Computational Modeling of Mood and Feeling from Emotion

Robert P. Marinier (rmarinie@umich.edu)
Department of Electrical Engineering and Computer Science, 2260 Hayward Street
Ann Arbor, MI 48109 USA

John E. Laird (laird@umich.edu)
Department of Electrical Engineering and Computer Science, 2260 Hayward Street
Ann Arbor, MI 48109 USA

Abstract

We propose requirements for computational models that combine mood and emotion to create feeling and feeling intensity within an appraisal theory framework. Meeting these requirements involves solving many representational issues, such as determining the range of values for each appraisal dimension. We present functions that have been realized in a computational model that fulfill these requirements.

Keywords: Emotion; mood; feeling; intensity; appraisal theories; computational models.

Introduction

In appraisal theories of emotion, an agent evaluates a situation with respect to its goals, and that evaluation leads to an emotion (for an overview, see Roseman & Smith, 2001). The evaluation is hypothesized to take place along multiple dimensions, such as goal relevance (is this situation important to my goals?), goal conduciveness (is this situation good or bad for my goals?), causality (who caused the situation?), control (can I change the situation?) and so on. Appraisal theories differ primarily in which evaluation dimensions they propose, although there is significant overlap between theories.

Existing computational models of appraisal theories of emotion emphasize the calculation of appraisals. However, comprehensive models of emotion require many more components that are "downstream" from the calculation of appraisals, and existing models of these components are either simplified or missing altogether. This paper attempts to fill in two important components: the interaction between mood, emotion, and feeling, and the calculation of feeling intensity. Most existing appraisal models (computational or otherwise) do not make clear distinctions between emotion (the immediate result of the situation), mood (a summary of recent situations), and feelings (what the agent actually perceives). Even when they do, they do not provide detailed computational theories of how these components interact. Moreover, existing models of feeling intensity fail to take into account the complexities of complete appraisal models.

Since relevant human data is scarce, in this paper we propose a set of criteria to guide the creation of these models, partially based on existing attempts (e.g. Neil Reilly, 2006). These criteria incorporate both qualitative behavior results as well as computational requirements. We present a computational model based on these criteria.

Background

We are working with Scherer's (2001) appraisal theory, which is distinguished by its completeness. Most appraisal theories contain six to eight appraisal dimensions, but claim that there are probably more to be found. Scherer's sixteen dimensions (of which we currently model eleven; see Table 1) provide a richer platform for exploring computational models. Existing computational models tend to be based on a small number of dimensions (e.g. Gratch & Marsella (2004) has five dimensions).

Table 1: Modeled appraisal dimensions with ranges.

Suddenness [0,1]	Unpredictability [0,1]
Goal Relevance [0,1]	Discrepancy from Expectation [0,1]
Intrinsic Pleasantness [-1,1]	Outcome Probability [0,1]
Conduciveness [-1,1]	Causal Agent [self, other, nature]
Control [-1,1]	Causal Motive [intentional, negligence, chance]
Power [-1,1]	

The collection of current values for each appraisal dimension is called an *appraisal frame* (Gratch & Marsella, 2004). Since appraisal theories hypothesize a fixed mapping from appraisal values to emotions, an appraisal frame essentially defines the current emotion.

In some theories of emotion (e.g. James, 1890; Damasio, 1994), there is a distinction between the emotion as generated by cognition and the agent's perception of that emotion. In these theories, the emotion influences physiology, and the agent perceives these changes, which are called the agent's feelings. We theorize that feelings are also modulated by mood, which acts as a memory of recent emotions. Thus, feelings are determined by the agent's appraisal of the current situation (emotion) as well as recent events (mood).

Since emotion, mood and feeling interact; we chose to use the same representation for all three: an appraisal frame. That is, the agent generates an appraisal frame (its emotion), which interacts with another appraisal frame (its mood) to generate its perceived appraisal frame (its feeling) as shown in Figure 1. We label these appraisal frames because their structure is a collection of appraisal dimensions, not because the agent, via an appraisal process, directly sets the contents of them (the agent only directly sets the contents of the emotion frame). When necessary, we will distinguish among them by referring to the emotion frame, the mood frame, or the feeling frame. This contrasts with most other theories in which emotion, mood, and feeling are not distinguished by separate structures.

