Extending Identity Control Theory: Insights from Classifier Systems*

PETER J. BURKE
University of California, Riverside

Within identity control theory (ICT), identities control meaning and resources by bringing perceptions of these in the situation into alignment with references levels given in the identity standard. This article seeks to resolve three issues in ICT having to do with the source of the identity standard, the correspondence between identity standards and the identity relevant meanings perceived in the situation or environment, and the activation of identities. Classifier systems, as developed by John Holland, are inductive, flexible, rule-based, message-passing, adaptive systems that are able to learn, to fit in, and to adapt to various and changing environments. Classifier systems are introduced and are extended to incorporate the central components of the model of identity as held in ICT. In this manner, a new identity model is proposed that has inductive and adaptive capacities able to resolve the three issues identified. Unexpectedly, the new model also lays the groundwork for the possible resolution of a long-standing issue within the symbolic interaction framework concerning the origins of shared meanings.

INTRODUCTION

Within identity control theory (ICT), identities control meaning and resources by bringing perceptions of these in the situation into alignment with references levels given in the identity standard. There are three important issues that have remained unresolved in identity control theory and without their resolution the real link between identity and social structure remains glossed. It has always been recognized that identities operate within social structure, but the nature of the connection has remained obscure. These issues concern (1) the origin of identity standards; (2) the correspondence between the perceptual inputs of identity relevant meanings in the situation and the identity standard (i.e., the fact that the dimensions of meaning input to the identity system are the same dimensions of meaning that are calibrated into the identity standard so that they can be compared); and (3) the conditions that activate identities.

Past work has provided some movement toward resolving these issues; for example, with respect to the first issue on the origins of identity standards, Stryker (1980) has noted that identity standards arise in the societal expectations that are held for persons occupying particular positions in the social structure upon which positions role identities are based. Thus, one acquires an identity standard by being socialized into the expectations that exist for a person in that position. However, socialization is not defined within the current scope of ICT. Nor is it clear in the literature what is involved in socialization other than gaining some internal standard or set of standards that guide behavior or why, given two persons with similar backgrounds, one

---

*I wish to thank Jan Stets for her insights and helpful comments on earlier drafts of this paper. Direct correspondence to: Peter J. Burke, Department of Sociology, University of California, Riverside, CA 92521-0419; email: peter.burke@ucr.edu.

Sociological Theory 22:4 December 2004

will acquire one standard while the other may acquire a different standard. What is involved in the learning of identity standards? And how are standards set for new identities that evolve over time in society? Clearly, the answer of *socialization* requires additional theorizing.

Similarly, with respect to the third issue of the conditions that activate an identity, Stryker (1980) has pointed out that the salience of an identity (i.e., its likelihood of being invoked in a situation) is a function of the commitment one has to the identity (Stryker and Serpe 1982; Stryker 1980). Thus, a person with strong commitments to a role identity, by being tied to many others through that identity and by having a strong emotional tie to the identity, will be more likely to activate that identity. Again, however, knowing that an identity is more or less likely to be invoked in a situation does not tell how it is activated or exactly what the activating mechanism is.

To begin to resolve these issues, this article extends ICT by placing it within the context of Holland's (1975) model of classifier systems. A classifier system is a system of conditional rules that evolve adaptive strategies to survive and to become more "fit" within a (possibly changing) environment, much as an organism uses genetic mutation to adapt to an environment. In this way, a new model of an identity is created that selects goals from its perceptions that make it more "fit" and able to survive in its environment. Yet, having its own goals, it is able to maintain its own agency as a control system. Bringing these two systems together helps move toward a resolution of the aforementioned issues in ICT concerning where identity standards originate, why self-perceptions are relevant to those standards, and how identities come to be activated.

THE IDENTITY MODEL

Identities and Meaning

Identity control theory was developed within the framework of structural symbolic interaction (Stryker and Burke 2000; Stryker 1980). Following James's (1890) notion of multiple selves, ICT posits that each of these selves correspond to an identity (or role-identity) within the overall structure of the self. The self-concept has always been considered to be a set of meanings describing who one is (Stryker 1980; Blumer 1962; Mead 1934). Within the structural symbolic interactionist framework (Stryker 1980), rather than the overall global self-concept having meaning, each of the multiple identities is seen as having its own set of meanings. Based upon Osgood, Suci, and Tannenbaum's (1957) conception of meaning as a mediational response, Burke and Tully (1977) propose a method for measuring the meaning of identities. In so doing, they change the understanding that identities are meanings from a mere assumption of symbolic interactionism to a quantitatively researchable idea. This allows researchers to investigate the self as an active agent by investigating the consequences of having particular identities.

This investigation required a conceptualization of the link between identities and behaviors in a manner consistent with the developing theory of measurement of identities. Burke and Reitzes (1981) point out that identities are linked not to particular behaviors but to the meanings of behaviors and the effects of behavior on other meanings in the situation. In this sense, meaning is the metric relating behaviors and identities. Both identities and behaviors are to be measured in terms of their meanings on identical dimensions.

Taking the self as an active agent one step further, ICT (Stryker and Burke 2000; Burke 1997; Burke 1996), affect control theory (Smith-Lovin and Heise 1988), and self-verification theory (Swann 1983) all conceptualize identity processes as cybernetic control systems that regulate the perception of self-relevant meanings of actors in situations (see Figure 1). In this model, the perceived meanings of the self that are implied in the situation (reflected appraisals and self-appraisals) are compared with meanings held in the identity standard. This can be seen in Identity 2 on the lower right side of the figure. Discrepancies or differences between these two, as indicated by the comparator, change the behavior of the actor, which changes in the situation, which in turn change the perceptions. Such changes in **action → situation → perception** continue until the discrepancy between perceptions and the identity standard are eliminated. In this model, the behavioral choices of individuals are outcomes neither of identities alone (as an automation) nor of the situation alone (as an S-R model). Instead, behaviors are seen as outcomes of the difference between an identity (meanings contained in the internal standard) and perceptions of situation (perceptions of self-relevant meanings reflected in the situation). However, it would be a mistake to think that particular discrepancies lead to particular behaviors. Rather, behaviors are chosen for their consequences, i.e., reducing the discrepancy, and any

