Personality and Social Psychology Review
2002, Vol. 6, No. 4, 316–325

Personality as a Dynamical System: Emergence of Stability and Distinctiveness from Intra- and Interpersonal Interactions

Yuichi Shoda and Scott LeeTiernan
Department of Psychology
University of Washington

Walter Mischel
Department of Psychology
Columbia University

The implications of conceptualizing personality as a cognitive-affective processing system that functions as a parallel constraint satisfaction network are explored. Computer simulations show that from dynamic interactions among the units in such a network, a set of stable attractor states and functionally equivalent groups of situations emerge, such that IF exposed to situation group X, THEN the system settles in attractor Y. This conceptualization explicitly models the effect of situations on a given individual, and therefore can also be used to model the function of interpersonal systems. We demonstrate this possibility by modeling dyadic systems in which one partner's behavior becomes the situational input into the other partner's personality system, and vice versa. The results indicate that each member of the dyad will, in general, exhibit new attractor states. This suggests that the thoughts, affects, and behaviors that an individual typically experiences are a function not of that individual's personality system alone, but rather a function of the interpersonal system of which the individual is a part. Just as individuals have distinctive and stable IF-THEN signatures, so do interpersonal relationships. Understanding the structure of the cognitive-affective processing system of each relationship partner also should enable predictions of their distincitve relational signatures as emergent properties of the interpersonal system that develops.

How our thoughts and feelings come and go, or what William James captured with the phrase "stream of consciousness" (1890), has been a basic topic in psychology for almost a century. It is also a central feature of all living things: if what looks like an animal doesn't change its behavior in response to what happens to it, we may begin to wonder if it is even alive. Signs of life are seen in variability.

Yet, when considering personality and individual differences, the vital importance of change and flow must be reconciled with the notion of constancy, and with the assumption that each individual is characterized by stable and distinctive qualities. The person that we are in one moment is basically the same person in the next moment, even if our thoughts, feelings and behaviors vary substantially. Is there something in us that remains invariant through the changing stream of thoughts, feelings, and behaviors, and if so, how might one conceptualize such invariance?

A key to embracing such variations and invariance in personality structure, we believe, lies in reconceptualizing personality as a dynamical system (Mischel & Shoda, 1995). Stability is expected in the underlying structure that generates the thoughts, feelings, and behaviors which themselves change from one moment to the next. The stability of the underlying structure will be reflected in part in the overall average levels of different types of behavior, but it also will be seen most clearly in the way they change. As every good novelist knows, both the subtle texture of personality and its underlying dynamics may be revealed in the seeming inconsistencies evident in a person, and these may be even more informative than the apparent consistencies and overall average behavior tendencies.

The study was supported in part by the University of Washington Royalty Research Fund, Grant MH39349 and NRSA Predoctoral Fellowship from the National Institute of Mental Health, as well as a Microsoft Fellowship in computer science and related fields. We are grateful to Eliot Smith, Andrzej Nowak, and two anonymous reviewers for their helpful comments, and especially to Steve Read, for his encouragement of the work and comments on earlier drafts, from which the present article benefited greatly.

Parts of this research were presented at the annual meeting of the Society of Personality and Social Psychology, 1999 and 2000.

Requests for reprints should be sent to Yuichi Shoda or Scott LeeTiernan, Department of Psychology, University of Washington, Box 351525, Seattle, WA 98195–1525. E-mail: yshoda@u.washington.edu or slt@u.washington.edu

The Emerging Social Cognitive-Affective Conception of the Person

When we started looking for alternative conceptualizations of personality several years ago, we were struck by the fact that in the broader field of social cognition and the cognitive sciences in general, a conception was already emerging, if not explicitly, in the form of a set of implicit assumptions. This conception views the person as consisting of mental representations whose activation leads to the thoughts and feelings experienced and the behaviors displayed (e.g., Higgins, 1990). The cognitions and affects particularly relevant to personality include diverse content, or "cognitive-affective units" (Mischel & Shoda, 1995). These encompass the person's encoding or construal of the self and of situations, enduring goals, expectations and feeling states, as well as specific memories of the people and events that have been experienced, and a host of competencies and skills particularly important for self-regulation. These cognitive-affective units have been previously discussed as "person variables" on which individuals differ (Mischel, 1973). Some of these units or mental representations are more available and accessible, ready to become activated, while others are quite inaccessible (Higgins, 1996). Most importantly, the activation of cognitive affective units is not constant, but changes from one time to another, from one situation to another. The renewed attention on such dynamics was calling for a way to conceptualize coherence in the face of change, accounting for the predictable patterns that exist in the pattern of changes distinctively characterizing a person.

Central to this emerging conception of personality is the idea that mental representations of the psychological meaning of situations, representations of self, others, possible future events, goals, affects, beliefs, expectations, as well as behavioral alternatives are not isolated, but are interconnected. We proposed that for a given individual the likelihood that thought *A* leads to thought *B* and emotion *C* is guided by a network of associations among cognitions and affects available to that individual. Through this network, for example, thinking about a person can activate the memory of the thoughts and feelings associated with a particular event in the past, which in turn may lead to other memories and thoughts that may make us smile or cry (e.g., Andersen & Baum, 1994; Andersen & Cole, 1990). Individuals differ stably in this network of inter-connections or associations, and such differences constitute a major aspect of personality (Mischel & Shoda, 1995).

