

## CONSTRAINT, CONSTRUAL, AND COGNITIVE SCIENCE

Robert P. Abelson  
Yale University

Cognitive science has barely emerged as a discipline -- or an interdiscipline, or whatever it is -- and already it is having an identity crisis.

Within us and among us we have many competing identities. Two particular prototypic identities cause a very serious clash, and I would like to explicate this conflict and then explore some areas in which a fusion of identities seems possible. Consider the two-word name "cognitive science". It represents a hybridization of two different impulses. On the one hand, we want to study human and artificial cognition, the structure of mental representatives, the nature of mind. On the other hand, we want to be scientific, be principled, be exact. These two impulses are not necessarily incompatible, but given free rein they can develop what seems to be a diametric opposition.

The study of the knowledge in a mental system tends toward both naturalism and phenomenology. The mind needs to represent what is out there in the real world, and it needs to manipulate it for particular purposes. But the world is messy, and purposes are manifold. Models of mind, therefore, can become garrulous and intractable as they become more and more realistic. If one's emphasis is on science more than on cognition, however, the canons of hard science dictate a strategy of the isolation of idealized subsystems which can be modeled with elegant productive formalisms. Clarity and precision are highly prized, even at the expense of common sense realism. To caricature this tendency with a phrase from John Tukey (1969), the motto of the narrow hard scientist is, "Be exactly wrong, rather than approximately right".

The one tendency points inside the mind, to see what might be there. The other points outside the mind, to some formal system which can be logically manipulated (Kintsch et al., 1981). Neither camp grants the other a legitimate claim on cognitive science. One side says, "What you're doing may seem to be science, but it's got nothing to do with cognition." The other side says, "What you're doing may seem to be about cognition, but it's got nothing to do with science."

Superficially, it may seem that the trouble arises primarily because of the two-headed name cognitive science. I well remember the discussions of possible names, even though I never liked "cognitive science", the alternatives were worse; abominations like "epistology" or "representonomy".

But in any case, the conflict goes far deeper than the name itself. Indeed, the stylistic division is the same polarization than arises in all fields of science, as well as in art, in politics, in religion, in child rearing -- and in all spheres of human endeavor. Psychologist Silvan Tomkins (1965) characterizes this overriding conflict as that between characterologically left-wing and right-wing world views. The left-wing personality finds the sources of value and truth to lie within individuals, whose reactions to the world define what is important. The right-wing personality asserts that all human behavior is to be understood and judged according to rules or norms which exist independent of human reaction. A similar

distinction has been made by an unnamed but easily guessable colleague of mine, who claims that the major clashes in human affairs are between the "neats" and the "scruffies". The primary concern of the neat is that things should be orderly and predictable while the scruffy seeks the rough-and-tumble of life as it comes.

I am exaggerating slightly, but only slightly, in saying that the major disagreements within cognitive science are instantiations of a ubiquitous division between neat right-wing analysis and scruffy left-wing ideation. In truth there are some signs of an attempt to fuse or to compromise these two tendencies. Indeed, one could view the success of cognitive science as primarily dependent not upon the cooperation of linguistics, AI, psychology, etc., but rather, upon the union of clashing world views about the fundamental nature of mentation. Hopefully, we can be open minded and realistic about the important contents of thought at the same time we are principled, even elegant, in our characterizations of the forms of thought.

The fusion task is not easy. It is hard to neaten up a scruffy or scruffy up a neat. It is difficult to formalize aspects of human thought which are variable, disorderly, and seemingly irrational, or to build tightly principled models of realistic language processing in messy natural domains. Writings about cognitive science are beginning to show a recognition of the need for world-view unification, but the signs of strain are clear. Consider the following passage from a recent article by Frank Keil (1981) in Psychological Review, giving background for a discussion of his formalistic analysis of the concept of constraint:

"Constraints will be defined...as formal restrictions that limit the class of logically possible knowledge structures that can normally be used in a given cognitive domain." (p. 198).

Now, what is the word "normally" doing in a statement about logical possibility? Does it mean that something which is logically impossible can be used if conditions are not normal? This seems to require a cognitive hyperspace where the impossible is possible.

It is not my intention to disparage an author on the basis of a single statement infelicitously put. I think he was genuinely trying to come to grips with the reality that there is some boundary somewhere to the penetration of his formal constraint analysis into the viscissitudes of human affairs. But I use the example as symptomatic of one kind of approach to the cognitive science fusion problem: you start from a neat, right-wing point of view, but acknowledge some limited role for scruffy, left-wing orientations. The other type of approach is the obvious mirror: you start from the disorderly left-wing side and struggle to be neater about what you are doing. I prefer the latter approach to the former. I will tell you why, and then lay out the beginnings of such an approach.

The strategy of trying to move leftward from the right suffers from a seemingly permanent limitation on the kinds of content and process you are

willing to consider. If what really matters to you is the formal tractability of the domain of investigation, then your steps are likely to be small and timid. Recent history in several social and behavioral science areas makes this quite clear.

In cognitive anthropology, there was a great deal of fascination 25 years ago with the orderliness of systems for kinship terminology. Kin terms in different societies were found to be precisely describable by concatenations of the values on a handful of well-specified components such as sex and generation. Formal mini-models captured these regularities elegantly. Originally it was thought that this kind of componential semantics held great promise for the analysis of culture and language in general, but gradually it was realized that outside of kinship terms and pronoun systems, precious little else in the language of any society was ordered so neatly. Faith in tight componential analysis has largely been abandoned.

