

EMA: A computational model of appraisal dynamics

Stacy Marsella

University of Southern California
marsella@isi.usc.edu

Jonathan Gratch

University of Southern California
gratch@ict.usc.edu

Abstract

A computational model of emotion must explain both the rapid dynamics of some emotional reactions as well as the slower responses that follow deliberation. This is often addressed by positing multiple appraisal processes such as fast pattern directed vs. slower deliberative appraisals. In our view, this confuses appraisal with inference. Rather, we argue for a single and automatic appraisal process that operates over a person's interpretation of their relationship to the environment. Dynamics arise from perceptual and inferential processes operating on this interpretation (including deliberative and reactive processes). We illustrate this perspective through the computational modeling of a naturalistic emotional situation.

1 Introduction

Change is an inherent quality of emotion. Aroused by an unpleasant event, a person might explode into anger, fume at a slow boil or collapse into a depression. Once aroused, emotions influence our actions and judgments concerning the event, altering what Lazarus [1] calls the person-environment relationship. Changes to this relationship can induce new emotional responses, resulting in a cycle of change in the person's relation to the environment. These changes can be rapid, on the order of milliseconds, or unfold over days and weeks. In short, emotions are inherently dynamic, linked to both the world's dynamics and the dynamics of the individual's physiological, cognitive and behavioral processes.

A key challenge for any theory of emotion is to capture this dynamic emotional process. Over the last 50 years appraisal theories have been established as a leading theory of emotion. These theories argue that emotion arises from a person's interpretation of their overall relationship with the environment as characterized by a set of *appraisal variables* (e.g., is this event *desirable*, who *caused* it, what *power* do I have over its unfolding). To date, however, appraisal theories have largely focused on structural considerations (e.g., specifying the components of appraisal) [2]. Far less progress has been made in detailing the processes that underlie appraisal, with some notable exceptions. For example, Lazarus proposed a cyclical relationship between appraisal, coping and re-appraisal. At a finer grain, Scherer's sequential checking model argues for the sequencing of distinct appraisal checks (relevance, etc.) that determine the appraisal variables [3]. Smith and Kirby argues for a two process model of appraisal whereby *associative processing*, a memory-based process and *reasoning*, a slower and more deliberative process, operate in parallel [4]. Moors investigates automaticity of some appraisal processes [5]. Reisenzein proposed a model of appraisal that distinguishes between hardwired appraisal processes for some ap-

praisal dimensions and peripheral appraisals that can be deliberative or learned [2].

However, further progress is needed if appraisal theories are to provide a full account of emotional processes and their role in behavior. Specifically, theories must detail the processes by which each appraisal variable is determined, including the logical and temporal dependencies between these appraisal processes [4]. Additionally, the basic calculus of how the results of appraisal result in emotions of varying types, intensities and durations must be determined. Completing the cycle, the impact of emotions on coping responses and subsequent changes in the environment-person relationship must be detailed.

We see computational models of emotion as a powerful approach to concretizing and exploring the dynamic properties of appraisal. The construction of a computational model forces specific commitments about how the person-environment relationship is represented, how appraisals are performed on those representations, the role of perception, memory, interpretation and inference in appraisal, and the relationship between appraisals, emotions and coping responses. Often these commitments raise issues that are unforeseen at more abstract specifications of a theory. Further, once computationally realized, simulation allows the model to be systematically explored and manipulated, thereby generating predictions that can be further tested with human subjects. Indeed given the complexity of appraisal theories, exploring dynamic properties of a theory and contrasting alternative theories from a process-based perspective may arguably hinge on their computer realization.

This paper advocates a theoretical stance towards the problem of capturing emotional dynamic informed both by the appraisal theory of Smith and Lazarus [6] and our experience in realizing this theory in a computational model called EMA (**EM**otion and **Adaptation**). In our view, process theories of appraisal have con-




Figure 1

flated appraisal and inference. Rather, we argue for a single and automatic appraisal process that operates in parallel over a person's interpretation of their relationship to the environment. Dynamics arises from perceptual and inferential processes operating independently over this interpretation (including deliberative and reactive processes). We illustrate this perspective by modeling of a naturalistic emotional situation in EMA.

An example of emotion dynamics.

