Knowledge Acquisition as a Constructive Modeling Activity

Kenneth M. Ford
Institute for the Interdisciplinary Study of Human & Machine Cognition, University of West Florida, Pensacola, Florida 32514

Jeffrey M. Bradshaw
Computer Science, Research and Technology, Boeing Computer Services, P.O. Box 24346, M/S 7L-64, Seattle, Washington 98124

Jack R. Adams-Webber
Department of Psychology, Brock University, St. Catharines, Ontario, Canada L2S 3A1

Neil M. Agnew
Department of Psychology, York University, Room 145, Behavior Science Building, 4700 Keele Street, North York, Ontario, Canada M3J 1P3

Knowledge acquisition is a constructive modeling process, not simply a matter of "expertise transfer." Consistent with this perspective, we advocate knowledge acquisition practices and tools that facilitate active collaboration between expert and knowledge engineer, that exploit a serviceable theory in their application, and that support knowledge-based system development from a life-cycle perspective. A constructivist theory of knowledge is offered as a plausible theoretical foundation for knowledge acquisition and as an effective practical approach to the dynamics of modeling. In this view, human experts construct knowledge from their own personal experiences while interacting with their social constituencies (e.g., supervisors, colleagues, clients, patients) in their niche of expertise. Knowledge acquisition is presented as a cooperative enterprise in which the knowledge engineer and expert collaborate in constructing an explicit model of problem solving in a specific domain. From this perspective, the agenda for the knowledge acquisition research community includes developing tools and methods to aid experts in their efforts to express, elaborate, and improve their models of the domain. This functional view of expertise helps account for several problems that typically arise in practical knowledge acquisition projects, many of which stem directly from the inadequacies of representations used at various stages of system development. To counter these problems, we emphasize the use of mediating representations as a means of communication between expert and knowledge engineer, and intermediate representations to help bridge the gap between the mediating representations themselves, as well as between the mediating representations and a particular implementation formalism.

INTERNATIONAL JOURNAL OF INTELLIGENT SYSTEMS, VOL. 8, 9–32 (1993)

CCC 0884-8173/93/010009-24

I. INTRODUCTION

Most working knowledge engineers rarely consider explicitly the conceptual rationales of the methods and tools which they employ. Like experts in many other fields, they tend to rely on intuition and personal experience rather than theoretical considerations to guide them in their work. Indeed, many of the methods used in knowledge engineering are based on epistemological presuppositions that some practitioners would perhaps reject if they were to consider them explicitly. For example, the "mining analogy," which pervades so much of the knowledge acquisition literature, reflects underlying epistemological assumptions that are fundamentally at odds with much current research in cognitive science.[1] This analogy suggests that our eliciting knowledge from experts involves "mining those jewels of knowledge out of their heads one by one" (Ref. 2, p. 2). The underlying assumptions are that there exists some "gold standard" of knowledge and that the expert has captured a discrete (presumably large) part of the "reality" governing observed events in the domain.

We regard the mining analogy as misleading in several respects. Expertise is not a natural resource that can be harvested, transferred, or captured. Experts involved in knowledge acquisition are not restating a coherent body of knowledge that already exists in their minds; rather, they are engaged in a constructive modeling process, in the context of which formal representations are newly created and shaped.[3] Thus expertise is also more than mastery of some set of widely shared consensual beliefs of the kind that can be found in textbooks. In fact, we claim that the most significant aspects of experts' socially situated knowledge and skills are those of their own making, constructed out of personal experience with their social constituency (e.g., supervisors, colleagues, clients, patients). Over time, experts seem to develop and deploy a repertory of predictive hypotheses, or "rules of thumb," that constitute functional but fallible anticipations held with high confidence and uncertain validity. The greater their expertise, the further the experts' cognitive processes deviate from those of typical practitioners, and the greater the importance of their personally constructed knowledge.

In contrast to earlier simple notions of knowledge acquisition as "expertise transfer," some recent work in knowledge acquisition is frankly predicated on the specific theoretical assumption that the development of expert knowledge systems is, by definition, a constructive modeling activity, and not simply a matter of "information transfer" (e.g., Refs. 4–11). In general, constructivist theories of knowledge, and, in particular, Kelly's[12] personal construct theory, have served as the basis for several new approaches to the design and construction of automated (i.e., computer-based) knowledge acquisition tools. These include the Expertise Transfer System[13–15] (ETS), PLANET,[16] *Aquinas*,[17,18] FMS Aid,[19] Kitten,[20] Kriton,[21,22] KSSn/KRS,[23–26] *Nicod*,[27,28] and ICONKAT.[29] This constructivist knowledge acquisition paradigm is derived directly from an explicit theoretical framework that simultaneously supports the conceptual models of both the domain expert and the knowledge engineer.

The main sections of this article are concerned with the following related topics: Section II stresses the importance of theory to tool-builders and tool-users and suggests that a constructivist theoretical stance may aid those engaged in the knowledge acquisition process; Section III offers a view of knowledge in which expertise is seen as personally constructed and socially situated; Section IV characterizes knowledge acquisition as a collaborative modeling process and discusses modeling from a constructivist perspective; Section V addresses the implications of our constructive modeling approach for tool-builders with particular emphasis on representational issues; and finally, Section VI provides a general summary.

II. WHAT'S IN A CONSTRUCTIVIST EPISTEMOLOGY?

As Lewin[30] observed, there is nothing quite so practical as a good theory. We expect a knowledge acquisition tool to prove useful to the extent that we have a serviceable theory to explain its basis of operation and delineate its range of application. For example, tool-makers can exploit theory as a basis for clarifying their underlying assumptions, and also as an infrastructure upon which to build integrated collections of tools and techniques. Tool-users, on the other hand, need a robust theory to serve as the conceptual rationale for the principled application of their tools. An operator's manual alone is not sufficient. For these reasons, we advocate a *theory-based* approach to the development of knowledge acquisition tools.

Personal construct theory, as formulated by Kelly[12,31,32] and elaborated by Adams-Webber,[33,34] incorporates a formal model of the organization of human cognitive processes that provides a comprehensive and systematic foundation for addressing central epistemological issues in knowledge acquisition (cf. Ref. 4). Specifically, Kelly's principle of "constructive alternativism" asserts that "reality" does not reveal itself to us directly, but rather is subject to as many different constructions as we are able to invent. Thus any given event is open to a variety of alternative interpretations. This does not mean, however, that one interpretation is as good as any other. On the contrary, different ways of construing the same event can be evaluated by comparing them systematically in terms of their relative predictive utility.[35] It is highly likely that some interpretations of an event will prove more useful than others for anticipating similar events in the future.

The basic units of analysis in Kellyan theory are interrelated dimensions termed "personal constructs," which are viewed as templets that a person "creates and then attempts to fit over the realities of which the world is composed" (Ref. 12, p. 8). We employ our networks of personal constructs to forecast events, and later to evaluate the utility of our forecasts. This does not mean that the same event ever actually recurs, but rather that we use our personal constructs to represent perceived similarities and differences among events, and then organize these representations into coherent patterns or "schemata" within the framework of which we are able to detect certain recurrent themes in our experience over time, and then feed these representations forward in the form of expectations about future events (cf. Ref. 36). With the passage of time, the perception of new events constitutes an ongoing validational process which serves to confirm or disconfirm many of our anticipations. As a result, our constructs may undergo continuous, progressive change as they are revised in the course of experience. Specific changes in either the structure or content of our personal construct systems occur primarily in response to predictive failures (cf. Ref. 34).

Kelly[12] argues that every construct has a specific "range of convenience," which comprises "all those things to which the user would find its application useful." Accordingly, the range of convenience of each construct defines its extension in terms of a single aspect of a limited domain of events.[37] On the other hand, a particular construct seldom, if ever, stands alone in our experience, because it is usually deployed together with one or more other related constructs in interpreting and predicting events. Indeed, a necessary condition for organized thought is some degree of overlap between constructs in terms of their respective ranges of convenience.[38] This overlap, or intersection, between the extensions of our constructs enables us to formulate "hypotheses." That is, in interpreting an event we essentially categorize it in terms of one or more constructs, and then by reviewing our networks of related constructs (schemata), we can derive predictive inferences from our initial categorization. It is this predictive function of personal constructs that provides the logical rationale for Kelly's (Ref. 12, p. 46) assertion that "a person's processes are psychologically channelized by the ways in which he anticipates events." As elaborated by Ford[36] (p. 190):

> We humans frequently anticipate the occurrence or nonoccurrence of future events based on our willingness to project observed uniformities into the future. Thus, we continually glide from the past into the future with our previous experience preceding us—illuminating and organizing the manner in which subsequent events will be manifest to us.

