Emotional Gestalts:
Appraisal, Change, and the Dynamics of Affect

Paul Thagard *Josef Nerb*
University of Waterloo *University of Freiburg*

Running head: Emotional Gestalts

Address for correspondence: Paul Thagard, Philosophy Department, University of Waterloo, Waterloo, ON, Canada, N2L 3G1.

E-mail: pthagard@uwaterloo.ca.

April 18, 2002

Emotional Gestalts:
Appraisal, Change, and the Dynamics of Affect

Paul Thagard *Josef Nerb*
University of Waterloo *University of Freiburg*

Abstract. This paper interprets emotional change as a transition in a complex dynamical system. We argue that the appropriate kind of dynamical system is one that extends recent work on how neural networks can perform parallel constraint satisfaction. Parallel processes that integrate both cognitive and affective constraints can give rise to states that we call *emotional gestalts,* and transitions can be understood as *emotional gestalt shifts*. We describe computational models that simulate such phenomena in ways that show how dynamical and gestalt metaphors can be given a concrete realization

Introduction

One Tuesday morning, Professor Gordon Atwood strode energetically into his department office at Yukon University. It was a great day – the sun was shining, his family had been good natured at breakfast, and he was finished teaching for the semester. He cheerily greeted the department secretary and went to his mailbox, eagerly noticing the long-awaited letter from the *Journal of Cognitive Chemistry*. After opening the letter, Gordon's heart sank as he read the dreaded words: "We regret to inform you that, based on the reviewers' reports, we are unable to accept your submission." As he scanned the reviews, sadness turned to anger when Gordon realized that one of the reviewers had totally failed to understand his paper, while the other had rejected it because it did not sufficiently cite the reviewer's own work. Stomping out of the department office, Gordon chided the secretary for running out of printer paper yet again.

After an unproductive morning largely spent surfing the Web, Gordon met his wife for lunch. She reminded him that he had recently had two articles accepted for publication, and the rejected article could always be published elsewhere. They chatted about their daughter’s recent triumph in a ballet school performance, and made plans to see a musical that weekend. The Thai curry was delicious, and the coffee was strong. Gordon returned to work and happily updated his article for submission to a different journal.

Everyone has had days like this, with transitions between different moods and emotions. A psychological theory of affect must explain both how different emotional states arise and how one state can be replaced by another one that is qualitatively very different. Affect is a natural subject for a dynamical theory that emphasizes the flow of thought and the complex interactions of emotion and cognition. Our aim in this paper is to develop such a theory of the emergence and alteration of emotional states.

We proceed first by interpreting emotional change as a transition in a complex dynamical system. This metaphorical interpretation, however, is limited in its explanatory power without concrete specification of the structures and mechanisms that can give rise to emotions and emotional change. We argue that the appropriate kind of dynamical system is one that extends recent work on how neural networks can perform parallel constraint satisfaction. Parallel processes that integrate both cognitive and affective constraints can give rise to states that we call *emotional gestalts,* and transitions such as those experienced by Gordon can be understood as *emotional gestalt shifts*. Finally, we describe computational models that simulate such phenomena in ways that show how dynamical and gestalt metaphors can be given a concrete realization.

Emotion as a Dynamical System

At its simplest, psychological theory postulates causal relations between mental properties and behavior, for example that there is a personality trait of extraversion that leads people who have it to talk frequently with other people. Since the rise of cognitive science in the 1950s, psychological theories have increasingly postulated representational structures and computational processes that operate on the structures to produce behavior. More recently, some psychologists and philosophers have proposed that psychological theories should be analogous to theories of complex dynamical systems that have been increasingly popular in physics, biology, and other sciences (see, for example, Port and van Gelder, 1995; Thelen and Smith, 1994).

Thagard (1996, ch. 11) described how dynamical systems theory can be applied to psychological phenomena by means of the following explanation schema:

Explanation Target: Why do people have **stable** but **unpredictable patterns of behavior?**

Explanatory Pattern

Human thought is describable by a set of **variables**.

These variables are governed by a set of nonlinear **equations**.

These **equations** establish a **state space** that has **attractors**.

The system described by the **equations** is **chaotic**.

The existence of the **attractors** explains **stable patterns** of behavior.

Multiple attractors explain abrupt **phase transitions.**

The **chaotic** nature of the system explains why **behavior** is **unpredictable**.

It is easy to apply this explanation schema to emotions. We want to be able to explain both why people have ongoing emotions and moods, and also how they can sometimes make dramatic transitions to different emotional states. Hypothetically, we might identify a set of variables that describe environmental, bodily, and mental factors. Equations that describe the causal relations among those variables would clearly be nonlinear, in that they would require specification of complex feedback relations between the different factors. The system described by the equations would undoubtedly be chaotic, in the sense that small changes to the value of some variables could lead to very large changes in the overall system: it only took one event to dramatically change Gordon's mood. On the other hand, the emotional dynamic system does have some stability, as people maintain a cheerful or terrible mood over long periods. This stability exists because the system has a tendency to evolve into a small number of general states called attractors, and the shift from one mood to another can be described as the shift from one attractor to another.

