14 Representing Knowledge Development

Gordon H. Bower
Stanford University

My task is to evaluate critically the chapters by Katherine Nelson, Lee Gregg, Steven Kosslyn, and Dorothea and Herbert Simon. The chapters are interesting, informative, and contain much that is praiseworthy. I have relatively few criticisms, most of which are very minor. I hope the contributors will take my remarks in the understanding of my assigned role here.

COMMENTARY ON NELSON'S CHAPTER

Nelson has chosen a good issue to investigate – the child's developing knowledge of everyday events. She probes for this knowledge by asking children to tell her what happens at lunchtime at their day care center and during similar routines. She then examines the sequential character of the child's descriptions of the activities.

I found Nelson's paper, particularly the opening remarks, full of very interesting theoretical questions and distinctions, but somewhat short on useful methods for collecting valid data to decide these questions. I can agree with most of what she says. Yes, of course, the knowledge a child acquires about routine daily activities serves as a background against which he learns to talk about those activities and learns to refer to objects and characteristics of such situations in their absence. And yes, many of these activities have a prototypic serial order to them. And yes, we can readily believe that older or more experienced children will have learned more about the actions, conditions, and objects involved in particular activities. In proving this point, it was unfortunate that in

Study II Nelson confounded age with time in residence at the day care center; nonetheless, I am willing to accept her conclusion that children who have been at the center longer know more about its lunchtime routines.

However, all these conclusions would have come as no surprise either to William James or John B. Watson, to mention a few of our august forefathers. The main difference between Nelson and, say, John B. Watson is in their vocabularies. Watson, of course, used the language of stimulus–response habit sequences to talk about routines, whereas Nelson uses the vocabulary of "scripts" that have variable slots that can be instantiated with details of a specific situation.

It is therapeutic to pause on occasion to ask what we have bought in switching to the frame or script formulations and away from the habit sequence vocabulary. I believe the main advantages of frames or scripts are their abstractness, the notions of variable slots, semantic constraints on what are acceptable fillers, and the passing along of variable bindings from one slot to the next in the frame. To put it concretely, it is difficult to state in S–R terms the metarule that in a restaurant scene, for example, the food you eat is usually the same food you ordered as well as the food you pay for. It is not clear when children learn such constraints – Nelson's MacDonald's interviews are silent on the point – but I suppose that appropriate test probes might elicit such information, even from 4-year-old children.

Let me discuss methods for studying scripts. Nelson has chosen to assess children's knowledge of scripts by the interview technique. She asks, "What happens when you eat at MacDonald's?," and follows that with general, then more specific probes. Finally she tabulates common responses and reports commonalty measures. There are several problems with this methodology. The first major problem is that little children do not know how to talk very well; they cannot always plan a lengthy recitation so that it all comes out straight. A second problem is that in normal conversation, even with a child, the answer you get to a question depends very much on what your listener assumes you want. Thus, when I ask you to "Describe what happens at a typical dinner," you will try to make *general* statements rather than instantiating specific details of specific events. If generalities are asked for, then that is what you get from a comprehending adult – you will not often get details of specific instances. This conversational rule might explain the absence of detail in Nelson's children's descriptions of the central event of eating, such as the type of food they ate, how they ate it, and so on. Such details may be deleted because they vary a lot, so that no instance defines the class of lunches. Notice, however, that hamburgers are mentioned for MacDonald's, where such food is invariable. Similarly pizzas and ice cream would be named for pizza parlors and ice cream parlors.

A problem with the free description technique is that the experimenter must rely on order of mention of events as the main indicator of serial knowledge in memory. But order of mention is neither a necessary nor a sufficient condition for inferring a temporally ordered memory structure. Consider some counterexamples. A baseball game unfolds in time. Yet if you read newspaper accounts of games – or listen to people describe a game they have seen – they almost never describe the events of the game inning by inning. They are more likely to mention events in their order of importance in determining the outcome. The sportswriter deliberately renders selected highlights, although his data base – his memory record, if you will – is a temporally ordered set of events. As another counterexample, people's knowledge of their house or apartment is not temporally organized, but if they are asked to describe their house, as did Linde and Labov (1975), people give a temporally ordered tour around the rooms of the house, usually starting with entering the front door. In this case, order of mention depends on spatial contiguity of the objects, as well as where the narrator begins his tour and how he proceeds through the house. Such examples only emphasize that order of mention of actions in a free description is neither necessary nor sufficient evidence for conclusions about a temporal-serial memory structure.