The representation of feeling as an appraisal frame is most likely a simplification because feelings are the perception of the physiological reactions to the combination of mood and emotion. Nonetheless, whatever structure is produced will be the basis for intensity, and the analysis we develop below should apply to that structure. Non-computational ideas regarding such structure have been proposed in, e.g., Lambie & Marcel 2002.

Given a feeling frame, we can also calculate the intensity of that feeling. Intensity gives the agent an indication of how important a feeling is, and thus helps determine to what extent the feeling should influence behavior.

Within this theory, two questions must be answered:

1. How does the emotion appraisal frame combine with mood to produce the feeling appraisal frame?
2. Given a feeling appraisal frame, how is the intensity of that feeling calculated?

In the creation of the theory, we have tried to rely on existing work and data. However, for the level of detail required in a computational model, such prior work is limited. Thus, we are faced with numerous decisions where there is little or no guidance from the literature. In these cases, we have tried to choose the simplest alternative (recognizing that "simplest" can be a subjective concept); that is, we are applying Occam's Razor. Our long-term strategy is to see where these simple assumptions fall short, which will indicate where additional complexity is required. Thus, the theory we present is likely oversimplified, but it provides a starting point for future work.

Mood

Emotion is based on the agent's appraisal of the current situation independent of any historical context. To avoid wild fluctuations in feeling, historical context is necessary, but this context should be biased toward those evaluations that are temporally relevant. Mood provides this historical context of recent emotions. Thus, we make the simple assumption that the mood combines with the current emotion to form the feeling that the agent perceives; more complex models are possible, of course.

To the extent that mood is physiological in nature, there are some phenomena that can guide our model. In the undoing effect (Fredrickson & Levenson 1998), physiological changes due to negative emotions return to baseline (the natural state for some positive emotions) more quickly when followed by a positive emotion. One possible interpretation of this is that the mood "chases" the emotion (i.e. the mood tries to change to match the state defined by the emotion), but will still decay on its own if left alone.

In our model, the mood starts out neutral (i.e. all zero values). To model the influence of emotion on mood, the mood "moves" towards the emotion each time step. In the current model, we have adopted a simple approach where the mood moves x% (our current experimental value is 10%) of the distance along each dimension towards the emotion in each cycle. Additionally, the system decays mood by y% (experimental value is 1%) each cycle; thus, if there were no influence of emotion, mood would eventually become neutral. This model summarized in Figure 1.

Combining Mood and Emotion to form Feeling

In general, the relationship between appraisal frames may be complex with interactions among multiple dimensions. However, we have no reason to assume this, so instead we assume that each appraisal dimension in a frame influences only the corresponding dimension in the other frame.

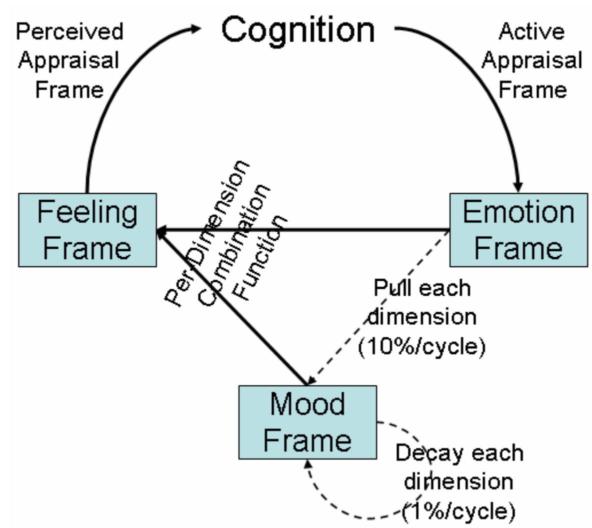

Figure 1: An emotion frame influences and combines with the mood frame to produce the feeling frame, which is perceived by the agent.

Before we can discuss how mood and emotion combine to create feelings, we must discuss the nature of the appraisal dimensions and their values that make up the frames.

Value Ranges for Appraisal Dimensions

One of the challenges in combining appraisal frames is that not all dimensions are on the same scale. Some dimensions are continuous numeric dimensions, but others are categorical (see Table 1).