Furthermore, each unit is potentially connected to every other unit in the network, and each pair of units is characterized by a distinct and stable strength of association between them. Called *recurrent* networks, one of the most notable properties of such networks is that they settle into a set of activation patterns to satisfy multiple simultaneous constraints represented by the patterns and strengths of connections among the units in the network. The use of a recurrent, or parallel constraint satisfaction, network is consistent with models of human information processing in the broader cognitive sciences, including analogical reasoning (Spellman & Holyoak, 1992), attitude change (Spellman, Ullman, & Holyoak, 1993), explanatory coherence (Read & Marcus-Newhall, 1993; Thagard, 1989), dissonance reduction (Read & Miller, 1994; Shultz & Lepper, 1996), and impression formation and dispositional inference (Kashima & Kerekes, 1994; Kunda & Thagard, 1996; Read & Miller, 1993).

The Cognitive-Affective Personality System (CAPS)

Figure 1 illustrates the basic structure of the system that in this conceptualization underlies changing thoughts and feelings, and stable differences among the patterns of change, summarized in a highly simplified, schematic outline (Mischel & Shoda, 1995). The personality system contains a number of cognitive and affective units, connected to one another to form networks of associations that distinctively characterize each individual. When the individual encounters a situation, a subset of these units (the feature detection units) becomes activated, depending on the configuration of features present in the situation. These activated units in turn activate other cognitions and affects (processing units), following the stable associative links in the individual's cognitive-affective processing system (or CAPS) network. The figure is, of course, a greatly simplified view of the rich system of interconnections among the cognitive and affective units. We assume that each individual's characteristic network among units reflects culture and subculture (Mendoza-Denton, Shoda, Ayduk, & Mischel, 1999), as well as an individual's specific social learning and biological-genetic history.

Behavioral Expressions of Stable Individual Differences in the CAPS Network

First, we consider in some detail how such a system "works" to produce behaviors. Cognitions and affects become activated, either by the salient "psychologically active" element in the situation (Shoda et al., 1994) or internally by other activated cognitions and affects. Activation then propagates through an individual's unique network of associations, and ultimately activates thoughts or emotions the individual is aware of, or a behavior that is observable.

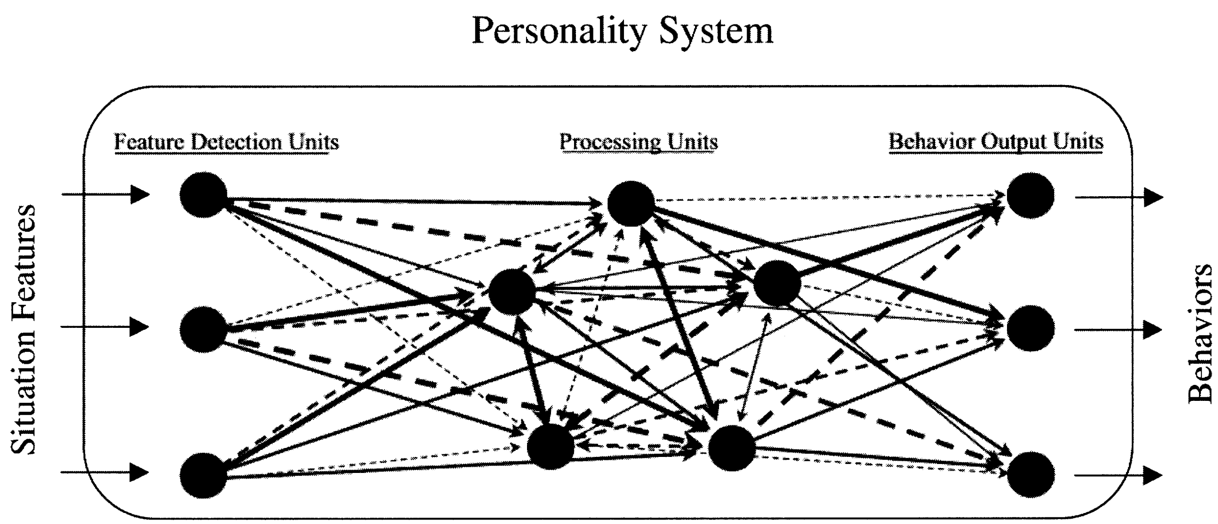

Figure 1. An Illustrative CAPS network. When a unit connected to another unit by a solid line becomes activated, it increases the activation of the second unit, in proportion to the thickness of the arrow. When a unit connected to another unit by a broken line becomes activated, it decreases the activation of the second unit, in proportion to the thickness of the arrow. The CAPS networks used in the simulations were larger than the network shown here, consisting of 5 input units, 10 processing units, and 5 output units each.

For example, suppose in situation 1, a person is sitting on an airplane, and a man with a ponytail sits next to her. Some of the feature detector units may become activated (left side of Figure 1) upon seeing the man's ponytail, which might in turn activate some of the processing units (middle section of Figure 1), such as conflicted feelings associated with the times when her ex-hippie father, who had a ponytail, came to her school on parents' day. This may even activate other processing units such as a specific memory of her friends' teasing her about her father's long hair, and for not having a TV in the converted school bus that was their home. These feelings and memories in turn may influence her response to the present situation, leading to her feeling angry and hurt when a flight attendant accidentally skips her row when serving lunch.

What if the same personality system is in a different situation, situation 2? Suppose this time the man who sits next to her has a big mustache which reminds her of her uncle, who had a farm and many children, and made her feel welcome whenever she came for a visit, thus activating cognitions and feelings that did not become activated in situation 1. The result may be that the processing units that become activated in this situation are distinctively different from those that become activated in situation 1. As a result, she may be understanding and even somewhat jovial when the flight attendant skips her row.

