PROCEEDINGS OF THE INTERNATIONAL CONFERENCE
ON FIFTH GENERATION COMPUTER SYSTEMS 1984,
edited by ICOT. © ICOT, 1984

QUALITY ISSUES IN NEW GENERATION COMPUTING

Henry J. Kohoutek

Product Assurance Manager, Hewlett-Packard Company
Fort Collins Systems Division
3404 East Harmony Road, Fort Collins, Colorado, 80525 U.S.A.

ABSTRACT

The next generation of computer products and the emerging commercialization of Artificial Intelligence require a reassessment of current quality assurance know-how and technology for its ability and readiness to support and contribute to this new phase in the computer industry.

After defining quality as a multiattribute product characteristic, the author reviews quality issues that are associated with knowledge, its representation, storage and processing, and also with computing hardware and human interfaces. These issues are discussed within the framework of current quality assurance concepts. This introductory and elementary review indicates that assurance sciences are generally adequate for developing needed quality strategies and it identifies a few areas of expected breakthroughs. The emphasis of this paper is on applications rather than theory; it deliberately avoids issues related directly to application domain facts, assumptions and problems.

1. INTRODUCTION

A paradox we often face is the apparent contradiction between the great triumphs and the equally dramatic failures of the human mind. The same organism that routinely solves inferential problems that are too subtle and complex for the most powerful computers, often makes errors in the simplest of judgements about everyday events (Nisbett and Ross 1980).

If we accept John R. Anderson's (Anderson 1983) three points of evidence for the plausibility of a unitary theory of mind:

- the unity of human cognition,
- the belief that all higher cognitive processes are different manifestations of the same underlying system, and
- the claim that production systems are computationally universal (i.e. capable of modeling all cognitive activity),

with the manifested ambition to mechanize the mind's major abilities, we can recognize the potential risks of computational and reasoning failures associated with current knowledge systems. If our evidence shows that the reasons for erroneous human judgements makes them more traceable to violations of inferential rules and misapplications of problem-solving tools by the Piagetian adult than to the largely motivational errors of the Freudian adult, then we can see the potential benefits of assurance science tools and methodologies for improving the quality of Artificial Intelligence (AI) based products by making the occurance of these errors less probable.

Because AI products are beginning to prove themselves in the marketplace, the reasons for applying Quality Assurance (QA) and Control know-how may also be highly pragmatic.

2. DEFINITION OF QUALITY

In this study, "quality" is understood to be a part of a product's value-in-use. It is distinct from performance which is another contributor to the value-in-use. It is also distinct from, but influences, the value-in-exchange via price (Weiner 1983, Page 1968).

Quality is perceived as a multiattribute product characteristic which can be expressed by a generalized, overall rating which is based on multidimensional measurements that reflect rank orderings of preferences and their relative importance (Monroe and Petroshius 1973). This definition of quality allows us to integrate the customers' expectations (objective and subjective, real and perceived) and market trends with state-of-the-art assurance know-how to identify sets of performance and acceptance requirements. These requirements constrain the product design and manufacturing goals by paradigms, architectures and implementation alternatives. This definition is also suitable for studies in consumer behavior (Kassarjian and Robertson 1981) and market research.

The most successful attempts today in developing a categorical structure of the notion of computer product quality reflect consumer behavior studies and observations along with the practical experience of the assurance specialist. These attempts identify four major quality attributes:

- fitness for use,
- correctness of results,
- reliability and
- maintainability.

The economical, technical and psychological aspects of these major attributes have a significant impact on the customer's satisfaction and must be considered when developing quality strategies for the new generation of computers.

Because the categorical structure of Artificial Intelligence or Knowledge System products is not yet fully developed, despite many brilliant attempts by Newell (Newell 1982) or Frederick Hayes-Roth (Hayes-

Roth 1983) for example, we will try to address the quality issues in an intuitive sequence: 1) knowledge, 2) problem solving methods, 3) human interface, 4) hardware, and 5) the complete knowledge system.