Compare the kinds of perceptual transitions that were identified by the gestalt psychologists. When you see a Necker cube or a duck-rabbit, you see more than just the lines that make up the figure. The cube flips back and forth as you see it in different gestalts, and the duck-rabbit appears to you as a duck or as a rabbit but not both. Attending to different aspects of the drawing produces a gestalt shift in which you move from one configuration to the other. In the language of dynamical systems theory, the perceptual system has two attractor states, and the gestalt shift involves a phase transition from one attractor to the other. Analogously, we might think of an emotional state as a gestalt that emerges from a complex of interacting environmental, bodily, and cognitive variables, and think of emotional change as a kind of gestalt shift.

So much for metaphor. If we want a scientific rather than a literary explanation of emotional and perceptual phenomena, we need to flesh out the dynamical systems explanation schema by specifying the variables and the equations that relate them. We can then use computer models to simulate the behavior of the mathematical system to determine how well it models the complex behaviors of the psychological system under investigation. In particular, we want to know whether the dynamical theory can explain both stability and change in the psychological system, for example both gestalts and gestalt shifts. We will now describe how theories of thinking as parallel constraint satisfaction can provide the desired explanations. These theories are implemented by connectionist (artificial neural network) models with variables and equations that define the behavior of dynamical systems.

Thinking as Parallel Constraint Satisfaction

Thinking has traditionally been conceived as a serial process in which rules are applied to mental representations to generate new representations. In contrast, many kinds of thinking can be usefully understood as simultaneously involving numerous representations that constrain each other, with processes that operate in parallel to maximize constraint satisfaction. For example, Kunda and Thagard (1996) explained how people form impressions of other people from stereotypes by proposing that representations of different stereotypical features, traits and behaviors interact to produce an overall impression or gestalt. Your overall impression of someone you encounter as a male, Canadian, irritable professor will depend on the simultaneous interaction of your representations for each of these concepts. Read, Vanman, and Miller (1997) have reviewed how recent parallel constraint satisfaction theories have illuminated the relevance of traditional gestalt principles to social psychology. Many cognitive phenomena such as analogy and categorization can be understood in terms of parallel constraint satisfaction (see, for example, Holyoak and Spellman, 1993; Thagard, 2000).

What takes this approach beyond the metaphorical is the availability of computational models that show how parallel constraint satisfaction can be performed by artificial neural networks. Representations such as concepts and propositions can be modeled by units, which are artificial neurons with variables that indicate the degree of activation (plausibility or degree of acceptance) of each unit. The constraints between representations can be represented by excitatory and inhibitory links between units; for example, the positive constraint between the concepts *professor* and *intelligent* can be captured by an excitatory link between the corresponding units, while the negative constraint between *professor* and *rich* can be captured by an inhibitory link. Parallel constraint satisfaction is performed by means of equations that specify how the activation of each unit is updated as a function of the activations of the units to which it is linked and the strength (positive or negative) of those links. Updating usually leads such systems into stable states in which activations of all units cease to change; this is called settling. However, the systems are chaotic in that slight changes to inputs can lead the system to settle into a different state. For example, it is easy to model the Necker cube by a neural network in which each unit stands for a hypothesis about which feature of the cube is front or back (Rumelhart and McClelland, 1986, vol. 2, p. 10). Slight changes to input about a particular feature being front or back can produce a flip in the overall state of the network. Thus artificial neural networks that implement parallel constraint satisfaction can naturally model gestalts and gestalt shifts. We will now describe how parallel constraint satisfaction and connectionist modeling can be extended to emotional thinking.

Emotional Gestalts: Theory

Extending parallel constraint satisfaction to emotion requires representations and mechanisms that go beyond those required for purely cognitive coherence. In addition to representations of propositions, concepts, goals, and actions, we need representations of emotional states such as happiness, sadness, surprise, and anger. Moreover, the cognitive representations need to be associated with emotional states, most generally with positive and negative evaluations or *valences* (Bower, 1981; Lewin, 1951). For example, Gordon Atwood's belief that his paper has been rejected has negative valence and is associated with sadness, whereas his belief that his daughter is an excellent dancer has positive valence and is associated with pride. Like beliefs, concepts can also have positive valences, for example ones representing success and sunshine, while other concepts have negative valences, for example ones representing death and disease.

In purely cognitive coherence, representations are connected by positive and negative constraints representing the extent to which they fit with each other. Similarly, representations can be connected by affective constraints. Some of these are intrinsic to the person that has the representations: most children, for example, attach an intrinsic positive valence to *candy.* In other cases the valence is acquired by associations, as when a child acquires a positive valence for *grocery store* because grocery stores sell candy.

In purely cognitive coherence models, representations are accepted or rejected on the basis of whether doing so helps to satisfy the most constraints, where a positive constraint is satisfied if the representations it connects are both accepted or both rejected. A judgment of emotional coherence requires not only inferences about the acceptance and rejection of representations, but also about their valence. Gordon is forced to accept the conclusion that his paper will not appear in the journal, and he attaches negative valence to this proposition. Emotional coherence requires not only the holistic process of determining how to best satisfy all the cognitive constraints, but also the simultaneous assessment of valences for all relevant representations.