Another hazard of the verbal description method is that a child may name an action or event without fully understanding its proper role in the social act. For instance, a child may mention "Daddy gives money to MacDonald's lady," and may place this act in the correct temporal position. Yet the child may understand nothing of the significance of the giving of money. He may not know that it is an exchange, a fulfillment of an obligation under an implicit social contract, that the amount of money depends on the food ordered, and so on. All of these are testable bits of knowledge. Any of them might be missing even though the child can describe the event of "giving money." The point is that one must be cautious not to attribute more knowledge to children than they have.

There are other ways to probe for knowledge of an event sequence besides free description. One way is to use forced-choice recognition testing. This can be used for identifying likely *objects* or props in a script ("Is your day care lunch more likely to include milk or coke?"); for identifying the *sequence* of canonical events ("Do you nap before or after you eat?"); and for identifying *actors* and their standard *roles* in the scene ("Who's likely to call you to day care lunch, the teacher or your mother?"). Such questions could elicit considerable information that a child has stored about lunch at the day care center. Nelson's "specific probing" technique comes close to such forced-recognition methods.

What do we do with the information so obtained? How do we decide which information, elicited by what means, is in the "lunch script," versus what is ancillary and can be reconstructed from specific instantiations of the general script? This seems to be a nearly unanswerable question. Of the several indexes possible – free description, probed recall, recognition, and enactment – there is no good argument for claiming that information gained by one method is in the memory script whereas the remaining knowledge revealed by a more sensitive method is not in the script. When different retrieval methods provide slightly different answers, we must assume that the methods differ in their ability to tap into the memory that is there as well as differing in the way they get children to talk about what they know.

Having noted some alternative methods for studying scripts, I end with a few cautions, because it is not clear that scripts will solve all our problems regarding the representation of actions. First, it is not clear in Nelson's paper or in Schank and Abelson's (1977) book at what level of abstraction a memory script is to be described. For example, suppose a text is about someone going to visit a particular cardiologist. How do I describe that in memory? Is it an instance of the "visit cardiologist" script, or the "visit doctor" script (including clinical psychologists), or the "visit professional" script (including lawyers), or the "visit a person" script, or a "go to location *X* and talk to person *Y*" script? It's all right to say that these are connected in a subset hierarchy and that properties true about scripts at one level are also true about subordinate scripts in the tree; but when we fill in the slots of the new memory structure set up to encode the visit to the cardiologist, from which scripts (and at which levels) are copy tokens pulled?

A second problem is that it is not clear how many contextual variations are stored under one script header. Schank and Abelson (1977) introduce the notion of different tracks through a class of scripts. Thus, I have tracks for dining at MacDonalds, at a summer camp, a cafeteria, a church picnic, a Moroccan bazaar, and so on to the limits of my experience. The place itself tells me which track to call up from memory. But if this is so, then my knowledge is beginning to sound just like a vast collection of specific memories about eating occasions, with a context-dependent specification of roles, props, reasons, typical foods, and action sequences. Where is the conceptual economy of the general script formalism? Another tough problem is to see how context modulates and alters details of a script. To take an example, Nelson had her subjects playact a script in a toy-world of MacDonalds with dolls. How does the toy-ness of that pretend world get passed along as a context to modulate this instantiation of the child's real world restaurant script? The child expects the hamburgers to be plastic and without smell; he doesn't expect to pay real money, and so on. The issue is how we permit a context like "toy world" selectively to cancel certain aspects or expectations of the script but not others.

Despite these reservations, I find the script idea attractive, and I believe Nelson has made some fine, initial explorations into probing for action--knowledge in children. I wish her well with her further efforts.

COMMENTARY ON GREGG'S CHAPTER

Gregg's chapter concerns the child's learning of spatial concepts and the development of exocentric representations. His paper was not available to me when these remarks were prepared, so I have only my memory of his speech to rely on. I generally thought that Gregg's task domain was an interesting one, full of possibilities for research. He has 4- to 5-year-old children learning to guide a toy turtle to a specific destination by remote commands (button presses that rotate the turtle clockwise or counterclockwise, or move it forward). The task requires the child to place himself "on the turtle's head," taking its perspective when deciding whether to move it left or right. Presumably, the child could carry out the instructions if he or she was there in place of the turtle; the problem is therefore one of mapping the turtle's spatial layout onto a set of buttons.

