constituent elements in emotional communication: specifically, those neural entities involved in displays of emotion in senders as well as preattunements to those displays in receivers.

Neural entities cannot, however, be the critical constituent "parts" in the case of higher-level social, cognitive, and moral emotions. Biologically based emotions may be constituent elements in higher-level emotions: for example, biologically based prosocial emotions underlying attachment and love may be involved in social emotions, and biologically based emotions underlying exploration and curiosity may be involved in cognitive emotions, and both of these may be involved in moral emotions (Buck 1999). These biological systems may serve to provide the affective "fire" underlying such powerful higher-level emotions as pride, envy, jealousy, guilt, shame, awe, dread, resentment, humiliation, and gratitude (Buck 2004). However, other considerations – comparative gain and loss, fairness and equity, social norms and roles – must arguably be constituent elements in higher-level emotions.

I would like to applaud and expand upon one contention in Lewis's account: that neither fully articulated emotions nor fully articulated cognitions can exist in isolation from one another. I fully agree, and suggest that this is the case with the concept of "motivation" as well. In their fully articulated forms, emotions imply cognitions imply motives imply emotions, and so on (Buck 1985).

Emotion is remarkable in its relevance to phenomena at widely different levels of analysis, literally from atoms, molecules, and genes to social, cultural, and historical phenomena. The DS approach helps us to sort out the ingredients of this self-organizing and dynamic stew in a systematic way, by identifying the constituent elements of ingredients, specifying how they interact with other ingredients, and – critically – assigning the ingredients to the correct level of analysis.

Emotion theory is about more than affect and cognition: Taking triggers and actions into account

Charles S. Carver

Department of Psychology, University of Miami, Coral Gables, FL 33124-0751. **CCarver@miami.edu**
http://www.psy.miami.edu/faculty/ccarver/

Abstract: Understanding how emotions emerge is difficult without determining what characteristic of the trigger actually triggers them. Knowing whether emotional experiences self-stabilize is difficult without remembering what other processes are set in play as the emotion emerges. It is not clear either that positive feedback is required for the emergence of emotion or that an attractor model captures well what is happening when an emotion arises.

Lewis introduces the target article as an effort to create a bridge between emotion theory and neuroscience. The bridge is narrower than the introduction implies, however. The target article is concerned with the processes by which emotions, once triggered, emerge as full experiences, via emotional influence on cognition and cognitive influence on emotions. The article focuses less on facets of the puzzle that interest me most, however: the intrinsic meaning of emotions and their functional (action-related) properties (Carver 2001; 2004; Carver & Scheier 1998).

I care particularly about two elements the target article downplays. The first is the properties of the trigger – the event evaluated (appraised) as being important to the self. The second is the behavioral function of the emotion. The target article says little about either of these. In my view the analysis thereby loses some of its potency, because the largest share of the functionality of emotions is left out of the discussion.

What triggers emotional experiences? The target article is remarkably abstruse on this key question (sect. 3.3.1). My answer would be that emotional experiences are triggered by events that facilitate or impede either the attainment of a desired condition or the avoidance of an undesired condition. I have characterized facilitating and impeding in terms of rates of progress (Carver & Scheier 1998), though there are other ways to conceptualize them (e.g., moving in the desired direction versus its opposite). A looming object means an approaching impact (undesired). A slow driver willfully obstructing your way means violation of your entitled (desired) place in the world (positive affects are disregarded in this commentary, but see Carver 2003). I believe emotions begin in the (subcortical) registering of such changes.

This is where Lewis picks up the story. Lewis presents a dynamic model of the emotional episode (the rising of emotion from trigger to complete experience), in which positive feedback transforms a minimal reaction into a larger one, and negative feedback then limits the reaction's growth, stabilizing it at a level representing an attractor for that emotional state. I consider the two phases of this depiction in turn.

The idea that initial affect biases subsequent perceptual processing, yielding intensification of the affect (a positive feedback cycle), would creatively account for what brings the affect noticeably off baseline. There is another way to look at this flow of events, however. The event Lewis uses as his example – Mr. Smart's suddenly noticing a slow-moving car in front of him – is sufficiently abrupt in registering that it takes time for the various responses that constitute the emotional reaction to catch up with the perception (it is not just an obstacle but a suddenly appearing one). Further, if affect depends on time (as I believe it does; Carver & Scheier 1998), a 5-second impediment is a smaller provocation to Mr. Smart than a 5-minute impediment. Thus, even without biased information search and confirmation, the mere passing of time creates a steady increase in the trigger's potency, by increasing the extent of the perturbation. Although positive feedback may occur, it is not needed to account for an increase in emotion from baseline.

