KNOWLEDGE ACQUISITION: THE CONTINUUM LINKING MACHINE LEARNING AND EXPERTISE TRANSFER

Brian R Gaines

Knowledge Science Institute
Department of Computer Science, University of Calgary
Calgary, Alberta, Canada T2N 1N4
(gaines@cpsc.ucalgary.ca, (403) 220-5901)

ABSTRACT

There is currently a major paradigm split in knowledge acquisition research and practice between techniques for the transfer of existing knowledge from human experts through "expertise transfer" and those for the creation of new expertise through empirical induction or "machine learning". There is, however, a fundamental relation between the two paradigms in that existing expertise was at some time derived through empirical induction. There is also continuity between the two paradigms in that existing expertise may be partial, erroneous, and of various forms, such that it cannot completely replace empirical induction but may serve to to guide and expedite it. This paper provides a quantitative framework for knowledge acquisition that encompasses machine learning and expertise transfer as related paradigms within a single spectrum, shows that the development of practical knowledge-based systems generally lies between these extremes, and suggests system architectures that combine the underlying paradigms.

INTRODUCTION

Research and practice in advanced information systems has developed major technologies for "knowledge acquisition" without very strongly founded concepts of the nature of "knowledge" itself. In the economic literature the notion that "knowledge" in general is an abstract term for a concrete economic entity similar in status to other "goods" was extensively developed by Machlup (1980, 1982, 1984), and has found increasing acceptance in both professional (Machlup & Mansfield 1983) and popular literature (Sveiby & Lloyd 1987). It underlies the powerful metaphor that Hayes-Roth (1984) used to explain the significance of expert systems in the early 1980s, that:

> "The power is in the knowledge. Mining, molding, assembling and refining expert knowledge are the main problems."

There are strong analogies between the classical resource management and extraction industries and the new knowledge management and extraction industries which serve to guide us in developing knowledge-based industries.

There are also differences which we do not yet fully comprehend and which may mislead us if we stretch the analogies too far. One of these is the essential human role in knowledge production and application as a by-product of our social existence. The knowledge sources that we mine are already partially processed by prior human activity and will continue to be further processed in use. There are parallels in classical resource extraction and manufacturing, but the ongoing dynamics of knowledge processing dominate knowledge-based industries in such a way as to typify them. Knowledge is an essentially dynamic resource which requires active management to maintain its existence.

This paper focuses on the analogy underlying Hayes Roth's metaphor, attempting to develop a framework for knowledge acquisition that encompasses the partial processing and dynamic aspects of knowledge resources. In particular, the framework developed encompasses knowledge acquisition through both machine learning and expertise transfer.

EXPERTISE TRANSFER AND EMPIRICAL INDUCTION

There is currently a major paradigm split in knowledge acquisition research and practice between techniques for the transfer of existing knowledge from human experts and those for the creation of new expertise through empirical induction or "machine learning" (Gaines & Boose 1988, Boose & Gaines 1988). There is, however, a fundamental relation between the two paradigms in that existing expertise was at some time derived through empirical induction. There is also continuity between the two paradigms in that existing expertise may be partial, erroneous, and of various forms, such that it cannot completely replace empirical induction but may serve to to guide and expedite it (Gaines 1987a).

Figure 1 makes these relations explicit:

At the top, some "world" acts as a source from which data may be acquired.

On the left the acquisition is by a person who models the data and becomes an "expert" about some aspect of the world. In developing a knowledge-based system such experts may be interviewed using expertise transfer tools (Boose 1986) and their "knowledge" (model of part the world) may be transferred to a computer-based knowledge base.

On the right the acquisition loads a computer-based database. In developing a knowledge-based system such a database may be modeled using empirical induction tools (Quinlan 1986) and the resultant "knowledge" (model of part the world) may be transferred to a computer-based knowledge base.

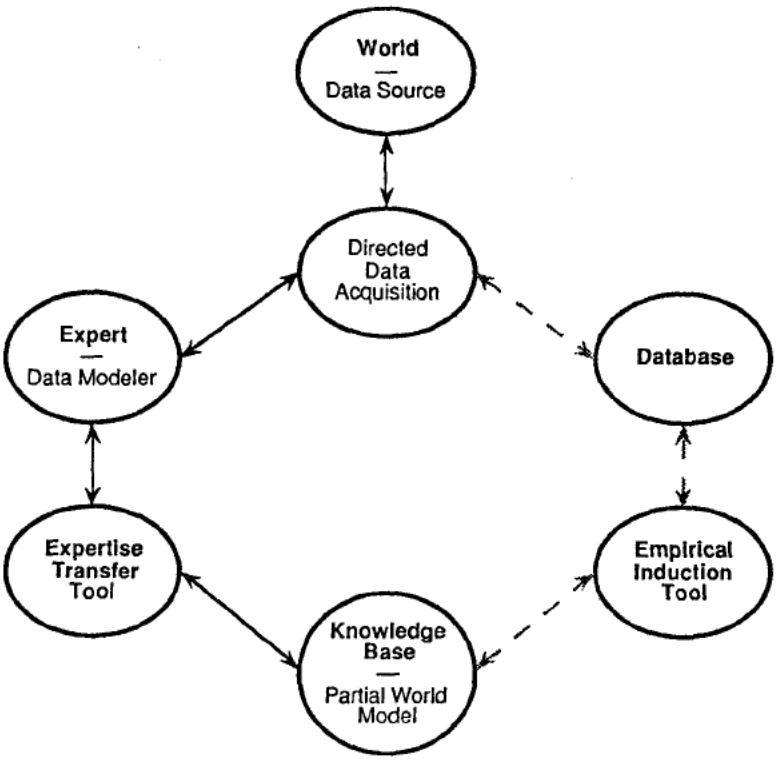

Figure 1 Expertise transfer and empirical induction

Figure 1 neatly captures the relation between expertise transfer and empirical induction. However, it over-emphasizes the separation between the two approaches:

On the left, the "expert" rarely makes available pure knowledge, that is, a minimal correct model of part of the world. Experts provide a mixture of relevant and irrelevant viewpoints, correct and incorrect modeling rules, case histories with partially known outcomes, and so on. The current generation of "expertise transfer" and "knowledge support" systems, such as Aquinas (Boose & Bradshaw 1987) and KSSO (Gaines & Shaw 1987) provide a wide range of tools to cope with, integrate, and clean up, this variety of forms and qualities of knowledge.