A puzzling result from the first study is that the initial orientation of the turtle with respect to the child produced no differences in the percentage of problems solved without errors. We are left to wonder why. Perhaps it is simply because the task of taking the turtle's perspective is equally difficult regardless of the turtle's orientation with respect to the child. Or perhaps the null result came about because the children inserted extra loops and amusing maneuvers into the turtle manipulations (which were counted as "errors"), and having such fun was more rewarding than getting the turtle directly to its stated goal. The low level of "correct performance" in the first experiments suggests such an interpretation. Later experiments, introducing the horn as an explicit error signal, shaped the children's behavior somewhat more toward the experimenter's criterion of "good behavior."

It would have been informative to have an error analysis of the protocols rather than a report of overall success percentages. I would like to see a breakdown of errors by the type of move, number of steps from beginning to end in the solution sequence, and so on. I suspect that many "errors" were caused by "overshooting," by the child is turning the turtle farther in one direction than the straightline path to the goal. These analytic measures were not reported in Gregg's chapter.

Putting colored earmuffs corresponding to the color of the button on the button on the turtle's head was an ingenious method of converting the problem of perspective orientation into one of color matching. The child merely has to note which ear color is closer to the goal (if the turtle's head is pointed 90° or 270° away from the goal) and press the button of that color to orient the turtle correctly. This "redundant cue" training did not transfer to the task without the earmuff cues. Specific transfer effects really cannot be assessed in this design, however, because there is no control group (one not receiving the color-cue training), and there is a general practice effect across the entire experiment.

Gregg emphasized in his chapter that he was reporting preliminary observations from pilot studies. This is an accurate assessment of the reliability of his findings. Nonetheless, I believe he has developed an interesting task with which he and his colleagues will be able to design experiments to study reliably children's developing spatial perspective abilities. We can wish him well with his continuing investigations in this domain.

COMMENTS ON KOSSLYN'S CHAPTER

I liked Kosslyn's chapter because it is highly organized and well written, and the arguments are laid out in an orderly fashion. I believe that his reasearch provides most impressive evidence regarding the characteristics of our sensory information store and the processes that operate upon it. He is also to be commended for developing an explicit simulation model that appears to capture the major qualitative trends of his data. The existence of a concrete model serves as a good target for focusing scientific efforts; it also eases the job of a commentator, since the bolder and more definite the target, the easier it is to pick away at. Having given proper priority to my very favorable impressions of Kosslyn's paper, I mention a few reservations.

The paper may be divided into four sections. First are some general remarks about how to proceed in theoretical research. Second is a review of Kosslyn's experimental findings on critical issues surrounding the representation of sensory information in memory. Third is the computer simulation model. Fourth are some experiments on children's handling of sensory information.

First, I simply did not understand the reference to teleology and its relevance to developmental psychology; I do not see how final goal states can be a reasonable explanation of cognitive development. Specifically, I do not believe that a child's knowledge of adult cognitive processes is what causes him to develop them; instead, a lot of social modeling and interactive learning in an educational environment would appear to be largely responsible.

Second, laying down in advance and following a "decision tree" regarding critical issues in theory development (as Kosslyn advocates) is nice if and when you can manage it. But, alas, very little scientific research that I know of proceeds with such surefooted mowing down of thoughtful alternatives. My impression is that we mostly bumble and bungle along, thrashing about for some new idea or experiment to keep ourselves busy and out of mischief. I suspect Kosslyn's decision tree is, in fact, a reconstruction of a poorly ordered research history, done sometime after the fact and after considerable research had shown that his four issues were indeed of primary importance. But only he can evaluate the truth of that.

In the second section of his paper, Kosslyn adduces evidence for the view that mental imagery is not epiphenomenal but contains real information. Let us concede at the outset that organisms – people among them – store descriptions of sensory information, of patterns, and that this includes spatial measures such as the approximate size of familiar objects and the size and location of parts relative to the whole object (e.g., the size and location of the nose relative to the whole human face). We may suppose that this normative size and location information is attached directly to the concept nodes in memory networks that represent common objects such as rabbits, collies, and German shepherd's ears. It is an interesting intellectual exercise to see how much of Kosslyn's data can be handled by a processor working over such a semantic network, where imagery is not essentially involved (Pylyshyn, 1973). On this theory, the internal TV is tuned on simply for the amusement of the executive monitor but it has no real function. (Although I don't believe this thesis [i.e., I think Kosslyn is right], the Devil's Advocate role helps to clarify critical from noncritical experiments. Kosslyn and I have played out this scenario several times before to sharpen his arguments.)