Kelly's is not the only model of human representational processes that implies that their primary function is the anticipation of events. Several other psychologists have emphasized the anticipatory nature of all cognitive processes, including Bartlett, Dewey, Neisser, and Piaget, among others. In his logical analysis of the reflex arc concept, Dewey[39] established that, although the nominal stimulus can often be identified as a "physical" event external to the person (i.e., it can, in principle, be described adequately by a physicist), the functional stimulus (i.e., what needs to be explained by a psychologist) is constituted by the anticipatory processes of the person. As Kelly notes, "Dewey emphasized the anticipatory nature of behavior and the person's use of hypotheses in thinking" (Ref. 12, p. 129).

Heidbrieder[40] points out that Dewey's insight led his student Watson[41] to conclude that the scientific investigation of sensation and perception, let alone thinking, is impossible because the cognitive processes of the person are never directly accessible to external observers. This dogma became part of the philosophical groundwork for the development of a radical behaviorism which, in its most extreme form, included an attempt to explain all human behavior—even "verbal behavior"—in terms of external events and the organism's responses as related to its past history of inputs.[42] Nonetheless, as Deese concisely puts the case,[43] psychology is concerned not only with overt behavior, but also with our everyday conscious acts of perceiving, remembering, and anticipating events (cf. Ref. 44).

For example, Bartlett[45] argues that both perception and memory involve not only the registration of sensory patterns, but also the construction of these sensory data into something having significance that goes beyond their sensory character. He referred specifically to the process of connecting a given stimulus pattern with some preformed setting or "schema" (a term borrowed from Head[46]) as an "effort after meaning." Bartlett also suggested that, although a stimulus array may possess "reactive significance" at the level of reflex responses, as soon as the reacting persons become aware of the material with which their reactions deal, there is "meaning." In this sense, even the most elementary perceptual processes involve inferential constructions that go beyond the given sensory data. For example, according to Bartlett, any perceived similarity between events must depend on active schemata that lead to the grouping together of items of input which possess a welter of diverse sensory characteristics.

From a Kellyan perspective, Bartlett's most important contribution was to further our understanding of how we can recognize a given event as the "same" or as "different" from that which we had anticipated. For example, let us suppose that a person recognizes a currently perceived event as the "same" as the one that she observed on a previous occasion. Since this kind of recognition is frequently highly detailed, there must be some way in which specific information is preserved in the perceiving system from the first to the second occasion. The traditional solution to this problem assumed that recognition of the "same" event on a subsequent occasion requires the reexcitation of a specific "trace" or comparison with a preserved "copy" of previous sensory input. As Asch,[47] among others, points out, this so-called "solution" still leaves open the question of how the present stimulus input makes contact with the "correct" trace or copy without its prior recognition, which is exactly what needs to be explained in the first place. It was Bartlett's crucial insight that "in all cases recognizing is rendered possible by the carrying over of orientation or attitude from the original presentation to the representation" (Ref. 45, p. 193).

It also seems clear that, in order for us to recognize the "same" event on a second occasion, the new sensory data (input) must exert some control over the perception of similarity. That is, there must be some common properties in the two stimulus patterns that the processes of cognition are prepared to seize upon and elaborate (cf. Ref. 48). Thus, even if perception on each occasion involves inferential constructions, the input information itself must also play a role in accurate recognition. In short, whenever there is repeated perception of the "same" event, the stimulus patterns that activate sensory processes are presumed to have something in common.

Nonetheless, as Bartlett's analysis reveals, something more is required. Following Bartlett, Neisser[49] contends that stimuli do not simply impose their impressions on a passive receptor. For instance, we are able to "see" an object only after an elaborate process of construction, which typically makes use of both the available stimulus material and "traces" of previous acts of construction. It follows, according to Neisser,[48,49] that the whole conception of structured cognitive processes is fundamentally different from that of a simple response sequence. Neisser[49] also adopts Bartlett's suggestion[45] that experience leads to a gradual building up of cognitive structures that are nonspecific, but organized representations of a great number of individual acts of construction, and proposed further that a cognitive system stores information about its own constructive processes rather than the products of those constructions. That is, the information that is retained consists of traces of similar acts of construction, and it is organized in ways that correspond to the structure of those acts. These cognitive structures (schemata) control the fate of information that is to be stored, and are themselves information of the same kind. Thus they are integral parts of all of our memories, and they also provide articulate patterns (anticipations) into which new material can be assimilated.

In a similar vein, Piaget[50] refers to this sort of anticipatory schema as a "gestalt with a history." More precisely, he submits that[51] (pp. 86–87):

> Perception itself does not consist in a mere recording of sensorial data, but includes an active organization in which decisions and preinferences intervene and which is due to the influence of perception as such on this schematization of actions or of operations.

From the standpoint of the entire cognitive system, the activity of each schema can be viewed as "the part (i.e., the sector of activity or functioning sector) played by a substructure in relation to the functioning of the total structure and, by extension, the action of the total functioning on the functioning of the substructure" (Ref. 52, p. 165).

An important epistemological issue is that of how our anticipatory schemata become adapted to our environment. To ignore this question would leave us with the position that any possible interpretation of an event is just as useful as any alternative interpretation. Mancuso and Adams-Webber[35] point out that the problem of cognitive adaptation can be viewed as essentially a matter of convenience in anticipating events. Thus it can be related directly to Kelly's assumption that any construct, or subsystem of interrelated constructs (schema), has a limited "range of convenience" which comprises all those events to which an individual would find its application predictively useful. The range of convenience of a construct, or subsystem of interrelated constructs (schema), will by definition delimit the specific search space that is relevant to evaluating that construct or subsystem in terms of its predictive efficiency. If this were not the case, adaptation could not take place. That is, if our representations were not specifically anticipatory in the sense that they are open to "relevant" positive and negative feedback from events, then cognitive development could not be constrained in any way by whatever environmental parameters govern the pattern of sensory input (cf. Ref. 49).

As Agnew and Brown[1] point out (p. 154):

> Kelly (1955) early recognized that individuals must possess mechanisms that automatically restrict their range of attention. He postulated that a construct, or a hierarchy or network of constructs, bounds our anticipations of particular experience, and selects abstractions from possible worlds, large or small, to serve the anticipations. Our constructs reflect our bounded rationality by limiting the number of events addressed, and by operating within a restricted or manageable frame of reference.

Thus Kelly resolves in "functional" terms the problem of reducing the search space of a problem to manageable size in that, as Agnew and Brown note,[1] the search space of any problem is automatically constrained by the range of convenience of the constructs that we apply to it (see Refs. 33, 36, 53). Nonetheless, this "solution" provides us with precious little guidance concerning the "pragmatic" issue of what particular constructs should we attempt to apply to a given event. From a Kellyan standpoint, this question is entailed in his fundamental tenet of "constructive alternativism," which implies that events are, in principle, subject to as many alternative ways of construing them as we ourselves can invent (cf. Ref. 33). This "larger" problem is also identified by Agnew and Brown: "if our knowledge relies on robust feedforward mechanisms, and highly selected abstracted feedback, then much of such knowledge must be highly fallible" (Ref. 54, p. 21).

Mischel[55] raises the related issue of how any of our anticipations can ever be invalidated if we evaluate all "feedback" from the environment in terms of the same set of constructs (i.e., "schemata") that we originally used to formulate those anticipations? Warren notes that Mischel's question is hardly specific to personal construct theory[56] (p. 11):

> Taking a more general view, I consider the point Mischel raises here to be a basic problem for all psychological theories which attempt to take perception seriously. It is a matter of the veridicality of perception or construction and how it is checked by the perceiver. It crops up under headings like "the Selectivity of Perception" or the "Transformation of Information Input." All theories using the concept of "hypotheses" or "expectation" run into this issue sooner or later.

