Theory & Psychology, 10 (1), 117-125 (2000). <Pre-final text>

Generative Theories in Cognitive Psychology

Gerhard Strube
University of Freiburg, Germany[1]

Modern cognitive psychology has strictly adhered to the experimental methodology of the natural sciences. Often, contributions in T&P have addressed shortcomings and possible remedies of this predominant approach and its emphasis on 'effects'. My comment will contrast this approach with the generative theories (cognitive simulation models) developed in cognitive science about thirty years ago and still not widely accepted in psychology. I'm going to characterize these generative theories, and discuss their weaknesses and their advantages over the usual way of theorizing in cognitive psychology. I hope to convince at least some readers that (i) in order to proceed in this manner, you need not to buy a ready-made 'cognitive architecture', and (ii) that this approach results in a much more rigorous theorizing (although still well controllable as a scientific endeavor).

A flourishing field of science

Cognitive psychology has probably been the fastest growing sub-field of scientific psychology, and one that has almost exclusively opted for a natural science approach. Its standard methodology consists in experimentation and statistical techniques for data analysis. Much has been written in *Theory & Psychology* and elsewhere, concerning the pros and cons of that methodology, its relation and possible amalgamation with qualitative or even hermeneutic techniques, and on the drawbacks of rigidly and sometimes even senselessly adhering to the usual methods of statistical hypothesis testing. My feeling (which I have sometimes shared with authors when reviewing articles) is that, in spite of all this valuable work, the impact of these articles will remain minor because (i) most of these articles discuss negative examples without giving positive ones, (ii) a well-reflected application of statistical models is more difficult indeed than just repeating what has always been done, and (iii), even editors of prestigious journals are sometimes happier with a 'standard' analysis even when its underlying model (say, GLM) can be shown to be inadequate for the data under scrutiny.

---

[1] Address for correspondence: Prof. Dr. Gerhard Strube, IIG / Univ. Freiburg, Friedrichstr. 50, D-79098 Freiburg, Germany. Email: strube@cognition.iig.uni-freiburg.de Homepage: http://www.iig.uni-freiburg.de/cognition

However, my present topic is different, concerning theory. Acknowledging that experimental research in the information processing paradigm has been tremendously successful during the past fifty years, I still have the uneasy feeling that much of cognitive psychology has remained below a level of theorizing that it could have easily attained. By this I do not mean the intended theory-abstinence occasionally practised by editors of experimental journals (I once co-authored an article for *Memory & Cognition* which was accepted on the condition that practically the whole theoretical part – the most interesting one in our opinion – was cancelled), rather I mean that too much experimental research in psychology focuses on isolated effects.

A psychology of effects?

I remember an incident from my student years: One hour of lectures on 'learning and memory' I attended closed with a short treatment of the reminiscence effect, only to be followed two days later (at the start of the next lecture hour) by an equally short treatment of the warm-up effect – without ever bothering to discuss their different causes, nor even the conditions where each of these effects is to be expected.

This is, of course, just anecdotal evidence. But consider the following example: If you have taught cognitive psychology for over two decades like me, you will be acquainted with experimental evidence for the usefulness of information like headlines or illustrations in text understanding. Almost every textbook cites Bransford and Johnson (1972, or 1973), where participants received an intentionally obtuse paragraph, which could only be understood with the help of a headline, or a drawing depicting the whole scene. These studies had been done within a research program that emphasized constructive aspects and the role of prior knowledge in cognition, and indeed the authors were able to show that the additional information had to be presented before the text in order to be helpful. Textbook authors, however, seem so keen to present the effect that they forget that this experiment is by no means evidence for any *general* importance of advance information in text understanding: the effect could well be limited to intentionally cryptic texts.

Let me follow this up with a last negative example. In their milestone textbook on psycholinguistics, Clark and Clark (1977) listed no less than fourteen syntactic or semantic strategies in text understanding for which there was solid experimental evidence. Each of these was discussed, with examples given, but there was absolutely no integration – indeed, it would have turned out that no integration was possible, and the experimental materials and paradigms were not compatible at all. This truly reflected the state of the art at that time, when research had not yet tried to arrive at working models of parsing and interpretation.