In short, the network of associations among specific cognitions and affects (i.e., personality structure) that characterizes the person can be invariant across situations, but its behavioral output is expected to vary greatly, and predictably, from one situation to another. The cognitions and affects that are activated change from one time to the next, but the relations between the cognitions and affects activated at one time and those activated next are assumed to reflect the stable personality structure of the individual (Mischel & Shoda, 1995; Shoda & Mischel, 1998). In this case, we would predict that for this person, IF in situation 1, THEN conflicted and angry, but IF in situation 2, THEN warm and friendly. In other words, the behavioral expression of this system is in the set of IF-THEN relations that characterize how the person's behavior varies from one situation to the next, constituting the person's "behavioral signature of personality" (Shoda et al. 1994).

The Role of Internal Situations and the Stream of Cognitive-Affective States

The simplified illustration, however, does not address a key aspect of the model: the pattern of connections among the processing units is assumed to form a recurrent, rather than a strictly feedforward, network (e.g., McClelland & Rumelhart, 1986). Because of the multiple feedback loops present in such a network, "downstream" units can activate "upstream" units, generating a flow of thoughts, feelings, and even behaviors without necessarily requiring an outside stimulus. That is, not only the external input, but also the thoughts and feelings at a given point influence what happens next in the system. The result is something that might resemble a "stream of consciousness." Will there be any stable and distinctive pattern discernible in the ever-changing states of such a system? Or, will it keep changing in essentially random, unpredictable ways?

Network simulations. To address this question, we conducted a computer simulation of personality dynamics, using the basic architecture reported more fully in Shoda and Mischel (1998). The computer simulation can only capture an extremely small subset of the characteristics of the actual CAPS system that generates people's experiences and behavior. Nevertheless, a system that implements such a subset is complex enough so that a computer simulation is useful, or even necessary, to provide insight into the system's functioning and characteristics.

A hypothetical personality was simulated in a 20-unit connectionist network consisting of 5 situation feature detection input units, 10 cognitive-affective processing units, and 5 behavior output units (similar to the one depicted in Figure 1, except the network used in the simulation used a larger number of units). Information travels unidirectionally from the feature detection units into the 10 processing units. The processing units themselves formed a fully interconnected recurrent network, containing many feedback loops. Information then travels from the processing units to the behavior, or output, units. Connection weights for each distinct pair of units were set for each simulated personality by randomly assigning a value drawn from a continuous distribution, ranging from –1 to 1. The connection weights were fixed thereafter, so that in this simulation no learning took place to alter the CAPS network. (In real life, of course, experiences, particularly intense or prolonged exposure to critical events, can change the connection weights, which in the CAPS theory represent *personality change.*)

Figure 1 illustrates the structure of associations among the units, with the lines connecting each pair of these units representing the associations among them. Some of the associations are positive (shown by a solid line), indicating that when one unit becomes activated, it will increase the activation of the other unit to which it is connected. Other associations are negative (shown by a dashed line), indicating that when one unit becomes activated, it will suppress the activation of the other unit to which it is connected. Thicker lines indicate stronger connections, either positive or negative.

To understand the operation of such a system, we "stimulated" the network by temporarily increasing the activation of a subset of the feature detection units. This is analogous to when a person is exposed to a social stimulus. A stimulus was assumed to have a characteristic set of features, which activated corresponding feature detection units in the CAPS network. Activation was assumed to spread from the activated units to other units in the network. In this case, activation would spread from the stimulated feature detector units to the cognitive-affective processing units, to the behavioral output units.

Such spreading activation is not unlike the biological system that inspired the original connectionist modeling. However, in a biological system the activation levels of all the units change simultaneously, and more or less continuously. To simulate simultaneous and continuous changes in conventional computers, which operate serially (i.e., executing one operation at a time, albeit extremely rapidly), a series of small adjustments are made in the activation levels of each unit, implemented in many "cycles" of adjustment. Specifically, in one cycle, the inputs to each unit from all units to which it is connected are summed, and added to its previous level of activation minus a slow decay, such that 98% of its activation in the previous cycle "carries over" to the next cycle (i.e., without any input into a unit, its activation value was 98% of its value in the previous cycle).[1] Then, the adjustment process starts all over again in a new cycle. It has been shown (Anderson, Silverstein, Ritz, & Jones, 1977; Hopfield, 1982, 1984) that after a number of updating cycles these types of networks settle into characteristic stable states, known as *attractors*. In our case, the attractor states represent the characteristic "states of mind" (e.g., a set of beliefs, affective states, etc.) activated in our simulated personality, in response to different combinations of feature detector inputs. In the hypothetical example of an airplane passenger mentioned earlier, the attractor states in the simulation correspond to her states of resentment and security, in response to two different situations.