From a constructivist standpoint, by definition, there can be no independence of the thing cognized from the cognition of it.[57] As Mischel puts the case, "since experience is not 'given' but is constructed by us according to rules that we prescribe for it, what we know is always things as they appear to us, never things in themselves" (Ref. 58, p. 18). It follows that all of our experience is constituted by our own constructions. Our anticipations are themselves constructions of the same sort, only projected toward future events. As Agnew and Brown[1] point out, it follows that the "feedback" in terms of which we evaluate our anticipations is also constructed by ourselves, and does not necessarily reveal the "real" nature of events as they are in-themselves, that is, independently of our own construing. Consequently, the problem of how closely our representations correspond with "events-in-themselves" simply cannot arise from a constructivist perspective.[57,58]

If, as Agnew and Brown put it, "reality does not directly reveal itself to us," (Ref. 1, p. 6) how can we evaluate the adequacy of our knowledge? Logically, we cannot "step outside" of the framework of the representations that we ourselves have constructed in order to compare them directly with external events. Kelly[32] was reaching toward a possible resolution of this fundamental epistemological issue when he suggested a strictly pragmatic approach to assessing the adequacy of our representations. He proposed specifically that they should be evaluated in terms of their basic function, which is anticipation.

None of the explicit assumptions of personal construct theory logically precludes the possibility that the underlying structure of reality may someday become fully intelligible to us. As Agnew and Brown note, "Kelly's model does not rule out the possibility of isomorphism between subjective criteria, on the one hand, and domain structure, on the other" (Ref. 1, p. 19). He did maintain, however, that the only currently available criterion for evaluating the adequacy of our construing is its predictive utility with respect to our own experience. Insofar as the principle that regulates cognitive processes lies in the "mind" and not in external events, it consists of our intention to bring about a correspondence between our future experience and certain of our anticipatory representations (cf. Ref. 35). Thus our confidence in our "knowledge" tends to be enhanced by new experience that is evaluated as consistent with our anticipations. As Warren explains[56] (p. 11):

> [The] criterion for a person's assessment of the outcome of his anticipations [is] the internal consistency of the personal constructions within the person's construction system . . . truth becomes a matter of coherence within a system rather than of correspondence with reality.

On the one hand, we are not in a position to specify the relationship between events and our own representations. On the other hand, we are ready to assume that, as we improve our capacity to anticipate events, the overall pattern of our experience will gradually become more coherent. Moreover, there is no specific reason for us to suspect that our construing will not continue to accommodate to whatever (unknown) parameters define "reality." Indeed, it might be the case that, as Agnew and Brown suggest, "reality plays an indirect and approximate editing role for some of our perceptions and beliefs" (Ref. 54, p. 17). Thus we agree with them that "Kelly's theory can provide for an optimism that some knowledge, through time and through intra- and inter-individual winnowing, achieves increased 'external' and general validity, knowledge that represents more than disposable cultural myths, or highly local or personal empirical or symbolic fabrications" (Ref. 1, p. 11).

These considerations may help us specifically to explain how so many of our anticipatory representations (hypotheses) can be functionally useful, despite our being unable to determine their so-called "objective" truth status. That is, even if all our current representations are of indeterminate validity with respect to their degree of correspondence with an independent "reality" underlying events, they can still prove useful for anticipating new possibilities as we persevere in our efforts to improve the range of convenience of our representations, and to explore still unknown potentials of human experience. According to Kelly, the issue is not whether any of our current hypotheses are true or false, but rather the pragmatic question of which of them might be the most useful axes of reference for charting alternative courses of action in terms of their anticipated consequences, and then making sense of feedback from future experience.[59]

III. WHAT'S IN AN EXPERT? (OR, WHERE'S THE EXPERTISE?)

As noted earlier, many workers in knowledge acquisition presume that there exists some "gold standard" or reality of knowledge and that experts discover various parts of this existent knowledge. Some of them assume, for example, that two or more well known experts may disagree because they have access to different chunks of the "total knowledge." However, from a constructivist perspective, we would expect that experts in the same domain are likely to agree about much of their knowledge (i.e., widely shared consensual beliefs) and yet each of them might rely also to a considerable extent on a unique fund of personal experience. Thus a critical task in knowledge acquisition research is the development of adequate tools and techniques for the purpose of assisting the knowledge engineer and expert in their task of collaboratively building a domain model. This modeling activity can make explicit the valuable personally constructed experience that experts frequently use, but are often unable to articulate.

For instance, it has been reported widely that when domain experts are asked to explain how they reach a given conclusion, they often construct plausible lines of reasoning having little correspondence with their actual problem-solving methods.[60] Waterman has described a particularly troublesome knowledge engineering paradox,[61] that is, the more competent domain experts become, the less able they are to describe the knowledge they use to solve problems. It appears that in many domains, the "experts" are largely unable to communicate that specific knowledge that makes them experts. This sort of discussion raises an interesting question: If experts cannot verbalize it and knowledge engineers cannot find it, we might wonder—where's the expertise?

We have suggested above that whatever resides inside experts it is not veridical pictures or maps of reality. But if objectivity is dead[62]—what replaces it? Some authors[1,4,63] suggest that knowledge can be viewed as functional but fallible constructions not of reality writ large but of experience. In this sense then an expert is perceived to possess more functional representations than nonexperts. Such representations may be more functional on various dimensions such as: simplicity, coherence with other treasured representations of a given constituency, and the degree of correlation with an institutionalized power base (e.g., science, church, military), among others. Notice that the expertise does not reside in the expert per se but in the expert-in-context. In brief, expertise is socially situated. Not only have we lost an external (reality) reference for expertise, but we have lost an individual reference as well. The minimum unit of analysis of all perception by expert and nonexpert alike is not the individual, but rather is the individual in context, or the expert in concert with his or her constituency.[63]

According to the above perspective, the rational–empirical knowledge engineer is in trouble on at least two counts: first, for assuming that there exists an objective criteria (reality) by which to judge expertise; and second, for acting as if expertise resides in, and can be extracted from, the expert. In contrast, a knowledge engineer with a constructivist orientation would be looking for expertise in terms of functional (but fallible) interactions between the expert and his/her social context. For example, some physicians are deemed to be "experts" not necessarily because they "possess" more valid medical information than their colleagues, but rather because they are *perceived* to be experts (for a variety of reasons) by their medical constituency. They are experts because their interactions with their patients and colleagues are perceived to be more functional than those of others. Notice, this perspective suggests that expertise is a quasistable state resulting from the selection of an expert by a constituency. The expert's constructions or procedures (i.e., mental models) need not be valid, in a rational–empirical sense, they need only be functional in helping the constituencies manage their uncertainty, just, for example, as all kinds of "invalid" past medical practice (when seen from the vantage of current medical belief) have done. This suggests a kind of natural selection of experts (and their constructions) to service the current needs and criteria of the constituency network.

Such a functional view of expertise helps address several problems that typically arise in practical knowledge acquisition projects. For instance, it accounts for why rational–empirical knowledge engineers often encounter so much trouble in their attempts to "mine" expertise out of domain experts. These knowledge engineers are looking for the wrong thing in the wrong place. Their efforts should instead have a twin focus, concentrating first on modeling the expert's personally constructed knowledge, and second on the current functional selection criteria of the constituency in relation to the trial and error gambits of their domain experts. In particular, they should be looking for what constitutes a functional solution not only in terms of formal domain content, but also in terms of knowledge that can help explain this expert's selection by the constituency network. We suggest that this could account for much of the difficulty in the knowledge acquisition process—there may be little expertise of the kind being sought—however, there may be a great deal of highly functional, if highly fallible, expertise located in the goodness of fit between the selected expert and the practical needs of their constituencies.[63]

As noted, we look at expertise in terms of expert-in-context, where the unit of analysis is an interaction between a constituency and the selected expert. At an individual level of analysis, experts manage their dynamic and painful representational incoherences through cognitive/affective trial and error gambits conducted within a shifting person-in-context flow of experience. Given that two potential experts meet the minimum qualification of perceived constituency membership—and so are in competition for a given network niche—which one is selected to occupy the niche will be mainly a function of the goodness-of-fit between the potential expert's current constructions and the quasistable selection criteria of the constituency.