Rational decision theory has until recently had a tight hold on the views of economists and some psychologists of the way people made decisions. The typical rational decision model specifies a set of uncertain outcomes, with each of which is associated a probability and a utility. Choices among ensembles of outcomes are then said to be predictable on the basis of a strict composition rule on the probabilities and utilities. The only trouble is, the behavior of human subjects overwhelmingly disobeys the predictions of the models, no matter how hard the axioms try. There have been numerous attempts to rescue the general framework, including the clever strategy of training subjects to obey the rational model following initial demonstrations of deviation from it. I would recommend this device also to people promoting competence models of syntax in the face of incompetent performers, except that I cannot as a psychologist bring myself to believe that it tells us anything about human psychology. In any case, there are some many violations of rational decision theory that it is a clear failure as a descriptive or explanatory psychological model. Only an approach that deals directly with observed decision phenomena (for example, the work of Tversky and Kahneman (1980)) has a chance of success. (For fuller reviews of this field, see Einhorn and Hogarth (1981), March (1978), and Abelson & Levi (Note 1)).

Other examples of excessive faith in the unaided power of formalisms to subdue the beast of psychological explanation could be adduced from within experimental psychology itself. A good case from some years back is provided by stochastic learning models (Bush & Mosteller, 1955), which were extremely rich as mathematical objects, but turned out to have applicability to a very small range of problems, indeed. Models like this were part of the "bottom up" tradition of doing science within experimental psychology, the belief that by starting with very tightly controlled, limited, and isolated laboratory phenomena, one could gradually explicate the operation of the whole organism. This tradition is of course still strongly honored by many experimental psychologists, but I think that those psychologists interested in cognitive science have largely departed from that tradition, at least in its most extreme form. In the service of studying more important and more general phenomena than those falling within the formal

boundaries of mini-models, cognitive science psychologists have been willing to use messier stimulus materials and at least contemplate non-laboratory methodologies. The way is not easy, and there is much anguishing. That, I claim, is the price of trying to move leftward from a right-wing starting point.

Linguists, by and large, are farther away from a cognitive science fusion than are cognitive psychologists. The belief that formal semantic analysis will prove central to the study of human cognition suffers from the touching self-delusion that what is elegant must perforce be true and general. Intense study of quantification and truth conditions because they provide a convenient intersection of logic and language will not prove any more generally informative about the huge range of potential uses of language than the anthropological analysis of kinship terms told us about culture and language. On top of that, there is the highly restrictive tradition of defining the user of language as a redundant if not defective transducer of the information to be found in the linguistic corpus itself. There is no room in this tradition for the human as inventor and changer and social transmitter of linguistic forms, and of contents to which those forms refer. To try to understand cognition by a formal analysis of language seems to me like trying to understand baseball by an analysis of the physics of what happens when an idealized bat strikes an idealized baseball. One might learn a lot about possible trajectories of the ball, but there is no way in the world one could ever understand what is meant by a double play or a run or an inning, much less the concept of winning the World Series. These are human rule systems invented on top of the physical possibilities of the batted ball, just as there are human rule systems invented on top of the structural possibilities of linguistic forms. One can never infer the rule systems from a study of the forms alone.

Well, not I have stated a strong preference against trying to move leftward from the right. What about the other? What are the difficulties in starting from the scruffy side and moving toward the neat? The obvious advantage is that one has the option of letting the problem area itself, rather than the available methodology, guide us about what is important. The obstacle, of course, is that we may not know how to attack the important problems. More likely, we may think we know how to proceed, but other people may find our methods sloppy. We may have to face accusations of being ad hoc, and scientifically unprincipled, and other awful things.

Sometimes we worry about such matters ourselves. There is a neat person struggling to get into every scruffy person (just as there is a scruffy person struggling to get out of every neat person). What is required is that we act on our worries, that we try to take the criticisms seriously and see what can be done about them. The messy intuitions and theories, albeit they concern very general and important problems (God bless 'em), need to be articulated and developed in a more orderly way.

I will take the work of the Yale Artificial Intelligence Project, and in particular, the programmatic statements in the Schank and Abelson (1977) book as point of departure. The Yale point

of view is quintessentially scruffy, and has been criticized accordingly. No matter that scripts and plans and goals and themes are psychologically reasonable, and that computer programs using such concepts are operational at the frontier of realistic processing of natural language, nevertheless, it is said that the system of concepts is not formal enough.

The make a system more formal is to define its concepts more precisely, and to have them enter into general predictions and explanations according to a set of principles, preferable a small elegant set. Let me first address the question of the definition of concepts. The Yale group deals with knowledge structures such as scripts and plans, it may seem at first that these are pretty amorphous entities. What counts, say, as a script or as clearly not a script? How can you tell? When we gave examples such as the restaurant script, and cognitive psychological experiments (e.g., Bower, Black & Turner, 1979; Galambos & Rips, 1979; Graesser, 1981) forthwith used verbal stimuli from the restaurant situation, along with doctor visits and laundromat activities and so on, it may have produced misleading impressions that the intention was to define scripts solely by waving at passing examples, or perhaps by writing down a definitive list of 111 of them, or worse, by allowing any damn thing to be a script just by calling it that.

I say that these impressions are misleading because in fact we have become acutely aware (Schank, Note 2) that scripts have been loosely used. The original intention was not at all to create a haven for loose concepts; in fact, scripts (among other knowledge structures) were very tightly defined, by a set of interdependent constraints. Indeed, if a knowledge structure is proposed as crucial in the top-down processing of certain inputs, then clearly it must embody of us to leave these constraints largely implicit, rather than spelling them out systematically.