The result can be an emotional gestalt consisting of a cognitively and emotionally coherent group of representations with degrees of acceptance and valences that maximally satisfy the constraints that connect them. Of course, circumstances may prevent the achievement of coherence, for example when there is support for inconsistent beliefs or when competing actions are associated with conflicting valences. But normally the process of parallel constraint satisfaction will lead a person to acquire a set of stable acceptances of propositions and concepts that apply to the current situation, as well as a set of valences that indicate the emotional value of those representations. Particular emotions such as happiness and sadness can then emerge from an assessment of the extent to which the current situation satisfies a person's ongoing goals.

Emotional gestalt shifts occur when changes in representations and their valences generate a new array of acceptances and valences that maximize constraint satisfaction differently from before. When Gordon reads his letter from the journal, he has to shift the cognitive status of the proposition that his article will be published in it from accepted to rejected. Through parallel constraint satisfaction, this shift may alter the acceptance status of other propositions, for example ones representing whether he will get promoted or receive a good raise this year. In addition, there may be a shift in valences attached to his representations of the article, the journal, or even his whole career.

The theory of emotional coherence is highly compatible with the appraisal theory of emotion, according to which emotions are elicited by evaluations of events and situations (Scherer, Schorr, and Johnstone, 2001). Appraisal theorists have explained in great detail how different cognitive evaluations produce many different emotions. But they have been vague about the specific mechanisms that generate evaluations and elicit emotions. We propose that appraisal is a process of parallel constraint satisfaction involving representations that have valences as well as degrees of acceptance, and we will describe in the next section neural network models that are capable of performing some kinds of appraisal. According to appraisal theory, changes in emotion such as those that occur in psychotherapy are due to changes in how individuals evaluate their situations; the theory of emotional coherence gives a more specific account of how cognitive therapy works, by the therapist introducing new evidence and reforming coherence relations in ways that change emotional appraisals (Thagard, 2000, pp. 209). Let us now consider more concretely how processes of parallel constraint satisfaction can produce emotional gestalts and affective change.

Emotional Gestalts: Models

Computational models are an indispensable tool for developing psychological theories and for evaluating their explanatory and predictive power. The view of thinking as parallel constraint satisfaction did not precede computational modeling, but rather arose because of the development of artificial neural network (connectionist) models (Rumelhart and McClelland, 1986). It is useful conceptually to distinguish among the following: (1) a psychological theory of the structures and processes that produce a kind of thinking, (2) a computational model that rigorously spells out the structures and processes in terms of data structures and algorithms, and (3) a running computer program that uses a particular programming language to implement the model and serves to test the power and empirical relevance of the model and theory. In practice, however, the development of computer programs can be a major part of the creation of new models and theories.

Computer simulations are crucial for theories of complex dynamical systems. Metaphorical use of dynamical system terminology (nonlinear systems, chaos, attractors, self-organization, emergence, etc.) can be useful for reconceptualizing complex phenomena, but a deep understanding of the phenomena requires mathematical specification of the key variables that produce them and of the relations among them. Once this specification is available, it is rarely possible to infer the behavior of the system characterized by mathematical deduction or hand simulations, so computer simulations are needed for determining the implications of the mathematical model. As Vallacher, Read, and Nowak (this issue) emphasize, dynamical research in psychology needs computer modeling as a complement to empirical research.

Fortunately, implemented artificial neural network models of emotional cognition are already available. We will describe how two models, HOTCO and ITERA, provide partial realizations of emotional gestalts, and indicate some directions for future work. Computer simulation of emotional or “hot” cognition originated with the work of

Abelson (1963) on rationalization, but current models use artificial neural networks to integrate cognition and emotion.

HOTCO

The model HOTCO (for “hot coherence”) incorporates previous computational models of analogy, decision making, hypothesis evaluation, and impression formation (Holyoak and Thagard, 1989; Kunda and Thagard, 1996; Thagard, 1992; Thagard and Millgram, 1995). All of these models perform parallel constraint satisfaction using artificial neurons called units to correspond to mental representations such as concepts and propositions; excitatory and inhibitory links between these units correspond to constraints between representations. As described earlier, algorithms update the activations of the units until the network of units settles, that is, all units have achieved stable activations.

HOTCO differs from previous cognitive models, however, in that each unit has a variable for the valence of a unit as well as for its activation, as well as algorithms for updating valences as well as activations. In addition, a special unit called VALENCE, whose valence is always the maximum value of 1, provides input to units corresponding to representations with an inherent emotional value. In a standard connectionist algorithm, the activation of a unit j, a_j, is updated according to the following equation:

$$a_j(t+1) = a_j(t)(1-d) + net_j(max - a_j(t)) \text{ if } net_j > 0,$$

$$\text{otherwise } net_j(a_j(t) - min).$$

Here d is a decay parameter (say .05) that decrements each unit at every cycle, min is a minimum activation (-1), max is maximum activation (1). Based on the weight w_{ij} between each unit i and j, we can calculate net_j, the net input to a unit, by:

$$net_j = \sum_i w_{ij} a_i(t).$$

HOTCO uses similar equations to update valences. The valence v_j of a unit u_j is the sum of the results of multiplying, for all units u_i to which it is linked, the activation of u_i times the valence of u_i, times the weight of the link between u_i and u_j. The actual equation used in HOTCO to update the valence v_j of unit j is similar to the equation for updating activations:

$$v_j(t+1) = v_j(t)(1-d) + net_j(max- v_j(t)) \text{ if } net_j > 0,$$

$$net_j(v_j(t) - min) \text{ otherwise.}$$

Here d is a decay parameter (say .05) that decrements each unit at every cycle, min is a minimum valence (-1), max is maximum valence (1). Based on the weight w_{ij} between each unit i and j, we can calculate net_j, the net valence input to a unit, by:

$$net_j = \sum_i w_{ij} v_i(t) a_i(t).$$

Updating valences is just like updating activations plus the inclusion of a multiplicative factor for valences. What these equations do is ensure that the valence of a unit, corresponding to the emotional value of a representation, increases if there is positive valence flowing into it from other active units.

In the original version of HOTCO, activations could influence valences but not vice versa (Thagard, 2000). This reflects the normatively appropriate strategy that beliefs should affect emotions but not the reverse. In real life, however, people are subject to wishful thinking and motivated inference where what they want influences what they believe to be true (Kunda, 1990). Support for the influence of desires on beliefs comes from recent findings in behavioral decision theory concerning cases where the outcomes of decisions or the targets of judgments are affectively rich (Finucane, Alhakami, Slovic, and Johnson, 2000; Rottenstreich & Hsee, 2001). These findings contradict both expected utility theory and prospect theory, which assume that probability and utilities are independent. Forgas (1995) reviews the role of affective states in social judgments, explaining motivated processing effects in his Affect Infusion Model (AIM), which specifies how and when the valence of a mood state may infuse judgments.

Accordingly, HOTCO 2 allows units to have their activations influenced by both input activations and input valences (Thagard, forthcoming). The basic equation for updating activations is the same as the standard one, but the net input to a unit's activation is defined by a combination of activations and valences:

$$net_j = \sum_i w_{ij} a_i(t) + \sum_i w_{ij} v_i(t) a_i(t).$$

The first summation is the standard one that makes the input to a unit's activation a function of the activations of all the units linked to it, while the second is the valence net input defined above. Together they make acceptance depend in part on desirability. These equations are implemented in a LISP program available on the Web at http://cogsci.uwaterloo.ca.

HOTCO has been used to simulate various psychological phenomena, including trust and juror reasoning. Figure 1 shows an abstract network in which a combination of emotion input and evidence input generate mood by means of an overall assessment of coherence. In the case of Gordon, the new evidence that his paper has been rejected produces a temporary reappraisal of his situation and himself, dramatically altering his mood.

<< INSERT FIGURE 1 HERE >>

For social psychologists, the most interesting HOTCO 2 simulation involves studies of stereotype activation reported by Sinclair and Kunda (1999). They found that participants who were praised by a Black individual tended to inhibit the negative Black stereotype, while participants who were criticized by a Black individual tended to apply the negative stereotype to him and rate him as incompetent. According to Sinclair and Kunda, the participants motivation to protect their positive views of themselves caused them either to suppress or to activate the negative Black stereotype.

Figure 2 shows the structure of a simplified HOTCO 2 simulation of aspects of the experiment in which praise and criticism produced very different evaluations of the individual who provided them (Thagard, forthcoming). Without any evidence input that the evaluation is good or bad, the program finds equally acceptable the claims that the evaluator is competent or incompetent. However, a positive evaluation combines with the motivation for self-enhancement to generate positive judgments of the evaluator and blacks, while a negative evaluation combines with self-enhancement to generate negative judgments of the evaluator and blacks. In the simulation shown in figure 2, the positive valence of the *I am good* unit supports activation of the accurate-evaluation unit, which activates the competent-manager unit and suppresses the black stereotype. HOTCO 2 thus shows how thinking can be biased by emotional attachment to goals such as self-enhancement

<< INSERT FIGURE 2 HERE >>

It should be evident from this example and the previous discussion that HOTCO models a dynamical system. The variables are activations and valences of the various units, and the equations that define a state space for the system are the equations for updating activations and valences stated above. The equations are nonlinear in the mathematical sense that the variables are related multiplicatively, and the system is also nonlinear in the metaphorical sense that there are many interactive feedback influences. The HOTCO network settles into very different states depending on changes in the input representing a positive or negative evaluation. Similarly, the experimental participants developed a very different emotional gestalt depending on whether the Black professional praised or criticized them.

Although HOTCO can model some interesting psychological phenomena, it is far from providing a complete account of the range of emotions discussed by appraisal theory. HOTCO does not differentiate between specific positive or negative emotions such as sadness and anger directed at particular objects or situations, although it does compute a kind of global happiness and sadness based on measures of constraint satisfaction (Thagard, 2000). The main specific emotion that HOTCO models is surprise, which is generated when units undergo a substantial change in degree of activation. However, another recent model allows for greater differentiation of emotions.