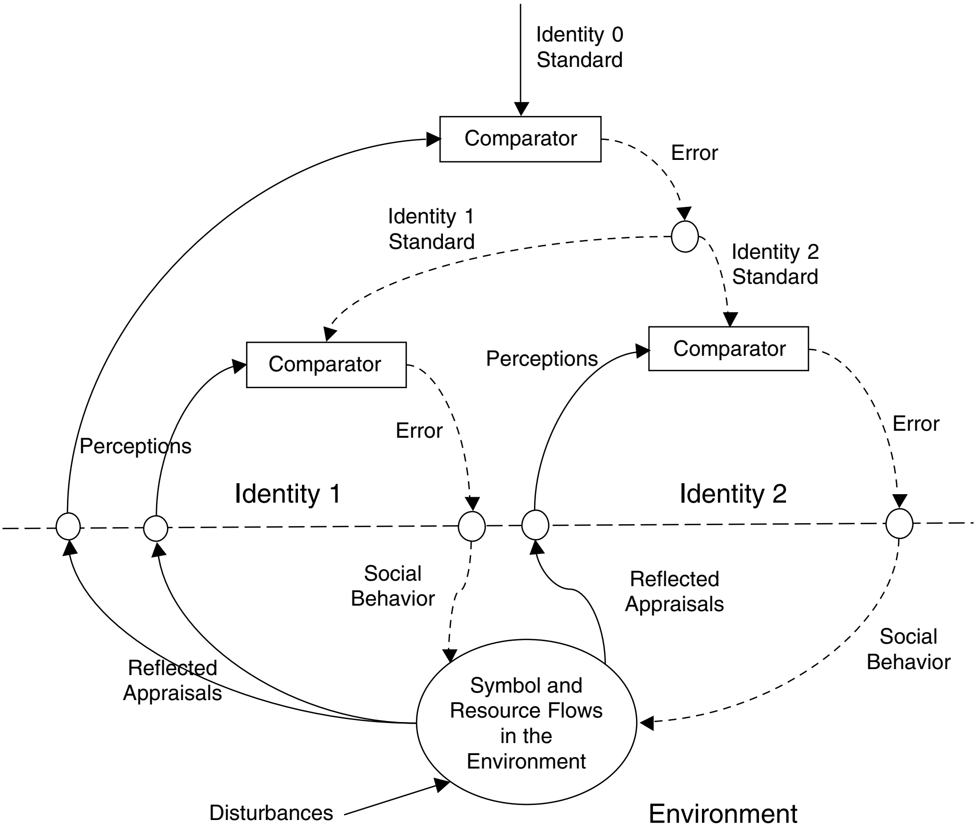

Figure 1. Identity model showing two levels and two identities.

behavior that has this result could be used, for it is the result that is important and not the means by which it is accomplished.

Recent developments of ICT have involved a broadening of the notion of *meaning* (Freese and Burke 1994). Traditionally, symbolic interactionists have viewed meaning in terms of symbols, that is, substitute stimuli that lead to shared, arbitrary, socially conventionalized responses to stimuli. Indeed, by focusing almost exclusively on shared symbols (as opposed to signs), the objects and processes that are symbolized and the nonshared aspects of social life have been ignored almost completely. Signs, sometimes called *natural signs* (Lindesmith and Strauss 1956), are also substitute stimuli that have been learned but responses to which may not be shared with others. Signs are not used in communication (not being shared) but act as signals, telling us about the world around us (resources). While signs have been largely ignored in the research literature, Lindesmith and Strauss (1956:53) point out, "The world . . . is full of [signs], and behavior is largely to be accounted for by responses to them." Continuing on, they say, "Natural signs are not 'natural' in the sense that they occur only in nature. They also may be human artifacts, as when the psychologist sets up a sequence of buzz—food in a dog's experience, or when lines in a spectrum are taken as evidence of the presence of certain elements. Similarly, the click of a Geiger counter is a natural sign of the passage of [a radioactive particle]" (1956:56). For this reason, it becomes important to focus not only on symbols but also on the things that are symbolized and to focus as well on signs and the things that are signified (Freese and Burke 1994).

Identities and Resources

Following the work of Freese (Freese and Burke 1994; Burke and Freese 1989; Freese 1988), *resources* may be defined as processes that sustain a system of interaction (loosely, individuals, groups, organizations, and so forth). This is in contrast to the usual notion held in sociology and social psychology that views resources as scarce, valued, consumable possessions. Defining resources in terms of their effects on a system of interaction (to sustain it) means that anything that acts to sustain a system of interaction is a resource, even if it is not recognized as such.

As a consequence, I am dealing with resource functions, and I view the interaction processes of any social system in terms of connected resource flows (e.g., flows of money, power, oil, air, love) (cf. Freese 1988). From this perspective, food and money would be resources insofar as they are used to sustain individuals in interaction, but love and attention also would be resources insofar as they too sustain persons in interaction. So too would be an opportunity that is taken, a sequence of activity that restores health, or a change in the sequencing of information in a database. Signs allow us to manipulate the world because they signal the current state of objects and processes in the environment. We can reach for a pen to write with or a glass to drink from; we can push a car that will not start or chop down a tree for fire. Taking a biosocial approach, Freese and Burke (1994) suggest the possibility that symbols evolved from signs in order to facilitate the control of resources by dealing with *potential resources*, i.e., resources that need to be symbolized because they exist in the future or otherwise are not present in the situation. Thus, underlying the control of symbolic meanings is the control of resources to sustain the biophysical aspects of persons and interaction.

In this expanded view of identities, it is suggested that not only do identities monitor (perceive) symbolic meanings arising in interaction and compare them with the meanings of the identity standard, but identities also monitor various resources by comparing sign meanings of these with additional meanings held in the identity standard. Thus, each identity standard contains meanings about the way that signs and symbols are "supposed" to be in a situation,[1] and if there are disturbances that modify resource and symbol flows so that perceptions of these meanings no longer match the meanings held in the standard, then action is taken to modify the situation and to bring the perceived resource and symbol flows back into alignment with the standard.

An identity thus monitors and controls both the symbols and resources associated with that identity. The identity standard therefore must contain information not only about the symbolic meanings of who one is (in that identity) but also about the sign meanings of the resources that are to be managed directly by the identity. As an example, consider a person with a parent identity. The parent must monitor *resources* to provide food and security, sufficient love, warmth, and attention for the child, as well as intellectual stimulation and a host of other things needed to support the child. In addition, the parent may have to monitor *symbols* of authority and power, as well as all of the *symbols* needed to plan, organize, and assure potential resources (that are not present in the interactive setting).

Social Structure, Resources, and Identities

ICT asserts a link between micro and macro processes in defining and understanding social action. Specifically, this link is identified, in part, in the actions (outputs) of identities that modify the local (situational) transfer and transformation of resources and symbols (micro), whose large-scale flows and coordination define social structure (macro). Structural symbolic interactionism also assumes the existence of social structure as an integrated and coordinated flow of symbols and resources, which serves as a "field" in which identities as agents carry out their actions (Burke 2004). Larger social structural characteristics of the individual (e.g., age, gender, race) both constitute locations in the large-scale flow of resources and symbols and provide individual variation in the contents of the identity standards that guide the local behavior of individual identities (Stets and Burke 1996).

The concept of managing resources, from a structural symbolic interactionist perspective, can be understood best in the context of sustaining a given social system in which the role (attached to the role identity) is embedded. The identity specifies choices a person may make with regard to resources. By serving as an active agent with respect to these resources (and meanings), the identity serves to sustain the social structure within which it is embedded. This is the other part of the link between micro and macro perspectives on interaction. From an ICT perspective that incorporates the concept of resources, the structure of the self remains stable over time to the extent that the individual is successful in managing the resources necessary to sustain and to maintain the social system in which the role and identity are embedded. Of course, neither the identity nor the social structure is maintained perfectly, since there continue to be disturbances and changing conditions. Adaptation to the new conditions must be part of the normal operation of the identity processes, and these adaptations and changes then are reflected in changes in the social structure.

---

[1] "Supposed" to be, not as a moral imperative, but definitionally, because these meanings are part of the standard toward which perceived meanings are adjusted.

Issues

The control model on which identities operate works well when the standard is known and fixed and when the perception mechanisms are in place. Each of these assumptions, however, may be problematic. With respect to the question of where the identity standard comes from, Powers (1973) suggests that standards for a control system are the outputs of a hierarchically higher-level control system. While such an answer ultimately leads to a problem of infinite regress, for our purposes it only means that identity standards are variable to the extent that the higher-level control system changes its outputs in response to a discrepancy between its inputs and its standard. Stability and change in identities are thus aspects of the more complete identity system as suggested in the whole of Figure 1, showing a higher-level control system (Identity 0) and two lower-level systems (Identities 1 and 2) (Burke 2004; Cast, Stets, and Burke 1999; Burke and Cast 1997). Nevertheless, this does not tell where the standard came from originally—i.e., why is there a higher-level control system, and why are the content meanings of the standard as they are? This is the first issue that a more complete identity theory must answer.

