Proceedings of the IASTED International Conference on Expert Systems, Artificial Intelligence, and Neural Networks, pp. 44-47, 1996

Concerning a General Framework for the Development of Intelligent Systems

Dan Ventura
Tony R. Martinez

Computer Science Department, Brigham Young University, Provo, Utah 84602
e-mail: dan@axon.cs.byu.edu, martinez@cs.byu.edu

Key Words: Cognition, Hybrid Systems, Connectionism, Symbolism

Abstract

There exists on-going debate between Connectionism and Symbolism as to the nature of and approaches to cognition. Many viewpoints exist and various issues seen as important have been raised. This paper suggests that a combination of these methodologies will lead to a better overall model. The paper reviews and assimilates the opinions and viewpoints of these diverse fields and provides a cohesive list of issues thought to be critical to the modeling of intelligence. Further, this list results in a framework for the development of a general, unified theory of cognition.

1. Introduction

Developing artificial models of cognition is a challenging field that has roots in such diverse fields as cognitive psychology, computer science, linguistics, neurophysiology, and mathematics. Due at least in part to this fact, vastly differing theories for explaining and/or imitating cognitive capabilities have been developed including Connectionism [1] [2], Symbolic Artificial Intelligence [3], Fuzzy Logic [4], Machine Learning [5] and Genetic Algorithms [6], to name a few. Each of these approaches to modeling cognition (or, more accurately, some aspect thereof) has various strengths and weaknesses that are becoming better understood, and each claims some fairly impressive successes as well as some disheartening failures.

For the most part, relations among the various paradigms have been indifferent at best and often have been almost hostile. Though this is true of all the approaches mentioned above to some extent, the majority of the contrasts, comparisons, and conflicts are embodied in the Symbolism vs. Connectionism debate that has raged continuously over the last ten years.

However, some researchers have begun to recognize the apparent complimentary nature of these two paradigms. They believe that neither approach alone can produce adequate models of cognition but that in combination richer models with greater functionality can be realized. Efforts in this direction have variously been referred to as Hybrid Systems, High-Level Connectionism, Symbolic Connectionism, Symbol Processing Connectionist Systems and the like. Though work in this area has been on-going since the resurgence of Connectionism in the late eighties [7] [8], a proliferation of new books on the subject ([9] [10] [11] [12], for example) indicates its increasing importance.

To date, the field of Hybrid Systems has produced some interesting applications [9] [10] [11] [12] [13], and some ground work has been laid concerning the proper integration of various approaches to modeling cognition; however, the critical next step must be the development of a formal theory for the integration of Connectionism with Symbolism and thus for a general theory of (artificial) intelligence. A promising approach, in our opinion, is to view Connectionism and Symbolism as complimentary approaches and to attempt to extract from each the primitives that embody them. Another way of looking at this is to realize that both Connectionism and Symbolism can be modeled by a Turing machine, and therefore our job is to identify a simple subset of Turing-computable functions that embody the primitives necessary for a model of intelligence.

This paper is an attempt to outline a framework for the development of such a general theory. The approach we take is to first review the differences and similarities, and the strengths and weaknesses of Connectionism and Symbolism. Second, a list of critical issues pertaining to intelligence is discussed. Since everyone has a differing opinion of what the critical issues are, we are attempting to distill the many ideas that *could* be considered into the essence of what *must* be considered. Third, a brief discussion of how these issues can guide in the development of a general theory of intelligence is presented.

2. Connectionism and Symbolism

Whatever their differences may be, Connectionism and Symbolism both share the assumption that at some level cognition can be functionally modeled as a computational process. Other than this extremely basic commonality, there is little else involving the issue of Connectionism and Symbolism that enjoys any kind of majority consent. Some researchers argue that Connectionism is fatally flawed as a model of cognition [14] or that while Connectionism may play a minor though important role, Symbolism represents the heart of any realistic model [15]. On the other hand, cases are presented for the desirability of Connectionism over Symbolism [16] and that the eventual dominance of Connectionism is only a matter of time [17]. Some see the two as complimentary or dualistic in nature and are optimistic that their combination will bear productive fruit [18] [19] [20] [21] [22], while others argue for their equivalence and cite performance issues as the main concern [23] [10]. Still others are more pessimistic as how intelligent a Connectionist-Symbolic hybrid will ever be [24], and some claim that the entire idea of modeling cognition as computation is hopeless [25].