For the numeric dimensions, most existing computational models use the range [0, 1] (e.g. Gratch & Marsella, 2004). The implication is that the 0 end of the range is less intense than the 1 end of the range. For some dimensions, this is true: an event with Suddenness 1 would be considered more sudden that an event with Suddenness 0. For some dimensions, though, being at the "low" end could be just as intense as being at the "high" end. For example, if I pass an exam, I will appraise this as high Conduciveness and have a strong positive feeling. However, if I fail the exam, I will appraise this as very low Conduciveness, (i.e. highly unconducive) and will experience a strong negative feeling. Thus, for these dimensions we use the range [-1, 1] – that is, values near zero (e.g. not very conducive or very unconducive) would have a low impact on feeling, but values near the extremes (e.g. very conducive or very unconducive) would have high impact on feeling. Table 1 shows the dimensions with their ranges.

For categorical dimensions, we use a numeric representation for the purposes of combination. Causal Agent and Causal Motive can each take on three values: Self, Other, Nature; and Intentional, Chance, Negligence, respectively. Our approach is to convert these categorical values into mutually exclusive features, each with its own numeric value in the range [0, 1]. Thus, the original Causal Agent feature is expanded into three features: Causal-Agent-Self, Causal-Agent-Other, and Causal-Agent-Nature. For the emotion frame, the selected value gets 1 and the others get 0. For example, if the value of Causal Agent is nature, then the dimension Causal-Agent-Nature gets a value of 1 while Causal-Agent-Self and Causal-Agent-Other get 0. The values for these dimensions are now numeric and are treated like other numeric values so that the mood tracks recent historical values for these dimensions. The feeling value is then the combination of these dimensions from the frames, just like the other dimensions. However, after combination, multiple categorical values can be non-zero, representing confusion about which is the true value. In these cases, the agent perceives the categorical value of the dimension with the highest numeric value. Thus, if Causal-Agent-Self = .4, Causal-Agent-Other = .7, and Causal-Agent-Nature = .2, the agent would perceive Causal Agent = Other.

Criteria for the Combination Function

There are many options for combining the values of mood and emotion to produce a feeling; we introduce several criteria below that such a combination function should meet. Simple combination functions such as averaging or multiplication have been shown to be inadequate, as our criteria will illustrate. Existing work (Neal Reilly 1996, 2006) has already provided some relevant criteria; however, that work has been done at the more abstract level of emotions of the same kind (e.g. Joy .3 and Joy .2). Our theory is defined at a lower level, that of individual appraisal dimensions and their "intensity" (e.g. Suddenness .3 and Suddenness .2). However, the criteria defined for these higher-level models still apply at the lower level, because the criteria are about how to combine intensities of the same kind, and have little to do with what the kinds are.

As mentioned earlier, we assume the dimensions are independent, so our combination function takes as input a particular dimension from the mood and emotion frames to produce the corresponding dimension of the feeling frame.

We begin by noting that we want to avoid a large range of inputs from mapping onto a small range of outputs because then the agent will not be able to distinguish between those inputs, and thus will not be able to form diverse responses. This criterion is subjective.

1) Distinguishability of inputs: Large input ranges should have large output ranges. Capping of extreme values may be necessary, but it should have minimal impact.

Next, we consider constraints from prior work: when combining values of the same sign, the result should be further from zero than the input with the largest magnitude, but less than or equal to the sum of the inputs (Neal Reilly 1996, 2006). The intuition is that the values should build on each other, but the combination should not be more than the parts. For example, if the mood's Suddenness value is .3 and the emotion's Suddenness value is .5, the feeling's Suddenness value should be at least .5 but no more than .8.

For values of opposite signs, the result should be closer to zero than the maximum magnitude, but be at least the sum of the inputs. Furthermore, the result should be further from zero than the sum of the results. The intuition is that the smaller value is dragging down the larger value, but the amount of the reduction should be no more than the magnitude of the smaller value. For example, if mood's Conduciveness is .3 and emotion's Conduciveness is -.5, the result should be between -.5 and -.2.

We can state the above by defining the combining function *C*, which has inputs vemotion and vmood:

2) Limited range: $C(v_{emotion}, v_{mood})$ should be between the input with the maximum magnitude and the sum of the inputs.