The second step in Lewis's analysis is that the emotional experience then stabilizes. Does it? Intuition suggests that people whose entitlement was violated can stay in a state of anger for an extended period if nothing changes the situation. Even when put aside, the anger can be re-evoked fairly readily by a reminder of the event. Moreover, the intensity of the anger seems roughly constant (again, absent situational changes), fitting the level of the perceived violation and extent of entitlement. These intuitions are consistent with an attractor model wherein the person bounces quickly to a level of anger and stays there. But this view leaves several disconcerting questions dangling.

Why would stabilizing feedback arise to keep the emotional reaction stable? Further, if the emotion is then in an attractor, why should it change? The target article leaves Mr. Smart stewing in anger. In reality, that is not how such an episode typically ends. It often ends with Mr. Smart acting to reassert his violated sense of entitlement (road rage sometimes leads to violence). Alternatively, he may do something to symbolically reinstate his entitlement (cut someone else off in traffic later, yell at his wife). Another possibility (cognitive, rather than behavioral) is that Mr. Smart can remind himself that he is a particularly saintly person who endures travail with equanimity. Yet another possibility is that he may decide his entitlements are doomed to failure (cf. Carver 2004).

In all except the last denouement, Mr. Smart acts to change the situation so that his goals are being better met. He reinstates his entitled position, or he diminishes its relevance to core values he is embodying. Only in the last case does that not happen. The last case entails giving up, anger changing to depression, disengagement from the goal (Carver & Scheier 1998; Klinger 1975; Nesse 2000). In all the other cases, though, Mr. Smart moves himself toward a desired goal, and in so doing he reduces the negative emotion.

Return, then, to the question of how self-amplification might give way to self-stabilization. I believe the answer lies in the evoking of action aimed at removing the trigger. This point brings me

https://doi.org/10.1017/S0140525X05260041 Published online by Cambridge University Press

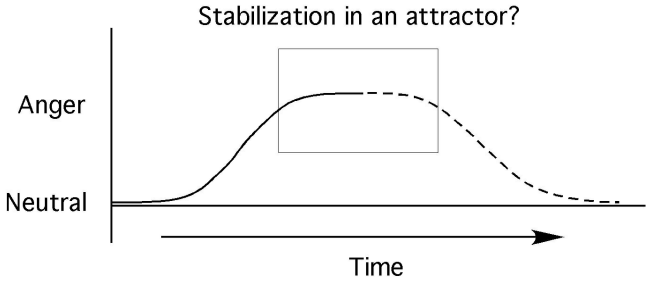


Figure 1 (Carver). Affect across time: Stabilization in an attractor, or gradual countering of a perturbation?

back to my second emotion-theory interest: the behavioral (and eventually affective) consequences of emotion. Behavioral consequences address the emotion's source (removing the obstruction, or reorganizing one's goal system to diminish its importance). When these behaviors are successful, the emotion diminishes (thus, the behavioral consequences have emotional consequences). Toward the end of section 3.3.3, Lewis brings up the possibility that functions pertaining to action play some role in stabilization. That seems far too little too late, however. Functions related to action are critical here.

Indeed, this view leads to skepticism that self-stabilization actually occurs. If affect emerges with registration of the violation, action tendencies emerge simultaneously to counter the violation (a point Lewis makes in the neuroscience part of the article, sect. 5.1). When those action tendencies yield perceived results, the anger diminishes. What appears to be stabilization may actually be the affect-countering effect of the actions (see Figure 1). Because the action often requires time to be fully effective, the emotion may cease to rise, yet fail to display immediate reduction, creating the illusion that stabilizing forces are acting to maintain it at that level. In this case, however, two directional forces are at work (one pushing emotion higher, the other dampening it) rather than a stabilizing force. To interpret this situation as a negative feedback loop maintaining the emotion at that level seems very misleading.