In the case of some narrow technical domains the constituency criteria and the potential expert's representations are relatively easy to specify and compare. Whereas in most domains of sufficient importance to warrant the construction of an expert system (e.g., medicine), "political" and nonrational criteria will weigh more heavily, and require different knowledge engineering strategies to locate the personally constructed and socially situated expertise.[63] We agree in substance with Clancey[64] that representations are context dependent, and that many cognitions are constructed "on the fly." The life span of a representation, or of expertise, depends on its functional fit to context. Thus the survival of expertise, like the survival of species, is at the mercy of their fit to context.

IV. WHAT'S IN A MODEL?

As noted earlier, according to Lewin[30] there is nothing more practical than a good theory. We recognize the pragmatic value of a theory when it provides explicit predictions that can be checked against careful measurements. This level of sophistication, however, usually is reached near the end of the theory-building enterprise. We frequently overlook how vitally useful even a vague or metaphorical model can be in helping us initially to identify and structure a potentially rich problem domain. Even a metaphoric model can afford a rich and expanding basis for communication and debate. For example, Bohr's metaphoric representation of the structure of the atom served as an intuitively based model which attracted a host of other investigators. They engaged his model as a shared metastructure, or representation of the problem domain, and explored its implications. Initially, a useful model should provide a linked set of explicit, as well as implicit, assumptions that engage and sustain both scholarly debate and technical innovation. This ill-defined process of trial and error exploration eventually may illuminate implicit assumptions, some of which will be found wanting and, consequently, either modified or jettisoned. This may help us to avoid major failures in later stages of implementation.

In a constructive approach to knowledge acquisition, experts and knowledge engineers cooperatively build models comprising explicit representations of problem-solving processes for a particular domain. These external models are largely based on the expert's internal mental "model" of how to successfully interact with the domain (including his/her constituency). Thus the product emerging from the knowledge acquisition process is essentially a model of a model. We are not implying that these models will necessarily constitute the basis of the performance system (as in a model-based reasoning system). Neither are we claiming that these models represent in any but a very abstract sense what resides "in the head of the expert." These models are valuable because they can provide rich descriptions of domain knowledge independent of any

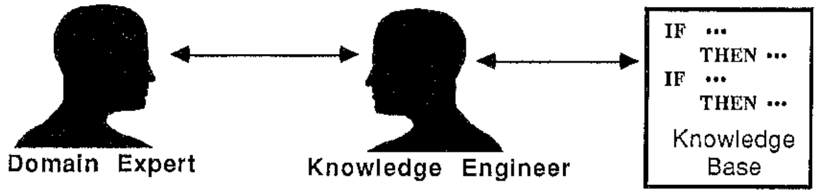

Figure 1. A traditional depiction of the knowledge acquisition process.

particular implementation formalism. Furthermore, they can serve as a basis for communication between expert and knowledge engineer.

From a constructivist perspective, a model is not a "picture" of the problem, but rather a device for the attainment or formulation of knowledge about it.[65] Indeed, sometimes the most important outcome of the modeling process may not be the model itself, but rather the insight we gain as we struggle to articulate, structure, critically evaluate, and agree to it.[66] By the same token, the value of a particular knowledge acquisition effort derives not simply from a final "correct" representation of the problem, but additionally from our success in framing the activity as a self-correcting enterprise that can subject any part of the model to critical scrutiny, including our background assumptions. From this standpoint, the crucial question for knowledge engineers is not "How do we know the model is correct?" (every model is, to some extent, an oversimplification); but alternatively, "How useful is the model (and the modeling process) as a means of facilitating our understanding of the domain?"

Unfortunately, the emphasis given to rapid prototyping in traditional accounts of knowledge acquisition, together with the faulty notion that "the production of working code is the most important result of work done," can lead to the premature encoding of knowledge in an implementation formalism associated with a specific performance environment[5](see Fig. 1). Johnson has described some of the problems with this approach[67] (pp. 179–180):

> When an expert system shell is the final destination for the expert's knowledge it is not uncommon for there to be no independent statement of the knowledge other than the rule base and some glossaries in the help information of the system. While this state of affairs may be quite acceptable for small scale applications . . . it is likely to be quite unacceptable in large-scale applications. For where many are gathered together to build an expert system, the team will need access to a statement of the problem and, as the project progresses, to the emergent knowledge not yet expressed [or expressible] in the concepts of the final implementation language . . . [T]he problem is the familiar one of software engineering where mistakes made early show up late and are thus more expensive to correct. Thus for software engineering in general (and KBS in particular) flaws in the requirements analysis (cf. elicitation stage) may not be apparent until a substantial prototype is in the field.

In brief, time invested in modeling—exploring and illuminating the domain expert's implicit assumptions (as well as our own) about the structure and dynamics of a problem domain is time saved in endless retrofitting, or failure, at the implementation stage. The problems of premature encoding of knowledge in implementation-driven representations have spurred efforts to develop other representations that more adequately support the early stages of conceptual modeling. We call these mediating representations (see Sec. V.A).

V. WHAT SHOULD BE IN A KNOWLEDGE ACQUISITION TOOL?

The *modeling* perspective implies the need for active participation by both experts and knowledge engineers in the creation of knowledge bases. It is precisely this kind of collaboration that automated knowledge acquisition tools should be designed to support. This support, however, should not be provided in an undisciplined, ad hoc manner. Thus, as discussed in Sec. II, we advocate a *theory-based* approach to the development of knowledge acquisition tools.

There is also the practical issue of how the model is to be maintained during its gradual evolution through numerous cycles of refinement. Knowledge acquisition does not culminate at some arbitrary point in development, but rather extends throughout the life of the system. It follows that modeling tools must facilitate the gradual evolution of the model through numerous cycles of refinement. Therefore they should support at least the following four facets: (1) elicitation and model construction, (2) analysis and refinement of the model, (3) maintenance of the knowledge base in the resultant performance system, and (4) model elaboration as part of an explanation capability. Advances in knowledge acquisition theory and methodology have led to the development of a new generation of knowledge acquisition tools (e.g., ICONKAT, DDUCKS, and KSSn/KRS) guided by this theoretical and methodological framework.[68]

Ideally, our representations and tools should support a smooth evolution of the model from an easily communicated, relatively unconstrained, conceptual statement of the problem to an unambiguous specification of design for the performance system. This suggests a requirement that automated knowledge acquisition tools accommodate the changes in representation that may accompany successive stages in model construction: from vague mental models to increasingly refined and explicit conceptual models via elicitation and analysis techniques, and eventually, from these highly elaborated models to an operational knowledge base via formalization and implementation procedures.[69]

Consistent with this objective, some researchers have found it useful to distinguish representations according to the roles that they play in the knowledge acquisition process. While some representations are directly executable in the performance system (e.g., production rules embodied in an implementation-specific syntax), others cannot be directly executed, but are useful because they serve as a medium of communication between expert and knowledge engineer (*mediating representations*, such as repertory grids[12] and concept maps[70]). Likewise, still other representations (*intermediate representations*) have been invented to help bridge the gap between mediating and executable representations. Mediating and intermediate representations are discussed in Secs. V.A and

V.B, respectively. Unfortunately, the selection of appropriate representations for knowledge acquisition is often far from straightforward. An important issue in the choice of a knowledge representation is the tradeoff between "acquirability" (i.e., ease of use by humans) and the computational expressiveness of knowledge representations. In Sec. V.C we discuss some implications of this tradeoff.

A. The Role of Mediating Representations

In the most general sense, we can think of a representation as being a set of conventions for describing some aspect of the world.[71] For example, we can think of alphabets, mathematical symbols, musical notation, and engineering drawings as being different forms of representation, each tailored for some specific purpose. There are often many logically equivalent methods by which the same information can be represented. Traditionally, the choice of a knowledge representation formalism has been based upon considerations of representational adequacy, inferential adequacy, and inferential efficiency.[72] Most knowledge representation research has focused on these problems. In contrast, our own work is more concerned with the acquisitional efficiency of alternative knowledge representations. Wielinga *et al.*[11] present a number of compelling arguments for making a clear distinction between knowledge-level conceptual models and implementation-focused design models in the knowledge-based system development process (see also Schreiber *et al.*[73]). We argue further that considerations of human efficiency far outweigh considerations of machine efficiency for complex modeling problems.