It is not my main intention today to remedy this neglect for scripts or any other specific type of knowledge structure, but rather to make clear my general view of the role that constraints play in the process of understanding text. However, having raised the issue, it is useful to begin by indicating what constrains script structures. Related remarks apply to related types of structures such as MOPs (Schank, 1980) and metascripts (Abelson, 1981).

The casual definition of a script is "a stereotyped sequence of events familiar to the individual". Implicit in this definition are two powerful sources of constraint. One is the notion of an event sequence, which implies the causal chaining of enablements and results for physical events and of initiations and reasons for mental events. Causal chains are highly ruleful, and many of those rules have been spelled out explicitly (Schank, 1975; Schank and Abelson, 1977, Ch.2). The other constraint generator comes from the ideas of stereotypy and familiarity. That an event sequence is stereotyped implies the absence of fortuitous events. Also, for events to be often repeated implies that there is some set of standard individual and institutional goals which gives rise to the repetition. Furthermore,

there are almost certainly subgoals, each of which defines a scene involving a transaction between particular role players in a certain physical setting, using given props.

At the scene level of a script, therefore, there run in parallel four networks of coherences: Those governing the transactions, the role players, the physical settings, and the props. None of these entities can enter into sequences arbitrarily. Scene transitions between one physical setting and another, for example, follow the topographical rules of familiar environments. One does not step off the airplane directly into a swimming pool, or go from the doctor's waiting room into the kitchen. Role players remain from scene to scene except when somebody makes a purposive and expected entry or exit. It does not require belaboring these coherences in full detail to realize the enormous degree of constraint thus imposed on input relevant to any given script.

Perhaps one of the things that disguised the high order of systematicity of scripts was that some of the computer programs that used them, such as SAM (Cullingford, Note 3), were written in a way that did not insure against ad hoc violations of some of the constraints. A prankish programmer could perfectly well prepare an expected event sequence wherein the customer ate the check and gave the food to the cashier, thus thwarting their mutual goals, and nothing in the Script Applier would protest illegitimacy of such expectations. Of course they would turn out by experience to be useless in matching realistic inputs, but that is a very weak way to recognize absurdity. (And it is still vulnerable to the possibility of prankish inputs). Later programs such as POLITICS (Carbonell, 1978) did not, by the way, suffer the same degree of vulnerability, but this whole issue has not been treated as explicitly as it might be.

Why is any of this important? Well, there may be some people who feel it is not important, that there are more compelling issues for language AI to worry about. But it bothers me that the concept of structural constraint seems to have been coopted by the neats, when all the while the scruffy Yale programs are based very heavily on a whole series of implicit constraints.

Let us look more closely at some general issues pertaining to the idea of constraint. In the original formulation of information theory and communication theory, the structural constraints on the communicative elements were presumed mutually accessible to the sender and receiver of messages. They each knew the redundancies of letter strings or phoneme strings, and this consensus was the basis for an analysis of the information content of messages. In effect, one could ignore most of the properties of the receiver, and concentrate the analysis on the properties of the stimulus ensemble.

Nowadays the emphasis in cognitive science is on chunks of meaning, and one cannot generate meaning simply by higher-order approximations to the structure of low-level stimulus elements. The idea that the set of possible messages is very much constrained is still a powerful idea, but at least two drastic changes are necessary in applying an information theoretic type of analysis to higher-level meaning elements, say,

sentences, rather than to low-level stimulus elements such as letters. For one thing, the number of possible elements is infinite rather than finite. For another, there is no guarantee at all that the receiver of messages adequately comprehends the structure of contingencies between sentences that can possibly be generated by the message source.

People in the Chomskyan tradition writing about constraints in knowledge structures do not usually distinguish between structural constraint intrinsic to the stimulus ensemble and structural constraint characterizing the receiver's construal of the stimulus possibilities. The former focuses the analysis structure strictly on analysis of language, completely defining the psychology of the receiver out of independent existence. This is a very right-wing attitude, in the sense I have previously discussed. People are said to be the way they are because of immutable external regularities. There is little interest in studying learning, or human error, or individual differences in intelligence.

Furthermore, there is total disregard of cultural shaping of knowledge structures. That is, even in cases where there is a structural match between the semantics of the language and the corresponding mental representation in a particular domain, this match may have been produced by a process of cultural invention rather than by the inevitable emergence of a natural truth. Much social knowledge pertains to what anthropologists (cf. D'Andrade, Note 4) call constitutive rule systems, extensive networks of how to construe, how to behave in, and even how to feel about culturally defined situations. The nexus of rules defining the meaning of marriage is one example. Other examples of cultural rule systems are mental illness, senior citizenship, golf, and sexual harassment. It seems to me perfectly obvious that there is no foreordained meaning for any of these domains, or a thousand others, rather, a meaning which evolves under the pressure of social, political and economic motives and experiences. I belabor the banal here because of recently renewed claims that to know knowledge, one only need know semantics.

Having thus gored the one-horned ox, let me try to lay out a balanced view of one aspect of the interplay between mental representatives and stimulus structures. I will place the argument in the context of text understanding.