ITERA

Nerb and Spada (2001) present a computational account of how media information about environmental problems influences cognition and emotion. When people hear about an environmental accident, they may respond with a variety of emotions such as sadness and anger. Following the appraisal theory of emotions, Nerb and Spada hypothesized that a negative event will lead to sadness if it is caused by situational forces outside of anyone's control. But an environmental accident will lead to anger if someone is responsible for the negative event. If people see themselves as responsible for the negative event, then they feel shame; but if people see themselves as responsible for a positive event, they feel pride. Nerb and Spada (2001) show how determinants of responsibility such as agency, controllability of the cause, motive of the agent, and knowledge about possible negative consequences can be incorporated into a coherence network called ITERA (Intuitive Thinking in Environmental Risk Appraisal).

ITERA is an extension of the impression-formation model of Kunda and Thagard (1996). The main innovation in ITERA is the addition of units corresponding to emotions such as anger and sadness, as shown in figure 3. ITERA is given input concerning whether or not an accident was observed to involve damage, human agency, controllability, and other factors. It then predicts a reaction of sadness or anger depending on their overall coherence with the observed features of the accident. This reaction can be thought of as a kind of emotional gestalt that summarizes all the available information.

<< INSERT FIGURE 3 HERE >>

In three studies investigating the evaluation of environmental accidents, Nerb and Spada (2001) compared the performance of ITERA with people's reactions. In their empirical experiments and their simulations experiments, the authors varied the degree of responsibility for an accident by manipulating the determining factors of damage, human agency, controllability, higher goal, and knowledge (cf. Weiner, 1995). In the domain of environmental problems, a higher goal may be given if the accident happened as a consequence of an action that was meant to achieve something that had a positive outcome, such as an increase in the overall benefit for society. Knowledge reflects whether agents knew in advance that there is a possible contingency between their action and threats to the environment. In the ITERA model, information that is known about the accident has a special status. Nodes representing such information have links to one of the two special nodes, OBSERVED+ or OBSERVED-; these special nodes have fixed maximal or minimal activation. A link to OBSERVED+ reflects that the corresponding determinant is given, a link to OBSERVED- means that it is not given. No link to one of the specials nodes represents that nothing is known about the determinant. By manipulating links to theses nodes, different experimental settings were realized; see figure 3 for a concrete example of an accident scenario.

Because the determinant-emotion links are bidirectional in ITERA, there is a feedback relationship between determinants and emotions. These feedback relations allow for interactions among the activations of nodes for cognitive determinants. Hence the final activation of a node can be different across simulation experiments in which the inputs to other nodes are varied, so that the model predicts that people will construe an aspect of a situation differently depending on other aspects of the situation. Nerb and Spada called these types of model predictions *coherence effects*. For instance, ITERA predicts that manipulating the controllability node will produce coherent patterns for human agency, higher goal, and knowledge. In ITERA, a controllable cause produces high activation values for human agency and knowledge but low activation for higher goal, whereas an uncontrollable cause leads to low activation values for human agency and knowledge but high activation for higher goal. Note that there are no direct links between the cognitive determinants. The bidirectional spreading of activation within the parallel constraint satisfaction network of ITERA is sufficient for producing this coherent covariation of the cognitive determinants.

Overall, ITERA produced a good fit for the model predictions with data for anger and the intention to boycott the transgressor. In particular, the predicted coherence pattern among appraisal criteria were confirmed by empirical evidence (see Nerb, Spada, & Lay, 2001, for more empirical evidence for the model). Such interaction effects among appraisal criteria are compatible with existing appraisal theories and are also supported by recent empirical findings (Lazarus, 1991; Lerner and Keltner, 2000, 2001). Interactions among appraisal criteria tend to be ignored by other types of computational models and by non-computational appraisal models. For instance, rule-based appraisal models that realize the cognition-emotion relationship as a set of if-then associations do not capture the interaction effects among appraisal criteria (e.g., Scherer, 1993). ITERA accounts for such affective coherence effects among appraisal criteria by using bidirectional links for the cognition-emotion relationship.

Unlike HOTCO, ITERA does not incorporate variables for valence as well as activation, and it lacks algorithms for global calculations of coherence. But it is more psychologically realistic with respect to the differentiation of particular emotions such as sadness and anger, and Nerb is now working on a synthesis of HOTCO and ITERA intended to combine the best features of both. Another connectionist model of emotional cognition has been produced by Mischel and Shoda (1995), as part of their cognitive-affective system of personality.

There are many other possibilities for future developments of computational models of the dynamics of emotion and cognition. Wagar and Thagard (forthcoming) describe a much more neurologically realistic model of the interactions between emotion and cognition. It is more realistic than HOTCO both with respect to individual neurons and with respect to the anatomical organization of the brain. The new model uses distributed representations, which spread information over multiple artificial neurons, rather than the localist ones in HOTCO and ITERA, which use a single neuron to stand for a concept or proposition. In addition, the artificial neurons in Wagar and Thagard's new model behave by spiking in the manner of real neurons, rather than simply spreading activation. Moreover, they are organized in modules corresponding to human neuroanatomy, including the hippocampus, neocortex, amygdala, and nucleus accumbens. The result is intended to be a model that captures much more of the dynamic activities of the brain than previous connectionist models of emotional cognition. Wagar and Thagard (forthcoming) simulate some of the fascinating phenomena discussed by Damasio (1994), especially the decision-making deficits found in patients such as Phineas Gage who have brain damage that disrupts the flow of information between areas responsible for reasoning (the neocortex) and emotions (the amygdala).