In our simulations, attractor states were determined by stimulating the network with a variety of stimuli, analogous to encountering a variety of situations, and letting the network "settle" after exposure to each stimulus. Specifically, we created 100 "situations", each of which corresponded to a unique combination of values of the five feature detector units, with values chosen

---

[1] The activation level of each unit was updated following the "squashing" function (e.g., McClelland & Rumelhart, 1986; Shultz & Lepper, 1996), which scaled the effect of net input into each unit in proportion to the distance left to the ceiling (when activation < 0) or to the floor (when activation > 0) of the possible range of activation, which was set from -1 to +1. Specifically,
when $net_i \geq 0$,

$$a_i(t+1) = d \cdot a_i(t) + net_i \cdot \{ceiling - a_i(t)\}$$

when $net_i < 0$,

$$a_i(t+1) = d \cdot a_i(t) + net_i \cdot \{a_i(t) - floor\}$$

Net input into unit i was computed as: $net_i = c * \sum_{i \neq j} (w_{ji} * f(a_j(t)))$, where is the activation level of unit j at time t and is the weight of the connection between the unit j and unit i. The weights were symmetrical, so that $w_{ij} = w_{ji}$. $f(a_j(t)) = \frac{1}{1 + e^{(-\lambda * (a_j(t) - \theta))}}$, which is a sigmoid function of $a_j(t)$ over -1 to +1. Constants are set as follows: c = 0, d = .98, ceiling = 1.0, floor = -1.0, λ = 10, θ = 0.

randomly from a continuous distribution ranging from −1 to 1. "Situation" 1 may be represented by the feature detection unit values −.23, .41, .09, −.35, .17, for example, while "situation" 2 may be represented by .81, −.63, −.14, .31, −.28. After stimulation by each input pattern, activation levels of units in each network were adjusted until no change in activation was detected from one cycle to the next (within the limit of the precision level available to the programming language used). The network reached such stable states with a median of 215 cycles. After the network settled, the activation values for the ten units were recorded.

The attractor states for each network were then identified by submitting the final activation values of the system resulting from each of the 100 inputs to the Partitioning Around Medoids (PAM) clustering algorithm[2], developed by Kaufman & Rousseeuw (1990) and available as an add-on package for the freely available statistical language *R*. The clustering algorithm determined the dissimilarity (indexed by Euclidean distance) among the 100 final states, each of which was characterized by the set of 10 activation values of the processing nodes. We then identified distinctive clusters of final states using a criterion that required a minimum level of dissimilarity between every member of a cluster and members of other clusters. This clustering algorithm required the user to specify the number of clusters. Therefore the data were analyzed with this clustering algorithm 99 times, with the requested number of clusters ranging from 2 to 100, producing 99 sets of cluster solutions. For each solution, we determined the number of clusters that met our criterion for being "isolated"[3], and defined the optimal solution as that with the maximum number of distinctive clusters. We then considered those clusters as attractors. For the simulated person shown in Figure 2, for example, a maximum number of distinct clusters were obtained with a 4-cluster solution (i.e., extracting more than 4 clusters did not increase the number of clusters that met the criterion for being distinctive).

To visualize the dynamic change in the network as it moved from the initial "exposure" to each situation, to an attractor or settled state, we conceptualized the state of the system at any time as a point in a 10-dimensional space, following the "brain state in a box" representation (e.g., Anderson, 1977; Golden, 1986). At any

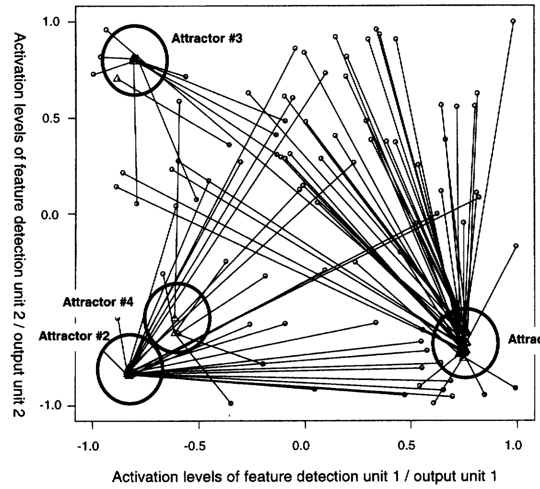

Figure 2. The attractor states of the CAPS network for one simulated person. The network was "exposed" to 100 distinctive "situations." The location of each small circle represents the initial activation levels of the first two feature detection units after exposure to each situation. The x-coordinate of a small circle corresponds to the initial activation level of the first feature detection unit, and the y-coordinate, that of the second feature detection unit. The location of each small triangle represents the final activation levels of the first two behavior output units, after exposure to each situation, and after the network has settled into a stable state. Note because the figure plots the initial activation of only 2 of the 5 input units, *starting* points that appear close in the figure may differ in the other 3 units. Thus two starting points that are close in this figure may lead to different attractors.

point in the updating cycle, the state of the network is defined by a set of 10 values, representing each of the 10 behavior output unit's activation levels. That set of 10 values can be thought of as the coordinates of a location in a 10-dimensional space. As the states change, the set of 10 values will change, which, in this framework, would correspond to a different location in the 10-dimensional space. Thus, the change from one time to the next in the system can be thought of as a point moving through the 10-dimensional space.

Drawing a 10-dimensional space is not possible, so we chose two of the processing units to illustrate the settling process. Figure 2 summarizes the system's response after being exposed to each of the 100 situations. In the figure we plotted the final activation levels of these two output units on the X-Y plane. The x-coordinate of each small triangle at one end of each line corresponds to the final activation level of the first output unit, while the y-coordinate corresponds to the final activation level of the second output unit. The initial states after the system was exposed to each situation are shown in the many small circles on the other end of the lines, with each circle representing the initial values for the first two of the five feature detection input units, superimposed on the same X-Y plane that is used to show the final states of the output units. The x-coordinate of each small circle corresponds to the initial activation level

[2]Downloadable at http://cran.r-project.org/src/contrib/PACKAGES.html#cluster.