A first impression is that one cannot think of how the results could have come out other than they did, given the way Kosslyn's basic experiments were run. Kosslyn induces his subjects to engage in an elaborate game of pretense or role playing. They play a perceptual game of "as if." Consider, for example, his investigation of the visual angle of the "mind's eye." Kosslyn says to his subject: "Suppose there's a cow standing over by the wall. Now, walk up to it until it fills your whole visual field, then stop. Okay, fine. Now suppose there's a skunk over there. Walk up to it until it fills your field," and so on. If you're a compliant subject, how can you possibly act sensibly on such instructions? Only by accessing your stored metric information about lengths of cows (say they're 6 ft long), and then walking up to the wall until you can just barely see a six-foot length of wall within a foveal ring. You play the same charade for skunks, elephants, and the like. Of course, when Kosslyn plots the distance you stop from the wall, it will increase linearly with the length of the animal. Yet, you never had to image anything to do the job. You would give the same data if Kosslyn told you to "image" a mythical object called a "Glunk" with only the information that it was 10 ft long.

Kosslyn answers this objection by noting that we assume a sensory magnitude component when we suppose the subject can judge a six-foot length on the wall. In other words, my objection begs the very abilities it presumes to question.

Next, consider Kosslyn's data that the time needed to scan from one object to another in an image is a function of the objects' physical distance. If you read Kosslyn's instructions carefully, they make it clear that the person is not simply to say whether the second object is in the picture; rather, he is to "scan over" to it and push the button after he arrives. The obvious demand characteristic of such instructions is that the subject should output times that are linearly related to the stored distance between the two objects. You would get the same data if I simply told you to output a response time proportional to the time it takes to drive from any point to another. But is it necessary to assume imagery is involved? Kosslyn could have the subject change the size of the "subjective map" or the speed of the "scanner," but this should simply result in a linear transformation of the response times. As an incidental comment, I suppose Kosslyn is aware that the time required to shift the gaze from one spot to another in physical space is determined by the latency of a saccad, which is nearly constant and independent of the physical distance between points (within limits).

Kosslyn would reply in three ways. First, I again beg the question of magnitude information when I assume a subject can (mysteriously) output an analog reaction time proportional to a distance. Second, several of his distance experiments did not instruct the subject so explicitly on scanning, so the compliance with demand characteristics should be less troublesome. Third, a map with 10 or so objects in it has 45 interpoint distances. It is inconceivable that subjects would compute, store, and use all such interpoint distances to guide their later reaction times. One needs a simpler representation of a map and something like a "scanner" to get Kosslyn's results. I become more convinced he is right.

Having examined some of the experimental evidence Kosslyn presents, let us turn to the simulation model of image generation that he and Schwartz have worked out. I found it a most intriguing model, well worth careful study. First, I like their idea of a visual short-term memory, where an image needs to be continually refreshed from long-term memory or it will fade out. This concept allows them to specify storage capacity in terms of the number of holdable chunks, where chunks may be objects or parts of objects. This clearly implies a tradeoff between the number of objects and the amount of detail that can be held active at one time in visual short-term memory. What Kosslyn does not elaborate on here are the conditions under which the processing of external stimuli will interfere with material in visual short-term memory. Experiments by Brooks (1968), Byrne (1975), Kroll (1975), and Segal and Fusella (1970) are clearly relevant to this issue of modality-specific interference, but it is not obvious how Kosslyn's model would deal with these differing results. That is something requiring further development.

I liked Kosslyn's way of representing perceptual information in terms of polar coordinates, the (R, Θ) pairs for points on the contours of an object. Although he doesn't mention it, I suppose one really has to store only the coordinates of inflection points, since these convey the most information about shape; a simple point connector can then fill in straight lines between the inflection points to draw a reasonable contour (see Attneave, 1954). As Kosslyn notes, the (R, Θ) format makes it easy to expand or contract images or to rotate them in the picture plane. What this representation lacks, of course, is the third dimension in which real objects exist. That requires another Θ value for each point. The model knows only a flat, two-dimensional world, whereas a real brain knows about projective geometry, knows how to relate an object's image size to its distance away in space, knows how to rotate an image in the depth plane, knows about occlusion, foregrounding, haze cues to depth, and so on. But again, these are issues for further development of the model.