Much remains to be done to understand long term emotional change. Psychotherapy can take months or even years to change the emotional tendencies of an individual by helping people to revise their cognitions and emotions. A marriage that begins with mutual love can dissolve into anger and even hatred. The terrorist leader Osama bin Laden began with a dislike for Americans that eventually turned into an intense hatred. A dynamical theory of emotion should deal with long term emotional changes as well as the sudden transitions discussed in this paper.

Another direction for future research on the dynamics of emotion is to expand models to capture group interactions. Thagard (2000, ch. 7) developed a theory of consensus resulting from communication between members of a scientific community who disagree about what theory provides the best explanation of the available evidence. He has run computer simulations in which up to 60 scientists, each of whom makes coherence-based decisions, reach agreement on disputed scientific issues. Future work will develop simulations in which group decision making involves emotions as well as evidence. This will require extension of existing models of emotional coherence and consensus to apply to group decisions in which individuals communicate not only their beliefs and goals but also their emotional evaluations of competing options.

Conclusion

Although the theoretical vocabulary and computational modeling described in this paper may seem unfamiliar to social psychologists, the basic ideas are similar to those of some of the classic theories in the field. Festinger’s (1957) notion of cognitive dissonance can be accounted for in terms of parallel constraint satisfaction (Shultz and Lepper, 1996); and it is possible with HOTCO to incorporate the affective dimension that later theorists found to be crucial to dissonance (Cooper and Fazio, 1984). Kunda and Thagard (1996) describe how parallel constraint satisfaction can be used to understand ideas about impression formation that Solomon Asch developed from Gestalt psychology.

Emotional coherence and the kind of computational model discussed in this paper can also be applied to McGuire's (1999) dynamic model of thought systems. McGuire's theory describes thought systems as consisting of propositions with two attributes: desirability (an evaluation dimension) and likelihood of occurrence (an expectancy dimension). These dimensions express how much the content of a proposition is liked and how much it is believed. The thought system is seen as dynamic so that a change that is directly induced in one part of the system results in compensatory adjustments in remote parts of the system. Together those assumptions imply that a person's evaluative and expectancy judgments on a topic will tend to be assimilated toward one another. McGuire (1999, p. 190) postulates "that causality flows in both directions, reflecting both a 'wishful thinking' tendency such that people bring their expectations in line with their desires, and a 'rationalization' tendency such that they bring desires in line with expectations." Emotional coherence models such as HOTCO provide mechanisms for both these tendencies, as activations corresponding to likelihood interact with valences corresponding to desirability.

Some proponents of a dynamical systems approach to understanding the mind have seen it as a radical alternative to symbolic and connectionist models that have been predominant in cognitive science (van Gelder, 1995). Our discussion has shown that this is a mistake: connectionist systems are one important kind of dynamical system, and they are arguably the ones best suited for explaining psychological phenomena, including ones involving emotion.

In sum, computational models based on parallel constraint satisfaction, with constraints and representations that are affective as well as cognitive, have much to contribute to a dynamical perspective on social psychology. We have shown how connectionist systems such as HOTCO and ITERA can be used to study the dynamics of emotional cognition at a concrete and not just a metaphorical level. In particular, they can model the generation and shifting of emotional gestalts.

Acknowledgments. This research was supported by a Natural Sciences and Engineering Research Council of Canada grant to the first author, and by a fellowship from the Alexander von Humboldt-Stiftung to the second author. We are grateful to Stephen Read, Robin Vallacher, Eliot Smith, and an anonymous referee for helpful comments on an earlier draft.

References

Abelson, R. (1963). Computer simulation of "hot" cognition. In S. Tomkins and S. Messick (Eds.), *Computer simulation of personality* (pp. 277-298). New York: John Wiley and Sons.

Bower, G. H. (1981). Mood and memory. *American Psychologist, 36*, 129-148.

Cooper, J., & Fazio, R. H. (1984). A new look at dissonance theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 17). New York: Academic Press.

Damasio, A. R. (1994). *Descartes' error*. New York: G. P. Putnam's Sons.

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford: Stanford University Press.

Finucane, M. L., Alhakami, A. S., Slovic, P., & Johnson, S. M. (2000). The affect heuristic in judgements of risks and benefit. *Behavioral Decision Making, 13*, 1-17.

Forgas, J.P. (1995). Mood and judgment: The Affect Infusion Model (AIM). *Psychological Bulletin*, *117*, 39-66.

Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology, 44*, 265-315.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295-355.

Kunda, Z. (1990). The case for motivated inference. *Psychological Bulletin, 108*, 480-498.

Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review, 103*, 284-308.

Lazarus, R. S. (1991). Cognition and motivation in emotion. *American Psychologist, 46,* 352-267.