Consider an individual who is exposed to a string of language, presented one chunk at a time, say, sentence by sentence. Considerable constraint will be imposed by the general context surrounding the presentation of the string, say, whether it is a story or a piece of conversation, and what nature of the topic and style of presentation. The local context operative at a particular place within the string will exercise further constraint. Is there a way to conceptualize a measure of the degree of structuredness at any given place in the presentation of the string? Further, is there a way to think about structure such that it is a joint property of the stimulus string and the interpretive machinery of the understander?

I propose a characterization relating the structuredness of a context to the constraint in the stimulus string and something I will call the construal function on the part of the understander. The constraint in the stimulus string can be expressed by the distribution of probabilities  $P(i)$  of occurrence over all potential next stimulus chunks. If a few inputs are moderately likely and all others are of very low probability, the stimulus contains more structure than if likelihood is fairly evenly distributed over a large set of possibilities. This is as in standard information theory.

But the understander may not extract from the stimulus the available structure. The individual has expectations of what sorts of things may occur next. If something which is highly expected occurs, it is difficult to process. In general, we may imagine that there is a distribution of measures of processing facility  $F(i)$  over all possible next stimulus chunks. Under various different construals of what the stimulus string is about, thus what expectations are appropriate, the distribution of facility measures will be different. Perhaps processing facility could be operationalized as the speed with which a given chunk can be comprehended, or perhaps in some more subtle way, but in any case the understander is conceptualized as having the capacity to prepare for coming inputs by making a differential allocation of facility measures  $F(i)$ . This is a much more realistic view of the understander than assuming that he has no expectations at all, thus relatively equal facility in accepting all inputs, or only a single dominant expectation. Artificial intelligence programs that work heavily top-down always in effect smear their expectations over a domain of related possibilities. A good image for this emaring tendency is Chuck Reiger's concept of the "expectancy cloud".

The average value of processing facility is the sum of cross-products of stimulus probability  $P(i)$  and facility measure  $F(i)$  over all possible input chunks. This average facility will be high or low depending on four things: (1) the general simplicity of the context; (2) the general facility of the understander; (these two factors can jointly be characterized by the mean of the  $F$ 's unweighted by stimulus probability); (3) the degree of predictability inherent in the stimulus string (which could be indexed by an uncertainty measure on the  $P$ 's); and (4) the match between the  $F$ 's and the  $P$ 's that is, the extent to which the construal by the understander appropriately allocates her preparations in the direction of inputs which are relatively likely to occur. Under an assumption of a fixed unweighted variance of the  $F$ 's, it is easy to show that the average proportionality relationship between the  $P$ 's and the  $F$ 's, high facility attaching to relatively high likelihood, and low facility attaching to relatively low likelihood.

The match between what is expected and what might occur should not automatically be presumed, either according to some cosmic principle of innate resonance between the individual and the environment, or on the basis of a gradual learning of, and accomodation to contextual contingencies. There are at least four reasons for this. First, it is very possible, indeed frequent, for people to misconstrue situations and have a whole series of misguided expectations.

Misconstrual tendencies are very interesting to social psychologists and there has been a good deal of recent research on stereotyping, on misleading first impressions, on the effects of inappropriate but salient schemata, and on the insensitivity of false constructions to empirical evidence (cf. Nisbert & Ross, 1980).

A second reason not to presume that construals reflect stimulus constraints is that people are generally extremely slow to pick up the contingency structure in novel input materials, if they ever get it at all. Contingencies are especially problematic when multidimensional dependencies are involved. It is clear from classical two-alternative guessing situations that people are very good at learning the simplest zero-order structure, that is, the relative frequencies of two different alternatives. Although appropriate data do not to my knowledge exist, there is little reason, however, to suppose that people are adept at learning the zero-order structure within large numbers of alternatives. And it is very clear from so-called "cue validity" studies (e.g., Hammond & Summers, 1965) that there are sharp limitations on the learning of higher-order structures. When many independent cues are modest predictors of an outcome variable, people are unable to use all the cues, but settle instead for (somewhat fallible) use of three or four of them. In realistic stimulus domains where it is not at all clear how many cues of what sort there might be, the situation can be even worse. For example, in studies of how people judge whether someone else is lying or not (Krauss et al., 1976; Kraut, 1978), the facial, gestural and speech cues that judges employ to diagnose the liar overlap hardly at all with the set of cues that actually predict lying.

A third reason, related to the second, is that stimulus structure is usually dependent on the source of the string being communicated. Different communicators have different styles and different angles on what to say or write about a given topic. The understander usually will not have long enough experience with particular communicators to pick up their individual contingency structures, even if learning were very rapid.

Fourth, the construal function must be flexible. in operation, so that when there is a shift in the topic of the stimulus string, the understander can establish a new set of expectations. A lack of matching could come about if adjustments in construal were sluggish, lagging behind the stimulus. It is intuitively clear that there are both individual and situational differences in the rate of adjustment of construals. Part of the ordinary concept of intelligence, or perhaps quick-mindedness, is the ability of an individual to keep up with what is being read or said, especially if the point is rapidly shifting. Situations, too, may help or hinder quick reconstrual. There is a good deal of psychological literature on the phenomenon of perseverance, wherein a person's analysis of a problem area continues in a vein which has previously been successful, despite the introduction of new materials which make the old analysis outmoded (Luchins, 1942), or the presentation of evidence that past success was spurious (Ross, Lepper, & Hubbard, 1975).