Many people agree that cognition can and should be described at different levels and that care should be taken to compare and contrast only those explanations/models/theories that exist on the same level [23] [14] [16]. However, others feel that an explanation of cognition should not or can not be completely differentiated from the biological implementation of cognition that we know as the brain [22][24] [17].

Tables 1 and 2 attempt to assimilate these views from two different angles. One common thread that runs through any discussion of Connectionism vs. Symbolism, either implicitly or explicitly is the dual nature of the two. People refer to this duality in different ways, and table 1 presents the computational dichotomy that these two approaches represent. For example, Connectionist networks are characterized by being spatially distributed while Symbolic approaches can be thought of as being temporally distributed. In related fashion, Connectionist networks perform processing in parallel over distributed representations but Symbolic models process serially over local ones. Connectionist models perform statistical approximations over continuous numerical data, and Symbolic models apply concrete logical operations to discrete symbolic data. Finally, on one hand, Connectionism is associated with cognitive tasks that are referred to as low-level, subconceptual, unconscious, or subdeliberative; on the other, Symbolism is associated with tasks that may be described as high-level, conceptual, conscious, or deliberative.

These are all prototypical generalizations, of course, and as is pointed out in, for example [21], the boundary dividing Connectionist models and Symbolic models is ill-defined at best. In fact Oden suggests that models that should be taken seriously will almost always exist in this fuzzy boundary region.

Connectionism vs. Symbolism	
Space	Time
Parallel	Serial
Distributed	Localist
Statistical	Logical
Continuous	Discrete
Numerical	Symbolic
Low-level	High-level
Subconceptual	Conceptual
Unconscious	Conscious
Subdeliberativc	Deliberative

Table 1. The Connectionist-Symbolic Dichotomy

The second approach to looking at this is to examine the relative strengths and weaknesses of the two approaches. People see this particular debate in many different lights. There is somewhat of a consensus as to what the respective strengths and weaknesses are. However much disagreement exists over their relative importance. Table 2 is a compilation of general strengths and weaknesses for Connectionism and Symbolism. Again, this is an attempt to assimilate as many viewpoints as possible, though doubtless some will be excluded or treated unsatisfactorily.

Attributes of Connectionist networks usually cited as strong points include *robustness*, the ability to *fine tune* knowledge through experience, automatic *acquisition of knowledge*, *fault tolerance*, and *adaptivity*. Robustness is the ability to respond appropriately to inputs that are noisy, novel, or unanticipated. Fine tuning of knowledge expresses the idea that the system can slightly alter its representations to account for new inputs without drastically affecting the representations gained from past experiences. Automatic knowledge acquisition is the ability to extract knowledge from the environment. Fault tolerance allows for graceful degradation of performance as the system fails. Finally, adaptivity deals both with a wide range of applicability of the model as well as with the ability of a given instantiation of the model to change with its environment.

Connectionism	Symbolism
Strengths	
robustness	handle complex structures
fine tuning	domain
knowledge acquisition	explanation
fault tolerant	serial computation
adaptive	
inherent parallelism	
Weaknesses	
difficult to interpret	brittleness
slow learning	symbol grounding
homogeneous stucture	knowledge acquisition
not naturally serial	narrow domain
scaling	scaling

Table 2. Strengths and Weaknesses of Connectionism and Symbolism

Symbolism, on the other hand, is usually considered to have as strengths the ability to represent and manipulate *complex structures*, *rapid learning* of concepts, *explanation*, and the ability to naturally perform *serial* computation. Complex structures are defined recursively as a set of atomic symbols together with those molecular symbols that can be created by combining atomic symbols according to a set of syntactic rules. These structures possess a semantic structure that is closely related to their syntactic one. Rapid learning refers to the ability to quickly and drastically change representations. Explanation refers to the ability of a system to explain its decision in understandable terms. Some people hold that the explanation of a system's behavior (not the behavior itself) is the real object of interest. Serial computation is argued to be necessary for some kinds of cognition including ordered problem solving, logical reasoning, generalized thinking involving variable binding, and certain types of structural aspects of language [20].