Another issue is that we should avoid going out of scale if possible. This can happen with middle values combined with a strict sum (e.g. .6 and .6). Values can always be capped, but capping middle values means the agent will be unable to distinguish among a large set of possible inputs, which violates our first criterion. Thus, our next criterion is that the combination should not be linear Neal Reilly (2006). While $C(.5, .5)$ should be much less than 1, $C(.1, .1)$ can be very close to .2. The intuition is that low-intensity events can result in a moderate intensity reaction, but moderate-intensity events should not result in extreme intensity reactions. That is:

3) Non-linear: For small inputs, *C* is nearly additive, but for large inputs, *C* is closer to a max. Put another way, for small values the derivative of *C* can be close to 1, but for large values, the derivative of *C* should be closer to 0.

We also identify several properties that enforce symmetry on the function. These properties do not result from any intuition or data, but rather represent reasonable first guesses given the lack of information. That is, these are default assumptions and not hard constraints. We would be satisfied with a theory that violated these criteria so long as the theory recognized the implications of the bias. For example, here may be some basis for a positivity bias (Diener & Diener 1996), but it is not clear whether such a bias belongs in the combination function or in the processes that generate the emotion frame.

4) Symmetry around 0: $C(x, 0) = C(0, x) = x$. If the mood or emotion input is 0, then the other input dominates. If they are both zero, then the result should be zero.

5) Symmetry of opposite values: $C(x, -x) = 0$. The mood and emotion can cancel each other out.

6) Symmetry of all values: $C(x, y) = C(y, x)$. The mood and emotion have equal influence on the feeling.

The Combination Function

As a starting point, we will use Neal Reilly's (2006) proposed function for combining intensity values of the same kind, and then modifying it as necessary to meet our criteria:

$$I = 0.1 \cdot \log_2 \sum_{em} 2^{10 \cdot em}$$

This function was designed to deal only with positive values. For most of those values, the function meets criterion 2 (limited range) and 3 (non-linear). The log combination ensures that the result is at least the max value, but no more than the sum. Further, the derivative of the log is near 1 for small values, but decreases for larger values. For example, $I(.1, .1) = .2$, but $I(.5, .5) = .6$.

Perhaps surprisingly, the function fails criterion 2 (limited range) at the lower extreme ($I(0, .1) = .15$), 4 (symmetry around 0; $I(0, .1) = .15$). The function does fulfill criterion 6 (symmetry of all values) for positive values. Criterion 5 (symmetry of opposite values) does not really apply since the function does not deal with negative values.

The problems with this function can be fixed. To deal with negative values (criterion 5), we introduce a Sign function and absolute values. The absolute values allow us to work with the magnitudes of the inputs, while the Sign function allows us to restore the signs that were removed by the absolute values. To do this, we break the function into two parts: the sum part and the log part. The sum part treats the exponent as a magnitude, but applies the original sign before including the value in the sum (see function below).

To center the function at 0 (criterion 4), we recognize that we need to end up taking the log of 1 (to get 0). If each input is 0, then the result of the exponent will be 1, and thus the sum part will be 2. To fix this, we subtract 1 from each magnitude of the sum (so the sum will be 0 for zero-valued inputs), and then add the Sign of the sum to the sum before taking the log (to maintain symmetry).

We originally chose b=e instead of 2 because the resulting values are less extreme near the edges of the input range, which helps meet criterion 1 (distinguishability of inputs). However, this function still fails criterion 1. The log scale of the function causes the result of an extreme input value and nearly any other input value of opposite sign to fall into a very narrow range. For example, $C(.9, -.1) = .89998$, whereas $C(.9, -.5) = .89816$ – nearly the same value. To fix this, we introduce a piecewise function that varies b depending on the inputs. If the signs are equal, then b=e. If the signs are opposite, then b=1.1, which spreads out the resulting values. For example, $C(.9, -.1) = -.85453$, whereas $C(.9, -.5) = .58561$.