So, is it better now? It is, at least in the area of human parsing (see Mitchell, 1994, for a more recent state-of-the-art review), but not universally as it should be. To the contrary, topics like priming or implicit memory have been researched by systematic experimental variation on about every dimension imaginable, largely without guidance from theories about the mental representations and processes involved. Research has been effect-driven rather than theory-driven.

Reasons for effect-driven research

A major historical reason for that may be the ideal of (a misconceived) physics so prominent in twentieth-century psychology. But it was Lewin (1927, p. 408) who cautioned that even the laws of physics have a context of permissible application.

Another possible reason may be found in the more 'Aristotelian' objective to find accurate descriptions for the 'facts' that turned out to be so elusive that even here statistical models were needed to allow for relevant generalizations. A well-known colleague and friend of mine, a respected experimenter in the area of memory, once told me that he intentionally refrains from any theorizing that is not strictly covered by experimental results. This opinion restricts theoretical sentences to careful generalizations over a series of experimental results, i.e., a 'robust' effect. Anything beyond that, e.g., a pair of hypotheses about mental representation in memory and processes of encoding and recall, is much too speculative for his taste.

Such restraint looks odd when compared to the bold speculations in theoretical physics and the wealth of theoretical entities and their often indirect and complicated relation to experimental data. Yet I have found the ideal of physics and a certain reluctance to theorizing to co-exist peacefully in quite a few colleagues.

An alternative: cognitive modeling

So why not be bold and see what may be gained? Psychologists also active in artificial intelligence, most notably Herb Simon and Alan Newell from Carnegie-Mellon University, started on computer programs for solving problems in the same way as people do (more a decade of research has been summarized in Newell & Simon, 1972). In particular, they pioneered production systems as a model of human cognition. Such systems employ knowledge in the form of if-then rules, a working memory, and an interpreter, i.e., a module organizing the flow of control. The current situation (represented in working memory) determines the subset of rules applicable, and the interpreter selects one according to certain heuristics. Taken all together, this type of program can be shown to do any computation.

From these early steps emerged a number of more structured production systems as 'cognitive architectures', with SOAR (Newell, 1990) and ACT (see Anderson & Lebiere, 1998, for the most recent version ACT-R) the most widely known among these. Both are strong in problem-solving and learning, and ACT includes a declarative memory along with the 'procedural' rule memory. All these systems claim to be models of human cognition at a functional level of description.

In the meantime, a lot of psychological phenomena have been modeled, and the predictions of the model tested in psychological experiments, with remarkable success, from adding numbers in pre-schoolers to list learning and the detection of analogies (for example, Anderson & Lebiere, 1998, contains a whole series of ACT-R applications).

Lots of other computational architectures have been developed as models of human cognition in general, or for some special cognitive task. Among those that are not, like the production systems, a kind of knowledge-based symbol-processing system, the artificial neural networks, both in their localist (Feldman & Ballard,

1982) and distributed varieties (Rumelhart & McClelland, 1986) have become the most popular.

Generative theories

What makes these computer simulations so interesting from a methodological perspective? I contend that it is their ability to generate the very phenomena they have been designed to explain, i.e., they are generative theories (Strube, in press).

A theoretical result in mathematics may consist in a proof that a certain problem can be solved, but without indicating how it could really be solved. A natural law may state that the speed of light is an upper bound on speed in our universe, and let us compute how much energy is needed to attain some velocity for an object of given mass, but it won't propel any object into space. But a generative grammar, for instance, is a theory of syntax that (combined with a very simple computer program) will actually generate grammatical sentences.

Likewise cognitive modeling, if successful, enables us to generate data that have the same characteristics as the empirical data, i.e., the model produces the experimental 'effects': error types as typically found in human problem solvers, computational complexity of tasks in good correlation with human reaction times, choice preferences like those of human deciders, etc.