Not surprisingly, the weaknesses of the respective approaches often seem to be the negation of a strength of the opposing method. For example, the difficulty in interpretation of Connectionist networks is in direct opposition to the strong explanation abilities of Symbolism. Likewise, Connectionism's slow learning rate, homogeneity of structure and lack of natural serial computation ability starkly contrast with Symbolism's strengths of use of domain knowledge, complex structure, and natural serial computation respectively. Similarly, Symbolism has its own weaknesses: brittleness, symbol grounding (or lack thereof), difficulty in acquiring knowledge, and narrow application domain (each system must be tailored to a specific domain). Brittleness results from a lack of robustness and fault tolerance, difficulty in acquiring knowledge contrasts with automatic knowledge acquisition, and a narrowness of domain results from lack of adaptivity. But robustness, fault tolerance, automatic knowledge acquisition, and adaptivity are exactly the strengths of Connectionism. The symbol grounding problem is interesting in that it does not directly oppose a strength of Connectionism, per se, (though Connectionism does not suffer from it and this may be thought of as a strength) and is treated in [15]. There are those, of course, who will point to one or another of one of the methods' weaknesses and argue that it is insurmountable and thus that the method is untenable. A good example of this is Fodor and Pylyshyn's denunciation of Connectionism on the basis that it lacks compositionality and systematicity (basically, the ability to represent complex structures and the existence of structure sensitive operations, respectively) [14].

3. Issues of Intelligence

Accepting the premise of Artificial Intelligence that cognition, at some level, can be modeled as computation and also taking the viewpoint that Connectionism and Symbolism represent a complementary dichotomy of computation, both poles of which explain certain aspects of intelligence, we would like to employ these assumptions to develop a framework for the development of a general theory of (artificial) intelligence. This framework will explicitly describe what such a theory must do in order to be a candidate for the explanation of cognition.

3.1 Formal primitives

A minimal set of formal primitives (a subset of the Turing computable functions) that are necessary and sufficient for modeling the building blocks of intelligence must be developed. Minimality is important for reasons of parsimony and elegance. Formality is a necessity for clarity, for testing hypotheses, and for implementation issues, among other things. Further, these primitives must adequately address the following issues: knowledge, complex symbol structure, symbol grounding, learning, robustness, fault tolerance, adaptivity, duality of intelligence, multiple levels of explanation, and scaling.

3.2 Knowledge

Possessing knowledge and rationally making use of it is a major component of intelligence and there are at least three facets to the knowledge problem: acquisition, representation, and explanation. Any cognitive system must be able to obtain information from its environment, internalize and manipulate that information, and respond appropriately to its environment.

3.2.1 Acquisition. A model of intelligence must provide a method for actively extracting pertinent information from the environment as well as allow for passive assimilation of information provided explicitly by an authority.

3.2.2 Representation. The model must internally represent knowledge in a way that is both useful to it and in a way that exhibits some kind of relationship with the external environment from which it was derived. This obviously relates to the issue of compositionality and to the issue of grounding.

3.2.3 Explanation. If for no other reason than for the practical one of trust, it is necessary that any intelligent system be capable of providing an explanation of its decisions. This is especially true of an artificial intelligence acting in any type of critical system (such as air traffic control or medicine). Some would claim that the explanation of a system's behavior is more important and more valuable (especially in the context of understanding intelligent behavior) than the behavior itself.

3.3 Complex symbols

For the staunch Classicists, the ability (or lack thereof) to represent and manipulate complex symbols is the crux of the matter. They believe that humans use symbol representations and manipulations in so many high-level cognitive tasks that this an absolute necessity. Connectionists are not convinced of this. It may certainly be argued that language syntax and semantics are based on this ability, and in fact, this ability is derived, for the most part, from a linguistic theory of mind. This concept of a complex symbolic representational system is so pervasive that any general model must either include provision for it (or a functional equivalent) or convincingly demonstrate it to be unnecessary.

3.3.1 Compositionality. This principle demands the ability to represent a theoretically infinite number of symbols (this is known as productivity) using a finite set of atomic symbols and a set of syntactic rules for their combination. This theoretical infinitude is limited in practice by finite resource constraints such as memory, time, etc. The important thing about these symbols is that their semantic interpretation is a semantic interpretation of their constituent parts along with the relationship that exists among these parts.

3.3.2 Systematicity. Systematicity complements compositionality in that it requires the existence of operations that are sensitive to syntactic structure. For example, in deductive logic a rule exists for transforming any structure of the form A $\wedge$ B to the form A. Because semantic structure is closely related to syntactic structure (by compositionality), this reflects the truism that if A and B exist (are true), then certainly A exists (is true).