The final function is shown below. A complete example showing mood and emotion frames combining to form a feeling frame is shown in Table 2.

$$C(v_{mood}, v_{emotion}) = 0.1 \cdot Sign(S) \cdot \log_b (|S + Sign(S)|)$$

$$\text{where } S = \sum_{v = v_{mood}, v_{emotion}} (Sign(v) \cdot (b^{10 \cdot |v|} - 1)$$

$$\text{and } Sign(v) = \begin{cases} 1 \text{ if } v \geq 0 \\ -1 \text{ else} \end{cases}$$

$$\text{and } b = \begin{cases} e \text{ if } Sign(v_{mood}) = Sign(v_{emotion}) \\ 1.1 \text{ else} \end{cases}$$

$$\text{If } C(v_{mood}, v_{emotion}) > 1 \text{ then } C(v_{mood}, v_{emotion}) = 1$$

$$\text{If } C(v_{mood}, v_{emotion}) < -1 \text{ then } C(v_{mood}, v_{emotion}) = -1$$

Calculating Intensity

Feeling intensity is important because it gives the agent a summary of how important this feeling is, and thus to what degree it should influence behavior. Feelings with low intensity are likely to be caused by less important events than feelings with high intensity.

We will combine the dimensions of the feeling frame to form a single intensity value. We cannot use the same "intensity" combination function that we derived above because, in this case, the input values are of different kinds, (i.e. Suddenness, Goal Relevance, Control, etc; see Table 1) and there are some different interactions among kinds that we want to capture. One of the challenges of computing a single numeric intensity from different kinds is that, as we mentioned earlier, not all dimensions are on the same scale. Additionally, some dimensions are continuous numeric dimensions and others are categorical (see Table 1). Once again, these challenges arise because of our assumption that feeling has an appraisal frame structure, and that structure allows for negative and categorical values. As noted above, the numeric representation for categorical values does not necessarily imply an intensity for each possible value, but rather a confusion over which value is correct. It is not clear that clarity should play a role in intensity, so for now, we ignore categorical dimensions in our current model of intensity.

Criteria

There are many ways to produce an intensity value from a frame, and although there is little theory or empirical evidence to guide us, there are three general criteria for an intensity function:

1) Limited range: Intensity should map onto [0,1]. This is common to most existing theories.

2) No dominant appraisal: No single appraisal value should dominate the intensity function; each should contribute to the result but no single value should determine the result. Furthermore, as intensity is unsigned, for intensity purposes, we are concerned with the magnitude of the appraisals and

not their signs. Previous intensity functions have tried multiplying the values together – for example, Gratch & Marsella (2004) multiply Desirability and Likelihood together. While this method might work well for simple appraisal models with small numbers of appraisal dimensions, it does not extend well to the large number of appraisal dimensions we have in our model. For example, there are many events that we would expect to be neither Intrinsically Pleasant nor Intrinsically Unpleasant, thus being 0 along this dimension. Using a simple multiplication model would lead to zero intensity no matter what the values of the other dimensions are.

3) Realization principle: Expected events should be less intense than unexpected events (Neal Reilly 2006). If an event is expected, then the intensity of that event when it occurs should be reduced. In Gratch & Marsella's (2004) function, the intensity is maximized when the event occurs (i.e. when Likelihood is 1). Neal Reilly's function for some emotions (he has variations for each) calculates intensity based on the change in Likelihood: I = *Desirability* x Δ*Likelihood*. That is, when the agent first realizes that an event is likely to occur (and thus it was unexpected), the intensity will be higher than when the (now expected) event actually occurs.

The Intensity Function

To construct our intensity function, we begin with the last criterion. In our model, Likelihood most closely maps onto Outcome Probability (OP). However, rather than computing the change in Outcome Probability, we instead rely on the value of Discrepancy from Expectation (DE). These dimensions together imply a change in likelihood. If outcome probability and discrepancy from expectation are both high, then the intensity should be high since expected outcomes were not met. Similarly, if outcome probability and discrepancy are both low, then intensity should be high again, because something that was considered unlikely actually happened. If outcome probability and discrepancy have opposite values, then intensity should be low. (because either a likely event occurred or an unlikely event did not occur). This leads us to the first part of our function:

$$I = (1 - OP)(1 - DE) + (OP \cdot DE)\ldots$$

This function has low values when Outcome Probability and Discrepancy are at opposite ends of their ranges (because each product will be a combination of a low and high value), and high values when they are at the same end (because one of the products will be the combination of two high values). For example, if Outcome Probability = .9 and Discrepancy = .1, then I = .18. Similarly, if Outcome Probability = .9 and Discrepancy = .9, then I = .82.