Lerner, J. S., & Keltner, D. (2000). Beyond valence: Toward a model of emotion-specific influences on judgement and choice. *Cognition and Emotion, 14*, 473-493.

Lerner, J. S., & Keltner, D. (2001). Fear, anger, and risk. *Journal of Personality and Social Psychology, 81*, 146-159.

Lewin, K. (1951). *Field theory in social science*. New York: Harper & Row.

McGuire, W. J. (1999). *Constructing social psychology: Creative and critical processes*. New York: Cambridge University Press.

Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and the invariance in personality structure. *Psychological Review, 102*, 246-268.

Nerb, J., & Spada, H. (2001). Evaluation of environmental problems: A coherence model of cognition and emotion. *Cognition and Emotion, 15*, 521-551.

Nerb, J., Spada, H., & Lay, K. (2001). Environmental risk in the media: Modeling the reactions of the audience. *Research in Social Problems and Public Policy, 9*, 57-85.

Port, R., & van Gelder, T. (Eds.). (1995). *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.

Read, S. J., Vanman, E. J., & Miller, L. C. (1997). Connectionist, parallel constraint satisfaction, and Gestalt principles: (Re)Introducing cognitive dynamics to social psychology. *Personality and Social Psychology Review, 1*, 26-53.

Rottenstreich, Y., & Hsee, C. K. (2001). Money, kisses, and electric shocks: On the affective psychology of risk. *Psychological Science, 12*, 185-190.

Rumelhart, D. E., & McClelland, J. L. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge MA: MIT Press/Bradford Books.

Scherer, K. R. (1993). Studying the emotion-antecedent appraisal process: An expert system approach. *Cognition and Emotion, 7*, 325-355.

Scherer, K. R., Schorr, A., & Johnstone, T. (2001). *Appraisal processes in emotion*. New York: Oxford University Press.

Shultz, T. R., & Lepper, M. R. (1996). Cognitive dissonance reduction as constraint satisfaction. *Psychological Review, 103*, 219-240.

Sinclair, L., & Kunda, Z. (1999). Reactions to a Black professional: Motivated inhibition and activation of conflicting stereotypes. *Journal of Personality and Social Psychology, 77*, 885-904.

Thagard, P. (1992). *Conceptual revolutions*. Princeton: Princeton University Press.

Thagard, P. (1996). *Mind: Introduction to cognitive science*. Cambridge, MA: MIT Press.

Thagard, P. (2000). *Coherence in thought and action*. Cambridge, MA: MIT Press.

Thagard, P. (forthcoming). Why wasn't O. J. convicted? Emotional coherence in legal inference. *Cognition and Emotion*.

Thagard, P., & Millgram, E. (1995). Inference to the best plan: A coherence theory of decision. In A. Ram & D. B. Leake (Eds.), *Goal-driven learning:* (pp. 439-454). Cambridge, MA: MIT Press.

Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.

Vallacher, R. R., Read, S. J., & Nowak, A. (this issue). The dynamical perspective in personality and social psychology. *Personality and Social Psychology Review*.

van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy, 92*, 345-381.

Wagar, B., & Thagard, P. (forthcoming). Using computational neuroscience to investigate the neural correlates of cognitive-affective integration during covert decision making. *Brain and Cognition*.

Weiner, B. (1995). Judgements of responsibility: A foundation for a theory of social conduct. New York, NY: Guilford Press.

Figure Captions

Figure 1. Mood changes affected by emotional coherence. From Thagard (2000) p. 209. Thick lines are valence links and thin lines are cognitive links. Valences spread from emotion input, while activations spreads from evidence input. Appraisals and moods are influenced by both valences and activations.

Figure 2. Evidential and valence associations leading to the motivated inhibition of the negative black stereotype. Solid lines are excitatory links and dashed lines are inhibitory. This figure and parts of the preceding paragraph are taken from Thagard (forthcoming).

Figure 3. ITERA network for emotional reactions to environmental accidents. Solid lines are excitatory links, and dashed lines are inhibitory links. This example represents a situation in which a media report states that there is damage caused by human agency that could not have been controlled. From Nerb and Spada (2001), p. 528.

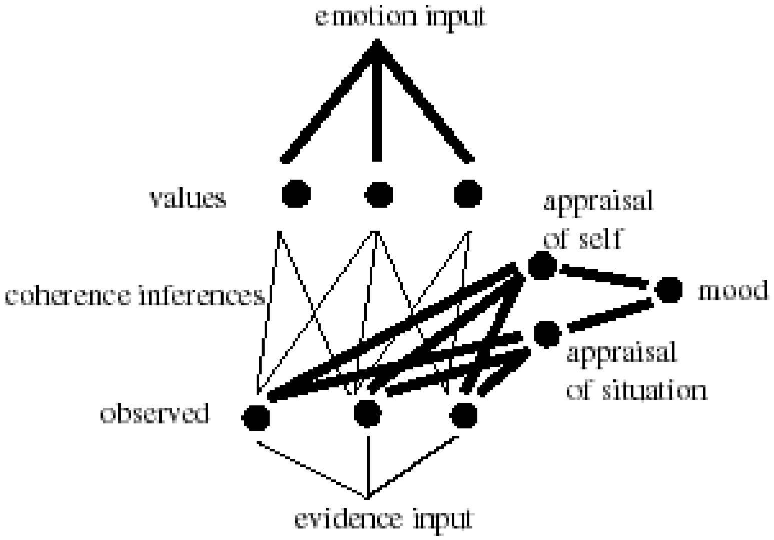
emotion input
values
appraisal of self
coherence inferences
mood
appraisal of situation
observed
evidence input