In short, the structure in personal construals need not match the structure of stimulus constraints, for several reasons. When there is a match, however, understanding is considerable facilitated. The example of script processing provides an especially clear case. An account of a highly scripted activity such as a visit to a doctor introduces very high stimulus constraint, because only a limited number of events have high probability of occurrence in the account. If the understander construes the account as indeed concerning a doctor visit (as opposed, say, to a chat with a professional colleague), then his relevant knowledge structure will highly constrain his expectations to a small set of events. Given a sufficient consensus on what sorts of things transpire in accounts of doctor visits, understanding will (on average) be very facile.

My discussion to this point has concealed an important aspect of the concepts of constraint and construal in text processing. There are really three different types of structural limitation on coherent stimuli and coherent expectations. Recall that we are supposing that the input string is received a chunk at a time, and that we are interested in the probability of occurrence and the processing facility associated with every possible chunk. For tangibility we may suppose that the chunks are sentences. Two somewhat different kinds of constraints are those applying within chunks, and those applying between chunks. Let us designate these respectively as combinatorial constraint and sequential constraint. A third kind I will call functional equivalence constraint, to which I will return shortly.

By within chunk, or combinatorial constraint, I refer to tendencies or rules for what linguistic components go with what. This would include all of syntax, semantic rules or "selection restrictions" about sensible meaning combinations, such as what actions require animate actors and what attributes are pertinent for what object classes, and also fragments of pragmatic knowledge that tell us what combinations are unlikely in the real world even though semantically possible, such as the Queen using obscenity or coal miners curtsying. In this category of constraint, it seems quite likely that rules characterizing stimulus structure would be generally well matched by rules characterizing construals. These aspects of constraint are widely appreciated and shared, and the reasons given above in support of mismatch tendencies are least likely to apply.

By between chunk, or sequential constraint, I mean tendencies for certain chunks to follow a given chunk or sequence of chunks. From a formal point of view, one might suppose that between chunks constraint is just another bundle of selection restrictions, or set of rules about what goes with what, and thus is just like within chunk constraint but operating on larger units. In the field of story understanding, the supposition has sometimes been made that a kind of syntax exists linking successive units, and so we have the notion of "story grammars" (cf. Rumelhart, 1977; Mandler & Johnson, 1977).

Story grammars may have some use as rough tools in restricted story domains, but they have I think rightly been criticized (e.g., Black & Wilensky, 1979) when they make overly strong claims. Stories are simply not grammatical in quite the same sense that sentences are. Nor are other bodies of text.

When a literate reader asserts that a short string of writing within a period at the end is not a sentence, the literate observer will almost always agree. But who is to say that a longer string of writing is not a paragraph, or not a story. Judgements here are somewhat fuzzy, as we are dealing not with inviolate rules but with general coherence tendencies. Pragmatics and stylistics become much more important, and the role of syntax dwindles. Construals can more readily be out of synch with stimulus constraints. To think otherwise, and to regard sequential constraint as merely an extension of combinatorial constraint, is to fall too much under the thrall of the bottom-up, start small, scientific strategy of the neats.

The third category, that I have called "functional equivalence" constraint, is rather different in nature. In the general conception I have sketched of the construal process, there is the clear unrealism of supposing that the organism allocates preparatory provisions among an infinite number of possible stimuli. Much more reasonable is the supposition that the understander prepares differentially for different categories of input content, in effect grouping stimuli into equivalence classes according to the functions they serve. Within each class, processing facility would be roughly constraint. This is a fairly strong form of subjectively imposed constraint, for it says that such-and-such stimuli are to be regarded as equivalent, and all processed with ease, whereas such-and-such other stimuli, constituting another equivalence class, will be processed with less facility, and so on. Now there is great opportunity for mismatching, depending on whether the understander does or does not carve the possibilities at different joints than the probability structure of the stimuli would. To maintain matching, it would be necessary for every stimulus within an equivalence class to be approximately equally probable. The understander, however, may lump possibilities together because they have comparable personal concern or affective significance, not because they are necessarily objectively substitutable.

These ideas may become clearer if I recapitulate and then amplify the model I am outlining. We are supposing an understander exposed to a stimulus string one chunk at a time. Strong structural constraints characterize the ensemble of stimulus strings which the given string instantiates, but we do not presuppose that the understander necessarily perfectly appreciates those constraints. Instead, the understander, in a construal process, or general expectational policy, imposes constraints of her own. Construals are from time to time altered as the stimulus string unfolds, but while in force, each given construal is a processing allocation which determines the particular degree of facility with which different possible inputs and sequences of inputs would be processed. A construal defines functional equivalence classes of stimuli, such that stimulus possibilities within class are processed alike.

The average facility of processing is maximized when there is a match between the probability structure of the stimulus ensemble and the facility structure of the construal. Furthermore, given some degree of match, the average facility increases with increasing structuredness. There are really three aspects of both types of structure: the zero-order structure partitioning the possibilities according

to how likely or expected they are, and the contingency structure limiting what goes with what, both within chunks and across chunks.

I have argued that matching between construal structure and stimulus constraint structure should not in general be presupposed, but that it is likely to occur in certain contexts. Within chunks, it is likely that syntactic constraints will be well matched. Across chunks, certain types of knowledge structures permit ready construals of well-structured realities. I have already mentioned scripts as one example.