One approach to studying emotion dynamics is to use assessment tools that ask a subject to imagine an evolving situation and to introspect on their emotional reactions [7]. This approach has in turn been leveraged to evaluate computer-based models of emotion [8].

However, the focus in this paper is different. Instead of a slowly evolving situation, our interest here is in an evocative situation that elicits a wide array of emotional responses over a short time period. Specifically, we will analyze and model a naturalistic emotion-invoking situation recorded fortuitously during one of our lab studies. We were videotaping actors at 30 frames per second as part of a study on gestures and postures. In the midst of instructing the actors, a pigeon unexpectedly flew in window. Figure 1 reveals

the reactions of one of the two actors. Although such unexpected, uncontrolled event makes a rigorous analysis problematic, this naturalistic setting serves well to illustrate the rapid dynamics that we would want to cover in a process model of appraisal.

In the video, the actor holding the umbrella goes through a sequence of behaviors that suggest the following *interpretation*:

- surprise at an unexpected event (frame 2),
- fear (12),
- an aggressive stance of self-protection (13-23),
- relaxation (29),
- concern for others (29-42), specifically for the bird that caused the initial negative reaction and, finally,
- an active helping strategy (62-80) combined with relaxed facial features and smiling suggestive of relief.

The sequence of behaviors that suggest that interpretation is as follows. By frame 2 (F2), the actor has begun to turn and orient toward the sound of the bird. Her eyebrows rise (F3 through F4). The eyebrows return to a more neutral level and the mouth begins to open by F8. The eyebrows lower and the jaw then drops during F11 and F12. In F13, she begins to grab the umbrella at the base, move the left foot back away from bird and starts to raise arms. She raises the umbrella (F14 through F22), shifting her

weight to her right, rear foot away from the bird. Her posture and grasp of the umbrella suggests she is prepared to ward off a presumed attack of the bird by whacking it with the umbrella. She continues her backward motion. Her motions slow and by F29 her left hand starts to let go of the umbrella and move towards her mouth. The umbrella is lowered in F34 and her left hand covers her mouth by frame F42. By F62 the backward motion stops (she moves approximately 6 feet) and the left hand begins to lower from covering her mouth. By F66, the actor begins to move forward and the hand lowers sufficiently to reveal relaxed facial features. In F72 through F80, the forward motion continues, the hand forms into a stop gesture and the face appears to be smiling (laughter and utterances expressing concern for the bird are also heard).

A seemingly identical sequence of reactions is visible in the other actor: raised eyebrows, lowered eyebrows and jaw drop, followed by expressions suggesting relief/amusement and compassion. But reactions also differ, for she becomes aware of the bird later, she is closer to the threat and certain responses are not facilitated by the instrumentality of the umbrella.

This rapid transition in the actor's expressive state and behaviors lasts 2.6 seconds. The expression of raised eyebrows often associated with surprise takes on the order 30-60 milliseconds and the expression of lowered eyebrows and lowered jaw often associated with anger and responses to threat takes on the order 300 milliseconds. Overall, this suggests an interesting progression from an appraisal of unexpectedness, to an assessment of personal significance, and finally an appraisal of significance to others (cf. Scherer's sequential checking, [3]). Tightly coupled with these appraisal dynamics from threat-to-self to threat-to-other and emotion dynamics of Fear/Anger to Compassion/Relief, there is a corresponding progression of coping responses from defend/attack to help.

Several factors can help us explain these dynamics. The process of appraising the situation in terms of its unexpectedness, congruency with the actor's concerns, etc. can have its own internal dynamics grounded in such factors as inferential demands that underlie the appraisal and/or potential logical ordering relations between steps of the appraisal. As Scherer's theory of appraisal checks argues, there may be a temporal course to the appraisal process. There are also perceptual processes and inferential processes that alter the actor's interpretation of the situation appraisal. Similarly there are the processes of forming a coping response or plan to deal with the event. Finally, the situation and subsequent re-appraisals occur due to changes as the event unfolds. This unfolding occurs both due to the actor's response of "arming herself" and moving away from the event as well as other entities becoming the focus (such as the threat to the bird).