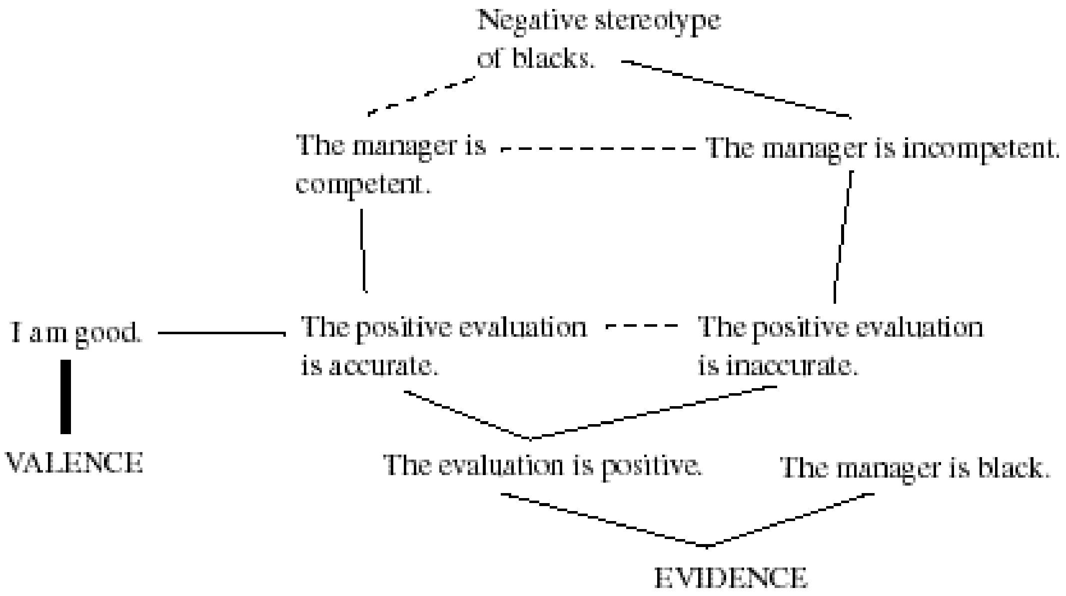
Negative stereotype of blacks.
The manager is competent.
The manager is incompetent.
I am good.
The positive evaluation is accurate.
The positive evaluation is inaccurate.
VALENCE
The evaluation is positive.
The manager is black.
EVIDENCE

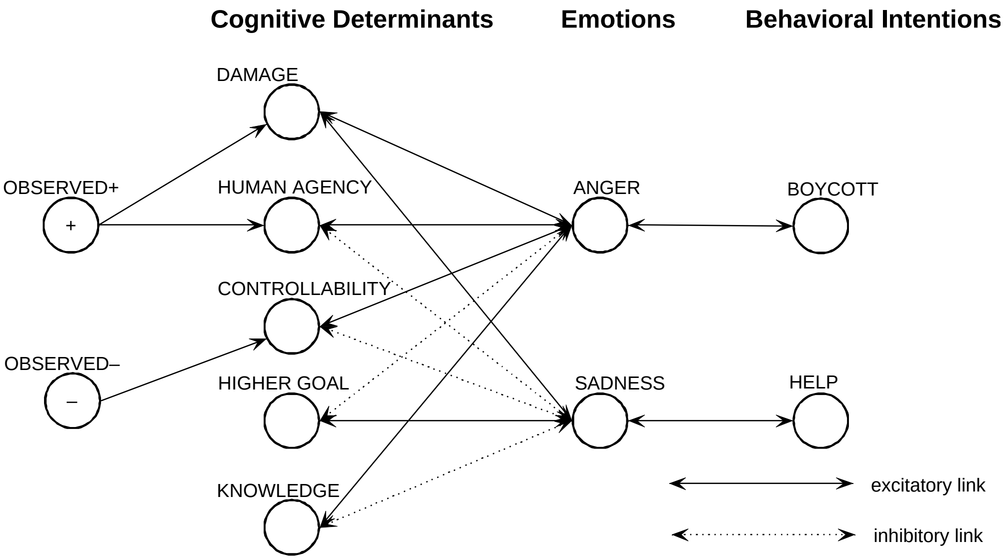
Cognitive Determinants
Emotions
Behavioral Intentions
DAMAGE
OBSERVED+
+
HUMAN AGENCY
ANGER
BOYCOTT
CONTROLLABILITY
OBSERVED–
–
HIGHER GOAL
SADNESS
HELP
KNOWLEDGE
excitatory link
inhibitory link