Another between-chunk example has to do with the role that knowledge of intentionality plays in story understanding. Stories tend to be summarized in terms of goals and plans of the main characters. Goals and successful plans tend to be well remembered. AI models of story understanding make a big point of tracking goals and goal fates. This was the case in Schank and Abelson (1977), and in other Yale models such as those of Wilensky (1978) and Carbonell (1978), and in a great many non-Yale models as well. Why is this? What is there about goals that makes them so important? From a scruffy common sense point of view, this question may not seem worth asking because the answer seems obvious. Goals underly most of the activities of people, and what interests people is hearing about things that other people do. One might amplify this intuition with the observation that intentionality is not observable, but must be inferred, and there is something especially intriguing in making inferences about people to explain why they do what they do. Another angle is that goals relate to emotions, as I will discuss later, and emotions are especially interesting.

A formalist would not be especially happy with such explanations. They essentially say that goals are interesting because everybody knows they are interesting; but no principled account is given. Well, it seems to me clear that a principled account can be given, and indeed, it is implicit in almost all analyses of the role of goals and plans, but it is usually not spelled out. Simply put, it is that intentionality is "inference rich"--it provides a high degree of stimulus constraint. Construals structured to match will confer great advantage to the understander. Intentional action introduces constraints of all three kinds: combinatorial, sequential and functional equivalence. Especially noteworthy is the possibility of remote sequential constraint, that is, the influence of somebody having a goal on his actions much later in time. In a story or a novel-- or in life-- it could be dozens or hundreds of pages-- or days-- before the major goals of individuals are actualized, yet that latent potentiality is present all the while. This constraint demon that goals unleash is a kind of inferential time bomb set to go off one knows not when. No Markov process behaves like this.

It is clear, therefore, that knowledge structures concerning goals are highly constraining, thus very important in the understanding process. Whether goals are the most constraining concept in texts about human activities is very hard to say, because we have no comparative constraint measures on different inference-rich concepts, averaged over all possible, or all available, or all experienced texts of a given type. I want to point out, however, that a good

guiding philosophy behind the choice of knowledge structures to investigate is to try to pick those which are both highly constraining and highly frequent, thereby being very useful for construal functions. Although we did not articulate it explicitly, this was indeed the rationale guiding the Schank & Abelson (1977) choice of scripts, plans, goals, and themes as the highest priority knowledge structure concepts to investigate. Far from being *ad hoc*, therefore, these concepts are very closely tied to ideas concerning constraint.

Thus far in my account, I have given mainly a way to talk carefully about the role of structural constraint in the process of understanding, with very little in the way of predictive principles or substantive claims. Let me now attempt some claims based on the construal concept.

I have said that construals are from time to time altered or updated during the course of understanding. I believe that what is remembered about a stimulus string is not the string itself, but the set of construals and reconstruals used in interpreting it. When a construal is active and inputs arrive which are not readily processed, that is, are unexpected in terms of the construal, then reconstrual becomes necessary. Memory for one's own construal structures, therefore, would include both the original construal and the reconstrual compelled by unforeseen stimulus events. Thus my proposal here is similar to Schank & Abelson's (1977) script pointer plus "weird list" idea, and to Graesser's (1981) "script pointer plus tag" hypothesis, except that it is more general in that the knowledge structure involved need not be a script, but could be anything. Indeed, it could be any combination of knowledge structures implicated in a construal which partially succeeds and then needs correction. (Among other recent treatments containing similar ideas are those of Lehnert (1979) and of Schank (Note 2)).

There are some memory data which fit nicely with this Construal Principle (e.g., Hastile, 1980; Spiro & Eposito, 1977; Graesser, 1981). But I am also aware that there are other data which are hard to explain by it (e.g., Anderson & Pichert, 1978), and that there are complications in applying it to data in general. Let that be a story for another occasion, however.

Reconstrual seems to me to be an especially interesting phenomenon. One class I would like to pursue arises when two incompatible construals compete with one another for processing dominance. By incompatible construals, I mean sets of expectations based on the other set. The more massive the preparations, the more serious the incompatibility. In cognitive terms, the massive preparatory part of a construal is the establishment of sequential constraints; thus incompatible construals involve opposing sequential constraints.

Characteristically associated with incompatible construals, under certain conditions, are affective experience. There are different types of incompatibility, as I shall spell out, and with each of these is associated a different variety of affect. These connections are compelling enough to serve as the basis for a theory of affect. Recently, there have arisen in psychology a number of systematic

taxonomies of affective states (deRivera, 1977; Roseman, 1979; Wilensky, Ortony & Collins, 1980) which set forth a number of disposing factors said to generate one or another affect. While highly evocative, the various schemes seem to lack a unifying principle common to their sets of affects, this I believe to be provided by the idea of incompatible construals.

My analysis owes much to the scheme devised by Roseman (1979), but is differently organized in order to get the benefit of the incompatibility principle. In conveying a preliminary version of my theory, I will continue to talk about construals, but I will depart freely from the text understanding paradigm to refer also to a behavioral situation paradigm where the prospective and actual events might happen or do happen to the individual involved, rather than his just imagining their occurrence for the characters in a story.

To avoid confusion, I should say what the theory is not about. It is not primarily about the pleasures and pains associated directly with physical sensations, either innately or thorough conditioning. Thus it is not primarily about sexual pleasure, or pleasurable tastes, or about fright, pain or disgust, or about the love or aversion for the people or objects associated with those pleasures or pains. Secondly, however, it may implicate pleasures or pains or other goal states, as will be explained. The theory is also not about the semantics of affect as coded in words or phrases capable of evoking associated emotions, such a pejorative adjectives applied to another person so as to stimulate dislike or the sunny lyrics of a love song. Rather, it concerns the emergence of an affective state as a consequence of the structural relationships in an ongoing situation: it is a structural theory of "on-line affect."