One objection to the computer model is that it does not address itself to one topic that has concerned imagery researchers for the past 10 years: Why does mental imagery improve memory? Why are pictures remembered better than words? Why are concrete words or sentences remembered better than abstract words or sentences? These are the issues Paivio (1971), myself, and many others have been concerned with. In his chapter, Kosslyn does not show how his model helps to clarify such issues. One way to augment the model to handle the benefits of imagery in paired-associateds learning is to assume that a semantic relation between two concrete concepts can be displayed on the internal TV screen and that this interactive scene can be stored in memory as a new data file. Thus, for the pair DOG–HAT, the system might store an image of a DOG wearing a HAT. A similar idea has been in the imagery and learning literature for 10 years. The problem is to deduce from that view the hypotheses needed to handle the myriad of findings in this area. For example, how would it explain the fact that paired-associates learning is helped by interactive imagery but not by imaging noninteracting objects standing side by side (see Bower, 1970)? Of course, the preliminary report of a new model should not be seriously faulted for the topics not immediately covered and explained by it. Perhaps a fairer attitude is to wait and see what develops from Kosslyn's active theoretical program.

I found Kosslyn's speculations at the end of his chapter interesting and worth pursuing. I agree with his conjecture that if we repeatedly query a subject's perceptual store (e.g., "Does a Volkswagen have ventwings?"), he will soon store the answer verbally. Perhaps this is a model of how children convert perceptual knowledge into verbal codes. I believe that people make such conversions only when someone explicitly asks for the implicit information. That is, I do not believe that children sit around spontaneously formulating propositions about their experience. They may be taught how to question others as well as themselves, but that would seem to be a kind of metaknowledge that I suspect most of us rarely use.

Having offered these criticisms of Kosslyn's paper, let me end by emphasizing again my highly positive reaction to his research and theoretical program. It is bold and imaginative. I believe my Devil's Advocacy has not really damaged his case. Kosslyn infers (properly, I believe) that the cumulative weight of his evidence supports his theoretical position. His theory also leads to several important developmental questions, as he illustrates at the end of his chapter. These developmental questions provide a rich lode to be mined by researchers in future years.

COMMENTS ON THE CHAPTER BY DOROTHEA AND HERBERT SIMON

In their paper on solving physics problems, the Simons characterize the performances of two problem solvers of quite different talents, and they indicate the similarities, as well as differences, in their approaches to problem solving. Specifically, they use production systems to model the sequences of equations the two subjects use while solving 19 problems involving the laws of motion. The subjects think aloud as they work. The protocols are then analyzed by content to decide which equation the subject is thinking of. The edited protocols are then reported in the Simons' Table 13.2, which contains the final data, comparing the solution paths of the two subjects to the authors' simulations.

My first comment is that there seem to be very few errors indeed, very little search, and few false starts that required backing up. The subjects seemed to have gone straight through the problems, varying greatly in the time they took to reach the solution. Because there are so few mistakes, one wonders how much we should be impressed by the correspondence between the production system model and the data. If the problems so constrain behavior that all the subject can do is plug into one formula or sequence of formulas, then it takes no strong scientific theory to predict that he will usually do just that. I could not tell how much uncertainty there was regarding the successful solution path for a given type of problem. Are there really more than two or three ways to get to most solutions? Once a theorist notices how someone characteristically handles problems of a given type, then he simply predicts that the subject will be consistent with later problems of that type. This does not seem to be a major theoretical accomplishment.

What parts of the problem-solving process are not modeled by the Simons' production system models? First, they have clearly ignored the language input part, the parsing device that identified variable names, relations among variables, and so on. Moreover, the model treats the problem formally, in terms of abstract variables only. I suspect that this is unrealistic. Studies of logical reasoning (with syllogisms) have repeatedly found that subjects are very much affected by the way the premises are stated as well as by the content of the premises. Logically equivalent paraphrases of the premises are not psychologically equivalent. Similar influences can be expected when novices solve kinematics problems. One aspect of the expert's talent probably is his ability to extract a relevant formal description from the specific contents of the problem.

A second problem is that the simulation does not predict the few mistakes there are – the incorrect substitutions, false starts, and subjects' detection of their false starts. For instance, the model for *S*2, the slower subject, cannot calculate something and then say "Oh no, that's not what I want." But their protocols probably include many such boo-boos. The simulation does not predict the ubiquitous checking and rechecking of the answers that both subjects show, nor does it deal with *S* 2's rehearsal and summarizing of what is known and what is wanted for the problem.