On the right, the "database" is not a complete or unselective collection of information about the world. It has been structured and collected by people for a purpose, and this structuring and collection already involves substantial expertise. Figure 1 emphasizes that the data used in both expertise and database formation is obtained by *directed* data acquisition. The provision of that direction is itself a task for experts.

Thus, an expert provides a mixture of knowledge, data, and garbage (incorrect or irrelevant knowledge and data), and so does a database. There is much in common between both the left- and right-hand paths in Figure 1—both expertise transfer and empirical induction involve the extraction of knowledge from partial mixes of knowledge, data, and garbage.

The following section provides a qualitative analysis of the various qualities of information between knowledge and data, and later sections provide quantitative measures of the resultant trade-off.

LEVELS OF EXPERTISE

Consider a practical knowledge-based system development in which both human experts and various sources of case histories are available, for example, Quinlan's development of thyroid diagnosis rules (Quinlan, Compton, Horn & Lazarus 1987). The experts might know:

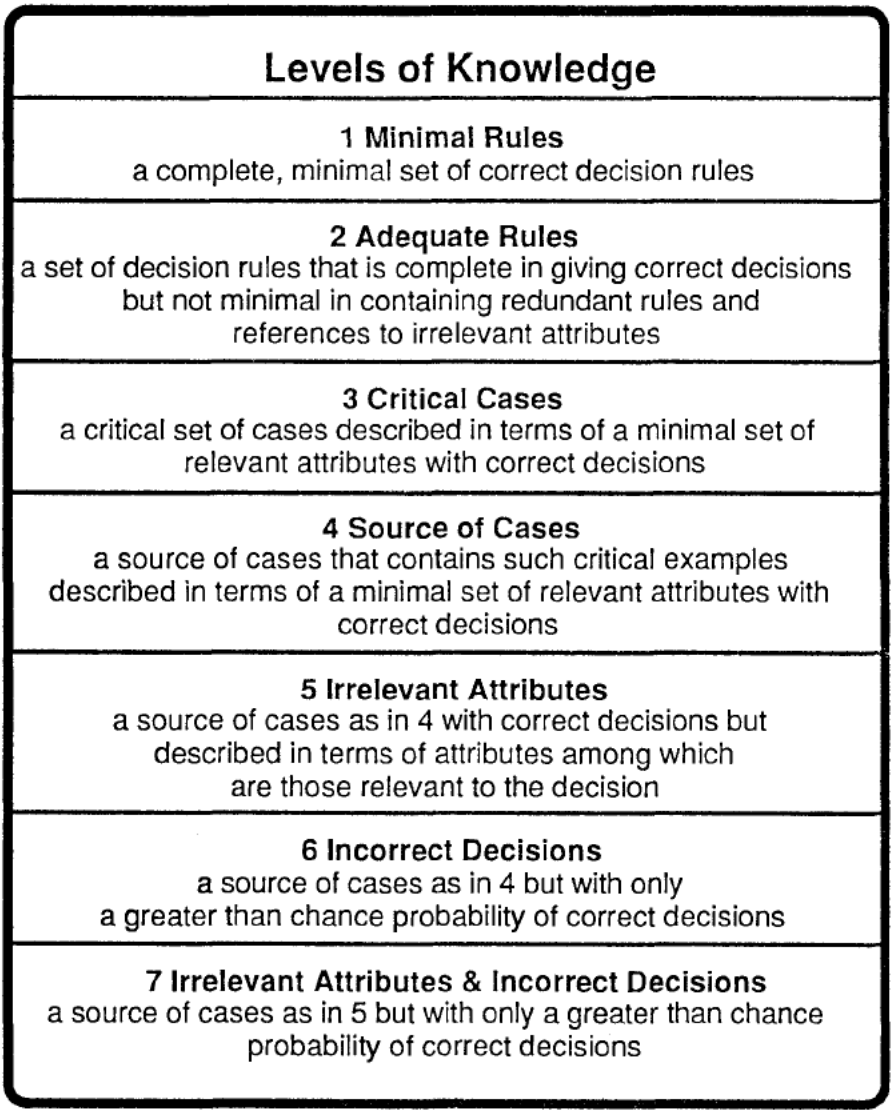

Figure 2 Levels of expert knowledge

This is a continuum of decreasing knowledge on the part of the human expert. It encompasses a range of situations met in practice, and it raises the question of how the amount of knowledge available from the expert affects the amount of data required for effective empirical induction. For case 1 no data is required for empirical induction since the correct answer is available. For case 7 the 'expert' has provided little except access to a source of data from which the correct answer might be derived. How much data is required for an optimal empirical induction procedure to derive 1 given 7, and how much less data is required for cases 2 through 6?

The following section report some studies that give quantitative answers to these questions. The studies are empirical rather than analytic so that it is not apparent yet how the answers generalize to arbitrary situations. However, they establish some base-line data which is interesting in its own right, a possible guide to practitioners of knowledge acquisition, and a test case for potential analytic estimates of the ratios involved.