The term *mediating representation* has various interpretations in the literature; however, we shall use it to "convey the sense of. . . coming to understand through the representation" (Johnson,[67] p. 184). A crucial requirement is that such mediating representations should be "easily readable by those who were not involved in the original development programme . . ." (Diaper,[74] p. 34). This is essential, since executable knowledge bases are seldom organized from the perspective of humans, but instead for the convenience of the representation and reasoning mechanisms of the performance environment. We may say, then, that the design of a mediating representation should be optimized for human understanding rather than for machine efficiency.

The choice of representation can have an enormous impact on human problem-solving performance.[75] As an example, consider the fact that numbers may be represented as Arabic numerals, Roman numerals, or as bits in computer memory. While all of these forms are logically equivalent, they are not the same in a practical sense. For example, it is much more efficient for a computer to multiply numbers represented as bits than as numeric symbols. Similarly, from a human perspective, it is easier to do multiplication with Arabic numerals than with Roman numerals.

Research on mediating representations aims to improve the knowledge acquisition process by developing and improving representational devices avail-

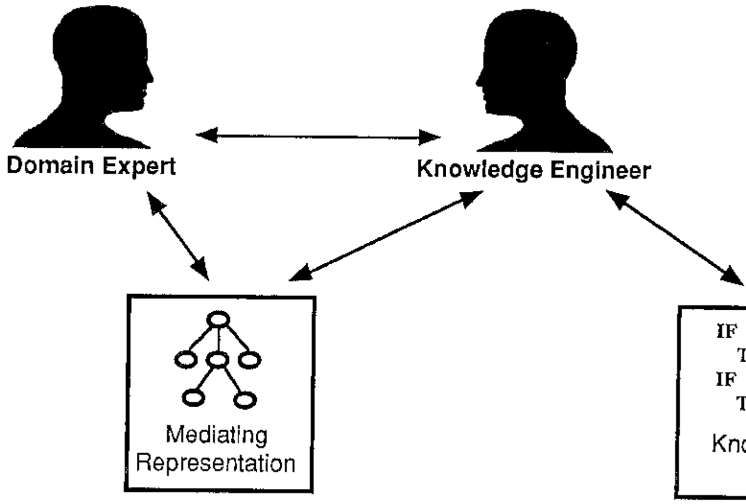

Figure 2. Mediating representations facilitate communication between domain expert and knowledge engineer.

able to the expert and knowledge engineer (see Fig. 2). A good mediating representation can facilitate modeling processes by providing a medium for experts to model their valuable, but difficult-to-articulate, knowledge in terms of an explicit external form. The mutual development of an external *cognitive artifact* supplementing the exchange of information between participants promotes and enriches communication, leading gradually to a shared understanding of the emerging conceptual model of the domain.[76] In this way, mediating representations enable domain experts and knowledge engineers to cooperatively build problem-solving models. In the later stages of system development, mediating representations may also facilitate maintenance and explanation by enabling both knowledge engineers and the system's eventual users to explore the conceptual domain model without resorting to low-level representations (e.g., C code, LISP, rules).

A number of automated knowledge acquisition tools are beginning to incorporate effective mediating representations. These tools tend to adopt one of two approaches. Either they contain interfaces that bear a close resemblance in appearance and procedure to the original manual task (e.g., cancer-therapy protocol forms in OPAL[9] and engineering notebooks in vmacs[77]) or they rely on some easily learned, generic knowledge representation form (e.g., concept maps and repertory grids in ICONKAT[29]).

B. The Role of Intermediate Representations

Since knowledge acquisition, like all modeling activities, is a process of iterative refinement, we would like to be able to map back and forth from the kinds of representations used in performance environments to mediating representations that are more useful for communication. For example, initial approaches to knowledge acquisition in ETS,[13–15] KSSO,[78] and *Nicod*[27,28] embodied procedures for transformation from repertory grids to rules. This was found to be a useful and productive step for knowledge engineers, particularly in the early prototyping phases of a project. Some kinds of information, however, could not be conveniently represented in simple repertory grids. Furthermore, this was essentially a one-way procedure—while the kinds of knowledge available in repertory grids could be transformed to rule sets, in most cases there was no natural mapping from rules back to grids.

Over time, the semantic gap between the representations used by knowledge acquisition tools and those typically associated with performance systems has widened dramatically. A distinguishing characteristic of some of these tools (e.g., ICONKAT,[29] DDUCKS,[68] and KSSn/KRS[23–26]) is the degree to which they promote the use of multiple perspectives on the same information and exemplify the push toward "informal" graphical and textual means of knowledge representation. As new mediating representations have enhanced the richness, complexity, and subtlety of the knowledge elicited by automated knowledge acquisition tools, researchers have defined a requirement for *intermediate representations* that can integrate diverse perspectives and help bridge the gulf between human participants and the implementation formalism required by the performance environment. In addition, intermediate representations facilitate the integration of knowledge acquisition and performance systems, allowing rapid feedback, dynamic analysis, and verification throughout the process of system development (e.g., Refs. 79, 80).

Figure 3 outlines a three-schemata architecture for knowledge acquisition tools, with mediating representations serving as an external schema, intermediate representations corresponding to a conceptual schema, and the knowledge base as an internal schema.* Views containing mediating representations are coupled to the underlying intermediate representation so that any changes made to one view may be reflected immediately in all related views. Knowledge analysis and performance tools may be similarly designed to exploit the integration of information at the intermediate level.

An intermediate knowledge representation is one "which only exists between flanking representations and is bound to them by clearly defined projection rules which map one representation to the next" (Johnson,[67] p. 184). The issue of mapping between representations is a troublesome one. For one thing, it is obvious that much of what can be modeled in mediating representations cannot be directly incorporated into the current commercial performance systems. Furthermore, since every transformation of knowlege is a reconstruction of that knowledge, we know that, even if logical equivalence as part of representational mapping is assured, we cannot assume practical or even conceptual equivalence. For these and other reasons, automated mapping between repre-

*Obvious similarities will be seen between our suggested architecture for knowledge acquisition tools and the proposed ANSI-SPARC three-schema model for data management.

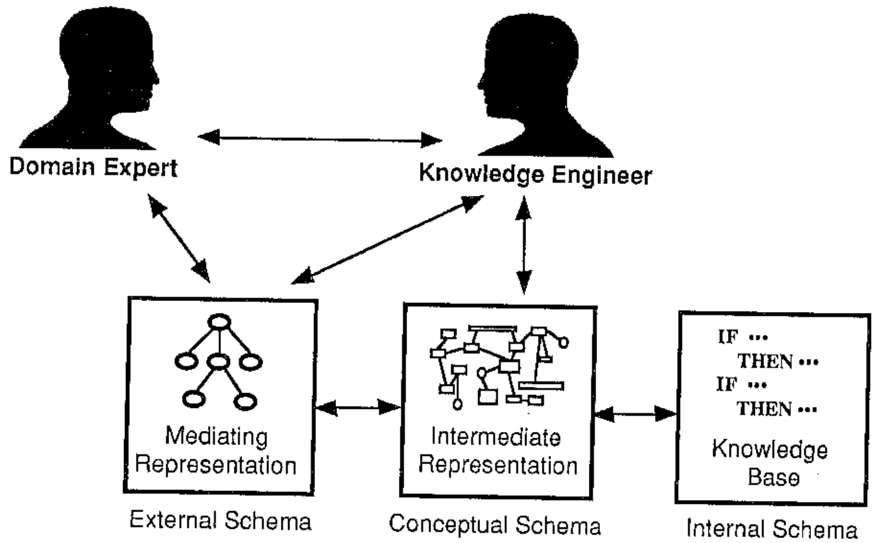

Figure 3. Three-schemata architecture for knowledge acquisition tools. Intermediate representations provide an integrating structure for the various mediating representations and can form a bridge to the knowledge base.

sentations will continue to be an issue, and some amount of manual mapping from one representation to another will remain common practice. Whether mapping is automatically assisted or manual, informal or formal, does not obviate the need for integrative, intermediate representations that are relatively independent of the constraints of the delivery environment.