The second issue stems from the need to explain one interesting feature of an identity as a control system. That is the fact that identity inputs (perceptions) are scaled in the same units and dimensions as the identity standard. Concretely, if a person's identity standard is calibrated in terms of a certain degree of, for example, powerfulness, the person observes the degree to which she or he is powerful in a situation and not the degree of liveliness or the degree of goodness (as examples) she or he seems to possess in the situation. The question is how do identities come to perceive and to monitor the things coded into the identity standard? The first and second issues clearly are related, since the origin of a standard may have something to do with the perception of meanings in a situation. For example, we know in terms of the way people operate that some of the standards are present by reason of instruction; we are told in our jobs, for example, what to do, how to do it, and how to know when it is done. Other standards emerge through an adaptive learning process as we become more familiar with a situation and what we need to do and to monitor in order to accomplish particular tasks. The questions posed in the first two issues are how do these standards come to exist and evolve, and how do we learn what to monitor in order to accomplish particular goals? ICT currently does not address these issues.

The third issue a more complete identity theory must address concerns the activation of identities. Again, the control system underlying the identity models assumes we are focusing on an active identity. The phrase "when an identity is activated" prefaces much of the work on identities because it is assumed that an identity must be activated before it can operate in the manner prescribed by the model. If there were only one identity, there would be no problem, since it could remain activated constantly. However, because people have several identities, there must be some mechanism that indicates when a particular identity is to be activated or deactivated. While the concept of *salience* (Stryker 1980) is used to indicate the likelihood of activation of an identity (more salient identities by definition are more likely to be activated), there is nothing in the concept of salience that tells what activates an identity. Similarly, the concept of commitment has been used in ICT to explain variations in salience (Stryker 1980), but this explanation concerns the likelihood of activation and not the specific situational conditions. The question raised is what activates one identity and deactivates another?

Because classifier systems incorporate ways of representing the environment (situation) within the system, as well as ways of adapting the system to the environment, and because the issues raised all deal with the source of internal representations (goals or standards) and the linking of perception with these internal representations (in terms of triggering mechanisms and correspondences), a possible answer to these three issues (and perhaps others) lies in modeling identity using a modified classifier system. Thus, at least in a preliminary way, we will be able to answer the questions where does the identity standard come from, what activates and deactivates identities, and why is the identity standard scaled in the same units as the input perceptions.

CLASSIFIER SYSTEMS

Classifier systems were introduced by Holland (1986, 1975; Holland et al. 1986) as general-purpose, passive, adaptive, learning algorithms. They monitor conditions and modify themselves to adapt to a constantly changing environment. Structurally, they are message-passing, rule-based systems in which large numbers of rules can be active simultaneously. As seen in Figure 2, there are four main parts or modules to the classifier system. The first is a *message board* that holds all messages to be processed. The second is a large set of *classifiers* or *rules*. All rules are in the form of condition/action pairs: for each condition there is an action. Each condition specifies the subset of messages that satisfy it, and each action specifies a message to be sent when the condition is satisfied. The heart of the system thus is concerned only with message processing as each classifier or rule scans the message board (simultaneously) for

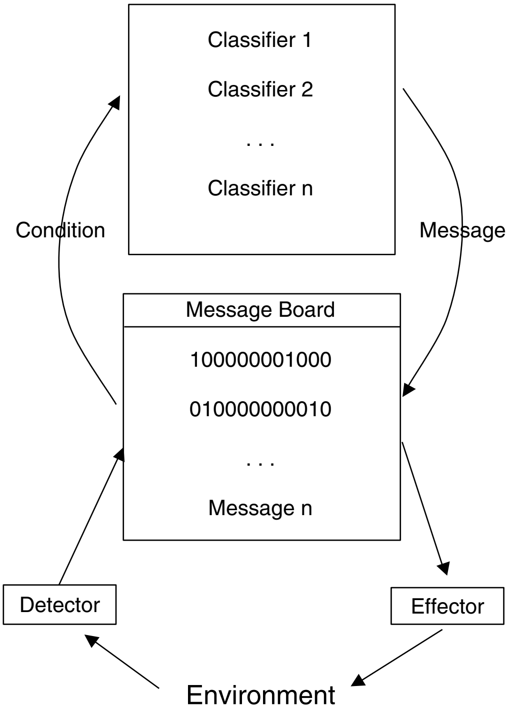

Figure 2. Heuristic model of classifier system.

messages that meet the conditional part of the rule. For each message that satisfies the conditional part of the rule, the classifier acts to post its action/message onto the message board.

Two additional modules interface between the system and its environment. The first is a *detector*, which sends messages to the message board representing conditions in the environment, such as whether a resource is present or not, whether it is located to the left or to the right, and so on. The second module is an *effector*, which causes the classifier system to take actions in the environment based upon messages contained on the message board. These actions might be to turn left or to turn right or to move forward, for example. There are two tasks that need to be performed to make the system adaptive within its environment. The first task is to rate the usefulness of existing rules for adaptation. This is the function of the bucket brigade algorithm, to be described later. The second task is to generate plausible new rules, and this is the function of the genetic-algorithm, also to be described later.

To make this abstract description a little more meaningful, I will describe an (overly) simple classifier system, adapted from Holland (1986), that can be thought of as existing in an environment of a flat space and moving about that space to come to rest next to an object that exists somewhere in that space. The classifier system will have the abilities to turn left, turn right, move forward (each in incremental units), and stop (do nothing).

The Message Board

The message board contains all the knowledge or information of the classifier system in the form of messages. In this initial simple classifier system, each message has the form of a binary string of zeros and ones, where each binary digit or "bit" represents the presence or absence of some attribute or characteristic; it is a piece of information about something. In more elaborate and complex classifier systems, each message could be in the form of a vector of real numbers rather than binary digits and thus could be able to represent more than the presence or absence of some attribute. The meaning of the "bits" that make up the messages or strings depends upon what type of message with which we are dealing. In this simple classifier system there are two types of messages: those that are placed on the board by the detector representing conditions in the environment and those that are placed on the board by individual classifiers.

In the example problem I am using, the *detector* puts a message on the board representing an object in the environment of the classifier system, perhaps some resource. The string of zeros and ones of the message then can be used to represent whether, for example, the object is large (1) or small (0), whether it is in the left field of vision (1) or not (0), whether it is in the right field of vision (1) or not (0), and whether the object is close (1) or distant (0).[2] The position of the 1 or 0 tells what attribute is present or absent (in this simple example). Thus, a 11011 as a message placed by the detector (initial "1") might represent a *large* object (second "1") in the *right* field of vision (the "1" in position four) and *close by* (the final "1"). Messages from the detector thus represent all that the classifier system "knows" about its environment. For any given application of a classifier system, what aspects of the environment the

---

[2] Note that placement of the object in the field of vision (left, center, or right) is a three-category variable represented as two binary variables where the omitted category (center) is represented by a 0 in "left" and a 0 in "right."

detector sees and thus what messages are placed on the message board need to be determined.[3]

Messages placed by *individual classifiers* are the second type of messages on the message board. The meaning of these messages depends upon whether they are read by the effector to take some action or by another classifier to post another message. The meanings for the effector lie in the actions initiated by the effector, which in turn depend upon the actions that the classifier system as a whole may take. Again, in the simple initial example, the classifier system can take three actions: turn right 20 degrees, turn left 20 degrees, and move one step forward, i.e., in the direction it currently is facing (or do nothing). Thus, a 100 might result in a turn to the right, a 010 in a turn to the left, a 001 in a step forward, and a 000 in no action. Again, for any given application of a classifier system, we must know what actions the classifier system (effector) can take and must assign these to the message strings such that there is a unique message for each type of action possibility.[4] However, what message string is assigned to what action does not matter, as the system learns the associations. This is much like any language where the actual words assigned to entities in the environment do not matter as long as we all agree on them. The meanings of messages responded to (as conditions) by other classifiers are also arbitrary and develop over time depending upon the classifiers that respond to each, the ultimate actions taken by the effector, and the consequences of those actions.