Basic Theoretical Assumptions

A central tenant in cognitive appraisal theories in general,

and Smith and Lazarus' work in particular, is that appraisal and coping center around a person's *interpretation* of their relationship with the environment. This interpretation is constructed by cognitive processes, summarized by appraisal variables and altered by coping responses. The goal of our work is to develop a process model of appraisal, realized as a computational process.

In our view, there are several key challenges for a computational model of appraisal. The model must explain both the rapid dynamics of some emotional reactions as well as the slower evolution of emotional responses that may follow some deliberation and inferences. The appraisal processes must in some way be consistent with both these reactive and deliberative responses. In addition, appraisal processes must operate over a range of phenomena spanning simple physical phenomena as well as complex social situations.

These challenges are often addressed by presuming multiple appraisal processes, for example fast pattern directed appraisal processes and slower more deliberative appraisals (e.g., [2, 4, 5]). Our theoretical stance on this issue differs considerably.

First and foremost, we clearly delineate between appraisal and inference. We argue that appraisal processes are always fast (reactive), parallel and unique in the sense that we postulate a single process for each appraisal variable. However, multiple other perceptual and cognitive processes perform inference (both fast and slow, both deliberative and reactive) over the interpretation of the person-environment relationship. As those inference processes change the interpretation, they indirectly trigger automatic reappraisal.

Thus, debates about which emotions have a cognitive or non-cognitive basis become moot. The relation between cognition and appraisal is that appraisals operate on the results of cognitive operations as well as any other operation that transforms the person's interpretation of their relationship to the environment.

A computational model of appraisal

EMA is a computational model that realizes this theoretical stance. We now sketch the basic outlines of the model (See [9] for a more complete description.) In general terms, we characterize a computational model as a process or processes operating on representations. A computational model of appraisal is a set of processes that interpret and manipulate a representation of the person-environment relationship.

In our view, a core requirement for the representation of the person-environment relation is that it support the derivation of appraisal variables. Moreover, as we argue that appraisal is fast and uniform, the representation must facilitate that assessment over the range of phenomena that induce emotional reactions.

To address those requirements, EMA uses a representation built on the causal representations developed for decision-theoretic planning augmented by explicit representation of intentions and beliefs (necessary for social attributions). The appraisal of relevance can be

expressed by a plan representation's ability to express the causal relationship between events and states. These causal representations are also critical for assessing causal attributions necessary for appraising blame or credit-worthiness. Appraisal variables of desirability and likelihood can be modeled by the decision-theoretic concepts of utility and probability. Explicit representations of intentions and beliefs are also critical for properly reasoning about causal attributions, as these involve reasoning if the causal agent intended or foresaw the consequences of their actions. The commitments to beliefs and intentions also play a role in modeling coping strategies, especially what is often called emotion-focused coping.

We call the agent's interpretation of its "agent-environment relationship" the *causal interpretation*. It provides a uniform, explicit representation of the agent's beliefs, desires, intentions, plans and probabilities that allows uniform, fast appraisal processes, regardless of differences in the underlying phenomena being appraised. In the terminology of Smith and Lazarus, the causal interpretation is a declarative representation of the current construal of the person-environment relationship.

Reactive and more deliberative processes map their results into this uniform representation. Architecturally, this is achieved by a blackboard-style model. The causal interpretation (corresponding to the agent's working memory) encodes the input, intermediate results and output of reasoning processes that mediate between the agent's goals and its physical and social environment (e.g., perception, planning, explanation, and natural language processing). At any point in time, the causal interpretation represents the agent's current view of the agent-environment relationship, which changes with further observation or inference. We treat appraisal as a set of feature detectors that map features of the causal interpretation into appraisal variables. For example, an effect that threatens a desired goal is assessed as a potential undesirable event.

Events are characterized in terms of appraisal variables via domain-independent functions that examine the syntactic structure of the causal interpretation:

- Perspective: viewpoint that the event judged
- Desirability: what is the utility (positive or negative) of the event if it comes to pass, from the perspective taken (e.g., does it causally advance or inhibit a state of some utility). The utility of a state may be intrinsic (agent X attributes utility Y to state Z) or derived (state Z is a precondition of a plan that, with some likelihood, will achieve an end with intrinsic utility).
- Likelihood: how probable is the outcome of the event, derived from the decision-theoretic plan.
- Causal attribution: who deserves credit/blame. This depends on what agent was responsible for executing the action, but may also involve considerations of intention, foreknowledge and coercion (see [10]).
- Temporal status: is this past, present, or future

- Controllability: can the outcome be altered by actions under control of the agent whose perspective is taken. This is derived by looking for actions in the causal interpretation that could establish or block some effect, and that are under the control of the agent whose perspective is being judged (i.e., agent X could execute the action).
- Changeability: can the outcome be altered by some other causal agent.