C. The Tradeoff between Acquirability and Computational Expressiveness

On what basis can we judge the probable efficacy of a proposed representation? How does one actually go about designing appropriate representations for a particular knowledge acquisition activity or tool? For a mediating representation perspective, the following criteria, derived from Johnson[67] and Winston[71] are important:

- Is the formalism sufficiently expressive?
- Does the formalism aid communication between the members of the development team?
- Does the formalism actually guide knowledge analysis in a significant way?
- Does it make the important things explicit, suppressing detail and keeping rarely used information out of sight, but still available when necessary?
- Does it expose natural constraints?
- Is it complete and concise, efficiently saying all that needs to be said?

If we could design a representation in conformance with the above "acquirability" criteria and also endow it with "Turing equivalent" computational

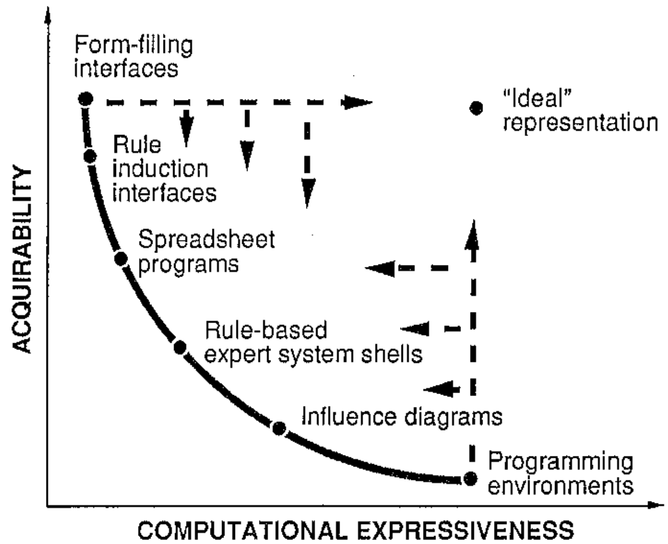

Figure 4. Tradeoff between acquirability and computational expressiveness. (Adapted from Gruber[81] and Webster.[82])

expressiveness, we would have achieved our ideal. Unfortunately, there appears to be an inevitable tradeoff between the acquirability and computational expressiveness of knowledge representations.[81,82] Figure 4 shows this tradeoff as a dark curved line. On one hand, programming languages are the epitome of computational expressiveness, but are not usable by those lacking special training (e.g., domain experts). On the other hand, form-filling interfaces (resembling the way a user normally enters information on paper) may be easy to acquire, but they tend to be rigid, and thus limited in their range of applicability to specific problems that the system designers have foreseen. The dotted arrows in Fig. 4 illustrate the knowledge acquisition tool designer's dilemma of trying to create an "ideal" representation that combines the naturalness of form-filling interfaces with the power and flexibility of a Turing machine. It seems that the more computationally powerful the representation, the more difficult it is to maintain a high level of acquirability. This is the same predicament as that faced by software engineering researchers in their attempts to achieve the goal of automatic programming.[83]

Knowledge acquisition tools do not eliminate the conflict between acquirability and computational expressiveness; however, they can act as a kind of "magnet" to help shift the curve (see Fig. 5). Through the use of knowledge acquisition tools, acquirable interfaces can become more powerful; and computationally powerful methods and representations can be more easily acquired and used.

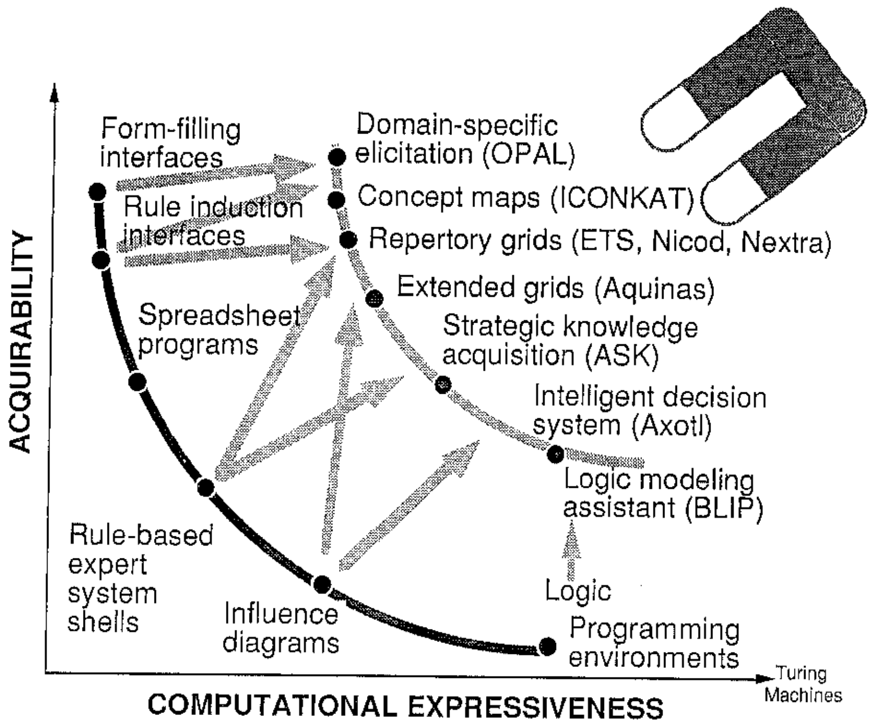

Figure 5. Knowledge acquisition tools can help make acquirable representations more powerful, and powerful representations more acquirable. (Adapted from Gruber.[81])

VI. SUMMARY

We have argued that knowledge acquisition is essentially a constructive modeling process, rather than simply a matter of "expertise transfer" or "knowledge capture." From this point of view, we have advocated the development of knowledge acquisition practices and tools that support active collaboration between experts and knowledge engineers in their efforts to cooperatively build useful domain models.

In Sec. II we suggested that an explicit theory of human cognitive processes, including perception, memory, representation, anticipation, and reasoning, can help us to simplify some of the central problems in this field. It also can provide an infrastructure upon which to build highly integrated hybrid knowledge acquisition tools in a principled way. Moreover, the same theory can also serve as a valuable reference and theoretical guide for the users of these tools. We proposed specifically that constructivist epistemology is a potentially rich set of ideas for those interested in developing computational models of human cognition and related systems for the assessment and representation of knowledge.

In Sec. III we hypothesized specifically that domain experts acquire their "expertise" not only from explicit knowledge of the sort found in textbooks (i.e., widely shared consensual beliefs), but also from their own private funds of personal experience essentially consisting of functional but fallible anticipations held with high confidence and uncertain validity, which, combined with their "book knowledge," make them expert practitioners. It would seem to follow that the greater their expertise, the further the experts' construct systems deviate from those of typical practitioners and the greater the importance of personally constructed knowledge. In short, experts have developed a socially situated expertise-in-context that in some important respects does not coincide with publicly available domain knowledge, rendering it extremely difficult for them to articulate their knowledge explicitly to either students and colleagues, or to knowledge engineers.

As previously noted, recent theoretical work in knowledge acquisition has emphasized that the creation of knowledge bases is a constructive modeling process, and not simply a matter of "expertise transfer" or "knowledge capture." In Sec. IV we discussed the modeling process, and described problems associated with premature expression of knowledge in a machine-oriented representation. A model-based description of the domain in a form that the user can intuitively understand has many advantages, the chief of which is that it can serve to mediate communication between developers, experts, and users of the system, helping all of them to understand and articulate the broader, higher-level problem context.

Ideally, knowledge acquisition tools should support the entire modeling life cycle, from initial conceptualization to eventual implementation of a knowledge-based system. Each phase of the life cycle, however, poses its own problems and has its own requirements. In Sec. V we noted that many of the problems associated with knowledge acquisition and maintenance stem directly from the inadequacies of the representations used at various stages in the development of knowledge-based systems. In order to surmount these problems, we have emphasized the deployment of mediating representations as a means of communication between expert and knowledge engineer; and intermediate representations to help bridge the gap between these mediating representations and a particular implementation formalism.

From this perspective, there are two intertwined critical tasks facing the knowledge acquisition research community. The first is to continue elaborating a variety of potential theoretical foundations for understanding the nature of expertise and the processes of modeling and representing knowledge. The second is to develop tools and techniques on the basis of these emerging theories.