Classifiers

I turn now to consider the individual classifiers or action-rules within the system. In a simple initial model we might have, say, 200 classifiers (action-rules). Each classifier functions individually, in parallel with the others,[5] to read messages on the message board and to place other messages on the board conditionally in response to the messages read. Each individual classifier as an action-rule has two parts, organized as an if-then clause: the condition (if) and the action or message (then). The conditional part is a pattern the classifier will recognize when it appears as part of a message already on the message board. When the condition part of the classifier is satisfied by the existence of an appropriate message on the message board, the classifier takes an action, i.e., posts its own message on the board. The conditional pattern is made up of three symbols: ones, zeros, and "don't cares," represented by "x." Thus, the conditional pattern 1xx10 would recognize any message that has a 1 in positions one and four and a 0 in position five, irrespective of the contents of the other positions. A classifier with this as its condition would place its message on the board any time the detector placed a message on the board indicating an object exists in the right field of vision but is not adjacent. Since the message is from the detector, there is a 1 in position one; since the object is in the right field of vision, there is a 1 in position four;

---

[3] This may seem to be a severe limitation, since it would appear that the programmer has to tell the classifier system exactly what to "see" in the environment. The beauty of the classifier system, however, lies in its inductive capacity to create higher-level concepts out of lower-level components. Thus, the combination of the low-level components of "small," "flying," "fast," and "black" might denote a crow, while "small," "flying," "fast," and "red" might denote a cardinal. Thus, if the perceptions are of sufficiently low-level components, all of the needed and useful perceptions would be learned inductively and would not be programmed from the start.

[4] Again, if the actions are sufficiently small, for example, governing the tension in each of the thousands of muscles in the human body, then larger "social behaviors" could be built inductively upon these micro-level actions.

[5] Because of the parallel operation, order is not relevant.

since the object is not adjacent, there is a 1 in position five. The classifier, then, would respond to either of the messages 11010 or 10010 by posting its own message.

Each classifier is thus a conditional rule in a system of conditional rules. Each functions in the manner of "if condition x then take action (post message) y." Thus, if our classifier system is set up to find an object (either large or small) in its environment and then to move adjacent to the object, we might have the following classifiers: 1x00x/100, 1x1xx/100, 1xx1x/010, 1x000/001, 1x001/000 where the string to the left of the "/" represents the conditional part of the classifier and the string to the right of the "/" represents the action part of the classifier, or the message posted. In this example the first (left-most) position of the conditional part is a tag representing a message from the environment via the detector (1) or not (0). The remaining positions are shown in Table 1. Thus, our five classifiers act as the following rules: if there is no object in sight (1x00x), turn left (action 100); if the object is in the left field of vision (1x1xx), turn left (action 100, to center it); if the object is in the right field of vision (1xx1x), turn right (010); if the object is centered visually and not adjacent (1x000), move forward (action 001); if the object is centered visually and adjacent (1x001), take no action (action 000). Some thought will show that these five classifiers thus would solve the problem of moving the classifier system adjacent to an object in the environment wherever the object was and wherever the classifier system was.

Table 1. Example Coding of the Classifier Messages and Actions

A. Example Coding of Messages

Position	Meaning
1	1 = message placed by detector; 0 = message placed by classifier
2	1 = object is large; 0 = object is small
3	1 = object in left field of vision; 0 = object not in left field of vision
4	1 = object in right field of vision; 0 = object not in right field of vision
5	1 = object is adjacent; 0 = object is not adjacent
...	Other codes

B. Example of Coding Actions

Position	Meaning
1	1 = turn to the left; 0 = no turn to the left
2	1 = turn to the right; 0 = no turn to the right
3	1 = move forward; 0 = no move forward
...	Other codes

However, the initial rules may not be fully known (or may be totally unknown) but must be learned by the system as it adapts to the environment in which it finds itself. Indeed, in most applications, the initial rules are random. The classifier system must be able to learn and to adapt over time, even as the environment changes. And its learning must be fairly rapid and efficient. I next show how the model does this.

The Bucket Brigade

There are two tasks that need to be performed to make the system adaptive. The first task, accomplished by the bucket brigade algorithm, is to rate the usefulness of existing rules, eliminating those that are not useful. The second task is to generate plausible new rules to replace the useless rules that have been dropped, and this is the function of the genetic-algorithm described in the next section (Holland 1986).

As just mentioned, each classifier posts a message when its conditional part is met. This is not quite true. Rather, each classifier whose conditional part is met makes a *bid* to have its message posted. The size of the bid the classifier makes is a function of a number of factors, but most important is its *strength*. Strength, as we shall see, is an indicator of how well the classifier is doing—how useful its messages are for solving the problem. When a classifier makes a bid, the size of that bid is subtracted from its strength (the cost of doing business). If the bid is high enough to outbid other competing classifiers, it is accepted, and the classifier posts its message on the board. In the simple example we are considering only the top two bids are accepted. Clearly, strong classifiers (ones that have been successful in the past) are able to make larger bids and therefore are more likely to have their messages posted.[6]

When a bid is accepted, the classifier puts its message on the board, and the amount of its bid is added to the strength of the classifier that placed the message on the board to which the current classifier responded.[7] If the message placed by the current classifier is found useful—that is, if it is responded to by another classifier making a bid—then the bid of that classifier is added to the strength of the present classifier. Or if it leads to a payoff from the environment (solution to the problem), then the amount of the payoff is added to the strength of the classifier whose action led to the payoff. Bids and payoffs thus are passed back along the chain of messages and add to the strength of those classifiers that post messages that other classifiers find useful. In this way, classifiers that do not lead to a payoff themselves but that are important for "setting the stage" for other classifiers to solve the problem (i.e., get the payoff) are rewarded properly.

On the other hand, if the present classifier places a message that no classifier bids to respond to or that leads to no payoff, then the classifier simply becomes weaker, as its bids are subtracted from its strength. If that continues, its bids will become smaller until it loses out in the bidding process. In this way, successful classifiers gain strength and are more likely to have their bids accepted because their bids will be larger.

From this discussion, it is clear that the payoff condition (solution to the problem) is the only point in the process where strength points are added to the system from the outside. The bidding process and the bucket brigade pass strength points from one

---

[6] An alternative formulation would have the probability of a bid being accepted to be a function of the magnitude of the bid. In that way, even weaker classifiers have a chance, though small, to have their bids accepted.