THE TRADE-OF BETWEEN KNOWLEDGE AND DATA

INDUCT (Gaines 1989) is part of a knowledge acquisition tool KSS0 (Gaines 1987b) which uses entity-attribute grid techniques (Shaw 1980, Shaw & Gaines 1983, Boose 1984) to elicit relevant attributes and critical entities from experts, and empirical induction based on these to build a knowledge base in terms of classes, objects, properties, values and methods (rules). The tool also accepts rules entered by experts and entity-attribute data from databases, so that its range of knowledge/data combinations encompasses all those listed in the previous section.

INDUCT is comparable in its approach and performance to other empirical induction algorithms (Quinlan 1987), having a particularly high noise rejection—technical details are given in the Appendix. However, we believe that the experimental results on the trade-offs between knowledge and data given below are fairly independent of the empirical induction algorithm used—they are systemic consequences of the quality of knowledge available rather than side-effects of a particular approach to inductive modeling.

Cendrowska's (1987) contact lens data (Figure 3) used in her exposition of PRISM has been taken as a starting point for our trade-off studies since it is well-defined, previously analyzed and results in a range of rules of varying complexity, some of which are supported by a high proportion of the data set and others of which are supported by only single cases. The deterministic, complete data set involves 24 cases described in terms of 3 binary attributes, 1 ternary attribute and 1 ternary decision attribute. PRISM gives a solution based on 9 rules, but using default logic correct solutions with only 6 rules are available.

On this data set INDUCT generates both the 9 rule and 6 rule solutions dependent on whether default logic is allowed (exception rules are given higher priority than partially-correct default rules). Figure 4 shows the 9 rule solution—the first number after each rule is the percentage correct on the test data—it would be lower if the data were noisy—the second number is the statistical significance of the rule in terms of the percentage probability that it has arisen by chance.

Age	Prescription	Astigmatism	Tear Production	Lens
young	myope	not astigmatic	reduced	none
young	myope	not astigmatic	normal	soft
young	myope	astigmatic	reduced	none
young	myope	astigmatic	normal	hard
young	hypermetrope	not astigmatic	reduced	none
young	hypermetrope	not astigmatic	normal	soft
young	hypermetrope	astigmatic	reduced	none
young	hypermetrope	astigmatic	normal	hard
pre-presbyopic	myope	not astigmatic	reduced	none
pre-presbyopic	myope	not astigmatic	normal	soft
pre-presbyopic	myope	astigmatic	reduced	none
pre-presbyopic	myope	astigmatic	normal	hard
pre-presbyopic	hypermetrope	not astigmatic	reduced	none
pre-presbyopic	hypermetrope	not astigmatic	normal	soft
pre-presbyopic	hypermetrope	astigmatic	reduced	none
pre-presbyopic	hypermetrope	astigmatic	normal	none
presbyopic	myope	not astigmatic	reduced	none
presbyopic	myope	not astigmatic	normal	none
presbyopic	myope	astigmatic	reduced	none
presbyopic	myope	astigmatic	normal	hard
presbyopic	hypermetrope	not astigmatic	reduced	none
presbyopic	hypermetrope	not astigmatic	normal	soft
presbyopic	hypermetrope	astigmatic	reduced	none
presbyopic	hypermetrope	astigmatic	normal	none

Fig.3 Cendrowska's (1987) contact lens data

prescription=hypermetrope & astigmatism=not astigmatic & tear production=normal -> lens recommendation=soft 100% 0.904%
age=young & astigmatism=not astigmatic & tear production=normal -> lens recommendation=soft 100% 4.34%
age=pre-presbyopic & astigmatism=not astigmatic & tear production=normal -> lens recommendation=soft 100% 4.34%

tear production=reduced -> lens recommendation=none 100% 0.355%
age=presbyopic & prescription=myope & astigmatism=not astigmatic -> lens recommendation=none 100% 39.1%
age=pre-presbyopic & prescription=hypermetrope & astigmatism=astigmatic -> lens recommendation=none 100% 39.1%
age=presbyopic & prescription=hypermetrope & astigmatism=astigmatic -> lens recommendation=none 100% 39.1%

prescription=myope & astigmatism=astigmatic & tear production=normal -> lens recommendation=hard 100% 0.463%
age=young & astigmatism=astigmatic & tear production=normal -> lens recommendation=hard 100% 2.78%

Fig.4 Decision rules derived from the dataset of Figure 3

This data set is used as a kernel from which to generate corrupted data with varying probabilities of errors and with varying numbers of irrelevant attributes as shown in Figure 5. The generator selects an entity at random from the 24 cases, randomly changes the decision according to a prescribed probability, and adds a prescribed number of irrelevant binary attributes with random values. The corrupted data sets are run incrementally on INDUCT to determine the minimum amount of data necessary for a correct solution. This quantity is itself a random variable and 10 data sets generated with the same parameters are run to obtain more robust estimates of the data requirement.

