (γ approaches 1), and ‘involvement and distance’ occurs if

$$\mu_{\tilde{I}}(u) \text{ and } \mu_{\tilde{D}}(u)$$

are almost fully compensatory (γ approaches 0).




Figure 3: Involvement-distance conflict of the robot with its user and the function of γ , which controls the angle in the trade-off outcome-space (gray) between the mean (solid line) and the minimum intensity (hatched) of distance (dotted) and involvement (dashed)

When a feature contributes equally to involvement (\tilde{I}) and distance (\tilde{D}), $\gamma = 1$, which means that the fuzzy_AND comes closest to the strict logical meaning of ‘and.’ When a feature contributes to only one set, $\gamma = 0$, which means that fuzzy_AND equals the mean $((0 + 1) / 2 = .5)$. In the trade-off outcome-space (gray areas of Figure 3) between mean and minimum intensity (of \tilde{I} or \tilde{D}), the value of γ that is chosen between 0 and 1 controls where the feature is placed and should be settled empirically.

In using γ , the robot can calculate the trade-off between the fuzzy sets that contribute to \tilde{I} and \tilde{D} for every combination of features, factor levels, or factors (ethics, aesthetics, epistemics, affordance, similarity, relevance, and valence).

Take notice that trade-offs take place for Involvement and Distance but not for ‘good’ against ‘bad,’ ‘beautiful’ against ‘ugly,’ etc. If a user is extremely bad, and not extremely good (a), Werners’ trade-off leads to the same outcome for Ethics as when a user is extremely good and not extremely bad (b), which is psychologically implausible. If, in addition, the Similarity between user and robot is based on their ethical features, situations (a) and (b) predict equal levels of robot-user similarity, which obviously is not the case. From a practical side, if a trade-off between two factor levels good vs. bad occurs, then where should the output go; to involvement or to distance? In other words, good, bad, ugly, etc. feed into involvement and/or distance but the trade-off only takes place between the latter two.

The trade-offs can be represented as an n -dimensional hypercube in which each factor level is a side or dimension (Figure 4) (Hayes and Mudge 1989). With 4 dimensions, a 4-D hypercube of values is created. In calculating the trade-off between three dimensions, their combination results into a 3-D cube, which can be reduced to a 2-D matrix, a 1-D array, and ultimately, a 0-D value that reflects the final involvement-distance judgment. This value should have predictive power for the level of satisfaction of the robot with its user. Satisfaction is then used to decide whether the robot wants to change its situation or not.

Combining the dimensions can be attained by a fuzzy_AND-operator, normal, or weighted mean. Although I expect from our empirical data that for most judgments, \tilde{I} and \tilde{D} occur in parallel, factors or factor levels can be combined into one encompassing judgment across dimensions (‘involvement’ or ‘distance’), according to threshold.

Figure 4 shows an example in which the factor levels of affordance, ethics, and valence form a $2 \times 2 \times 2$ hypercube of values ($n = 3$). Starting in the lower left corner, the values for *obstacle* and *aid* are combined into $(\tilde{I} \& \tilde{D})_{Aff}$ that is solely driven by affordance (the &-sign represents a combination operator such as γ). This output goes into the black node. Then *bad* is combined with *good* ($n = 2$), which updates the value in the black node: $(\tilde{I} \& \tilde{D})_{Eth \leftarrow Aff}$ is driven by ethics as preceded by affordance. The output for combining positively and negatively valenced features ($n =$

1) updates the black node to $(\tilde{I} \& \tilde{D})_{\text{Val} \leftarrow \text{Eth} \leftarrow \text{Aff}}$, which is driven by valence as preceded by ethics as well as affordance. This value should have predictive power for the level of satisfaction upon which situation selection is based.




Figure 4: Hypercube of values for affordance, ethics, and valence and their contribution to the involvement-distance trade-off. The serial order is interchangeable. The & symbol represents a combination operator of choice (e.g., γ)

Clearly, the hypercube of Figure 4 shows that there are six ways to calculate the overall judgment in the black node: Starting from the affordance nodes, the valence nodes, or the ethics nodes, which, in their turn, can be calculated in two ways. The values of the nodes for ethics, for instance, can be calculated by combining the nodes in the affordance-ethics dimension or those in the valence-ethics dimension. The different ways of calculation correspond to the order(s) in which the dimensions can be combined according to contextual rules or a given situation.