3.4 Grounding

This consideration is due to Harnad [15] and addresses the problem in pure symbol systems of the lack of *intrinsic* meaning in the system. Though such a system possesses a rich symbol structure, that structure is purely internal and completely arbitrary. It is not grounded to anything in reality. Harnad likens this problem to that of trying to learn Chinese as a first language with only Chinese literature as an instruction source. All the Chinese symbols are related to one another in complex ways, but without so much as a notion of language (remember we are learning this as a *first* language), there is no way to connect any of those symbols to anything real. Symbolic representation is practically worthless to a system trying to function in an environment in which none of its symbols are grounded. Harnad suggests that the solution to this quandary is to begin with an elementary set of nonsymbolic representations such as various input patterns that can be both discriminated and identified (classified) correctly. The classifications of these patterns will then be abstract symbols, but they will be grounded to the input patterns from the environment that produced them. These can then be further combined into more complex symbols as discussed above, and the grounded meanings of the atomic symbols will be inherited by any higher level symbols of which they are constituents (by the principle of compositionality).

3.5 Learning

Knowledge acquisition is one method for providing a system with its initial knowledge; other approaches such as inductive methods also may be employed. The ability to alter that knowledge and accordingly to alter behavior due to changes in the environment is what allows a system to intelligently function in a dynamic, realistic environment. This is the ability to learn. Any intelligent system should, over time, become more adept at performing tasks (not just in the sense of speed-up learning) and interacting with its environment through this process. Both the rapid assimilation of concepts necessary for immediately critical behavioral changes and the gradual refining of representations involved in capturing exceptional behaviors and concept drift must be supported.

3.6 Robustness, fault tolerance, and adaptivity

The arguments for these capabilities are almost platitudes. Independent of method it is obvious that the abilities to perform in noisy, novel or adverse conditions, to maintain reasonable levels of performance despite failures in the system (graceful degradation), and to adapt to a changing environment and/or unique types of problems are clearly desirable properties that any intelligence or model of intelligence will possess.

3.7 Duality

As discussed in section 2, there appears to be a dual nature to intelligence. It is multi-faceted and each of these has diametrically opposed poles that are both present (or at least seem to be) in intelligence. The most important of these dualities is briefly discussed below.

3.7.1 Continuous and discrete. Most of the real world is a continuum. Therefore it is critical that any model of intelligent behavior account for handling continuous data. On the other hand, we often process in discrete quantities as well, content to know that it is a cool day rather than that it is 63.7° F. Discretization, though far from satisfactory, is one approach to reconciling the two [26].

3.7.2 Numerical and symbolic. The argument here is related to the previous one. Clearly, we explicitly and consciously perform symbolic computation at least in language. Perhaps we perform some numerical computation in mathematical endeavors, though mathematics is by its very nature symbolic. More probably any numerical processing we perform occurs at an unconscious level, and there are numerous unconscious tasks that we perform (particularly of the pattern recognition sort) that can, so far, only be explained numerically.

3.7.3 Distributed and localist. A distributed system represents concepts as patterns over its units. A localist system represents a single concept with a single unit. Therefore a distributed system will be able to represent many more patterns than a localist one with similar resources. However, a fully distributed system typically can represent only a single concept at a time because all the units of the system are involved in the representation. On the other hand, a localist system can simultaneously represent as many concepts as it has units. The requirement here is to achieve the simultaneity of a localist system and the capacity of a distributed system. Other issues that could be addressed include those involving fault tolerance, performance issues, etc.

3.7.4 Parallel and Serial. Many processing tasks such as pattern matching, image recognition, memory recall, and natural language understanding require huge amounts of processing and yet they are performed by biological systems in a few hundred milliseconds. This suggests massive amounts of parallel computation. However, other tasks such as logic inference, problem solving, and the like have a very serial nature to them in which one computation cannot be accomplished until after another is complete. How do we reconcile two such computational approaches in the same model?

3.7.5 Statistics and logic. In a related vein, it is often necessary to perform a nearest match or a closest approximation or a degree of membership type operation. Certainly any completely novel experience will require such improvisation for the production of behavior. At other times, strict all or nothing matching or yes/no reasoning is required. Both the “soft” constraints of statistics derived from the environment (as per Connectionism) and the “hard” constraints of logic (as per Symbolism) are required.