3. KNOWLEDGE

If the human mind is the fundamental object to be modelled by the computational tools of Artificial Intelligence, then knowledge is the fundamental notion for AI's engineering application: Knowledge System Engineering. Knowledge can be viewed as a medium or specification of what symbol structures should be able to do. It represents a potential for generating actions.

3.1 Knowledge Levels

One of the working hypotheses of contemporary Knowledge Engineering assumes that expert system performance is dependent primarily on the size and quality of that system's knowledge. The factual part of our knowledge comes mostly in the form of scientific and technical data in three broad classes (Lide Jr. 1981):

- repeatable measurements on well-defined systems (Class A)
- observational data, often time or space dependent (Class B)
- statistical data (Class C).

Unfortunately, the quality of data preserved in scientific literature leaves much to be desired. This becomes apparent when data on a well-studied subject are systematically reviewed. For example, in the aforementioned article by David R. Lide, Jr., a three order-of-magnitude scatter of about 200 reported measurements of the thermal conductivity of copper as a function of temperature (in some temperature ranges), illustrates the pitfalls of relying on a single value retrieved from the literature.

Quality Assurance methodologies recommend using Class A data only from reliable sources such as the National Standard Reference Data System (NSRDS) (Brady and Wallenstein 1964) which is coordinated by the U.S. National Bureau of Standards (NBS). The NBS critically evaluates and maintains data from academic, industrial and government establishments.

The most effective way to assure the quality of Class B data is by careful maintenance and calibration of the measuring instruments prior to the data acquisition. Recording and preserving all auxiliary information required is also important.

Quality control of Class C data is a more difficult task, often hindered by disagreements in definitions and terminology. Also, Quality Assurance has not been a major issue in bibliographical services. The electronic revolution in data dissemination, the growth in the use of computer-based, on-line systems and the need for reliable data, will hopefully provide enough pressure to cause both format design and data quality assurance to become key factors in controlling Class C data reliability.

Resolving quality issues in the higher knowledge levels that are being engineered today (from definitions, taxonomies, discrete descriptions, constraints, empirical associations, perceptual structures, deductive methods, etc., up to search heuristics), presents many epistomological challenges. The question of the reliability of scientific knowledge is a serious intellectual issue. Once we have cast off the naive doctrines that all science is necessarily true, and that all true knowledge is necessarily scientific, we realize that epistemology is not just an academic philosophical discipline but a practical ground for decisions and actions (Ziman 1978). This leads us to understanding the dangers of

- cultural dependencies,
- historical inaccuracies,
- subjectivity,
- categorical imprecisions,
- hidden variables, and
- parascientism,

and to search for knowledge quality assurance tools and strategies. It also leads us to distinctions between knowledge and belief or introspection (Hintikka 1962).

The knowledge Quality Assurance task will consist mainly of evaluating its

- perceptual consensibility (unambiguity leading to consensus) (Ziman 1978),
- communicable consensuality (visibility and acceptability of perceived patterns) (Ziman 1978),
- credal probability (expectation determining) (Levi 1980), and
- methods of justification and validation (including extra-logicality) (Ziman 1978).

This evaluation should be made with the understanding that philosophers usually seek "certain" knowledge, while scientists remain content with "probable" knowledge (Gregory 1981), i.e. they accept frequent revisions and occasional dramatic changes in principles and premises.

Isaac Levi's insight that knowledge is a resource for inquiry and deliberation, as well as a standard for justifying choices between feasible options, the quality of which can be assessed in terms of credal probability, coherence, consistency and relevance (Levi 1980), can serve as a theoretical base for knowledge Quality Assurance studies. But on knowledge levels that are subjected to current systems engineering and its pragmatic interest, the key philosophers of knowledge (David Hume, Ludwig Wittgenstein, Rudolph Carnap, Bertrand Russell, Sir Karl Popper, Thomas Kuhn and others) offer only limited help.