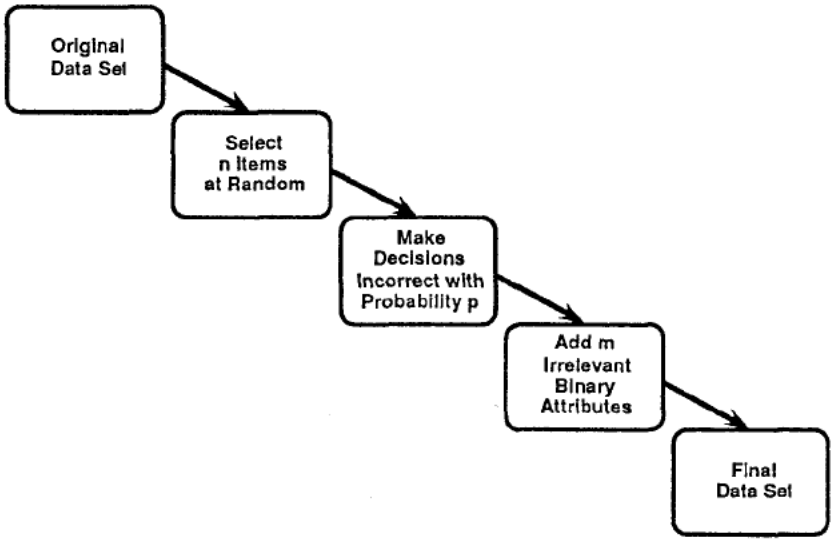

Fig.5 Dataset generation for empirical studies of knowledge/data tradeoff

Figure 6 shows the results obtained to date with a variation from 6 to over 1000 data items being needed to cover the spectrum of knowledge availability specified in cases 1 through 7 in Figure 2. It is interesting to note that very high levels of noise can be tolerated—the expert by no means has to be 100% correct, only better than chance. Noisy irrelevant attributes cause similar effects to noise in the decision—this seems to validate claims that knowledge acquisition tools targeted on eliciting relevant attributes are in themselves worthwhile. There is strong interaction between errors and irrelevancy—much more data is required to eliminate both than either alone.

Knowledge	*Data Required*	
	Mean	S.D.
Exact rules	6	0
Critical cases	18	0
Correct cases	90	43
10% Errors	123	49
25% Errors	326	159
1 Irrelevant Att.	160	77
2 Irrelevant Atts.	241	125
5 Irrelevant Atts.	641	352
10% Err. + 1 Irr.Att.	1970	1046

Fig.6 Data/knowledge tradeoff

COMBINING EXPERTISE TRANSFER AND EMPIRICAL INDUCTION

Current systems for expertise transfer generally also provide empirical induction tools for modeling cases and deriving rules from, for example, entity-attribute grid data. The studies reported in this paper give a formal framework within which the relative roles of these tools, and their interplay and integration, can be analyzed. The seven levels of expertise defined in Figure 2 are those we actually find in the "experts" typically available to us in the development of knowledge-based systems. Our "expertise transfer" methodologies give us problem-solving attributes and exemplary cases expressed in terms of those attributes. There are no guarantees that the attributes elicited are all necessary or that the cases elicited all have correct decisions. The results of the previous section show that we can eliminate irrelevant attributes and incorrect decisions providing:

(a) A sufficient set of relevant attributes for problem-solving has been elicited;

(b) an adequate number of cases has been elicited in relation to the number of irrelevant attributes and probability of error in exemplary cases.

Achieving condition (a) has been a major design requirement in even the earliest personal construct elicitation systems. Shaw's (1980) original design for a computer-based grid elicitation system emphasized features aimed at stimulating the user to add complementary attributes through feedback about those already elicited, and these have been further refined and developed in later systems.

The difficulty with achieving condition (b) is that estimates of the number of irrelevant attributes and the error rates in case data may be difficult to obtain. However, even coarse estimates make it possible to use studies such as that reported above to obtain lower bounds on data requirements.

What is interesting in considering requirement (b), however, is to note that additional case history data and improving estimates of error rates become available through the application of the knowledge-based system being developed. Whatever initial expertise is available can be used to prime an expert system that then, through use, acquires more case data on which to base improved empirical induction.

Additional cases that conform to, and hence confirm, existing rules present no problems. However, those which induce learning, that is a change in rules, also suggest that the decisions being recommended by the system are suspect. As in human behavior, the detection of anomalies between knowledge and data is critical to both action and learning.

I have previously suggested that rules based on empirical induction, such as those shown in Figure 4, should be coupled with anomaly detection rules that trigger "surprise" if their premises are true but their conclusions are false (Gaines 1988b). In such cases, further empirical induction on the enhanced data set will result in either the rule being changed or the new data rejected as in error. This is precisely the process that Stich and Nisbett (1984) have termed the expert's management of "reflective equilibrium" in their analysis of expertise using Goodman's (1973) model of induction.

The current generation of expert systems, based on taking static snapshots of expertise, will become obsolete very rapidly and present major problems of maintenance. We will be forced to move very rapidly toward what Steels (1986) has termed "second generation" systems able to learn from experience (Gaines 1988a). However, we have also to remember that the associated human experts will also be acquiring new knowledge from many sources not necessarily accessible to the knowledge-based system. Our system architectures must accomplish a long-term integration of expertise transfer and machine learning methodologies.

WHAT MAY BE ACQUIRED THROUGH COMBINED ET AND EI

The model of Figure 1 can be extended to encompass other aspects of knowledge acquisition, notably the roles of neural networks and text processing. The "directed data acquisition" at the top of the model is not readily emulated using current knowledge acquisition technologies. We may obtain labels for attributes from an experts and have them characterize cases in terms of these, but our methodologies do not capture the way in which the expert does this, the procedures for measuring the values of attributes. Thus, our knowledge-based systems cannot describe cases themselves but must question clients. In practice we assume that the clients themselves are already sufficiently expert to understand the expert's terminology and are able to supply valid characterizations of cases within the expert's conceptual framework. These are strong assumptions.

Figure 7 summarizes the levels and types of knowledge that we may regenerate given sufficient data based on lower levels of knowledge. A combination of expertise transfer and empirical induction enables us to regenerate the *prescriptive* knowledge of the expert in the form of decision-making models, the *descriptive* knowledge in the form of relevant attributes and their critical values, and the *ecological* knowledge in the form of cleaned-up case histories. What remains inaccessible is the role of the expert as an *interactive* transducer providing measurements of the world.