3.8 Multiple levels

The duality issues of section 3.7 all suggest that an explanation of multiple levels of cognition is required. In every case one half of the dual was associated with high-level or conscious tasks such as reasoning while the other half was associated with low-level or unconscious tasks such as perception. It is well accepted that there exist multiple levels of *explanation* for cognitive phenomena and that there are at least the two levels (high-level and low-level processes) of *extant* cognitive phenomena. Therefore, any model of intelligence may exist at or take into account multiple levels of description; however, a model of intelligence *must* account for at least the high and low levels implicated in the computational dichotomy of intelligence.

3.9 Scaling

Neither Connectionism nor Symbolism have really succeeded in scaling their proposed explanations of cognition to real-world levels. Not only is a theory that only works on toy problems of no practical use, it is soon of very little use intellectually as well, for no theory that cannot account for the handling of real-world problems can come close to modeling intelligence. We must eventually succeed in scaling solutions up to realistic problems or finally abandon our efforts to explain cognition.

3.10 What and how

Are the brain and the mind inextricably associated one with another or is the functionality of mind completely separable from the implementation of the brain? There are those that will argue both ways. Among those that hold to the separability view, some will argue for a top-down approach in which we must understand *what* is happening before we can understand *how* it is happening, while others will argue the opposite way. We are of the opinion that mind and brain cannot be wholly separated and thus some attempt must be made to understand both at once. However, whatever the viewpoint, both the *how* and the *what* must eventually be addressed. The difference comes in the order they are tackled.

4. Toward a General Theory

As the issues above are confronted and resolved, a general, unified theory of (artificial) intelligence will emerge, and some good preliminary efforts have been made. For example, both Arbib's Schema Theory [27] [28] and Michalski's Multistrategy Task-adaptive Learning [29] give good accounts of many of the aspects of knowledge, complex symbols and learning. Though it is possible that the emergence of a theory will come as a revolutionary breakthrough, it is more likely that it will evolve much more slowly as an eclectic result of such research performed in many different disciplines. For this to have a chance to succeed, not only must argumentation and hostility between various camps be avoided, but also interdisciplinary cooperation should be fostered to a greater extent. For this to succeed, as McKenna points out, "we must resist the temptation to critique one discipline by the specialized criteria of another" [17]. It should be reiterated that although the theory may evolve in an eclectic manner, the final product cannot be a hodge-podge of techniques, partial models and explanations. It must be cohesive, formalized, and elegant and it must address the issues discussed here whether by subsumption or refutation. Though the issues presented may not be sufficient for a general theory of cognition they are surely necessary.

References

[1] Rumelhart, David E., McClelland, James L. and the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, MIT Press, 1986.

[2] Wasserman, Philip D., *Advanced Methods in Neural Computing*, Van Nostrand Reinhold, 1993.

[3] Russel, Stuart and Norvig, Peter, *Artificial Intelligence: A Modern Approach*, Prentice-Hall, 1995.

[4] Zadeh, L. A., "Fuzzy Logic", *IEEE Computer*, 83-93, 1988.

[5] Shavlik, Jude W. and Diettrich, Thomas G. (eds.), *Readings in Machine Learning*, Morgan Kaufman, 1990.

[6] Goldberg, David E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, 1989.

[7] Barndan, John A. and Pollack, J. B. (eds.), *Advances in Connectionist and Neural Computation Theory 1, High-level Connectionist Models*, Ablex, 1991.

[8] Hinton, G. E. (ed.), *Connectionist Symbol Processing*, MIT Press, 1991.

[9] Goonatilake, Suran and Khebbal, Sukhdev (eds.), *Intelligent Hybrid Systems*, John Wiley & Sons, 1995.

[10] Honavar, Vasant, and Uhr, Leonard (eds.), *Artificial Intelligence and Neural Networks: Steps Toward Principled Integration*, Academic Press, 1994.

[11] Levine, Daniel S. and Aparicio, Manuel (eds.), *Neural Networks for Knowledge Representation and Inference*, Lawrence Erlbaum Associates, 1994.

[12] Sun, Ron and Bookman, Lawrence (eds.), *Computational Architectures Integrating Neural and Symbolic Processes: A Perspective on the State of the Art*, Kluwer Academic, 1995.