Each appraised event is mapped to an emotion instance of a type and intensity following the structural scheme proposed by Ortony et al. [11]. A simple activation-based focus of attention model computes a current emotional state based on most recently accessed emotion instances.

Another key aspect of EMA is that it includes a computational model of coping integrated with the appraisal process (according to Lazarus's theory). Coping determines how one responds to the appraised significance of events. Coping strategies are proposed to maintain desirable or overturn undesirable in-focus emotion instances. Coping strategies essentially work in the reverse direction of appraisal, identifying the precursors of emotion in the causal interpretation that should be maintained or altered (e.g., beliefs, desires, intentions and expectations). Strategies include:

- Action: select an action for execution
- Planning: form an intention to perform some act (the planner uses intentions to drive its plan generation)
- Seek instrumental support: ask someone that is in control of an outcome for help
- Procrastination: wait for an external event to change the current circumstances
- Positive reinterpretation: increase utility of positive side-effect of an act with a negative outcome
- Acceptance: drop a threatened intention
- Denial: lower the probability of a pending undesirable outcome
- Mental disengagement: lower utility of desired state
- Shift blame: shift responsibility for an action toward some other agent
- Seek/suppress information: form positive or negative intention to monitor a pending or unknown state

Strategies give input to the cognitive processes that actually execute these directives. For example, planful coping generates an intention to act, which in turn leads to the planning system to generate and execute a valid plan to accomplish this act. Alternatively, coping strategies might abandon the goal, lower the goal's importance, or re-assess who is to blame.

Not every strategy applies to a stressor (e.g., an agent cannot be problem directed if it is unaware of actions impacting the situation), but multiple strategies can apply. EMA proposes strategies in parallel but adopts them sequentially. A set of preferences resolve ties: e.g., EMA prefers problem directed strategies if control is appraised as high (take action, plan, seek information), procrastination if changeability is high, and emotion-focus strategies if control and changeability are low.

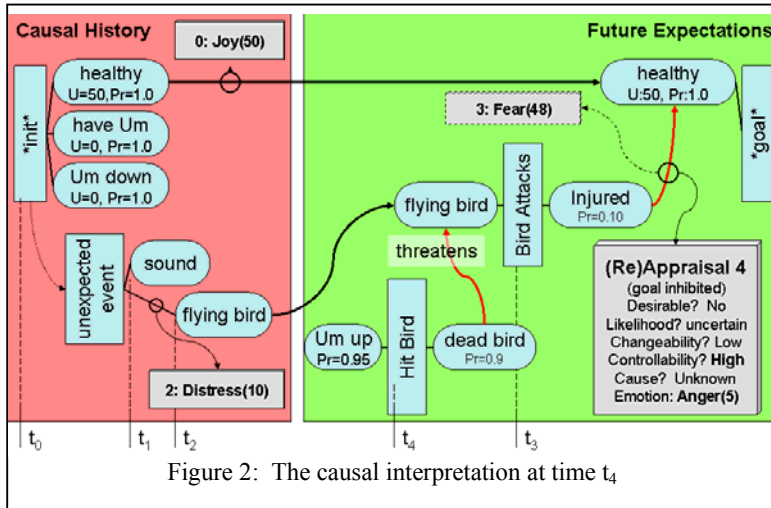


Figure 2: The causal interpretation at time t_4

In developing a computational model of coping, we have moved away from broad distinctions of problem-focused and emotion-focused strategies. Formally representing coping requires a crispness lacking from the problem-focused/emotion-focused distinction. In particular, much of what counts as problem-focused coping in the clinical literature is really inner-directed in an emotion-focused sense. For example, one might form an intention to achieve a desired state – and feel better as a consequence – without ever acting on the intention. Thus, by performing cognitive acts like planning, one can improve one's interpretation of circumstances without actually changing the physical environment.