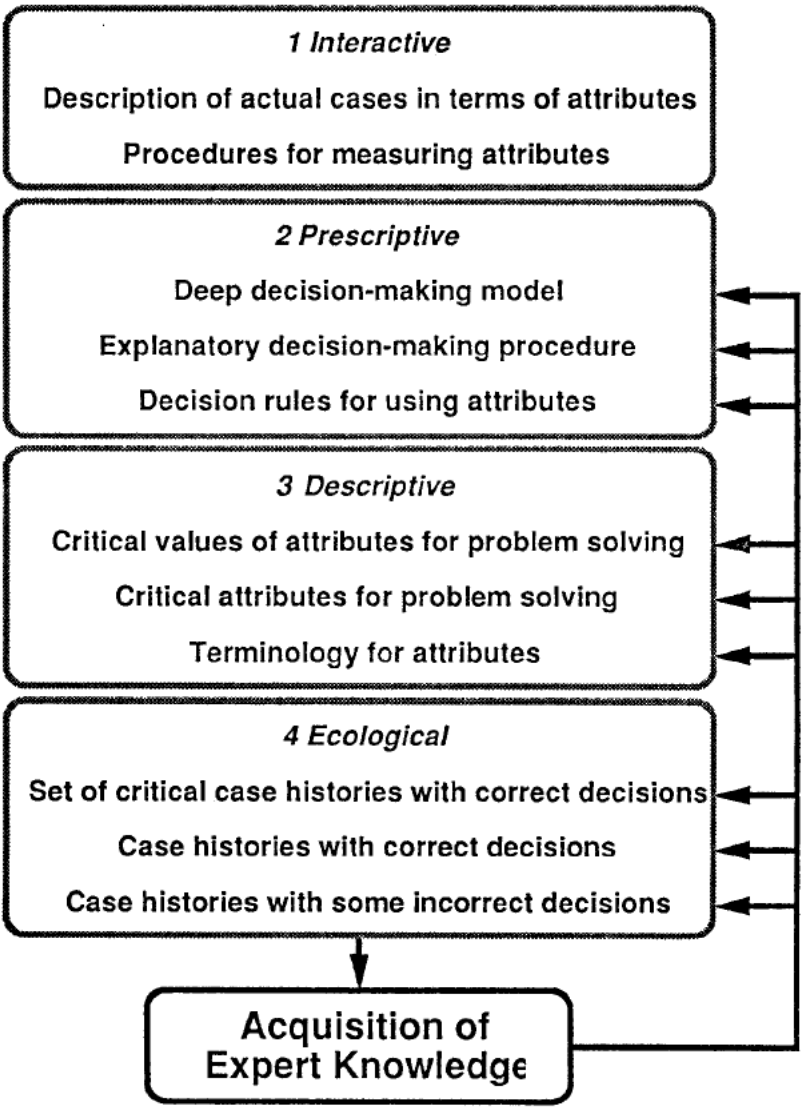

Fig.7 Regeneration of knowledge

Speculatively, we suggest that the role of neural network research within knowledge acquisition is to emulate the human capabilities of perceptual learning that enable us to form new cognitive constructs using the mass of available data to our sensations (Grossberg 1988, Eckmiller & Malsburg 1988). The lowest level of "directed data acquisition" in Figure 1 requires the formation of measurements, the recognition of new patterns, that provide values for attributes appropriate to problem solving.

This is not a new concept in terms of the direction of much research on neural networks. However, it is novel in the expertise transfer community because, so far, the emphasis has been on transferring the high level knowledge, not the perceptual capabilities. As knowledge-based systems develop it is reasonable to expect more emphasis of complete functionality, including the emulation of perceptual capabilities. It is reasonable to aim for these to be an enhancement of existing approaches rather than a complete change to radically different technologies.

Figure 8 shows Figure 1 extended with a neural network learning to emulate the directed data acquisition through perceptual learning. It also shows a second extension in that expertise is conventionally externalized and communicated through encoding in media such as text and diagrams. This provides yet another path into the knowledge base through tools that extract knowledge from such material.