[7] In the case where the classifier is responding to multiple messages, its bid is divided among the classifiers responsible for placing those messages.

classifier to another and represent an internal reallocation of strength. For the system as a whole to work, however, strength must come from the adaptation of the system to the environment in the sense of solving the "problems" the environment "poses." This is the place where the researcher needs to define what adapting to the environment or solving problems means for a given classifier system. For example, classifier systems that learn to play poker very successfully have been built; thus, winning at poker is adapting to the environment for this system (Smith 1980). A classifier system that regulates a natural gas pipeline distribution system also has been built, thus acting and "making decisions" in a way that previously only humans could. Maintaining a constant pressure and flow of natural gas in the event of disturbances such as leaks, changes in temperature, storage capacities, or changes in demand constitutes successful adaptation to the environment for this system (Goldberg 1983). In each of these examples, adaptation *success* is defined by the nature of the problem, and the payoff for success is provided when those success conditions are met.

The Genetic Algorithm

The other task that needs to be accomplished is that of discovering new and plausible rules, since there is no guarantee that the present rules will solve the problem or will be optimal. With the possibility of a changing environment, rules that work at one time may not be optimal or even work at another time. It is through the genetic algorithm that classifier systems learn and adapt to an environment. Individual classifiers that are successful (whose strength is built up through payoff and the bucket brigade) have an opportunity to reproduce themselves and to replace classifiers that have not been successful. Reproduction, however, is not accomplished by simply replicating successful classifiers, as this only would provide unnecessary duplication of existing rules. Rather, the conditional parts of successful classifiers are cross-linked to make new conditions.

For example, consider two successful classifiers (parents) with conditions 10011 (mom) and 11001 (pop). A cross-over point is selected randomly, for example, at position two. Then two new classifiers are created. One new classifier condition is defined consisting of the first two bits of the condition of the mom classifier and the last four from the condition of the pop classifier (yielding 10001), while a second new classifier condition is defined consisting of the first two bits from pop and the last four bits from mom (yielding 11011). The new classifiers do not replace the mom and pop but rather replace two "unsuccessful" classifiers (i.e., ones with low strength). Thus, successful classifiers are kept and continue to contribute to the success of the classifier system as a whole, while modifications of the successful classifiers are introduced that recombine, in new ways, previously successful patterns. These new patterns may or may not be successful themselves, but chances are in their favor since they are recombinations of successful patterns (Holland 1986). By allowing occasional mutation to occur in the recombination process (a 0 becomes a 1 or a 1 becomes a 0), additional rules can be generated that are close to successful rules but different in one respect or another that may prove important in exploring the rule space for useful rules.

Operation of the Classifier System

A classifier system begins by placing all the messages representing information about the environment from the detector into a list on the message board. The message list then is processed by all of the classifiers. Since the operation is essentially parallel, the order of processing does not matter. Each classifier scans the messages on the message list. If any message satisfies the conditional part of the classifier's rule, the classifier makes a bid to place its message (action part of the rule) on the list. As the bids are made, the amount of the bid is subtracted from the current strength of the classifier and is distributed among the classifiers whose messages to which the classifier in question is responding, by adding to their strength. The bids of all the classifiers are collected; some set of them (usually the highest bids, but possibly all bids or a weighted random selection) is accepted, and the action parts of the classifiers are posted as messages onto a new message list. Finally, the effector scans the new message list and initiates action for each classifier message, fulfilling the conditions for that action. At this point, the cycle starts again with the posting of new detector messages to the message list. After the first round, classifiers are responding to messages of classifiers from the prior round as well as detector messages from the current round.

Notice that the bucket brigade is built into the bidding system. The genetic algorithm, however, is not built into the previously mentioned execution cycle and must be added. The usual arrangement is to invoke the genetic algorithm on every cycle in which there is no action (since the system currently does not have an effective rule for the current situation) and periodically even when the system is taking action. In this way, the classifier system continues to explore potentially better rules even when it has some that work.

One aspect of classifier systems that may not be apparent from the description given thus far is the ability of classifiers as the intermediate processors of messages between the detector and effector to be coupled into clusters and hierarchies through the use of tags and conditions (Holland 1986). This coupling is not a formal part of the classifier system but is rather an emergent property of the system as defined. For example, one can define a classifier that attends to a specific set of classifiers. Thus, a pair of classifiers K1 and K2 that sends messages 1101 . . . and 1001 . . . , respectively, both will be responded to by a classifier whose conditional part is 1x0x. . . . Also, one classifier can initiate a complex sequence with a single message that is responded to by several classifiers, each of whose messages are responded to by other classifiers, and so forth. The clusters and hierarchies of related classifiers (rules) become "chunks" in larger, emergent organizations of classifiers.

And because of the parallelism of classifier systems, clusters of co-active rules can be used to define categories, objects, and processes inductively. Indeed, induction is one of the prime capabilities of the classifier system and figures heavily into its capacity to adapt to new and changing environments. The relations between these inductively "discovered" categories are carried by classifiers that couple the larger chunks together (Holland 1986). Such coupling may occur through a classifier that responds to the presence of two or more "chunks" or a classifier that initiates two or more "chunks." In this way, very complex adaptive responses to heterogeneous and changing environments are made possible. Indeed, as (Holland 1986) has pointed out, such a classifier system is computationally complete in the sense that, in theory, it can solve any problem that can be solved by an algorithm. Nothing is said, however, about how quickly such a solution might be found.

COMBINING ICT AND CLASSIFIER SYSTEMS

To begin to incorporate ICT into a classifier system, we need to make some correspondences between the parts and functions of each. Some are very simple. The *input* of the identity system corresponds to the *detector* of a classifier system: both bring information to the system from the environment. Similarly, the *output* of the identity system corresponds to the action of the *effector* of the classifier system, which may affect the environment. The *comparator* and *identity standard* of the identity model, however, have no counterparts in the classifier system and must be added.

Action (output) in the identity system is a function of the difference between the input and the standard and continues until there is no difference. In the identity system, there is a goal, which is to match the input to the standard. Action in the classifier system is a function of a match between a condition (input from the environment by the detector or generated by a classifier) and the condition part of a classifier whose action part will activate the effector. The action takes place as the result of a similarity, not a difference, and there is no goal state—it is more like a stimulus-response model: when a condition is satisfied (stimulus), a certain action is taken (response). We need to modify the basic classifier system to make it more like the identity model with its own standard (goal) and thus able to take the "purposive" action of making the perceptions match the standard.

Hence, the first modification is to allow the classifier system to act as a control system. This means that there must be some internal standard or goal condition and a comparator, so the system as a whole can be made to act to match that standard with inputs from the environment. To implement this aspect, I add a *standard* and a *comparator* to the classifier system. This standard is a message string that represents some state of the environment that might be "observed" by the detector, and, indeed, it normally would be taken from the messages placed on the board by the detector and would be stored in a separate place. In this way, a direct comparison between the goal condition (standard) and the other messages placed on the board by the detector (perceptions) is possible for the reason that the goal condition is a message just like the others. Because everything the system knows is either wired into the basic programming or is input from the environment, it thus makes sense that the standard is also an input from the environment in this manner and that it represents some state of the environment that is or might be realized through action of the system as a whole.[8] In this model, then, *socialization* is the acquisition of internal standards that, along with perceptions, will guide behavior. This and other modifications to the classifier system to create an identity model are illustrated in Figure 3.

The question still remains as to which of such messages might qualify to be a goal state. A minimal condition is what the detector perceives at the point in time at which a payoff occurs. The classifier system thus would "remember" that state of the environment when the payoff occurs, and that state (as the goal) represents those conditions under which a payoff occurs.

To complete the transformation to a control system, I also add a *comparator* that functions to compare the goal message with other messages placed on the board by the detector. The comparator operates to scan the message board and to compare each *detector message* (input) with the *goal message* (standard). The result of each comparison, however, is not just to place another message on the board representing the difference between the goal condition and the current perception (though it does this), and the difference pattern (message) represents a new condition to which the

---

[8] Actually, there is a third category of knowledge that is neither wired in nor gotten directly from the environment. This knowledge is the set of successful classifier rules, which consist of emergent patterns and relationships among patterns that have come to represent successful actions, perceptions, and the relationships between them. A standard taken from this category could represent a more advanced and abstract goal condition.