We express our appreciation to Miroslav Benda, John Boose, Guy Boy, Kathleen Bradshaw, John Brennan, Larry Bunch, Alberto Cañas, Bill Clancey, John Coffey, Tom Cowin, Steve Dobbs, Jim Fulton, Brian Gaines, Pete and Cindy Holm, Earl Hunt, Oscar and Sharon Kipersztok, Cathy Kitto, Joe Koszarek, Tim Lethbridge, Allen Matsumoto, Louis Montoya, Joseph Novak, Thom Nguyen, Steve Poltrock, Bob Schneble, Doug Schuler, Kish Sharma, Mildred Shaw, Dave Shema, Doug Skuce, Howard Stahl, Tony White, Bruce Wilson, Jeff Yerkes, and Debra Zarley for their contributions and support.

References

1. N.M. Agnew and J.L. Brown, "Foundations for a theory of knowing: I. Construing reality," *Can. Psychol.* **30,** 152–167 (1989).
2. E.A. Feigenbaum and P. McCorduck, *The Fifth Generation,* Addison-Wesley, New York, 1983.
3. W.J. Clancey, "The frame of reference problem in cognitive modeling," in *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society,* Ann Arbor, Michigan, 1989.
4. K.M. Ford and J.R. Adams-Webber, "Knowledge acquisition and constructivist epistemology," in *The Psychology of Expertise: Cognitive Research and Empirical AI,* R.R. Hoffman, Ed., Springer-Verlag, New York, 1992, pp. 121–136.
5. J.M. Bradshaw and J.H. Boose, "Knowledge acquisition as CASE for knowledge-based systems," presented at the Third International Workshop on Computer-Aided Software Engineering (CASE-89), London, July, 1989.
6. W.J. Clancey, "Implications of the system-model-operator metaphor for knowledge acquisition," in *Knowledge Acquisition for Knowledge-Based Systems,* H. Motoda, R. Mizoguchi, J.H. Boose, and B.R. Gaines, Eds., IOS Press, Amsterdam, 1990.
7. L.A. Cox, "Knowledge acquisition for model building," *Int. J. Intell. Syst.* **8,** 91–103 (1993) (this issue). Also in *Knowledge Acquisition as Modeling,* K.M. Ford and J.M. Bradshaw, Eds., Wiley, New York, 1993, pp. 91–103.
8. T.R. Gruber, "Justification-based knowledge acquisition," in *Knowledge Acquisition for Knowledge-Based Systems,* H. Motoda, R. Mizoguchi, J.H. Boose, and B.R. Gaines, Eds., IOS Press, Amsterdam, 1990.
9. M.A. Musen, "Generation of Model-Based Knowledge Acquisition Tools for Clinical-Trial Advice Systems, Ph. D. dissertation, Computer Science Department, Stanford University, 1988; also available as Reports Nos. KSL 88-06 and STAN-CS-88-1194.
10. M.L.G. Shaw and J.B. Woodward, "Modeling expert knowledge," *Knowl. Acquis.* **2,** 179–206 (1990).
11. B.J. Wielinga, J.M. Akkermans, G. Schreiber, and J.R. Balder, "A knowledge acquisition perspective on knowledge-level models," in *Proceedings of the Fourth Knowledge Acquisition for Knowledge-Based Systems Workshop,* J.H. Boose and B.R. Gaines, Eds., SRDG Publications, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada T2N 1N4, 1989, pp. 36.1–36.22.
12. G.A. Kelly, *The Psychology of Personal Constructs,* Norton, New York, 1955.
13. J.H. Boose, "Personal construct theory and the transfer of human expertise," in *Proceedings of the National Conference on Artificial Intelligence,* Austin, Texas, 1984, AAAI, Menlo Park, CA, 1984, pp. 27–33.
14. J.H. Boose, *Expertise Transfer for Expert System Design,* Elsevier, New York, 1986.
15. J.H. Boose, "A knowledge acquisition program for expert systems based on Personal Construct Psychology," *Int. J. Man-Mach. Stud.* **23,** 495–525 (1985).
16. B.R. Gaines and M.L.G. Shaw, "Interactive elicitation of knowledge from experts," *Future Comput. Syst.* **1,** 151–190 (1986).
17. J.H. Boose and J.M. Bradshaw, "Expertise transfer and complex problems: Using Aquinas as a knowledge-acquisition workbench for knowledge-based systems," *Int. J. Man-Mach. Stud.* **26,** 3–28 (1987). Also in *Knowledge Acquisition Tools for Expert Systems,* J.H. Boose and B.R. Gaines, Eds., Academic, London, 1987, pp. 39–64.
18. J.H. Boose, J.M. Bradshaw, C.M. Kitto, and D.B. Shema, "From ETS to *Aquinas:* Six years of knowledge acquisition tool development," in *Proceedings of the Fourth Knowledge Acquisition for Knowledge-Based Systems Workshop,* J.H. Boose and B.R. Gaines, Eds., SRDG Publications, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada T2N 1N4, 1989, pp. 5.1–5.17.
19. C. Garg-Janardan and G. Salvendy, "A conceptual framework for knowledge elicita-

tion," in *First Knowledge Acquisition for Knowledge-Based Systems Workshop*, 1986, Part 3 [*Int. J. Man-Mach. Stud.* **26,** 521–532 (1987)]. Also in *Knowledge-Based Systems Vol. 1: Knowledge Acquisition for Knowledge-Based Systems*, J.H. Boose and B.R. Gaines, Eds., Academic, New York, 1988, pp. 119–130.
20. M.L.G. Shaw and B.R. Gaines, "Techniques for knowledge acquisition and transfer," *Int. J. Man-Mach. Stud.* **27,** 251–280 (1987).
21. J. Diederich, I. Ruhmann, and M. May, "KRITON: A knowledge acquisition tool for expert systems," *Int. J. Man-Mach. Stud.* **26,** 29–40 (1987).
22. J. Diederich, M. Linster, I. Ruhmann, and T. Uthmann, "A methodology for integrating knowledge acquisition techniques," in *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, T. Addis, J.H. Boose, and B.R. Gaines, Eds., University of Reading, Reading, England, 1987, pp. E4.1–E4.11.
23. B.R. Gaines, "An overview of knowledge acquisition and transfer," *Int. J. Man-Mach. Stud.* **26,** 453–472 (1987). Also in *Knowledge-Based Systems Vol. 1: Knowledge Acquisition for Knowledge-Based Systems*, J.H. Boose and B.R. Gaines, Eds., Academic, New York, 1988, pp. 3–22.
24. B.R. Gaines, "How do experts acquire expertise?" in *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, T. Addis, J.H. Boose, and B.R. Gaines, Eds., University of Reading, Reading, England, 1987, pp. B2.1–B2.11.
25. B.R. Gaines, "Semantics and Implementation of Open Architecture Knowledge Representation Servers." Department of Computer Science, University of Calgary, Alberta, Canada, 1990 (unpublished).
26. B.R. Gaines, "An Open-Architecture Knowledge Support System," Department of Computer Science, University of Calgary, Alberta, Canada, 1990 (unpublished).
27. K.M. Ford, "An Approach to the Automated Acquisition of Production Rules from Repertory Grid Data," Ph.D. dissertation, Tulane University, 1987.
28. K.M. Ford, F.E. Petry, J.R. Adams-Webber, and P.J. Chang, "An approach to knowledge acquisition based on the structure of personal construct systems," *IEEE Trans. Know. Data Eng.* **KDE-3,** 78–88 (1991).
29. K.M. Ford, H. Stahl, J.R. Adams-Webber, A.J. Cañas, J. Novak, and J.C. Jones, "ICONKAT: Integrated Constructivist Knowledge Acquisition Tool," *Knowl. Acquis.* **3,** 215–236 (1991).
30. K. Lewin, *A Dynamic Theory of Personality*, McGraw-Hill, New York, 1935.
31. G.A. Kelly, "A mathematical approach to psychology," in *Clinical Psychology and Personality: The Selected Papers of George Kelly*, B.A. Maher, Ed., Wiley, New York, 1969, pp. 7–45.
32. G.A. Kelly, "A brief introduction to personal construct theory," in *Perspectives in Personal Construct Theory*, D. Bannister, Ed., Academic, London, 1970, pp. 1–29.
33. J.R. Adams-Webber, *Personal Construct Theory: Concepts and Applications*, Wiley, New York, 1979.
34. J.R. Adams-Webber, "Kelly's pragmatic constructivism," *Can. Psychol.* **30,** 190–193 (1989).
35. J.C. Mancuso and J.R. Adams-Webber, Eds., "Anticipation as a constructive process," in *The Construing Person*, Praeger, New York, 1982, pp. 8–32.
36. K.M. Ford, "A constructivist view of the frame problem in artificial intelligence," *Can. Psychol.* **30,** 188–190 (1989).
37. J.C. Mancuso and B.N. Eimer, "Fitting things into sorts: The range corollary," in *The Construing Person*, J.C. Mancuso and J.R. Adams-Webber, Eds., Praeger, New York, 1982, pp. 130–151.
38. J.R. Adams-Webber, "Actual structure and potential chaos," in *Perspectives in Personal Construct Theory*, D. Bannister, Ed., Academic, London, 1970, pp. 30–45.
39. J. Dewey, "The reflex arc concept in psychology," *Psychol. Rev.* **3,** 357–370 (1896).