In analyzing incompatible construals, we have to ask why more than one construal would ever be necessary. An obvious answer is that the ongoing construal leads to poor understanding, and must be changed. A more direct way to say this is that anticipation does not correspond to reality. What you imagine will happen does not in fact happen, and you must update your imaginings. If the update is incompatible with the previous construal, then an affective process will occur which is both a signal and a symptom of the activity of reconstrual (and which, incidentally, will be associated with high memorability for the event which precipitated the reconstrual).

Two conditions seem basic to the degree of incompatibility of construals. One is the range of possible cognitive chunks implicated by the two construals, the second is the discrepancy in the hedonic import of the two construals, whereby one is highly pleasurable or painful and the other is not. Inference-richness and hedonic import would seem in practice to cooccur, because one mainly makes extensive inferences about that which is personally consequential. But the two concepts are conceptually separable.

In any case, not every reconstrual involves compatibility, and many incompatibilities are quite trivial in extent and significance. Thus structural affect is not freely evoked by minor alterations from previous expectations. Thus if you mistook someone to say they were from Stanford

and it turned out they meant Stamford, (Connecticut), or if you thought that a session of the conference was in Room A and it turned out to be in Room B, those changes would not provoke affect (unless your mistake led to come commitment or consequence difficulty to undo).

It is instructive to consider systematically how the inference-richness of consequential alternative construals might vary and give rise to differing affective states. A useful rubric is the intentional action sequence, where a positive or negative outcome state is cognized by the individual along with an alternative outcome of opposite import. What is of interest is the point in the sequence at which the alternativity arises, thereby determining the depth of reconstrual which is necessary. There are different classes of cases, depending on whether two alternative construals are present only in imagination, or whether one is imagined and the other represents reality, or there are representations of two disparate realities because reality has changed.

Let us suppose a sequence in which a goal leads to some planned action which through some causal instrumentality determines an outcome. Consider first the case in which this sequence has progressed up to a certain point in reality, and then there are alternative imagined construals, one hedonically positive, the other negative, of the uncertain future course of the sequence. If goal, action, and causal instrumentality are fixed, but only the outcome is uncertain, there is a minimal range of content for the alternative construals to deal with. The associated affective experience can be characterized as SUSPENSE. If goal and action are fixed, but there are alternative causal instrumentalities each potentially controlling the outcome, the alternative construals are inferentially richer and the affect in general will be more elaborate. Perhaps there are alternative authorities who may become responsible for the outcome, as for example two judges who might be assigned to your legal case, one probably sympathetic, the other probably unsympathetic. The affect here is one of the pair of HOPE/FEAR, depending on whether the favorable or unfavorable construal is emphasized.

When only the goal is fixed, but two (or more) well defined and distinct action plans are construable, each with uncertain connection to the important outcome, the associated affective state may be characterized as AGONIZING. Then not even the goal is fixed, but incompatible possible goals can be clearly construed, the affective state is one of CONFLICT.

Consider next the case in which a particular sequence leading to a favorable outcome is construed in imagination, but reality forces an alternative construal in which the outcome is in fact negative. If goal, action, and causal instrumentality are fixed, but the real negative outcome differs from the imagined positive outcome, the state is one of DISAPPOINTMENT. If goal and action are fixed, but the real causal instrumentality producing a positive outcome differs from the imagined causal instrumentality producing a positive outcome, the affective state is one of FRUSTRATION or ANGER. In relation to

previously outlined cases, it can be seen that FRUSTRATION represents dashed HOPE, and DISAPPOINTMENT is negatively resolved SUSPENSE.

In a slightly different type of subcase, the negative reality has already occurred, but the individual imagines what might have been, by reconstructing the sequence starting at a particular point of departure. The idea "I shouldn't have done what I did; if only I had acted differently, things would have been different", corresponds to a state of EMBARRASSMENT or MORTIFICATION. If the recrimination goes all the way back to believing that one has pursued the wrong goal, then the affective state is one of GUILT.

Another set of subcases arises when there is an imagined negative outcome, but the actual outcome is positive. Without going through all the details, suffice to say that depending on the sequential point at which alternativity occurs, the respective affective states of LUCKINESS, GRATITUDE, PRIDE, AND RECTITUDE can be generated.

Finally, there is the case of incompatible construals which arise because one reality is suddenly replaced by another reality. This need have nothing to do with an intentional action sequence, because it can be outside of the individual's control. If the old reality was positive, and the new reality is no longer positive, the affect is SORROW. If the old reality was not negative, but the new reality is negative, the affect is DISTRESS. If the old reality was not positive, but the new reality is replaced by one which is not negative, a state of RELIEF is produced.

I have been perhaps somewhat scruffy in my presentation of this system of 16 affects (albeit I had earlier implied I would try to be neat). It was not my intention here, however, to be complete and well-disciplined, but only to lay out a particular direction of theory and research involving the role of construals in understanding, memory, and affective experience. The conception of a construal function as a system of subjective constraints which may or may not match objective stimulus constraints is, it seems to me, a very important conception. There is no reason why the idea of systems of constraint should be abandoned to cooption by the right wing within cognitive science, which presumes to investigate Mind without reference to minds. Instead, we need in cognitive science a fusion of left and right wings, of subjective and objective, of content and of formalism.