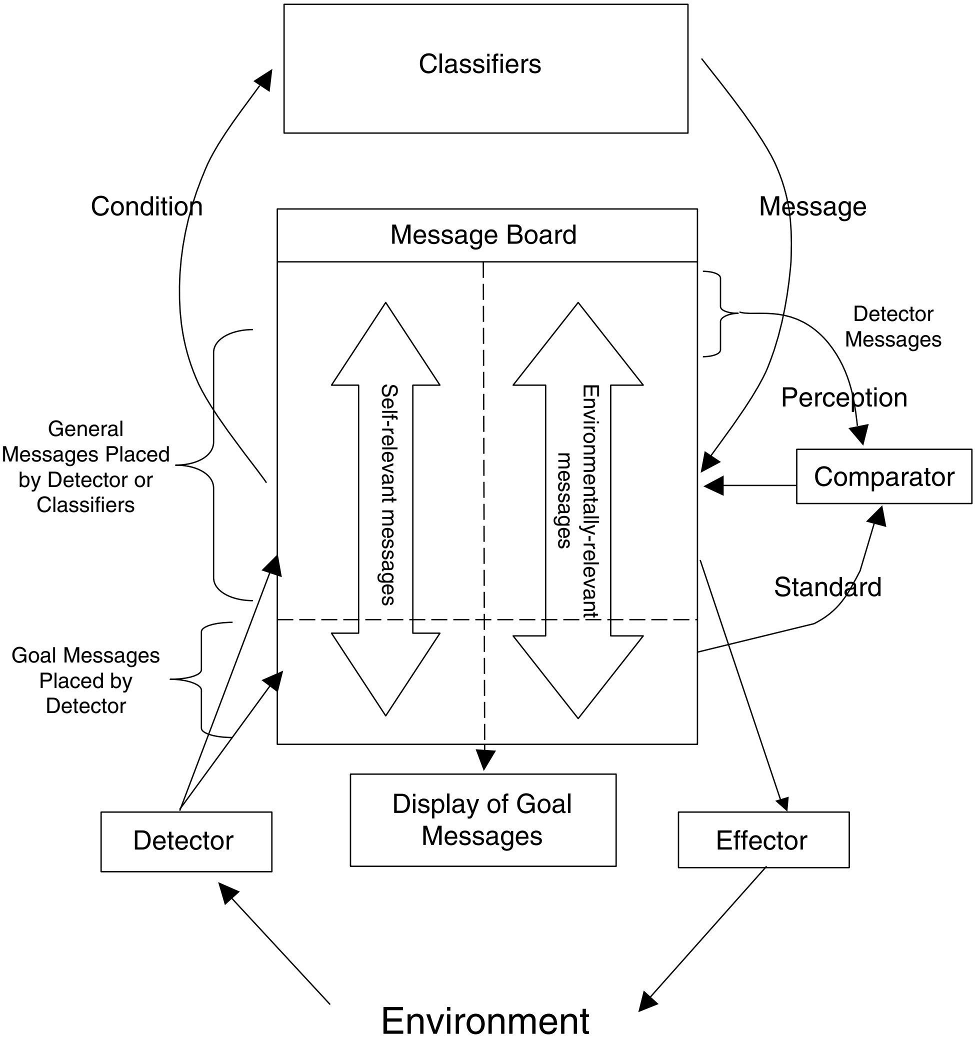

Figure 3. Conceptual model of identity as a classifier system.

classifiers may respond and may take action.[9] In addition, the comparator also contributes to the bucket brigade by providing a payoff. This payoff is given to those classifiers whose action brings the result of a comparison between the current input and the standard closer to a match. This acts to strengthen classifiers whose actions bring the system closer to the goal without waiting until the goal has been achieved fully for an external payoff to occur.[10]

Such closeness could be measured by counting the number of attributes that are the same between the input and the standard, by ignoring "don't care" attributes, and

---

[9] In this way, actions can be the result of differences between the perception and the standard as they are in the control system.

[10] The external payoff is paid into a resource pool from which the comparator draws its payoffs for achieving perceptions that are closer to the standard. The final action, which makes the perception equal to the standard, then yields a payoff from this internal pool, just as other adaptive but nonfinal actions. We might call such a resource pool *self-esteem* (cf. Cast and Burke 2002).

perhaps by assigning weights to particular attributes to indicate their importance. With the addition of the goal condition (standard), the comparator, and the intermediate reward mechanism to provide a "payoff" for steps taken toward the goal, we have transformed the classifier system from a reinforcement-based model into a control system by equating reinforcement with a match between the perception and the standard.

Thus far, I have assumed that there is only one payoff condition to become a standard or goal the system learns and then uses. Most systems, however, have a number of conditions in which a payoff will occur, and the goal system must be extended to accommodate these. Thus, each condition in which a payoff occurs is placed on a *goal board* as that condition is encountered. In addition, each *goal message* is evaluated for its closeness to current perceptions of the environment (ease of accomplishment) and the amount of payoff obtained when that condition is achieved (which must be remembered). Which of the goal states then becomes the current standard is then a function of its importance (amount of payoff) and closeness to current state (ease).[11] Other operators can be created to modify the list of goals on the goal board. Examples include (1) an operator that generalizes (if a large (1) or small (0) object provides a reward, then size is not important and can be ignored (x)); (2) a genetic algorithm that explores new goals through a mechanism of recombining existing goals; and (3) an operator that removes goals that no longer have a payoff (in a changing environment). In this way, identity standards (goals) become organized and ordered and are activated according to their importance and the closeness of the current situation to their achievement. In this sense, the situation does activate identity standards, though all are continually in a state of readiness for activation.

The second addition to the classifier system is to make it *reflexive* or able to recognize itself. Reflexivity is certainly an explicit part of the self as viewed in social psychology, though its incorporation into the identity model has remained somewhat implicit (Burke 1980). ICT does point out that identities monitor *self-relevant* meanings (as opposed to all meanings), and self-relevance cannot be understood without some degree of reflexivity (e.g., knowing what one's own goals are). Thus, ICT does incorporate reflexivity, though the mechanisms underlying its operation have not as yet been spelled out. In the classifier system, reflexivity is accomplished in part by having the detector place on the board messages that report not only the condition of the external environment but also messages that report on the state of the classifier system itself in that environment and any actions it has taken.

In a classifier system, there are tags on messages that distinguish between messages placed in the message list by the detector and messages placed by other classifiers (for example, see Table 1). It becomes straightforward by the use of appropriate tags, therefore, to distinguish messages placed by the detector concerning current states of the self from messages concerning current states of the nonself-environment.[12] By having messages in the message list concerning the current state of the self, i.e., the identity/classifier system, it is possible to evolve system rules that take into account not only the current environment but also the current state of the classifier system and the relationship between the two.

---

[11] Other conditions also might influence the importance of a particular goal. For example, a goal that needs to be achieved periodically (like eating) might become more important as a function of elapsed time.

[12] An interesting and important issue thus becomes what, perceptually, constitutes the self and distinguishes it from the nonself. Classifications of self and nonself cannot be absolute but must be conditional, as when I reach for a stick, the stick is a nonself object, but when I use the stick to point, the stick is then part of the self, and the thing to which I point is a nonself object.