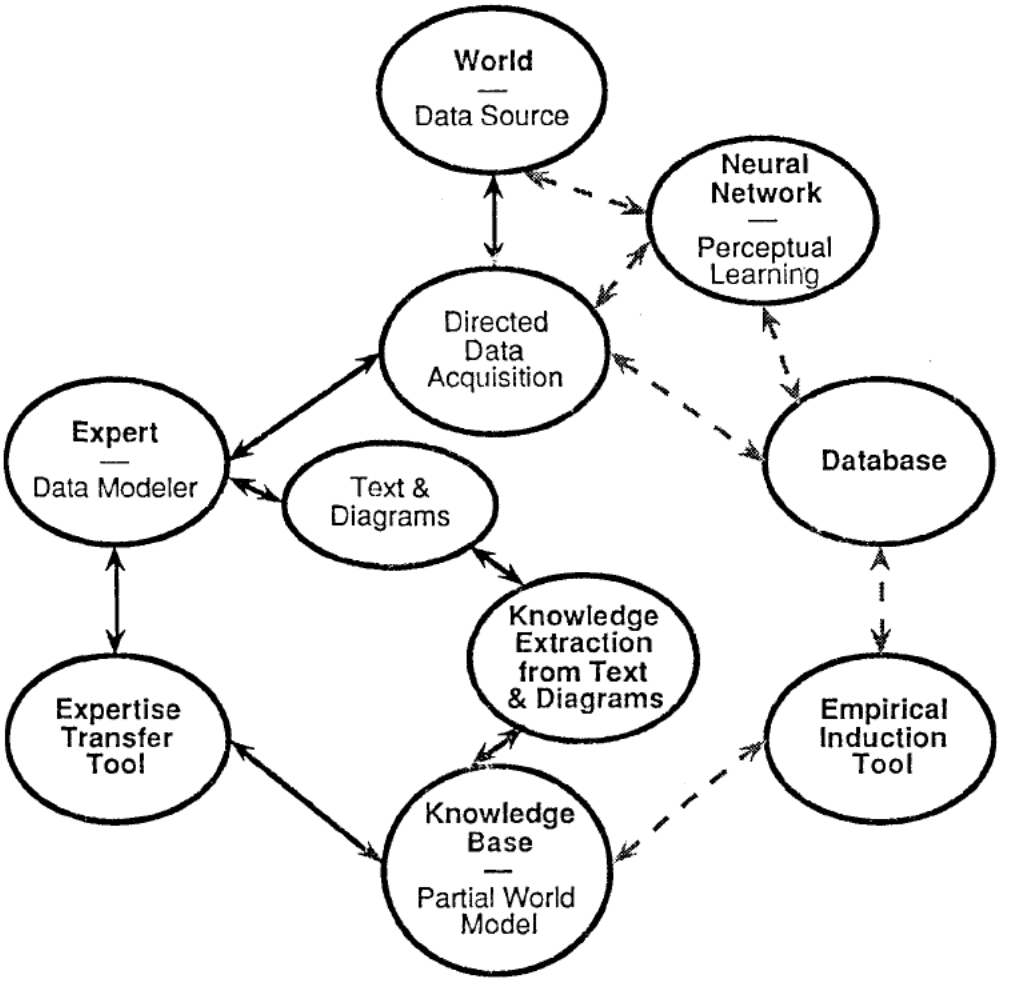

Fig.8 Integration of various forms of knowledge acquisition

A knowledge acquisition system with the architecture of Figure 8 would integrate the state of the art of four diverse artificial intelligence technologies. It is a reasonable objective for next generation research.

CONCLUSIONS

At one level the studies reported in this paper may be seen as legitimating current practice in research and development for knowledge acquisition systems. Most applicable systems are providing tools for both expertise transfer and empirical induction, and for using thse in closely coupled combinations. However, some of the tutorial literature presents them as almost competetive approaches rather than as extremes along a continuum of possibilities that match a continuum of types and levels of expertise.

In architectural terms, the analysis leading to Figure 8 provides a integrative structure for four of the major technologies supporting knowledge acquisition processes in artificial intelligence research. A practical system with such an architecture will involve many paths not shown in the figure which support the meta-tasks of managing the acquisition process, and it will involve multiple experts in different roles.

In practical terms, the studies reported in this paper provide a principled, quantitative approach to many of the different techniques being used for knowledge acquisition, from those based on expert interviews to those based on empirical induction. The trade-off data is a guide to practitioners as to the appropriate approach and data requirements in their situations—it needs testing with other data sets but that is now a matter of sheer number crunching.

In theoretical terms, the approach taken in this paper seems to offer the possibility of developing a quantitative science of knowledge in terms of the amount of data reduction that knowledge buys us when carrying out empirical induction. This can be seen as a reasonably principled economic evaluation of the knowledge.

Figure 9 shows the results of Figure 6 used to define a measure of knowledge in terms of the logarithm of the data required to achieve it. This gives a basis for an economic measure of the "value" of different forms and levels of knowledge if we assume, for example, that obtaining each item of data incurs equal cost. The studies in this paper are empirical using a particular situation. It should be possible to obtain some analytic results for more general cases.

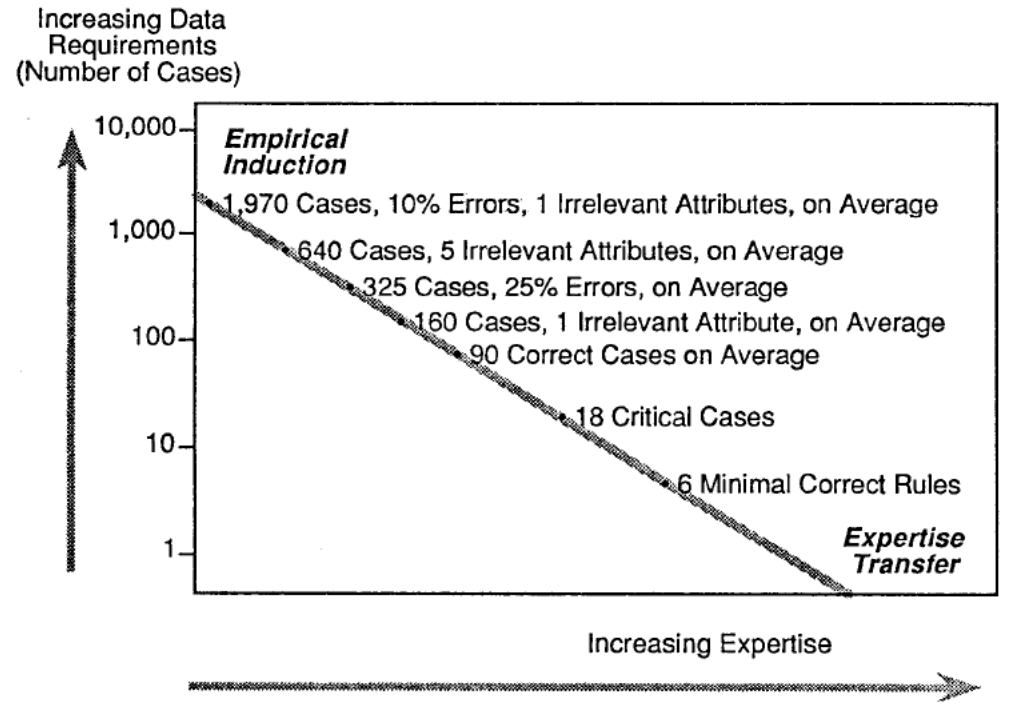

Fig.9 Knowledge/data tradeoff in expertise transfer and empirical induction

In design terms, the most important results obtained are that a single inductive algorithm is adequate to cover the complete spectrum of cases 1 through 7 of "levels of expertise" defined at the beginning of this paper, from rule simplification, to noise reduction and relevancy determination, and that it is possible to generate an integrated system that ranges across that spectrum, from the transfer of expertise from experts to the creation of equivalent expertise through empirical induction.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. I am grateful to many colleagues for discussions over the years that have influenced this paper. In particular I would like to thank John Boose, Jeff Bradshaw, Mildred Shaw, and Ross Quinlan, for access to their own research and many stimulating discussions.