Because the model does not deal with these issues or with memory searching, it does not predict the difficulty of the various problems or the time required to solve each. The solution time is not simply related to the number of equations evoked in the solution path.

As the Simons point out, the production system models provide a slightly misleading characterization of the relative skills of the expert and novice problem solvers. The model of the novice, $S2$, uses careful means–ends analyses for goal-directed planning, whereas the model of the expert, $S1$, "blindly" calculates new quantities from the givens, hoping to stumble on the desired variable. As a result, $S2$ often has fewer equations in her solution path than $S1$ has in his. She appears to be more efficient and planful, in the eyes of the model. The slack was taken up informally when the Simons claimed that the real differences were, first, that the expert calculates more swiftly and more confidently than the novice, and second, that the expert has better "physical intuitions" than the novice, thus helping him select the right equations. The Simons' data provide only exceedingly indirect evidence that the expert uses physical models. They also do not make clear how a model of the physical situation would call the production system to evoke appropriate equations. That interfacing would be interesting to see, but that is not what the chapter is about. One would like to have a generator of physics situations (like Kosslyn's procedures) that will turn on its TV screen and draw an internal model of the relevant variables and physical interrelationships described by the problem.

I have a few questions about production systems as models of psychological processes. A first question is this: Given a small sample of someone's behavior on a short list of kinematics problems (or any other kind), can any reasonably consistent behavior be reasonably well modeled by some production system? In other words, does the fact that the Simons developed a production system from each subject's solution path lend credibility to the production system method of theorizing? Second, would it not be proper to test a production system model by predicting what the subject will do on a new set of problems, either equivalent paraphrases of the prior problems or concatenations (embeddings) of them? Third, production system models are deterministic and therefore falsified by a single misprediction. If we use a criterion of falsifiability to advance our theories, will not all production system models quickly become rejected corpses along the roadsides of scientific progress? This reaction of mine probably reflects a fundamental difference in preference for deterministic versus probabilistic models.

I close my comments on the Simon and Simon paper by noting that they use the "thinking aloud" method for collecting data. At issue is whether such introspections validly report cognitive processes. Simon and Simon view their subjects' introspective reports as veridical indicators of when a production (equation) has been evoked and/or is being applied. They believe that the verbal report either precedes or occurs coincidentally with the theoretical process their simulation theory is about, namely, noticing the current conditions, evoking equations from memory, substituting values for variables, or performing the arithmetic operations. However, it may be the case that most problem-solving gymnastics occur outside of consciousness, so that only products of cognitive operations are available for report, and even these self-reports lag considerably behind, describing only those decisions that occurred several seconds previously.

I have recently been reading the literature on the accuracy (or rather inaccuracy!) of self-reports. I was particularly impressed by a paper by Nisbett and Wilson (1977) entitled "Telling more than we can know: Verbal reports on mental processes." They systematically marshall evidence for the view that people are often not aware of what they are doing, of how they have changed, of what is causing their actions, of what stimuli are controlling their feelings and behaviors. The main occasions when they happen to guess their controlling variables correctly is when the alternatives are few and conspicuous and when there is a cultural belief or causal schema which says that people's behavior in situations of type X are often reflections of this or that controlling variable.

Consider just a few examples of non-rule-governed behaviors in which introspections about causal events are false and misleading. Latané and Darley (1970) found that the biggest single variable inhibiting helping behavior from bystanders to a potential disaster (e.g., an epileptic seizure, a sickness, a robbery) was how many other bystanders heard or saw the same scene and remained passive. Yet, if you ask subjects who have been through the experiment, they almost never report that it was the presence of that passive audience that inhibited their helping. You get the same discounting from people who simply imagine the bystander apathy situation: They always underestimate the impact of social pressure. Another example: As part of an apparent consumer survey, women in a department store were asked to indicate their preference among many pairs of nylon stockings. It happened that the biggest factor in the entire experiment was the left-right positioning of the alternative (counterbalanced across types of stockings); the stocking presented on the subject's right was preferred almost four to one over that presented to her left, regardless of what it was. Yet, no subject identified position of the alternative stockings as the controlling variable. In fact, when asked that question directly, after having been shown the four-to-one preference, the women thought the questioner was crazy. Their causal schema supposed that people prefer stockings for their quality, texture, sheer appearance, and so on, not because of their left-right presentation position.