These two modifications (incorporating reflexivity and a standard or goal state) can be combined to provide the important ability that a particular self-state can become a goal of the system. In this way, not only can the system act on the current environment in response to the current environment; it also can act to create environments it has seen before or has been told about, and it can act to achieve goal states of the self by creating modifications to the environment that may not have been seen before. This greatly increases the inductive capacity of the system.

For example, start with a system that accidentally approaches a "food" object and thus obtains a payoff. If this happens often enough, the payoff and the bucket brigade algorithm will strengthen those classifiers (rules) that achieve the payoff. This would be accomplished faster if, the first time the system achieved the payoff, the environmental state at the time of the payoff (object visually centered and adjacent, for example) were used as a goal state that could provide intermediate rewards as it was approximated. This would more quickly strengthen classifiers that laid the groundwork for the ultimate payoff. Now, if the goal state were some condition of the self, for example, "possessing information about the location of food," then alternative behaviors might become available to attain this goal (such as asking someone, in addition to searching) that may make it easier to obtain food. This is in accord with Mead's (1934) suggestion that having a goal state of the self makes possible more purposeful control and organization by an individual of its own conduct.

The last set of modifications of the classifier system I wish to mention here allows the identity/classifier model to be more social. This can be accomplished through the use of a system of symbolic communication in which one classifier system can communicate with one or more other classifier systems. This is, perhaps, the most difficult and speculative of the modifications because it is not entirely clear exactly what needs to be incorporated into the system to accomplish this. What I have to say on this thus is limited and is only suggestive at this point.

Since a classifier system is already a message-based system, it makes sense to start with those messages as the building blocks. Like human symbolic communication, the coding of meaning into messages in a classifier system, as discussed already, is arbitrary and evolves over time only through the utility and functioning of those messages. The symbols (messages), and what they represent, are not fully predetermined; the genetic algorithm is always creating new messages. What counts is what the messages accomplish. The first modification to allow communication between separate classifier systems, then, is to have some means of passing messages from one classifier system to another.

One way to accomplish this is to have some sort of message display on each classifier system that would allow one or more messages to be "seen" (detected) by other classifier systems (and because of the reflexivity already built in through the first modification, they also would be detected by the classifier system itself in question). Thus, if classifier system A displays a message, the detectors of both classifier system A and classifier system B each can perceive the message and can place it or another message on its own respective message list.

What distinguishes symbols from signs (in addition to the arbitrariness in assigning meaning to symbols) is the sharing of meaning, that is, the response of each party to the display of a symbol is the "same" (Osgood, Suci, and Tannenbaum 1957). If we are to have symbolic communication between classifier systems, the responses of each of the classifier systems to the display of messages in some sense should be the same. By adding message displaying and detecting capabilities to the classifier systems, we have begun the process, but the question now becomes what else is needed so that responses will move toward “sameness”? By making the displayed message be a goal state (standard), then each classifier system that saw the displayed message and incorporated it as a standard would act to achieve that goal. In this way, the responses would be the same from the perspective of outcomes. However, the way in which each classifier system achieved the goal state may differ, but perhaps this is a difference that does not make a difference. Different persons given the instruction to “sit down” each will do so, yet the particular body movements to accomplish this act are not likely to be identical from one person to another. Similarly, if the displayed message, though not a goal state itself, functioned to activate a common goal state in each of the classifiers, we again would have the same responses.

The key, thus, is to have the communication aspect of the classifier concern goal states (at least initially) in order to function symbolically. This, of course, is not a new idea. Ferguson (1792:120–122), one of the Scottish moral philosophers noted that “the more general character of man’s inclinations . . . is not that of a blind propensity to the use of means, but instinctive imitation of an end, for the attainment of which he is left to discover and to choose, by his own observation and experience, the means that may prove effectual.”

I recognize that this is still not sufficient, however. Some people told to “sit down” will respond that they prefer to stand. It is clear that they understand the symbolic communication that is given to them, but they choose not to employ the goal state for themselves. Clearly, a goal state, when communicated, becomes only one of many possible goal states the person holds, and mechanisms for choosing among these must exist. The implementation of such in a classifier system needs to be worked out.

DISCUSSION

There is a dual purpose in modeling ICT as a classifier system. On the one hand, by formalizing ICT in a model we are forced to come to grips with all of the aspects of the theory that have not been well formulated but about which we need to be explicit in order to build the model. It is a very instructive way of proceeding. On the other hand, by making explicit the ideas and assumptions that are needed to create the model, issues that have not been resolved in ICT can be moved toward resolution. Consider again the three issues in ICT I raised at the beginning of the article.

The first issue raised with respect to the identity model concerned the source of the identity standard—the sets of meanings to be maintained by acting on the environment until the feedback meanings reflect (match) the meanings held in the standard. In the classifier model of an identity, one or more goal states have been added. The goal states are in the form of messages representing states of the environment that provide resources needed to sustain the system. These resourceful states are ones the detector has previously placed on the message list at the time of a “payoff.” The addition of mechanisms for specification, generalization, elimination, and creation (using a genetic algorithm) of goal messages then provided the means for the full exploration and organization of goal states or identity standards. In this way they could come to represent states of the environment or the self that have never been realized in the past, for example, a “possible self” (Markus and Nurius 1986) or an “envisioned” environment.

The second issue concerned the problematic character of the correspondence between the perceptual inputs of identity relevant meanings in the situation and the identity standard. This question can be framed in either of two ways: How is it that we are able to see "potency" (as a dimension of meaning) in a situation when "potency" is part of the identity standard? Alternatively, if we begin with the fact that we perceive "potency," how does that get to be part of the identity standard? By selecting goals from "successful" detector messages we have dealt with this second issue. The meanings encoded in the identity standard are the same meanings as encoded into the input from the environment because the standard is a perception, either actual or potential.

The third issue concerned the activation of identities and understanding how that might occur. In ICT, we usually characterize identities in terms of their standards. Being a shopkeeper, for example, implies that one controls the meanings and resources in the situation to bring them into line with the "expectations" or meanings held in the shopkeeper identity standard. In the combined classifier/identity model, the standard or goal is a message, though of a particular status, indicated, perhaps, by a special tag like the ones used to distinguish self-messages from environmental messages or detector messages from classifier messages.

Activation of an identity then amounts to *activation of a goal message*. When there is only one goal message, it is always active. With more than one goal message available, selection of the message to be made active would occur by a weighted combination of two factors. The first is which goal is closest to the current state of the environment as contained in the most recent detector message. The second is the amount of the external payoff that occurs when the goal is achieved.[13] These two factors are similar to the two factors of attitude accessibility and object evaluation that link attitudes and behaviors in attitude research (Fazio 1986). Thus, in modeling the classifier/identity system as I have, the three issues in ICT I raised at the beginning are answered. And while these answers may not be fully correct in their initial formulation, they at least become hypotheses that are subject to test and modification.

In modeling the classifier/identity system to deal with these issues, a potential answer has been provided to another issue concerning the source of consensus on meanings that makes symbols and symbolic interaction among persons possible. Resolving this issue would make it possible to begin to model interaction processes among several identities in order to study the nascent emergence of structure in social interaction. Recalling from the discussion on identities and meaning aforementioned, meaning is defined as a (mediational) response to a stimulus. Convergence in meanings thus would be manifest in a convergence of responses to a common stimulus. By modeling the communication of standards or goal messages as suggested already, a common response (achieving the goal state by matching perceptions with the standard brings each classifier/identity) is assured. Though how that response is achieved is not given, but this has never been an issue. For example, in training a rat to press a bar when the green light comes on, the experimenter does not concern himself or herself with the detail of whether the rat uses left or right paw, or nose or any other means to accomplish the bar press (which throws into doubt exactly what behavior is being reinforced!). Thus, while many of the details need to be worked out yet, in this way, with common goal states that arise under similar conditions, each of the classifier/identities will develop a "language" that can function as a language in the sense of having common meanings that are shared with other classifier/identities.