A final note: I am among those inclined to ignore the assumption that appraisal and emotion are distinct functions. How can appraising an event as having adverse implications for the self not imply negative affect? How can negative affect exist apart from registering (at some level, not necessarily conscious) that an event has adverse implications for the self? These seem two sides of the same coin.

I do not think abandoning the distinction renders emotion just another class of cognition, however. Valence, which is intrinsic to emotion, renders this class of experience distinctly different from others. Emotions differ from cognitions in other ways, too. The term emotion connotes physiological changes preparing the body to act. These changes are part of registering that the event has an adverse implication for the self, because adverse implications prompt behavioral responses. Such changes are not part of registering that an event constitutes a tree. This also makes emotions different from other experiences called cognition.

ACKNOWLEDGMENT

The author's work is supported by the National Cancer Institute (CA64710 and CA78995).

An intermediate level between the psychological and the neurobiological levels of descriptions of appraisal-emotion dynamics

Antonio Chella

Dipartimento di Ingegneria Informatica, Università di Palermo, Palermo, 90128, Italy. **chella@unipa.it** **http://www.csai.unipa.it/chella/**

Abstract: Conceptual space is proposed as an intermediate representation level between the psychological and the neurobiological levels of descriptions of appraisal and emotions. The main advantage of the proposed intermediate representation is that the appraisal and emotions dynamics are described by using the terms of geometry.

Lewis proposes two levels of description of appraisal and emotion dynamics. The higher, psychological level is characterized by perception, attention, evaluation, and reflection for the appraisal process, and by arousal, action tendency, and feeling tone for the emotion process (see Fig.1 of the target article). The lower, neurobiological level is characterized by the interaction among several parts and circuits of the brain.

An intermediate "conceptual" level of representation of appraisal and emotion is proposed and discussed, based on conceptual spaces (Gärdenfors 2000). A conceptual space is a geometric level of concept representation which is intermediate, in the sense of Jackendoff (1987), between the lower subsymbolic level characterized by descriptions in terms of dynamics of neural networks, as in the neurobiological level put forth by Lewis, and the higher level characterized by linguistic descriptions of emotion dynamics, as in the psychological level he describes.

As sketched below, the conceptual space level of representation has all the capabilities to describe the perception, attention, planning, and reflection processes discussed by Lewis as the basis of appraisal. Moreover, the conceptual space may be easily generalized in order to represent emotions.

The main advantage of this intermediate description is that the appraisal-emotion dynamics described by Lewis may be expressed in terms of geometry – that is, in terms of vectors, dimensions, geometrics operators, metric functions, and so forth. Geometric descriptions of cognitive processes are easy to model and to manipulate, as discussed in detail in Gärdenfors (2000); moreover, they may be immediately implemented in an artificial agent by standard geometric programming techniques.

A conceptual space is a metric space whose dimensions are related to the quantities processed by the agent sensors. Examples of dimensions could be color, pitch, volume, spatial coordinates. In any case, dimensions do not depend on any specific linguistic description: a generic conceptual space comes before any symbolic-propositional characterization of cognitive phenomena.

A *knoxel* (in analogy with *pixel*) is a point in the conceptual space that represents the epistemologically primitive perceptive element at the considered level of analysis. In an implemented robot vision system (Chella et al. 1997), in the case of static scenes, a knoxel corresponds to a *geon*-like three-dimensional geometric primitive (Biederman 1985). The agent itself is a knoxel in its conceptual space. Therefore, the perceived objects, like the agent itself, other agents, and the surrounding obstacles, are all reconstructed by means of geons and they correspond to suitable sets of knoxels in the agent's conceptual space.

Conceptual spaces may represent moving and interacting entities (Chella et al. 2000). Every knoxel now corresponds to a simple motion of a geon, expressed by adding suitable dimensions in the conceptual space that describe the variation in time of the knoxel. For example, consider the knoxel describing a rolling ball: the robot's dynamic conceptual space takes into account not only the shape and position of the ball, but also its *speed* and *acceleration* as added dimensions (Marr & Vaina 1982).

The example corresponds to a situation in the sense that the motions in the scene occur simultaneously; that is, they corre-

https://doi.org/10.1017/S0140525X05260041 Published online by Cambridge University Press