[3]The PAM algorithm identifies a user-specified number of clusters that minimize intra-cluster distances while maximizing inter-cluster distances. For each cluster, PAM computes an index of *separation*, the minimal distance between a member of the cluster and a member of another cluster. We considered a cluster to be distinct when its separation was 0.3 or greater (i.e., clusters whose members were at least a distance of 0.3 away from a member of another cluster, in a ten dimensional space where each dimension ranged from -1 to +1).

of the first feature detector unit, while the y-coordinate corresponds to the initial activation level of the second feature detector unit.

The distribution of the small circles on the X-Y plane corresponds to the locations of the input space we sampled. With 100 input vectors sampled, it still would not be fair to call it an exhaustive sample, but we can see most of the input space is represented to some degree. The graph shows that different feature constellations produce different end states for the same individual. But instead of the 100 different input patterns leading to 100 different end states, there were 4 clusters of end states, or attractors, shown in large circles. That is, one group of IFs lead to one THEN (i.e., a distinct attractor state), while another group of IFs led to a different THEN. Thus, for this simulated person, there were four distinct groups of IFs (input patterns) each leading to a distinct THEN.[4] These clusters may be thought of as corresponding to functional equivalence classes of situations that Gordon Allport hypothesized as an expression of an individual's personality, or "a neuropsychic system (peculiar to the individual) with the capacity to render many stimuli functionally equivalent, and to initiate and guide consistent (equivalent) forms of adaptive and expressive behavior" (1937, p.295).

Forty personalities were simulated in this manner, with each exposed to 100 unique feature detector unit input patterns. For each input pattern, the activation levels of the network units were adjusted until the network reached a stable state, and the final activation values for the 10 processing units were recorded. Using the clustering technique on these 10 values, for each simulated personality, we were able to determine the number and the size and location of the attractors for each of the simulated personalities. The 40 simulated personalities had, on average, 2.18 attractors of varying size and location.

Extending the Model to Address Interpersonal Systems: The Personality of a Dyad

An attractor state is a distinct pattern of activation in cognitive and affective units that is particularly stable; once a network is in that state, it is unlikely to change easily. As shown earlier, the personality system, modeled as a recurrent network, "settles" into one of its distinctive attractor states after encountering a situation that activates some of the cognitive-affective units. Furthermore, in response to different situations, the system may settle into different attractor states.

In the previous simulations, it was assumed that after an initial exposure to situations, the subsequent processing of the initial stimulus was carried out by the network in isolation, without further external inputs or constraints. Such a process may account for what might happen, for example, when an individual is watching a TV news program alone. As the individual encounters a news item, her mind is stimulated by it, and after a chain of associations, may settle into a state of recurrent thoughts, corresponding to one of her cognitive-affective network's attractor states. When the next topic is presented, the process repeats, this time settling into another attractor state, if the new item activates a different set of thoughts. Note that in this example, the situations (i.e., the news stories) are determined by factors outside the individual's influence. That is, the thoughts, feelings, and behaviors of the individual do not (under usual circumstances) influence the next item featured on the news.

Real-life, however, is replete with situations that are not determined independently of a person's actions, and indeed they tend to be the rule rather than the exception. For example, imagine a couple in a heated argument. In such a dyadic interaction, what one person says, the tone of voice, or even the quick glimpse of a facial expression, can significantly affect the other person's thoughts, feelings, and behavior. Furthermore, to the extent that the dyad represents a significant relationship, it would not be easy, or adaptive, to ignore one another. Every reaction of one individual may count, in determining the partner's cognitive, affective, and behavioral responses. Could the present framework be extended to model interactions in a dyadic system in which the response of one individual is closely coupled to the response of the other?

This theoretical possibility was explored in a review of recent empirical literature (Zayas, Shoda, & Ayduk, in press). In the present article, we address it computationally by simulating a dyadic system that combines two networks, each representing a person (Figure 3). Each individual CAPS network was constructed in exactly the same way as in the simulation described earlier and shown in Figure 1. In contrast to the previous simulations, however, this time pairs of networks were combined by connecting the behavior units of one network to the feature detector units of the other network and vice versa. In this manner, each individual network becomes part of a larger parallel constraint satisfaction system. This coupling of individual networks is conceptually akin to recent work on coupled dynamical systems (see Nowak & Vallacher, 1998), which modeled the properties of dyads such as relationship synchronization and the manner in which a dyad reaches equilibrium. Gottman and colleagues (e.g., Gottman, Swanson & Swanson, this issue) have

[4] In general, the larger the network, the larger the number of attractors. Thus in a more realistic network representation of an individual containing more than 10 processing units, the number of attractors, of course, should be much larger than is illustrated in the present simulation.

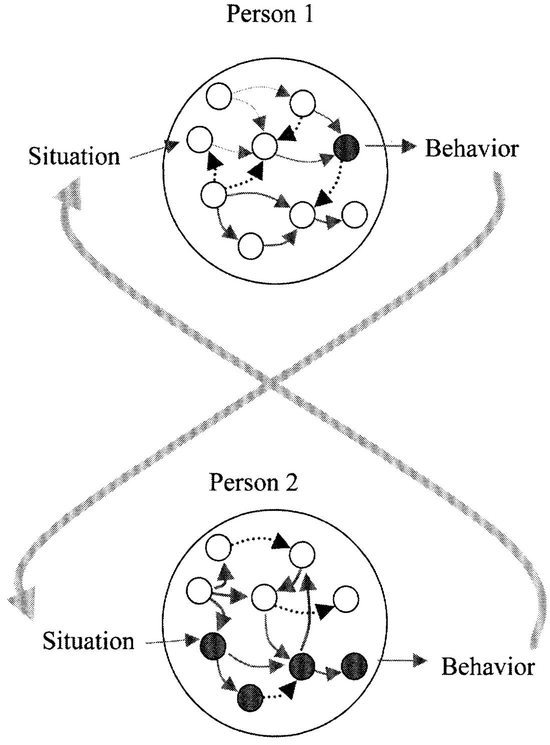

Figure 3. A dyadic system is formed in which the behaviors of one member activate the feature detection units of the other, and vice versa.