The advantage of this representation is that one hypercube contains all (combinations of) values, avoiding the calculation of variables with different numbers of dimensions. Moreover, different orders may lead to different outcomes, a situation that simulates the contextual sensitivity of human evaluative behavior. Table 1 provides a numerical illustration of this point for two 2-leveled factors (valence and affordance).

Table 1 shows that if the factor levels of valence are combined first, the resulting values for affordance are aid = .225 and obstacle = .35. The result for the involvement-distance trade-off, then, is $\tilde{I} \& \tilde{D} = .2562$ (with $\gamma = .5$). The picture changes if the factor levels of affordance are combined first. In that case, the values for valence are positive = .3 and negative = .275, so that the result for the involvement-distance trade-off is $\tilde{I} \& \tilde{D} = .2812$ ($\gamma = .5$).

The *fuzzy_AND* operator warrants that the order of input in the calculus has effect on the outcome of the involvement-distance trade-off. Yet, when all judgments run in parallel, *fuzzy_AND* is fruitful only if all dimensions are combined together. In that case, however, a conventional weighed mean also suffices. Thus, *fuzzy_AND* is a

sophisticated representation of the conflict between involvement and distance. In cases of parallel occurrence of judgments, a conventional weighed mean suffices.

Table 1: Example of values that enter an evaluative trade-off (upper panel) and how the order of calculation changes its outcomes ($\gamma = .5$) (lower panel).

	Positive valence	Negative valence
Aid	.6	.1
Obstacle	.2	.8

	Positive valence	Negative valence	$\tilde{I} \& \tilde{D}$
<i>Fuzzy_AND</i>	.3	.275	$\rightarrow .2812$
Aid	.225	.6	
Obstacle	.35	.8	
$\tilde{I} \& \tilde{D}$	\downarrow	.2562	

Discussion

To create an affect mechanism in robots, many AI systems focused on emotion description and generation through situational reasoning (e.g., ACRES - Frijda and Swagerman 1987; Affective Reasoner - Elliot 1993; SIES Dalkvist 1989; AbMal - O'Rorke and Ortony 1994; THUNDER - Reeves 1991). These systems, moreover, were focused on representing a limited number of emotions (e.g., fear, anger, and mistrust in PARRY- Cerf 1973; Colby 1981) or on the biological backgrounds (e.g., ACRES, also Thayer and Lane 2000). The current paper adds to these approaches as it is concerned with balancing approach and avoidance tendencies that co-occur – not the emotions or the affective reasoning per se. In this sense, the work could be used as the input side of models of affect regulation and coping (see Bosse, et al. 2008). It does so by accounting for an existing and well-tested theory (Interactive PEFiC) by means of a fuzzy sets approach.

In assuming that factors and factor levels can be represented by fuzzy sets, features of humans and robots can participate in multiple fuzzy sets that, from an affective point of view, are contradictory (e.g., a skilled user that raises admiration as well as mistrust). The *Fuzzy_AND*-operator γ (Werners 1988) is capable of balancing involving and distancing tendencies without violating Kolmogorov's axioms of probability (Hoorn and Konijn 2003). This is also important when sub sets should be regarded as proportional to one another as is the case with (dis)similarity ratings in which the intersection is proportionally related to the distinctive sets. Moreover, *Fuzzy_AND* allows for effects of prior evaluations on future evaluations, which is in line with the literature on, for example, judgmental biases through framing strategies.

Although the fuzzy hypercube representation of the I-PEFiC model seems promising, an open issue still is its validation with simulation data as compared with predicted overt behaviors, such as situation selection (Bosse et al., 2008). Patterns that are produced by the hypercube should be in line with I-PEFiC's explanations of human affective trade-off processes; a fruitful area of forthcoming research on our futuristic robot friends.