## REFERENCE NOTES

1. Abelson, R.P. & Levi, A. Decision-making and decision theory. Chapter in preparation for G. Lindzey & E. Aronson (Eds.), Handbook of social psychology. Reading, Mass.: Addison-Wesley, 1983.
2. Schank, R.C. Dynamic memory: A theory of learning in computers and people. Manuscript in preparation. Yale University, 1981.
3. Cullingford, R. Script application: Computer understanding of newspaper stories. Research Report 116, Dept. of Computer Science, Yale University, 1978.
4. D'Andrade, R. Cultural meaning systems. Paper presented at the SSRC conference on Culture and Cognition, May 1981.
5. Wilensky, R., Collins, A., & Orotny, A. A cognitive analysis of affect. Draft of paper. U. Cal. Berkeley, Nov., 1980.

## REFERENCES

- Abelson, R.P. Social psychology's rational man. In S.O. Benn & G.W. Mortimore (Eds.), Rationality and the social science. London: Routledge & Kegan Paul, 1976.
- Abelson, R.P. The psychological status of the script concept. American Psychologist, 1981, 36, 715-729.
- Anderson, R., & Pichert, J.W. Recall of previously unrecallable information following a shift in perspective. Journal of Verbal Learning and Verbal Behavior, 1978, 17, 1-12.
- Black, J.B. & Wilensky, R. An evaluation of story grammars. Cognitive Science, 1979, 3, 213-230.
- Bower, G.H., Black, J.B. & Turner, T.J. Scripts in memory for text. Cognitive Psychology, 1979, 11, 177-220.
- Bush, R.R. & Mosteller, F. Stochastic models for learning. New York: Wiley, 1955.
- Carbonell, J.G., Jr. POLITICS: Automated ideological reasoning. Cognitive Science, 1978, 2, 27-52.
- deRivera, J. A structural theory of the emotions. Psychological Issues, Monograph 40. N.Y.: International Universities Press, 1977.
- Einhorn, H. & Hogarth, R. Behavioral decision theory: Processes of judgement and choice. Annual Review of Psychology, 1981, 32, 53-88.
- Galambos, J. & Rips, L.J. The representation of events in memory. Paper presented to the Midwestern Psychological Association, May, 1979.
- Graesser, A. Prose comprehension beyond the word. New York: Springer-Verlag, 1981.
- Hammond, K.R. & Summers, D.A. Cognitive dependence on linear and non linear cues. Psychological Review, 1965, 72, 215-224.
- Hastie, R. Memory for behavioral information that confirms or contradicts a personality impression. In R. Hastie, T.M. Ostrom, E.B. Ebbesen, R.S. Wyer, D.L. Hamilton, D.E. Carlston (Eds.) Person memory: The cognitive basis of social perception. Hillsdale N.J.: Erlbaum, 1977.
- Keil, F.C. Constraints on knowledge and cognitive development. Psychological Review, 1981, 88, 197-227.
- Kintsch, W., Miller, J., & Polson, P. Problems of methodology in cognitive science. Paper for the Cognitive Science Symposium, Boulder, Col., July 1981
- Krauss, R.M., Geller, V., & Olson, C.T. Modalities and cause in the detection of deception. Paper presented at the American Psychological Association convention. Washington, D.C., 1976.
- Lehnert, W. Text processing effects and recall memory. Research Report 157, Dept. of Computer Science, Yale University, May 1979.
- Luchins, A. Mechanization in problem solving. Psychological Monographs, 1942, 54, Whole No. 248.
- Mandler, J.M. & Johnson, N.S. Remembrance of things parsed: Story structure and recall. Cognitive Psychology, 1977, 9, 111-151.
- March, J.G. Bounded rationality, imbiguity, and the engineering of choice. Bell Journal of Economics and Management Science, 1978, 9, 587-608.
- Nisbett, R. & Ross, L. Human inference: strategies and shortcomings of social judgement. Englewood Cliffs, N.J.: Prentice-Hall, 1980.
- Roseman, I. Cognitive aspects of emotion and emotional behavior. Paper delivered at the American Psychological Association convention. New York, 1979.
- Ross, L., Lepper, M.R. & Hubbard, M. Perserverance in self perception and social perception: Biased attributional processes in the debriefing paradigm. Journal of Personality and Social Psychology, 1975, 32, 880-892.
- Rumelhart, D.E. Understanding and summarizing brief stories. In D. LaBerge & S.J. Samuels (Eds.), Basic processes in reading: Perception and comprehension. Hillsdale, N.J.: Erlbaum, 1977.
- Schank, R.C. The structure of episodes in memory. In D.G. Bobrow & A. Collins (Eds.), Representation and understanding: studies in cognitive science. New York: Academic Press, 1975.
- Schank, R.C. Language and memory. Cognitive Science, 1980, 4, 243-284.
- Schank, R.C. and Abelson, R.P. Scripts, plans, goals and understanding. Hillsdale, N.J.: Erlbaum, 1977.
- Spiro, R.J. & Esposito, J. Superficial processing of explicit inferences in text. Technical Report 60, Center for the Study of Rading, Urbana, Ill., 1977.
- Tomkins, S. The psychology of being right -- and left. Trans-action, 1965, 3, 23-27
- Tukey, J.W. Analyzing data: Sanctification or detective work? American Psychologist, 1969, 24, 83-91.
- Tversky, A. & Kahneman, D. The framing of decisions and the psychology of choice. Science, 1980.
- Wilensky, R. Why John married Mary: Understanding stories involving recurring goals. Cognitive Science, 1978, 2, 235-266.