applied a dynamical model to understand the characteristics of interactions in married couples

We hypothesized that when two CAPS networks are connected to each other, configurations of thoughts and feelings that previously were not chronically activated may now become chronic. That is, when interlocked in a dyad, individuals in the dyad may be characterized by a new set of recurring thoughts, affects, and behaviors that they did not have when alone. This may occur if partner A's behavior activates certain thoughts and affects in partner B, and if in turn these thoughts are expressed in the behavior by partner B that led to partner A's behavior, thus forming a feedback loop that reinforces these new thoughts, affects, and behaviors. More specifically, we predicted that while each person's CAPS network has a set of characteristic attractor states, when they are interlocked in a dyadic system the individual CAPS network will not necessarily settle into the original set of attractors. Rather, they may settle into a new set of attractors. Such a positive feedback loop might underlie the "chemistry" of interpersonal relationships. The simulations to test these hypotheses followed these steps:

1. For each of the 40 simulated personalities, the attractor states of each CAPS network were identified by "stimulating" each, in isolation, with 100 stimulus patterns and allowing them to settle into a stable state after exposure to each stimulus pattern (as described earlier).

2. Twenty dyads were formed out of the original set of 40 simulated personalities. The two networks constituting each dyad were then "interlocked" by connecting person A's behavior units to person B's feature detector units, and vice versa. Specifically, behavior units #1, 2, 3, 4 and 5 of person A activated situation feature detectors #1, 2, 3, 4 and 5 of person B, respectively, while behavior units #1, 2, 3, 4 and 5 of person B activated situation feature detectors #1, 2, 3, 4 and 5 of person A, respectively.

3. Then the attractor states of these two interlocked networks were identified. We re-stimulated Person A's feature detector units with the same 100 input patterns used when her network was simulated alone. Because Persons A and B are "interlocked," the behavioral output of Person A in response to each stimulus pattern becomes the input into Person B, whose behavioral output in turn becomes the input to Person A, starting another cycle of interaction. The activation in person A's network spreads through to person B's network and back, creating not only intra-individual but also inter-individual feedback loops. The simulation continued over many cycles, until there was no change in the states of either person. The attractor states resulting from these simulations were compared with those resulting from being simulated alone (step 1). This procedure was repeated for Person B.

To test for the effect of the dyadic interaction, we compared the attractor states before and after the individual networks were paired in a dyad. Would each individual network settle into different attractor states now that it was part of the larger parallel constraint satisfaction network of the dyad? The answer seemed to be yes. As an illustration, compare Figure 4 to Figure 2. Figure 4 shows the attractors for the same simulated person as in Figure 2. But now the person is part of a dyadic system (see Figure 3). Note that before the network became part of a dyadic system it had four attractors, but when part of the dyadic system, it had three.[5] Furthermore, except for attractor #3, the new attractors were not the same as those shown in Figure 2.

These results are summarized in Table 1. In order to test if the difference in the distribution of attractors between when the network was stimulated alone and when it was embedded in a dyadic system might simply be due to chance, we computed χ^2 to compare the two distributions of attractors. For the results shown in Table 1, χ^2 was 191.3 ($df = 5$), with $p < .0001$. The effect size, indexed by Cohen's *e* (Cohen, 1977) which ranges from 0 (no effect) to 1 (maximum effect) was 0.96.

[5]As when the networks were simulated alone, the number of attractor states was derived from the final activation values of all 10 processing units, not just the two units shown in the figure.

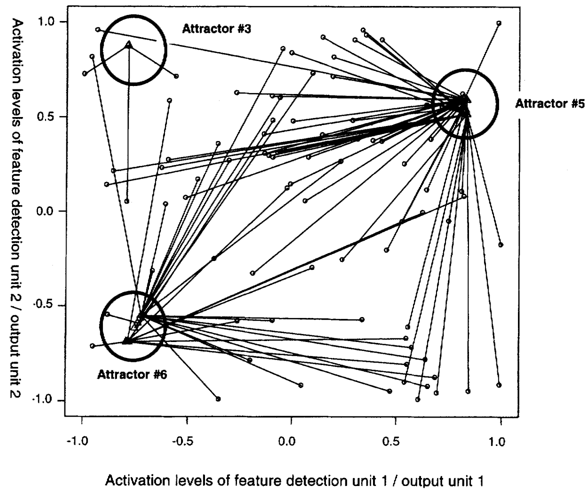

Figure 4. **The same CAPS network that produced the attractor states shown in Figure 2 were now embedded in a dyadic system with another CAPS network. The x-coordinate of a small circle corresponds to the initial activation level of the first feature detection unit, and the y-coordinate, that of the second feature detection unit. The location of each small triangle represents the final activation levels of a processing unit, after exposure to each situation, and after the network has settled into a stable state. The x-coordinate of a small triangle corresponds to the final activation level of one processing unit, and the y-coordinate, that of the second unit.**

Summary of Results Across All Dyads

Nineteen of the 20 dyads resulted in both members of the dyad producing at least one attractor both alone and when embedded in a dyad.