On the theoretical side, this means that we can say why and how these results are obtained because the model has been built according to certain principles of operation. (This holds for symbolic models, at least, but may be disputed for some types of connectionist models.) This kind of theories, then, is semantically richer: it tells us step by step how an effect is brought about, not only under which conditions it may be observed. And it gives the researcher an environment for as-if studies: how would performance be affected if certain resources were decreased, for instance? This method of 'cognitive lesioning' has already been explored by Carpenter and Just (1999). It is a special way of analyzing the interactions of specific structures and processes in a cognitive model, and their respective contributions to its overall performance, as advocated by Richman & Simon (1989), or Schneider (1988) for connectionist systems.

Encouraging cognitive modeling

So why does not everyone in psychology rush to make generative theories? Here is a short list of common objections and some answers to each of them.

Deterministic models are inadequate

Psychology has a well-established tradition of stochastic models. On the psychologically interesting functional level, many phenomena seem to be quite naturally described by recourse to probability distributions. This poses the fundamental problem of whether to conceive of the world as a probabilistic universe, or of observations and measurements in a probabilistic relations to underlying, (mostly) deterministic facts. Both ways have been followed up in psychology, especially the latter, which inspired theories of measurement and scaling, and is embodied in multivariate techniques like canonical analysis, LISREL, etc.

Computer programs, of course, are deterministic. (Still, there is room to include chance by means of pseudo-random algorithms, or true random values by online measurement of, e.g., thermic noise in a short wire.) But deterministic theories can be related to data by means of probabilistic models – the second way mentioned above. An example is the 'knowledge tracking' technique developed by Janetzko (1998), which relates conceptual networks, consisting of relations between concepts, to streams of behavioral data.

Cognitive models are not testable

This objection was heavily raised by Fridja (1967), and arguably was not quite unfounded at that time, now more than thirty years ago. Recently, Simon and Wallach (1999) have summarized six criteria of empirical adequacy for cognitive models as correspondences between a computational model's performance and the human behavior to be modeled:

- product correspondence (overall performance),
- correspondence of intermediate steps (processing in the model parallels separable stages in human processing),
- temporal correspondence (computational expense paralles reaction time patterns),
- error correspondence (same error patterns in model and in experimental data),
- correspondence of context dependency (comparable sensitivity to known external influences),
- learning correspondence (identical learning curves).

This impressive list shows that generative theories can be tested in a number of ways. What remains is the problem that many cognitive models are so complex that they cannot be tested in full. There are two ways out of this dilemma: (i) build your model on one of the well-established (and hence, already well-tested) cognitive architectures, or (ii), build a small-scale model for a specific cognitive function or task so that it can be tested extensively. A particularly nice example of fine cognitive modeling addresses the gaze control of readers (Reichle, Pollatsek, Fisher & Rayner, 1998).

There are just too many possible models

What still remains is the insight that indeed, by suitably choosing representational formats and computational processes, there may be many – infinitely many, in fact – combinations of representations and processes to produce some specific behavior.

But I hasten to add that this is nothing special to generative theories. The history of science demonstrates that our theories may be the best up to date (or even fall short of that), but may be upset tomorrow by others. The present state of psychological theory is that we typically have dozens of competing theories even for relatively small areas (think of emotion, for example). With generative theories, however, it becomes much easier to compare them and analyze how they bring about the phenomena.

Generative theories have two big assets, as compared to more traditional approaches: (i) they are sufficiently specified (or else the computer program would not run and generate the phenomena), whereas many traditional theories in psychology tend to become vague about uncertain details, or just ignore them. (ii) The very fact that a cognitive model has been implemented and is running on a computer testifies to its inherent logical consistency. Now go and try to prove consistency and completeness for a traditional theory!

Cognitive modeling is too hard

That is true, in a way. It is true that you must acquaint yourself with concepts of computation, architecture, representation and control, and techniques of computer programming – certainly not of greater difficulty or complexity than the statistical expertise quite usually acquired by psychologists, but, of course, an additional load. Again, it helps to build minimum models (for narrowly defined cognitive tasks), or have recourse to one of the ready-made architectures on the market: ACT-R, SOAR, EPIC, CAPS, or the recent tool-box for cognitive modeling, COGENT (Cooper & Fox, 1998). You may also be interested in connectionist models (McClelland, 1999).