40. E. Heidbrieder, *Seven Psychologies*, Appleton-Century-Crofts, New York, 1933.
41. J.B. Watson, *Psychology from the Standpoint of a Behaviorist*, Lippincott, Philadelphia, 1919.
42. B.F. Skinner, *Verbal Behavior*, Appleton-Century-Crofts, New York, 1957.
43. J. Deese, "Behavior and fact," *Am. Psychol.* **24**, 515–522 (1969).
44. F.I.M. Craik, "Will cognitivism bury experimental psychology?" *Can. Psychol.* **32**, 440–444 (1991).
45. F. Bartlett, *Remembering: A Study in Experimental and Social Psychology*, Cambridge University Press, London, 1932.
46. H. Head, *Studies in Neurology*, Froede, Hodder, & Stoughton, London, 1920.
47. S.E. Asch, "A reformulation of the problem of association," *Am. Psychol.* **24**, 92–102 (1969).
48. U. Neisser, *Cognitive Psychology*, Appleton-Century-Crofts, New York, 1967.
49. U. Neisser, *Cognition and Reality: Principles and Implications of Cognitive Psychology*, Freeman, San Francisco, 1976.
50. J. Piaget, *The Psychology of Intelligence*, Harcourt Brace, New York, 1960.
51. J. Piaget, *Psychology and Epistemology*, Viking, New York, 1971.
52. J. Piaget, "The problem of common mechanisms in the human sciences," *Hum. Context* **1**, 163–185 (1969).
53. K.M. Ford and P.J. Chang, "An approach to automated knowledge acquisition founded on personal construct theory," in *Advances in Artificial Intelligence Research*, JAI Press, Greenwich, CT, 1989, Vol. 1, pp. 83–132.
54. N.M. Agnew and J.L. Brown, "Foundations for a theory of knowing: II. Fallible but functional knowledge," *Can. Psychol.* **30**, 168–183 (1989).
55. T. Mischel, "Personal constructs, rules, and the logic of clinical activity," *Psychol. Rev.* **71**, 180–192 (1964).
56. N. Warren, "Constructs, rules and the explanation of behavior," presented at Symposium on Construct Theory and Repertory Grid Methodology, Brunel University, England, 1964.
57. M. Husain, "To what can one apply a construct?" in *Applications of Personal Construct Psychology*, J.R. Adams-Webber and J.C. Mancuso, Eds., Academic, New York, 1983, pp. 11–28.
58. T. Mischel, Ed., "Scientific and philosophical psychology: A historical introduction," in *Human Action: Conceptual and Empirical Issues*, Academic, New York, 1969, pp. 1–40.
59. J.R. Adams-Webber and J.C. Mancusco, Eds., "The pragmatic logic of personal construct psychology," in *Applications of Personal Construct Psychology*, Academic, New York, 1983, pp. 1–10.
60. P.E. Johnson, "What kind of expert should a system be?" *J. Med. Philos.* **8**, 77–97 (1983).
61. D.A. Waterman, *A Guide to Expert Systems*, Addison-Wesley, Reading, MA, 1986.
62. W.J. Clancey and J. Roschell, "Situated cognition: How representations are created and given meaning," presented at the AERA Symposium of Cognitive Theories of How the Nervous System Functions for Research and Practice in Education, Chicago, April 1991.
63. K.M. Ford and N.M. Agnew, "Expertise: Socially situated and personally constructed," In *Working Notes, AAAI Spring Symposium Series; Cognitive Aspects of Knowledge Acquisition*, Stanford University, March 1992, pp. 80–87.
64. W.J. Clancey, "Interactive control structures: Evidence for a compositional neural architecture," presented at the NATO Workshop on Emergence, Situatedness, Subsumption, and Symbol Grounding, Brussels, Belgium, April 1991.
65. A. Kaplan, *The Conduct of Inquiry*, Harper and Row, New York, 1963.
66. E.A. Moore and A.M. Agogino, "INFORM: An architecture for expert-directed knowledge acquisition," *Int. J. Man-Mach. Stud.* **26**, 213–230 (1987).

67. N.E. Johnson, "Mediating representations in knowledge elicitation," in *Knowledge Elicitation: Principles, Techniques and Applications,* D. Diaper, Ed., Wiley, New York, 1989.
68. J.M. Bradshaw, K.M. Ford, J.R. Adams-Webber, and J.H. Boose, "New approaches to constructivist knowledge acquisition tool development," *Int. J. Intell. Syst.,* **8**(2) (1993).
69. M.L.G. Shaw and J.B. Woodward, "Mental models in the knowledge acquisition process," in *Proceedings of the Fourth Knowledge Acquisition for Knowledge-Based Systems Workshop,* J.H. Boose and B.R. Gaines, Eds., SRDG Publications, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada T2N 1N4, 1989, pp. 29.1–29.24.
70. J.D. Novak and D.B. Gowin, *Learning How to Learn,* Cambridge University Press, New York, 1984.
71. P.H. Winston, *Artificial Intelligence,* second ed., Addison-Wesley, Reading, MA, 1984.
72. E. Rich and K. Knight, *Artificial Intelligence,* second ed., McGraw Hill, New York, 1991.
73. G. Schreiber, H. Akkermans, and B. Wielinga, "On problems with the knowledge level perspective," in *Proceedings of the Fifth Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems,* SRDG Publications, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, 1990, pp. 30.1–14.
74. D. Diaper, Ed., "Designing expert systems—from Dan to Beersheba," in *Knowledge Elicitation: Principles, Techniques and Applications,* Wiley, New York, 1989.
75. J.H. Larkin and H.A. Simon, "Why a diagram is (sometimes) worth ten thousand words," *Cogn. Sci.* **11,** 65–99 (1987).
76. D.A. Norman, *The Psychology of Everyday Things,* Basic Books, New York, 1988.
77. C. Sivard, M. Zweben, D. Cannon, F. Lakin, and L. Leifer, "Conservation of design knowledge," in *Proceedings of the 27th Aerospace Sciences Meeting of the American Institute of Aeronautics and Astronautics,* AIAA, New York, 1989.
78. B.R. Gaines and M.L.G. Shaw, "Induction of inference rules for expert systems," *Fuzzy Sets Syst.* **8,** 315–328 (1986).
79. D.B. Shema and J.H. Boose, "Refining problem-solving knowledge in repertory grids using a consultation mechanism," *Int. J. Man-Mach. Stud.* **29,** 447–460 (1988).
80. M. Linster and B.R. Gaines, "Supporting acquisition and performance in a hypermedia environment," presented at Terminology and Knowledge Engineering Workshop, Oct. 1990.
81. T.R. Gruber, *The Acquisition of Strategic Knowledge,* Academic, New York, 1989.
82. D.E. Webster, "Mapping the design information representation terrain," *Computer,* Dec., 8–23 (1988).
83. C. Rich and R.C. Waters, "Artificial intelligence and software engineering," in *AI in the 1980s and Beyond: An MIT Survey,* W.E.L. Grimson and R.S. Patil, Eds., MIT Press, Cambridge, MA, 1987.