To meet the first and second criteria, we notice that a simple function that allows each dimension to contribute is an average. A sum will not work because it would exceed the legal range as defined by the first criterion. To get magnitudes, we take the absolute values of those appraisals that can be negative. In general, one might expect that some dimensions contribute more than others do in the intensity calculation. In the absence of supporting data, however, we will assume all dimensions contribute equally. Thus, we normalize the dimensions with a [-1, 1] range.

We must now combine these two parts. Two obvious candidates are multiplication and averaging. We have chosen multiplication because it is consistent with the current approaches described earlier. Thus, we have:

$$I = [(1 - OP)(1 - DE) + (OP \cdot DE)] \cdot \frac{S + UP + \frac{|IP|}{2} + GR + Cond + \frac{|Ctrl|}{2} + \frac{|P|}{2}}{7}$$

Discussion and Results

The feeling intensity function is biased so that some classes of feelings are inherently more (or less) intense than others. For example, the class of feelings that Scherer's theory would label as Boredom/Indifference is composed of low values for most dimensions combined with high outcome probability and low discrepancy, resulting in low intensity. On the other hand, Scherer's Rage/Hot Anger feelings are composed of mostly high values, with high outcome probability and high discrepancy, resulting in high intensity. This is congruent with many circumplex models of emotion (see Yik et al 1999 for an overview), which also propose different intensities for different emotions. This may suggest a bridge between circumplex models and appraisal models.

The combination and intensity functions we presented can sometimes lead to unexpected results. Even though the combination function has a building effect (i.e. if the inputs have the same sign, the magnitude of the result will be at least as large as the magnitude of the largest input), the intensity of the feeling will not always be higher as a result. Given the way Outcome Probability and Discrepancy from Expectation influence intensity, even if both of those values go up, the intensity may actually go down. For example, suppose the Discrepancy and Outcome Probability for the feeling are both .1 (and assume all other dimensions were 1.0). This would lead to an intensity of .82. However, if both of these dimensions then increased to .2, the intensity would fall to .68.

Given the lack of relevant data on which to base our theory, our results are informal. First, we give a complete example showing the output of the combination and intensity functions and discuss its consequences. Then we show actual feeling intensity data from a complete system implemented in the Soar cognitive architecture that demonstrates the realization principle.

Table 2 shows a complete example of mood and emotion frames combining to create a feeling frame, along with the intensity of each frame. While we only discussed intensity for the feeling frame, it can be useful to apply the function to the other frames to aid our understanding of the system.

Figure 2 shows feeling intensity data excerpted from a trace of an agent implemented using the model we have described. As the figure shows, feeling intensity is maximized when the agent first realizes that it will achieve its goal, and is less when the agent actually achieves the goal. This is because going into the state where the realization occurs, the agent has a prediction which assumes that the goal completion is not imminent with some moderate probability. The realization that goal completion is indeed imminent violates this expectation. Thus, outcome probability was at least moderate, and discrepancy from expectation was high, leading to a higher intensity (assuming no major changes in the other appraisals). Following this, the agent now predicts that the goal will be accomplished with high probability, so when it is in fact accomplished, outcome probability is high and discrepancy is low, causing the intensity to be lower.

Table 2: An example combination of a mood and emotion frame to form a feeling frame. Approximate linguistic labels provided based on Scherer's (2001) modal emotions.

	Mood	Emotion	Feeling
Suddenness [0,1]	.235	0	.235
Unpredictability [0,1]	.400	.250	.419
Intrinsic-pleasantness [-1,1]	-.235	0	-.235
Goal-relevance [0,1]	.222	.750	.750
Causal-agent (self) [0,1]	0	0	0
Causal-agent (other) [0,1]	0	0	0
Causal-agent (nature) [0,1]	.660	1	1
Causal-motive (intentional) [0,1]	0	0	0
Causal-motive (chance) [0,1]	.660	1	1
Causal-motive (negligence) [0,1]	0	0	0
Outcome-probability [0,1]	.516	.750	.759
Discrepancy [0,1]	.326	.250	.362
Conduciveness [-1,1]	-.269	.500	.290
Control [-1,1]	-.141	.500	.402
Power [-1,1]	-.141	.500	.402
Label	anx-wor	ela-joy	ela-joy
Intensity	.088	.094	.127