To summarize, an agent's causal interpretation is equated with the output and intermediate results of processes that relate the agent to its physical and social environment. This configuration of beliefs, desires, plans, and intentions represents the agent's current view of the agent-environment relationship, which may subsequently change with further observation or inference. We treat appraisal as a mapping from domain-independent features of causal interpretation to individual appraisal variables. Multiple appraisals are aggregated into an overall emotional state that influences behavior. Coping directs control signals to auxiliary reasoning modules (i.e., planning, action selection, belief updates, etc.) to overturn or maintain features of the causal interpretation that lead to individual appraisals. For example, coping may resign the agent to a threat by abandoning the desire. The causal interpretation could be viewed as a representation of working memory (for those familiar with psychological theories) or as a blackboard (for those familiar with blackboard architectures).

Model of the Bird

The bird example can be readily encoded into EMA. Some of the dynamics of the situation arise from the world but others arise from the inferential processes of the model. Here we describe an encoding that produces the hypothesized emotional transitions that we derived

from our video analysis.

Our goal is not to definitively explain and reconstruct the inferences and emotions experienced by this actor, but rather to illustrate how such dynamic situations could be modeled by EMA. Many encodings are possible. We did, however, try to adopt what we felt were plausible inferences and beliefs based on post hoc analysis of the situation.

Figure 2 illustrates a snapshot of EMA's causal interpretation several steps into the situation. EMA makes discrete changes to this interpretation over time, as a result of perceptual and inferential updates. (The Soar cognitive architecture, upon which EMA is

based, assumes that updates occur once every 100 milliseconds.) The time stamps at the bottom (e.g., t_0) indicate the discrete time step in which elements are added or deleted from the causal interpretation. Vertical rectangles indicate actions including two degenerate actions **init** (whose effects generate the initial state of the simulation) and **goal** (whose preconditions correspond to the agent's goals). Ovals indicate predicates that describe the current beliefs of some feature of the world, including its likelihood of being true and its utility. For example, "have Um" indicates that the actor has an umbrella (Um) and believes this with certainty ($Pr=1.0$) and that this fact has no intrinsic value ($U=0.0$). Predicates are linked to actions (representing the actions preconditions or effects) or to other states via establishment links (i.e., this effect establishes a precondition for some other action) or threat links (i.e., this effect deletes a precondition for some other action). So, for example, the simulation began at time t_0 and the actor was healthy, holding an umbrella and the umbrella was lowered. Finally, 3D rectangles indicate appraisal frames. Each frame consists of number, indicating at which time step it was generated, a set of appraisal variables and an emotional label (due to space, the appraisal variables are omitted for all but the 4th frame).

To complete a model, one must provide several additional elements that sit outside the causal interpretation. EMA is provided with a plan library that consists of a set action definitions and set of recipes indicating how these actions could be combined. These recipes are used both for plan generation and plan recognition (as when inferring behavior of other agents based on observed actions and world states). EMA can also be provided with a set of simple stimulus-response rules that automatically trigger actions when certain world state configurations are perceived. Finally, EMA must be connected to some world simulation that defines how percepts change as a result of actions. We describe these additional elements in the course of describing the evolving situation.

At time t_0 the actor has the high utility goal of being healthy and this goal has been established in the initial state. This establishment is automatically appraised as desirable and certain, resulting in joy.¹

At time t_1 a sound is perceived, denoted in the causal interpretation by the predicate "sound". As no currently executing action predicted that a sound would occur, this is unexpected and EMA records this unknown event in the causal history. EMA supports partial observability of the world's state and this model we assume the bird cannot be seen unless the actor looks at it. To model this, we define a S-R rule that executes the "Look" action if a sound is perceived.

At time t_2 , EMA has perceived the flying bird. Typically, EMA is attached to a simulation environment that would maintain the world state and compute the observable effects of actions such as "look." For the model, we accomplish this through a domain-specific procedure that sets "flying bird" to be true 100 ms after looking at the sound. By definition in the model, "flying bird" has some negative utility for the action ($U=-10$). Thus, the effect of this action is appraised as undesirable and certain, leading to distress.