For each of these 19 dyads, we conducted the same analysis as described earlier for the example shown in Figures 2 and 4 and in Table 1. The mean χ^2 was 59.1 (mean df = 2.03), with a mean Cohen's *e* effect size index of .51. In 16 of the 19 dyads, at least one member of a dyad underwent changes in the distribution of attractors that were statistically significant at $p < .05$.

Concluding Thoughts

We have outlined and illustrated a conception of personality that allows a reconciliation of the human qualities of both change and stability and that indeed predicts them in the expressions of a personality system. This conception accounts for intraindividual variability in behavior, as well as for overall average individual differences in behavior tendencies, both enduring human characteristics (Mischel & Shoda, 1995). That is, although the particular thoughts and affects activated at a given moment change within a person, the internal organization of the cognitive-affective processing system itself remains relatively stable and invariant, at least in the short term, from situation to situation. The stable structural properties of the cognitive-affective processing system guide the dynamic activation within the network of particular cognitions and affects activated by a given situation. In turn, different sets of cognitions and affects lead to different behaviors, but to the extent that a person encounters situations with similar features, similar behavioral responses are generated. Key properties of the system were illustrated in computer simulations, in which individuals' characteristic cognitive-affective processing systems were modeled as a recurrent, parallel constraint satisfaction network. The results showed that an individual is likely to have a distinct set of attractor states, perhaps corresponding to recurrent thoughts, feelings, and behaviors, and functionally equivalent groups of situations emerge, such that IF encountering situation class X, THEN the system settles in attractor Y. The simulation illustrated how such aspects of personality can be an emergent quality, based on the stable and distinctive cognitive-affective network that characterizes an individual and his or her interaction with the environment. In addition, because this conceptualization of personality explicitly models the effect of situations on a given individual, it can also be used to model the function of interpersonal systems. We demonstrated this possibility by modeling dyadic systems in which one partner's behavior becomes the situational input into the other partner's personality system, and vice versa. The results indicated that each member of the dyad will, in general, exhibit new attractor states, suggesting that the thoughts, affects, and behaviors that an individual typically experiences are a function not of that individual's personality system alone, but rather a function of the *interpersonal system* of which the individual is a part.

There is much commonality between the present results and those shown by Gottman, Swanson, and Swanson (this issue). Both show that a model of dynamical interactions between pairs of individuals predict stable interaction patterns that reflect the

Table 1. *Attractors for the Network Shown in Figure 4, as a Function of Whether the Network was Isolated ("Alone") versus Embedded in a Dyadic System.*

	Attractor #1	Attractor #2	Attractor #3	Attractor #4	Attractor #5	Attractor #6
Alone	60	25	8	7	0	0
Embedded in a dyad	0	0	3	0	64	33

Note: Entries indicate the number of situations leading to each attractor. χ^2 (df = 5) = 191.3, $p < 10^{-38}$, Cohen's *e* = 0.96.

unique combination of the two people in the dyad, and that this is an emergent quality in that it is not a simple combination of the behavioral tendencies that exist in the two individuals alone. The present work and the Gottman et al. work complement each other in that the latter directly models the influence of one person on the other, with the possibility of empirically assessing the nature of these influences from observable data. The CAPS networks, on the other hand, focus on the intra-individual dynamics, that is, interactions among the cognitions and affects within an individual. Thus the present work can be thought of as modeling the intra-individual processes that may underlie the nature of the influence one person has on the other, which is the focus of the Gottman et al. model. By assessing the pattern of automatic associations among cognitions and affects that form individuals' CAPS networks to predict the nature of influence one person has on another if they form a dyad.

Furthermore, there are a number of questions that could be addressed by a relatively simple extension of the present work. For example, what would be the result if every possible combination of individuals were simulated? Would each distinctive pair result in a unique set of attractors, or would some individuals tend to result in repeated patterns of relationships, such as always getting into hostile and conflicted relationships across a range of partners? And if so, is it because the individual is relatively impervious to situational input, or is it because the individual manages to provoke a predictable pattern of reaction from others? Furthermore, real-life dyads' CAPS networks are related due to partner selection (e.g., Zayas et al., in press) and mutual influences over the course of a relationship, for example. What would be the effect of such relationships between the CAPS networks of partners in a dyad?

The simulation of the dyadic system illustrated here is one example of what may be possible with a model of personality that incorporates, rather than excludes, the effect of situations. It models how people respond cognitively, affectively, and behaviorally, to the important situations in which their lives are contextualized, as illustrated in dyadic relations with significant others. Such phenomena as the multiple "relational selves" that define identities (e.g. self as partner, self as daughter, self as mother) are increasingly recognized in social psychology and self-theories (e.g. Mischel & Morf, in press), although they are not routinely included in most current approaches to personality and individual differences. Our hope is that such an approach will allow phenomena like the relational nature of personality to be studied with increasing precision. Another fruitful line of research may be to investigate individual differences in the structure of these networks. For example, individuals may differ in the number and "depth" of attractors, as well as the complexity of their CAPS networks. Such differences may in turn allow one to predict the kinds of stabilities and constancies in observable behaviors that are likely to be expressed.