[13] Shibata, Takanori and Fukuda, Toshio, "Hierarchical Intelligent Control for Robotic Motion", *IEEE Transactions on Neural Networks*, vol. 5, no. 5, 823-32, 1994.

[14] Fodor, J. A. and Pylyshyn, Z.W., "Connectionism and Cognitive Architecture: A Critical Analysis", *Cognition*, vol 28, 3-71, 1988.

[15] Harnad, Stevan, "The Symbol Grounding Problem", *Physica D*, vol. 42, 335-46, 1990.

[16] Barnden, John A., "On Using Analogy to Reconcile Connections and Symbols", *Neural Networks for Knowledge Representation and Inference*, eds. Daniel S. Levine and Manuel Aparicio, Lawrence Erlbaum, 27-64, 1994.

[17] McKenna, Thomas M., "The Role of Interdisciplinary Research Involving Neuroscience in the Development of Intelligent Systems", *Artificial Intelligence and Neural Networks: Steps Toward Principled Integration*, eds. Vasant Honavar and Leonard Uhr, Academic Press, 75-92, 1994.

[18] Goonatilake, Suran and Khebbal, Sukhdev, "Intelligent Hybrid Systems: Issues, Classifications and Future Directions", *Intelligent Hybrid Systems*, eds. Suran Goontalilake and Sukhdev Khebbal, John Wiley & Sons, 1-20, 1995.

[19] Sun, Ron, "An Introduction: On Symbolic Processing in Neural Networks", *Computational Architectures Integrating Neural and Symbolic Processes: A Perspective on the State of the Art*, eds. Ron Sun and Lawrence A. Bookman, Kluwer Academic, 1-18, 1995.

[20] Boden, Margaret, "Horses of a Different Colour?", *Aritificial Intelligence and Neural Networks: Steps Toward Principled Integration*, eds. Vasant Honavar and Leonard Uhr, Academic Press, 3-19, 1994.

[21] Oden, Gregg C., "Why the Difference between Connectionism and Anything Else Is More Than You Might Think but Less Than You Might Hope", *Artificial Intelligence and Neural Networks: Steps Toward Principled Integration*, eds. Vasant Honavar and Leonard Uhr, Academic Press Inc., 93-103, 1994.

[22] Aparicio, Manuel and Levine, Daniel S., "Why are Neural Networks Relevant to Higher Cognitive Function?", *Neural Networks for Knowledge Representation and Inference*, eds. Daniel S. Levine and Manuel Aparicio, Lawrence Erlbaum, 1-26, 1994.

[23] Honavar, Vasant, "Symbolic Artificial Intelligence and Numeric Artificial Neural Networks: Towards a Resolution of the Dichotomy", *Computational Architectures Integrating Neural and Symbolic Processes: A Perspective on the State of the Art*, eds. Ron Sun and Lawrence A. Bookman, Kluwer Academic, 351-88,1995.

[24] Chandrasekaran, B. and Josephson, Susan G., "Architecture of Intelligence: The Problems and Current Approaches to Solutions", *Artificial Intelligence and Neural Networks: Steps Toward Principled Integration*, eds. Vasant Honavar and Leonard Uhr, Academic Press, 21-50,1994.

[25] Penrose, Roger, *Shadows of the Mind*, Oxford University Press, 1994.

[26] Ventura, Dan and Martinez, Tony, "An Empirical Comparison of Discretization Methods", *Proceedings of the Tenth International Symposium on Computer and Information Sciences*, 443-50, 1995.

[27] Arbib, Michael, "Schema Theory: Cooperative Computation for Brain Theory and Distributed AI", *Aritificial Intelligence and Neural Networks: Steps Toward Principled Integration*, eds. Vasant Honavar and Leonard Uhr, Academic Press, 51-74, 1994.

[28] Arbib, M. A., "Schema Theory", *The Encyclopedia of Artificial Intelligence*, 2nd edition, ed. S. Shapiro, Wiley-Interscience, 1427-43, 1992.

[29] Michalski, Ryszard S., "Toward a Unified Theory of Learning: Multistrategy Task-adaptive Learning", *Readings in Knowledge Acquisition and Learning*, eds. B. G. Buchanan and D. C. Wilkins, Morgan Kaufman, 1993.