APPENDIX: THE INDUCT ALGORITHM

It is relevant to the assessment of the results in this paper to understand the method used to derive decision rules from data, and this section outlines the INDUCT algorithm, linking it to previous approaches, clarifying enhancements made and their roles in the studies, and comparing performance with the best known methods.

Cendrowska (1987) has shown that empirical induction through decision trees and direct conversion to rules, even with pruning, leads to rule sets that test the values of irrelevant attributes and are much larger than is necessary. Her PRISM algorithm goes from entity-attribute data direct to rules but does not address the problems of noisy data or missing values. Quinlan (1987) has developed extended pruning techniques in a way that are effective in coping with noisy data but still involve decision tree production with problems of irrelevant attributes and missing values.

INDUCT extends the PRISM algorithm to control direct rule generation through statistical tests that are effective in dealing with both noisy data and missing values. Figure 10 shows the basis for these statistical tests. Given a universe of entities, E, a target predicate, Q, and a set of possible test predicates of the form, S, on entities in E, use them to construct a set of rules from which the target predicate may be inferred given the values of the test predicates.

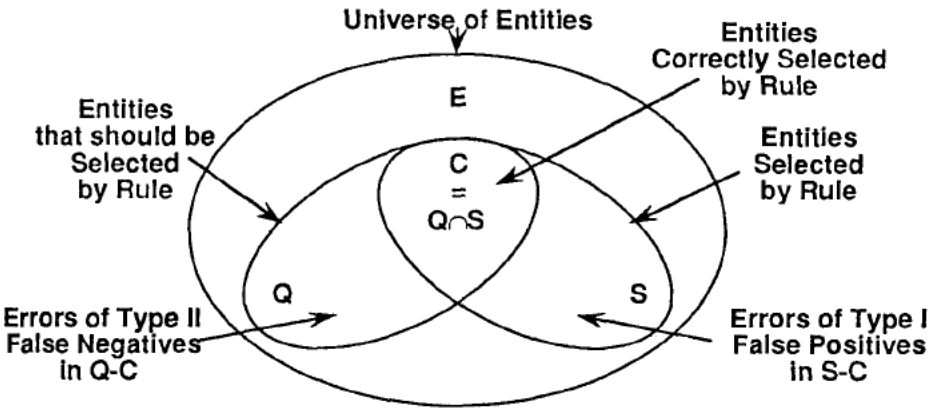

Fig.10 Problem of empirical induction

For the purposes of the statistical analysis the forms of S and Q do not matter. One may regard S as a *selector* choosing those e out of some subset of E for which to assert Q(e), and compare the selection process of the rule with that of random selection, asking "what is the probability that random selection of the same degree of generality would achieve the same accuracy or greater."

This probability is easily calculated, let Q be the relevant entities in E for which Q(e) holds, S be the selected entities in E for which S(e) holds, C be the correct entities in E for which both S(e) and Q(e) hold:

$$Q \equiv \{e: e \in E \wedge Q(e)\} \quad (1)$$

$$S \equiv \{e: e \in E \wedge S(e)\} \quad (2)$$

$$C \equiv \{e: e \in E \wedge S(e) \wedge Q(e)\} \quad (3)$$

Let the cardinalities of E, Q, S and C be e, q, s and c respectively. The probability of selecting an entity from E for which Q holds at random is:

$$p = q/e \quad (4)$$

The probability of selecting s and getting c or more correct at random is:

$$r = \sum_{i=c}^{s} {}^{s}C_i \, p^i \, (1-p)^{s-i} \quad (5)$$

The advantage of using r as a measure of the correctness of a rule is that is easily understood, as the probability that the rule could be this good at random, and that it involves no assumptions about the problem such as sampling distributions. Note that if c=s (all correct) then log(p) = s log(q/u) which seems to be the basis of 'information-theoretic' measures.

The probability that the performance of an isolated rule could be obtained by random selection is not adequate in itself to evaluate a rule that has itself been selected as best after a search through many possible rules. If n different rules have been searched then we may ask "what is the probability that the rule found by search achieved its performance by chance," that is the probability that a rule will have been found that in itself has a probability of being selected at random of r. This probability is:

$$t = 1 - (1-r)^n \quad (6)$$

As one would expect this approximates to nr if this term is substantially less than 1.
To obtain a value for t one needs to estimate n, the number of rules in the search space, and this requires information about the forms of possible rules. If the rules are conjunctions of value tests of attributes as illustrated in (3) then n_1 can be estimated for rules with one clause, n_2 for rules with two or less clauses, and so on. For example if there are m attributes the i'th of which can be tested in m_i possible ways, then:

$$n_1 = \sum_{i=1}^{m} m_i \quad (7)$$

$$n_2 = (n_1)^2 - \sum_{i=1}^{m} (m_i)^2 \quad (8)$$

Missing values are taken into account by assuming that they might have any value. When the selection of an entity is tested a missing value is assumed to have the required value for selection. In the statistics a selection based on missing values is allowed to contribute to false positives but not to correct positives. This has important consequence in knowledge acquisition since it allows the expert to enter conjunctive rules as if they were entities with missing values. INDUCT then generates the same or an equivalent smaller rule set. It is reasonable to test the consistency of an inductive procedure by requiring the rule-set produced by it to be 'fixed-point' if re-entered as data.