At time t_3 , EMA infers the bird will attack. This follows from a general plan recognition approach informed by the domain-specific plan recipes. This action has one precondition "bird-flying" that is established by the unexpected event in t_1 . It has one effect "injured" denoting that the actor will become unhealthy as a result of this action and this threatens the actor's goal of being healthy.² This threat relation is automatically appraised as undesirable, uncertain (as the effect may not occur), and uncontrollable (as there is no explicit way to respond to the threat, given the current causal interpretation). The result is Fear.

At time t_4 , EMA infers that there is a way to confront the threat to the actor's health: whack the bird with the umbrella. This follows from the general plan generation approach informed by the domain-specific recipes. The planner has determined that the bird's "plan" can be confronted by blocking its precondition. The planner further infers that the probability of the bird's "plan" succeeding is significantly reduced. Given that there is now an action in the causal interpretation under control of the actor and addressing the threat to health, the threat to health is automatically reappraised as controllable. The result is Anger.

We have modeled the scenario to this point, though the remainder is straightforward. The model must next infer that it must raise the umbrella to satisfy the "whack" plan and subsequently execute this initial

plan step. This plan, however, is overtaken by events as the bird becomes caught in another actor's hair, disabling the "attack" and possibly resulting in the undesirable state that the bird will be injured.

Concluding Remarks

Computational models of psychological phenomena are powerful research tools. The development of a model can help concretize theories, reveal shortcomings and can help derive predictions through simulations.

EMA provides a framework for exploring and explaining emotion dynamics. In particular, the simulation of the bird example, and the emotional dynamics it reveals, argues that the temporal characteristics of appraisal may be a by-product of other perceptual and cognitive processes. By modeling appraisal as a fast, uniform processes, EMA roots the temporal dynamics in those other processes that operate on the causal interpretation. Further, EMA's description of appraisal is economical, not requiring appeal to alternative fast and slow appraisal processes.

Acknowledgments

This work was supported by the HUMAINE Network

References

1. Lazarus, R., *Emotion and Adaptation*. 1991, NY: Oxford University Press.
2. Reisenzein, R., *Appraisal Processes Conceptualized from a Schema-Theoretic Perspective*, in *Appraisal Processes in Emotion: Theory, Methods, Research*, K.R. Scherer, A. Schorr, and T. Johnstone, Editors. 2001, Oxford Press. p.
3. Scherer, K.R., *Appraisal Considered as a Process of Multi-level Sequential Checking*, in *Appraisal Processes in Emotion: Theory, Methods, Research*, K.R. Scherer, A. Schorr, and T. Johnstone, Editors. 2001, Oxford University Press..
4. Smith, C.A. and L. Kirby, *Consequences require antecedents: Toward a process model of emotion elicitation*, in *Feeling and Thinking: The role of affect in social cognition*, J.P. Forgas, Editor. 2000, Cambridge University Press.
5. Moors, A., et al., *Unintentional processing of motivational valence*. *The Quarterly Journal of Exp. Psychology*, 2005.
6. Smith, C.A. and R. Lazarus, *Emotion and Adaptation*, in *Handbook of Personality: theory & research*, L.A. Pervin, Editor. 1990, Guilford Press: NY. p. 609-637.
7. Perrez, M. and M. Reicherts, *Stress, Coping, and Health*. 1992, Seattle, WA: Hogrefe and Huber Publishers.
8. Gratch, J. and S. Marsella, *Evaluating a computational model of emotion*. *Journal of Autonomous Agents and Multiagent Systems* (in press), 2005.
9. Gratch, J. and S. Marsella, *A domain independent framework for modeling emotion*. *Journal of Cognitive Systems Research*, 2004. 5(4): p. 269-306.
10. Mao, W. and J. Gratch. *Social Causality and Responsibility: Modeling and Evaluation*. in *International Working Conference on Intelligent Virtual Agents*. 2005. Kos, Greece.
11. Ortony, A., G. Clore, and A. Collins, *The Cognitive Structure of Emotions*. 1988: Cambridge University Press.

¹ Arguably, one has no emotions to being healthy but only reacts when there is a threat to health or they are unhealthy. EMA allows a utility distribution over predicates but we omit this distinction in our example for simplicity.

² In actuality, the effect of this action is to make "healthy" false (i.e., "injured" is shorthand for NOT(healthy)). The same holds for "dead bird" which denotes NOT(flying).