The depiction of a person condensed into a single still photo with an average of joy and sadness, hope and despondence, anger and conciliation may provide a useful composite for many purposes. But to begin to account for the individual's characteristic dynamics, one needs to examine the processes that underlie them, as expressed in the patterns of distinctive IF-THEN personality signatures. A glimpse into such processes comes from analyzing the structure of the person's cognitive-affective processing sytem (Shoda & Mischel, 1998). We propose that just as individuals have distinctive and stable IF-THEN signatures, so do interpersonal relationships (e.g., "This couple always starts arguing when life gets too good"). By understanding some of the structure of the cognitive-affective processing system of each relationship partner we also may become able to predict their distinctive relational signatures.

References

Allport, G. W. (1937). *Personality: A psychological interpretation.* New York: Holt, Rinehart, & Winston.

Andersen, S. M., & Baum, A. B. (1994). Transference in interpersonal relations: Inferences and affect based on significant-other representations. *Journal of Personality, 62*, 460–497.

Andersen, S. M., & Cole, S. W. (1990). "Do I know you?": The role of significant others in general social perception. *Journal of Personality and Social Psychology, 59*, 384–399.

Anderson, J. A. (1972). A simple neural network generating an interactive memory. *Mathematical Biosciences, 14*, 197–220.

Anderson, J. A. (1977). Neural models with cognitive implications. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading perception and comprehension* (pp. 27–90). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc..

Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: some applications of a neural model. *Psychological Review, 84*, 413–451.

Cohen, J. (1977). *Statistical power analysis for the behavioral sciences.* New York: Academic Press.

Golden, R. M. (1986). The 'brain-state-in-a-box' neural model is a gradient descent algorithm. *Journal of Mathematical Psychology, 30*, 73–80.

Gottman, J., Swanson, C., & Swanson, K. (2002). A general systems theory of marriage: Nonlinear difference equation modeling of marital interaction. *Personality and Social Psychology Review, 6*, 326–340.

Higgins, E. T. (1990). Personality, social psychology, and person-situation relations: Standards and knowledge activation as a common language. In L. A. Pervin (Ed.), *Handbook of personality: Theory and research* (pp. 301–338). New York: Guilford.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York: Guilford.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, 79*, 2554–2558.

Hopfield, J. J. (1984). Neurons with graded responses have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, 81*, 3088–3092.

James, W. (1890). *The Principles of Psychology*. New York: Holt.

Kashima, Y., & Kerekes, A. R. Z. (1994). A distributed memory model of averaging phenomena in person impression formation. *Journal of Experimental Social Psychology, 30*, 407–455.

Kaufman, L., & Rousseeuw, P. J. (1990). *Finding groups in data: An introduction to cluster analysis*. New York: Wiley.

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

Kohonen, T. (1972). Correlation matrix memories. *IEEE Transaction on Computers, C-21*, 353–359.

McClelland, J. L., & Rumelhart, D. E. (1986). A distributed model of human learning and memory. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructures of cognition: Vol. II. Psychological and biological models* (pp. 170–215). Cambridge, MA: MIT Press/Bradford Books.

Mendoza-Denton, R., Shoda, Y., Ayduk, O. N., & Mischel, W. (1999). Applying Cognitive-Affective Processing System (CAPS) Theory to cultural differences in social behavior. In W. J. Lonner, D. L. Dinnel, D. K. Forgays, & S. A. Hayes, (Eds.), *Merging past, present, and future in cross-cultural psychology: Selected papers from the Fourteenth International Congress of the International Association for Cross-Cultural Psychology* (pp. 205–217). Lisse, Netherlands: Swets & Zeitlinger.

Mischel, W. & Morf, C. (in press). The self as a psycho-social dynamic processing system: A meta-perspective on a century of the self in psychology. In M. Leary & J. Tangney (Eds.), *Handbook of Self and Identity*. New York: Guilford.

Mischel, W., & Peake, P. K. (1982). Beyond déjà vu in the search for cross-situational consistency. *Psychological Review, 89*, 730–755.

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

Nowak, A., & Vallacher, R.R. (1998). *Dynamical social psychology*. New York: Guilford.

Read, S. J., Jones, D. K., & Miller, L. C. (1990). Traits as goal-based categories: The importance of goals in the coherence of dispositional categories. *Journal of Personality and Social Psychology, 58*, 1048–1061.

Read, S. J., & Marcus-Newhall, A. (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology, 65*, 429–447.

Read, S. J., & Miller, L. C. (1994). Dissonance and balance in belief systems: The promise of parallel constraint satisfaction processes and connectionist modeling approaches. In R. C. Schank & E. Langer (Eds.), *Beliefs, reasoning, and decision making: Psycho-logic in honor of Bob Abelson* (pp. 209–235). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shoda, Y., & Mischel, W. (1998). Personality as a stable cognitive-affective activation network: Characteristic patterns of behavior variation emerge from a stable personality structure. In S. J. Read & L. C. Miller (Eds.), *Connectionist and PDP Models of Social Reasoning and Social Behavior* (pp. 175–208). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shoda, Y., Mischel, W., & Wright, J. C. (1994). Intra-individual stability in the organization and patterning of behavior: Incorporating psychological situations into the idiographic analysis of personality. *Journal of Personality and Social Psychology, 67*, 674–687.

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

Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology, 62*, 913–933.

Spellman, B. A., Ullman, J. B., & Holyoak, K. J. (1993). A coherence model of cognitive consistency: Dynamics of attitude change during the Persian Gulf war. *Journal of Social Issues, 49*, 147–165.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences, 12*, 435–467.

Zayas, V., Shoda, Y., & Ayduk, O. N. (in press). Personality in context: An interpersonal systems perspective. *Journal of Personality*.