INDUCT has been tested on a wide range of data sets in the literature together with many artificial data sets with known degrees of noise, missing values and irrelevant attributes, and found to perform consistently at least as well as the previously published best results. For example, with Quinlan's (1987) very noisy data sets, Prob-Disj (disjunction of three terms in three attributes, one irrelevant attribute, 10% noise), and Digits (7-segment display, 10% noise in attributes, Bayes optimal solution 26% errors) the comparisons are shown in Figure 11.

Results on Prob-Disj Data (3-term disjunct)

Data	Error Rate (Type I + Type II)	
	ID3 to Pruned Rules	INDUCT
Sample	4.2 rules	3 rules
Test1	10.0%	10.0%
Test2	10.0%	10.0%

Results on Digits Data

Data	Rules	Error Rate (Type I + Type II)	
		ID3 to Pruned Rules	INDUCT
Sample	10		31.05%
	16		26.8%
Test1	10		32.0%
	16	31.3%	29.0%
Test2	10		31.4%
	16	28.3%	27.8%

Fig.11 Comparison of INDUCT and optimally-pruned ID3

REFERENCES

Boose, J.H. (1984). Personal construct theory and the transfer of human expertise. **Proceedings AAAI-84**, 27-33. California: American Association for Artificial Intelligence.

Boose, J H. (1986) **Expertise Transfer for Expert System Design**. New York: Elsevier.

Boose, J.H. & Gaines, B.R., Eds. (1988). **Knowledge Acquisition Tools for Expert Systems**. London, Academic Press.

Boose, J.H. & Bradshaw, J.M. (1987) **Expertise** transfer and complex problems: using AQUINAS as a knowledge acquisition workbench for knowledge-based systems. **International Journal of Man-Machine Studies 26, 3-28.**

Cendrowska, J. (1987) An algorithm for inducing modular rules. **International Journal of Man-Machine Studies 27 (4), 349-370 (October).**

Eckmiller, R. & Malsburg, C., Eds. (1988). **Neural Computers**. Berlin: Springer.

Gaines, B.R. (1987a) An overview of knowledge acquisition and transfer. **International Journal of Man-Machine Studies 26** (4), 453-472 (April).
Gaines, B.R. (1987b). Rapid prototyping for expert systems. Oliff, M., Ed. **Proceedings of International Conference on Expert Systems and the Leading Edge in Production Planning and Control**. pp.213-241. University of South Carolina.
Gaines, B.R. (1988a). Second generation knowledge acquisition systems. **Proceedings of the Second European Workshop on Knowledge Acquisition for Knowledge-Based Systems (EKAW'88)**. GMD-Studien Nr.143. pp.17-1-17-14. Bonn, Germany: Gellsellschaft für Mathematik und Datenverarbeitung (June).
Gaines, B.R. (1988b). Positive feedback processes underlying the formation of expertise. **IEEE Transactions on Systems, Man & Cybernetics**, SMC-18(6), 1016-1020 (November).
Gaines, B.R. (1989) An Ounce of Knowledge is Worth a Ton of Data: Quantitative Studies of the Trade-Off between Expertise and Data based on Statistically Well-Founded Empirical Induction. **Proceedings of 6th International Workshop on Machine Learning**, Cornell University, Ithaca.
Gaines, B.R. & Boose, J.H., Eds. (1988). **Knowledge Acquisition for Knowledge-Based Systems**. London, Academic Press.
Gaines, B.R. & Shaw, M.L.G. (1987). Knowledge support systems. **ACM MCC-University Research Symposium**. Austin, Texas: MCC. pp.47-66.
Goodman, N. (1973). **Fact, Fiction and Forecast**. Indianapolis: Bobbs-Merrill.
Grossberg, S., Ed. (1988). **Neural Networks and Natural Intelligence**. Cambridge, Massachusetts: MIT Press.
Hayes-Roth, F. (1984). The industrialization of knowledge engineering. Reitman, W., Ed. **Artificial Intelligence Applications for Business**. pp. 159-177. Norwood, New Jersey: Ablex.
Machlup, F. (1980). **Knowledge and Knowledge Production**. Princeton University Press.
Machlup, F. (1982). **The Branches of Learning**. Princeton University Press.
Machlup, F. (1984). **The Economics of Information and Human Capital**. Princeton University Press.
Machlup, F. & Mansfield, U. (Eds.) (1983). **The Study of Information**. New York: Wiley.
Quinlan, J.R. (1986) Induction of decision trees. **Machine Learning** 1(1) 81-106 (March).
Quinlan, J.R. (1987) Simplifying decision trees. **International Journal of Man-Machine Studies 27** (3), 221-234 (September).
Quinlan, J.R., Compton, P.J., Horn, K.A. & Lazarus, L. (1987) Inductive knowledge acquisition: a case study. Quinlan, J.R. (Ed.), **Artificial Intelligence and Expert Systems**. pp.157-173. Sydney: Addison-Wesley.
Shaw, M.L.G. (1980). **On Becoming a Personal Scientist**. London: Academic Press.
Shaw, M.L.G. & Gaines, B.R. (1983). A computer aid to knowledge engineering. **Proceedings of British Computer Society Conference on Expert Systems**, 263-271 (December). Cambridge.
Steels, L. (1986). Second generation expert systems. Bramer, M.A. (Ed). **Research and Development in Expert Systems III**. pp.175-183. Cambridge: University Press.
Stich, S.P. & Nisbett, R.E. (1984). Expertise, justification and the psychology of inductive reasoning. Haskell, T.L., Ed. **The Authority of Experts**. pp. 226-241. Bloomington, Indiana: Indiana University Press.
Sveiby, K.E. & Lloyd, T. (1987). **Managing Knowhow**. London: Bloomsbury.