Simon and Simon may accede to these remarks about the inaccuracy of causal attributions in self-reports, but they would reply that the remarks are irrelevant to the manner in which they are using self-reports in their work. They use the person's report only as an indicator of which equation is in the current focus of attention, and they are not concerned here with modeling the subject's beliefs about what is causing different equations to come to mind in the order they do. In other words, they ascribe no causal validity to any attributions their subjects might have given, so the criticism simply does not apply to their work. This is probably correct in this instance.

FINAL COMMENTS

Before closing I want to share some thoughts with you about our popular conceptions of scientific criticism. This is the first time I have assumed the role of public discussant; I have found it a most disagreeable job. I dislike the role of examining someone's prized writing, looking for possible flaws or things to take issue with or disagree about. It is particularly distressing to do this with work of the quality of that by Nelson, Gregg, Kosslyn, and Simon and Simon. I have not enjoyed doing this hatchet job.

Why are critics enjoined to do this? It is the popular belief that Science and Truth (with capital S and T) are best advanced by the fine polishing of rough ideas before the hoary, deliberate sandblastings delivered by intellectual adversaries. This belief underlies the popular scientific debating style of challenging, questioning, probing, and seeking to find the weakest spot in someone's argument.

However, as I grow older and more mellow, I have become increasingly disenchanted with that adversarial view. It leads to devastating criticisms that are terribly destructive of the spirit, self-esteem, and enthusiasm of the parties involved as well as the onlookers. First and foremost, scientists are human beings whose productive work requires sincere praise and encouragement from peers in order for them to maintain enthusiasm for their scientific pursuits. Nothing kills our enthusiasm and joy in our work more than receiving a steady diet of nitpicking criticisms and put-downs. I believe it is much more beneficial to support one another and encourage one another's best efforts. From my perspective, what is good about a conference like this one is that it brings together a community of scholars, many of whom have been friends for years, and who respect and support one another. We depend on one another, not only for intellectual stimulation, but also for emotional support and reinforcement.

This is why the role of critic is so nihilistic and life denying. I propose that in the future we dispense with the critic and replace him or her with a Grand Celebrator, a Herald, a Cheerleader, who leaps up to proclaim the depth and originality of the participant's ideas and findings. Such a hornblower would surely make all of us – participants as well as onlookers – feel a lot better about ourselves and the professional life in which we toil.

To right the wrongs in just a small way, I thank the participants for the achivements of the intellect that have brought their work together in this volume. I thank them for sharing with us their new perspectives and insights into the topics. I have been impressed and inspired by the ideas, speculations, and new experiments reported here. We celebrate their achievements. They should feel good about themselves for what their dedication and work have contributed to knowledge.

REFERENCES

Attneave, F. Some informational aspects of visual perception. *Psychological Review,* 1954, *61,* 183–193.

Bower, G. H. Imagery as a relational organizer in associative learning. *Journal of Verbal Learning and Verbal Behavior,* 1970, *9,* 529–533.

Brooks, L. R. Spatial and verbal components of the act of recall. *Canadian Journal of Psychology,* 1968, *22,* 349–368.

Byrne, B. Item concreteness vs. spatial organization as predictors of visual imagery. *Memory and Cognition,* 1974, *2,* 53–59.

Kroll, N. E. A. Visual short-term memory. In D. Deutsch & J. A. Deutsch (Eds.), *Short-term memory.* New York: Academic Press, 1975.

Latané, B., & Darley, J. M. *The unresponsive bystander: Why doesn't he help?* New York: Appleton-Century-Crofts, 1970.

Linde, C., & Labov, W. Spatial networks as a site for the study of language and thought. *Language,* 1975, *51,* 924–939.

Nisbett, R. E., & Wilson, T. D. Telling more than we know: Verbal reports on mental processes. *Psychological Review,* 1977, *84,* 231–259.

Paivio, A. *Imagery and verbal processes.* New York: Holt, Rinehart & Winston, 1971.

Pylshyn, Z. W. What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin,* 1973, *80,* 1–24.

Schank, R., & Abelson, R. *Scripts, plans, goals, and understanding.* Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Segal, S. J., & Fusella, V. Influence of imaged pictures and sounds on detection of visual and auditory signals. *Journal of Experimental Psychology,* 1970, *83,* 458–464.