I don't want to commit myself to a 'cognitive architecture'

Due to the history of cognitive modeling shortly described above, many people identify generative theories in psychology with the descendants of the production systems, the so-called cognitive architectures. While I have already recommended those for their advantages of providing users with a head start on a basic, but already well-tested complex system, it is also true that in order to model a specific cognitive task, it might well be overkill to use them – especially if you happen to have doubts about the general adequacy of production rules as the basic building-blocks of human cognition.

The alternative is small-scale models. In Gigerenzer and Goldstein (1996), the authors present a 'cognitive algorithm' to model how people decide, when confronted with the names of two cities, which of these is bigger. As it turns out, their simple sequential process (if, for instance, you have heard about A, but not B, answer "A" and you're done) is superior to a fully rational decision under ecologically valid conditions. The process draws on knowledge about the relevance of cities' characteristics (e.g., you know the name, or you know it has a subway system) for the task.

Summing up

Space prohibits me from giving more examples of just how intriguing and fruitful the construction of generative theories in cognitive psychology can be. And I hope to have you encouraged to try and test this interesting and still not widespread approach. You need not 'buy' a given cognitive architecture to do it. And the theoretical advantages are many: fully specified, consistent theories that are open to extensive empirical testing and give insight into how behavior may actually be produced. What more could you want?

References

Anderson, J. R., & Lebiere, C. (1998). *Atomic components of thought.* Hillsdale, NJ: Erlbaum.

Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior*, *11*, 717-726.

Bransford, J. D., & Johnson, M. K. (1973). Considerations of some problems of comprehension. In W. G. Chase (Ed.), *Verbal information processing*. New York: Academic Press.

Carpenter, P. A., & Just, M. A. (1999). Computational modeling of high-level cognition versus hypothesis-testing. In R. J. Sternberg (Ed.), *The nature of cognition* (pp. 245-294). Cambridge, MA: MIT Press.

Clark, H. H., & Clark, E. V. (1977). *Psychology and language. An introduction to psycholinguistics.* New York: Harcourt Brace Jovanovich.

Cooper, R., & Fox, J. (1998). COGENT: a visual design environment for cognitive modeling. *Behavior Research Methods, Instruments & Computers*, *30*, 553-564.

Feldman, J. A., & Ballard, D. A. (1982). Connectionist models and their properties. *Cognitive Science, 6,* 205-254.

Fridja, N. H. (1967). Problems of computer simulation. *Behavioral Science*, *12*, 59-67.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103,* 650-669.

Janetzko, D. (1998). Tracking cognitive representations. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (pp. 1229). Mahwah, NJ: Erlbaum. *See also* http://cogweb.iig.uni-freiburg.de/KT

Lewin, K. (1927). Gesetz und Experiment in der Psychologie. *Symposion*, *1*, 375-421.

McClelland, J. L. (1999). Cognitive modeling: connectionist. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences* (pp. 137-141). Cambridge, MA: MIT Press.

Mitchell, D. C. (1994). Sentence parsing. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 375-409). San Diego: Academic Press.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Newell, A. & Simon, H. A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.

Reichle, E. D., Pollatsek, A., Fisher, D. L., & Rayner, K. (1998). Toward a model of eye movement control in reading. *Psychological Review*, *105*, 125-157.

Richman, H. B., & Simon, H. A. (1989). Context effects in letter perception: comparison of two theories. *Psychological Review*, *96*, 417-432.

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

Schneider, W. (1988). Sensitivity analysis in cognitive modeling. *Behavior Research Methods, Instruments & Computers*, *20*, 282-288.

Simon, H., & Wallach, D. (1999). Cognitive modeling in perspective. *Kognitionswissenschaft*, *8*, 1-4.

Strube, G. (in press). Cognitive modeling: research logic in cognitive science. In N. J. Smelser & P. B. Baltes (Eds.), *International encyclopedia of the social and behavioral sciences*. Amsterdam: Pergamon (Elsevier).