---

[13] Although costs traditionally have not been used in classifier systems, it may be useful to include costs as well as rewards (payoffs) in future models.

CONCLUSIONS

I began by raising three issues in ICT that, because the theory as it currently exists does not have the ability to address, have remained unresolved and somewhat unexplored. These issues concerned answers to the questions of where identity standards come from, how identities are activated in situations, and how the meanings and resources perceived in a situation by an activated identity are the very resources and meanings that are coded into the identity standard. ICT remains incomplete without being able to answer these questions. In addition, without an answer to these questions, the real nature of the link between identity and social structure remains obscure, and the real impact of one on the other cannot yet be understood or investigated.

I introduced the concept of a classifier system, a model initially formulated by Holland (1975), to create a system that is able to *learn inductively* to function within any environment and to *adapt* to any changes that might occur to that environment. By modifying the classifier system to include the basic elements of the identity model, a new model of an identity as an adaptive system with goals was created, which has the adaptive capacity of the classifier system but also has the goal-seeking, control-system attributes of the identity model. This new theoretical model not only serves to resolve the three issues that were raised but also takes steps toward the resolution of a longer-standing issue concerning the origins of shared meanings in the framework of symbolic interaction. By postulating that identity standards are set or are created in situations in which resources are encountered and are used by individuals, this new model suggests that identity standards are "goal states" that, when achieved, result in the reproduction of those resourceful situations. Through the manipulation of signs and symbols by identities, resources that sustain individuals and interaction are manipulated, transformed, and transferred. In this way, identities serve both an adaptive and control purpose for individuals and are intimately tied to the resources that sustain them.

REFERENCES

Blumer, H. 1962. "Society as Symbolic Interaction." Pp. 179–192 in *Human Behavior and Social Processes*, edited by A. M. Rose. Boston, MA: Houghton Mifflin Co.

Burke, P. J. 1980. "The Self: Measurement Implications from a Symbolic Interactionist Perspective." *Social Psychology Quarterly* 43:18–29.

———. 1996. "Social Identities and Psychosocial Stress." Pp. 141–174 in *Psychosocial Stress: Perspectives on Structure, Theory, Life Course, and Methods*, edited by H. B. Kaplan. Orlando, FL: Academic Press.

———. 1997. "An Identity Model for Network Exchange." *American Sociological Review* 62:134–150.

———. 2004. "Identities and Social Structure: The 2003 Cooley-Mead Award Address." *Social Psychology Quarterly* 67.

Burke, P. J. and A. D. Cast. 1997. "Stability and Change in the Gender Identities of Newly Married Couples." *Social Psychology Quarterly* 60:277–290.

Burke, P. J. and L. Freese. 1989. "Identity and Social Structure." Presented at the annual meeting of the American Sociological Association, August 18, San Francisco, CA.

Burke, P. J. and D. C. Reitzes. 1981. "The Link Between Identity and Role Performance." *Social Psychology Quarterly* 44:83–92.

Burke, P. J. and J. Tully. 1977. "The Measurement of Role/Identity." *Social Forces* 55:881–897.

Cast, A. D. and P. J. Burke. 2002. "A Theory of Self-Esteem." *Social Forces* 80:1041–1068.

Cast, A. D., J. E. Stets, and P. J. Burke. 1999. "Does the Self Conform to the Views of Others?" *Social Psychology Quarterly* 62:68–82.

Fazio, R. H. 1986. "How Do Attitudes Guide Behavior?" Pp. 204–243 in *Handbook of Motivation and Cognition: Foundations of Social Behavior*, edited by R. M. Sorrentino et al. New York: Guilford Press.

Ferguson, A. 1792. *Principles of Moral and Political Science: Being Chiefly a Retrospect of Lectures Delivered in the College of Edinburgh*, vol. 1. Printed for A. Strahan and T. Codell, London; and W. Creech, Edinburgh.

Freese, L. 1988. "The Very Thought of Resources." Presented at the annual meetings of the American Sociological Association, August 19, Atlanta, GA.

Freese, L. and P. J. Burke. 1994. "Persons, Identities, and Social Interaction." Pp. 1–24 in *Advances in Group Processes*, vol. 11, edited by B. Markovsky, K. Heimer, and J. O'Brien. Greenwich, CT: JAI Press.

Goldberg, D. E. 1983. "Computer-Aided Gas Pipeline Operation Using Genetic Algorithms and Rule Learning." Ph.D. Dissertation, University of Michigan, Petroleum Engineering, Ann Arbor.

Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.

———. 1986. "Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems." Pp. 593–623 in *Machine Learning: An Artificial Intelligence Approach*, edited by R. S. Michalski, J. G. Carbonell, and T. M. Mitchell. Los Altos, CA: Morgan Kaufmann Publishers Inc.

Holland, J. H., K. J. Holyoak, R. E. Nisbett, and P. R. Thagard. 1986. *Induction: Processes of Inference, Learning and Discovery*. Cambridge, MA: MIT Press.

James, W. 1890. *Principles of Psychology*. New York: Holt Rinehart and Winston.

Lindesmith, A. R. and A. L. Strauss. 1956. *Social Psychology*. New York: Holt Rinehart and Winston.

Markus, H. and P. Nurius. 1986. "Possible Selves: The Interface Between Motivation and the Self-Concept." Pp. 157–172 in *Self and Identity: Psychosocial Perspectives*, edited by K. Yardley and T. Honess. New York: Wiley.

Mead, G. H. 1934. *Mind, Self, and Society*. Chicago, IL: University of Chicago Press.

Osgood, C. E., G. J. Suci, and P. H. Tannenbaum. 1957. *The Measurement of Meaning*. Urbana: University of Illinois Press.

Powers, W. T. 1973. *Behavior: The Control of Perception*. Chicago, IL: Aldine Publishing.

Smith, S. F. 1980. "A Learning System Based on Genetic Algorithms." Ph.D. Dissertation, University of Pittsburgh, Computer Science, Pittsburgh, PA.

Smith-Lovin, L. and D. R. Heise. 1988. *Analyzing Social Interaction: Advances in Affect Control Theory*. New York: Gordon and Breach Science Publishers.

Stets, J. E. and P. J. Burke. 1996. "Gender, Control, and Interaction." *Social Psychology Quarterly* 59:193–220.

Stryker, S. 1980. *Symbolic Interactionism: A Social Structural Version*. Menlo Park, CA: Benjamin Cummings.

Stryker, S. and P. J. Burke. 2000. "The Past, Present, and Future of an Identity Theory." *Social Psychology Quarterly* 63:284–297.

Stryker, S. and R. T. Serpe. 1982. "Commitment, Identity Salience, and Role Behavior: A Theory and Research Example." Pp. 199–218 in *Personality, Roles, and Social Behavior*, edited by W. Ickes and E. S. Knowles. New York: Springer-Verlag.

Swann, W. B., Jr. 1983. "Self-Verification: Bringing Social Reality into Harmony with the Self." Pp. 33–66 in *Psychological Perspectives on the Self*, edited by J. Suls and A. Greenwald. Hillsdale, NJ: Erlbaum.