References

- Bosse, T., Pontier, M., and Treur, J. (2007). A dynamical system modelling approach to Gross' model of emotion regulation. In: Lewis, R.L., Polk, T.A., Laird, J.E. (Eds.), *Proc. 8th Int. Conf. on Cognitive Modeling, ICCM'07*. Taylor and Francis, pp. 187-192.
- Bosse, T., Hoorn, J.F., Pontier, M.A., and Siddiqui, G.F. (2008). *A Robot's Experience of Another Robot: Simulation*. In: Sloutsky, V., Love, B.C., and McRae, K. (Eds.), *Proc. 30th Int. Annual Conf. of the Cognitive Science Society, CogSci'08* (this volume).
- Cerf, V. (1973). *PARRY meets THE DOCTOR*. DATAMATION: 62-64.
- Colby, K.M. (1981). Modeling a paranoid mind. *The Behavioral and Brain Sciences 4*: 515-560.
- Davis, F. D., (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly 13* (3), 319-339.
- Dollard, J., and Miller, N.E. (1950). *Personality and Psychotherapy. An Analysis in Terms of Learning, Thinking, and Culture*. New York, NY: McGraw-Hill.
- Dyer, M.G. (1983). *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. MIT Press.
- Elliott, C. (1993). Using the affective reasoner to support social simulations. In *Proc. 13th Int. Joint Conf. Artificial Intelligence*, 194-200.
- Frijda, N., and Swagerman, J. (1987). Can computers feel? Theory and design of an emotional system. *Cognition and Emotion 1*(3): 235-257.
- Gratch, J. (2000). Modeling the Interplay Between Emotion and Decision-Making, *Proc. 9th Conf. on Computer Generated Forces and Behavioral Representation*.
- Gross, J.J. (2001). Emotion regulation in adulthood: Timing is everything. *Current Directions in Psychological Science, 10*, 6, 214-219.
- Hayes, J.P., and Mudge, T. (1989). Hypercube supercomputers. *Proc. of the IEEE, 77, 12*, 1829-1841.
- Hoorn, J.F., and Konijn, E.A. (2003). Perceiving and experiencing fictional characters: An integrative account. *Japanese Psychological Research 45*(4): 250-268.
- Hoorn, J.F., and Van Vugt, H.C. (2006). The role of social norm in user-engagement and appreciation of the Web interface agent Bonzi Buddy. *Lecture Notes in Artificial Intelligence (LNAI) 4133*: 456.
- Konijn, E.A., and Bushman, B.J. (2007). World leaders as movie characters? Perceptions of G.W. Bush, T. Blair, O. Bin Laden, and S. Hussein at the eve of Gulf War II. *Media Psychology, 9*: 157-177.
- Konijn, E.A., and Hoorn, J.F. (2005). Some like it bad. Testing a model for perceiving and experiencing fictional characters. *Media Psychology 7*(2): 107-144.
- Marsella, S., and Gratch, J. (2003). Modeling coping behavior in virtual humans: Don't worry, be happy. In *Proc. 2nd Int. Joint Conf. on Autonomous Agents and Multiagent Systems, AAMAS'03*. ACM Press, pp. 313-320.
- O'Rorke, P., and Ortony, A. (1994). Explaining emotions. *Cognitive Science 18*: 283-323.
- Rabi, A., and Aarabi, P. (2006). Face fusion: An automatic method for virtual plastic surgery. *Proc. Int. Conf. Information Fusion (Fusion '06)*. Available at http://www.apl.utoronto.ca/publication/i/rabi_fusion06.pdf
- Reeves, J.F. (1991). *Computational morality: A process model of belief conflict and resolution for story understanding*. Technical Report UCLA-AI-91-05, UCLA Artificial Intelligence Laboratory.
- Rollenhagen, C., and Dalkvist, J. (1989). *Cognitive contents in emotions: A content analysis of retrospective reports of emotional situations*. Technical report, Department of Psychology, University of Stockholm.
- Thayer, J.F., and Lane, R.D. (2000). A model of neurovisceral integration in emotion regulation and dysregulation. *J. of Affective Disorders 61*: 201-216.
- Van Vugt, H.C., Hoorn, J.F., Konijn, E.A., and De Bie Dimitriadou, A. (2006). Affective affordances: Improving interface character engagement through interaction. *Int. Journal of Human-Computer Studies 64*(9): 874-888.
- Van Vugt, H.C., Konijn, E.A., Hoorn, J.F., Eliëns, A., and Keur, I. 2007. Realism is not all! User engagement with task-related interface characters. *Interacting with Computers 19*(2) 267-280.
- Van Vugt, H.C., Konijn, E.A., Hoorn, J.F., Veldhuis, J. (2006). Why fat interface characters are better e-health advisors. *Lecture Notes in Artificial Intelligence (LNAI) 4133*: 1-13.
- Venkatesh, V.G.M.M., Davis, G.B., Davis, F.D., (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly 27* (3), 425-478.
- Werners, B.M. (1988). Aggregation models in mathematical programming. *Mitra*: 295-319.
- Zimmermann, H.J. (1994). *Fuzzy Set Theory - and its Applications*. Boston, MA: Kluwer-Nijhoff.

Acknowledgements

I am grateful to Elly A. Konijn, Stefan L. Frank, Matthijs Pontier, and Ghazanfar F. Siddiqui for their comments on earlier drafts of this paper.