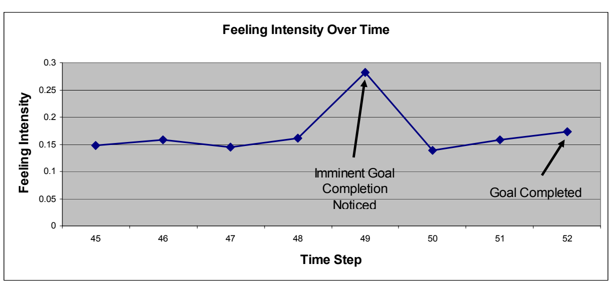

Figure 2: Feeling intensity is maximized when the agent realizes the goal will be completed, as opposed to when it actually completes.

Conclusion and Future Work

In conclusion, most existing computational models of emotion lack principled models of feeling intensity and how feeling arises from emotion and mood. Our contributions are three-fold. First, we proposed a concrete distinction between emotion, mood and feeling. This included a common representation for them (appraisal frames) and the possible range of values allowed for each appraisal dimension. Second, we listed criteria that models of feeling intensity and mood-emotion combination functions should fulfill, building on earlier criteria established by Neal Reilly (1996, 2006). Finally, we proposed specific mood-emotion combination and feeling intensity functions that fulfill those criteria, implementing them within an existing cognitive architecture.

Future work remains in discovering more criteria and alternative functions and in demonstrating that a model that fulfills these criteria has a functional advantage. Matching human data remains a holy grail of sorts for computational models of appraisal theories, and we will incorporate such data as the field makes progress. We also have a complete agent model, of which this work is only a part. We plan to demonstrate behavioral benefits of these theories for such a model in the near future.

Acknowledgments

The authors acknowledge the funding support of the DARPA "Biologically Inspired Cognitive Architecture" program under the Air Force Research Laboratory "Extending the Soar Cognitive Architecture" project award number FA8650-05-C-7253.

References

Damasio, A. (1994). Descartes' error: Emotion, reason, and the human brain. New York: Avon Books.

Diener, E. & Diener, C. (1996). Most people are happy. *Psychological Science*, Volume 7 Issue 3, 181-185.

Fredrickson, B. & Levenson, R. (1998). Positive emotions speed recovery from cardiovascular sequelae of negative emotions. *Cognition and Emotion*, 12(2):191-220.

Gratch, J. & Marsella, S. (2004). A Domain-independent framework for modeling emotion, *Journal of Cognitive Systems Research*, Volume 5, Issue 4, pp. 269-306.

James, W. (1884). What is an emotion? *Mind*, *9*, 188-205.

Lambie, J. A. & Marcel, A. J. (2002) Consciousness and the varieties of emotion experience: A theoretical framework. Psychological Review, Volume 109, No. 2, 219-259.

Neal Reilly, W. S. (1996). *Believable Social and Emotional Agents*. (Tech. Rep. CMU-CS-96-138). Pittsburgh, PA: Carnegie Mellon University.

Neal Reilly, W. S. (2006). Modeling what happens between emotional antecedents and emotional consequents. *Proceedings of the Eighteenth European Meeting on Cybernetics and Systems Research* (pp. 607-612). Vienna, Austria: Austrian Society for Cybernetic Studies.

Roseman, I. & Smith, C. A. (2001). Appraisal theory: Overview, Assumptions, Varieties, Controversies. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.) *Appraisal processes in emotion: Theory, Methods, Research* (pp. 3-19). New York and Oxford: Oxford University Press.

Scherer, K. R. (2001). Appraisal considered as a process of multi-level sequential checking. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.) *Appraisal processes in emotion: Theory, Methods, Research* (pp. 92-120). New York and Oxford: Oxford University Press.

Yik, M. S. M., Russell, J. A., & Barrett, L. F. (1999). Integrating four structures of current mood into a circumplex: Integration and beyond. Journal of Personality and Social Psychology, 77, 600-619.