3.2 Knowledge Representations

Because the goals of AI Research and Knowledge System Engineering are in program designs that exhibit intelligent behaviour, AI researchers have often taken a rather pragmatic approach to knowledge. They have focused on developing schemes to incorporate knowledge about the world into these programs (Barr and Feigenbaum 1981). These Knowledge Representation Techniques involve routines for manipulating specialized data structures to obtain intelligent inferences. The most frequently used formalisms are: state/space representation, formal logic, procedural representation, semantic nets, production systems, frames and special techniques for visual scene or speech representation.

In general, the representation of knowledge is a combination of data structures and interpretive procedures. The issues of form and notation have occupied most of the past discussion, but Anderson (Anderson 1983) recognized that the issue of what could or could not be done with a representation is more important. He has proposed his Tri-code Theory of Knowledge Representation which states that we need:

- linear structures to represent the order of a set of items or events;
- spatial structures to encode spatial configurations;
- abstract propositional structures to represent logical interpretations and to encode meaning.

The most important quality attribute of a knowledge representation is its adequacy. J. McCarthy and P.J. Hayes (McCarthy and Hayes 1981) have defined three kinds of adequacies:

- metaphysical, assuring a form that does not contradict the facts that interest us.
- epistemological, which guarantees that the representation can be used practically to express the facts one actually has about a given aspect of the world.
- heuristical, that assures the reasoning processes leading to a solution are expressible in the language.

Other attributes that allow quality assessment of knowledge representations are (Barr and Feigenbaum 1981):

- Modularity – allows adding, deleting or changing information independently.
- Uniformity – assures understandability by other parts of the system via rigid knowledge structures.
- Naturaliness – reflects the ease of expressing important kinds of knowledge.
- Understandability – the degree to which the system is understandable by humans.
- Modifiability – assures context independence and allows the meaning of a fact to be specified, when the fact is entered or removed, independently of the rest of the system.

Studies on semantic primitives, which deal with establishing representational vocabularies, lead Wilks (Wilks 1977) to such attributes as finitude, comprehensiveness, independence, noncircularity, and primitiveness.

Other quality related issues, such as those described in SIGART "Special Issues of Knowledge Representation" (Brachman and Smith 1980), represent almost overwhelming diversity. All this indicates the need for standardization as a base for effective Quality Assurance.

New Quality Assurance problems arise when changing the representations during problem solving. Because the two dimensions of representation space are Information Structure and Information Volume, the representation transformations must be evaluated in terms of isomorphism and homomorphism. In general, isomorphic transformations change the Information Structure of a representation while leaving the Information Volume fixed. Conversely, homomorphic transformations alter the Information Volume while leaving its structure unchanged (Korf 1980).

3.3 Knowledge Bases

Addressing the issues of knowledge quality and the quality of its representation will not free us from concerns about the quality of the databases where this knowledge will be stored. The conventional strategies for improving data base quality have culminated in fault tolerant designs. The problems associated with transaction failures in midstream, or with system failures, are being solved via recovery algorithms. The most difficult aspect of designing recovery algorithms (Bernstein, Goodman and Hadzilacos 1983) is deciding when to write updated pages into the stable database while also obeying the "Commit Rule" and the "Log-Ahead Rule" (Gray 1979). The impact of failure modes in which system components continue to run, but perform incorrectly, is minimized by redundant design (Pease, Shostak and Lamport 1980).

Knowledge base implementations for the next generation of computers are expected to require distributed and relational databases which will bring into focus some new quality issues:

- Protection against interference among subsets of the distributed data base. This is complicated by the environment of computer networks which have no central operating system. Current research tends toward making transactions atomic without affecting the transaction time.
- Protection against new levels of pattern sensitivity.
- Knowledge base transparency assurance in high structural complexity conditions (Warfield 1979).
- Knowledge maintenance quality controls.

Implementing quality strategies dealing with these new issues will be complicated by the problems of scale, speed, and system complexity. Methods for validating a system's knowledge base and evaluating its performance for the management of maintenance and refinement, are relatively difficult tasks. But the first successful steps in this direction have already been made (Politakis and Weiss 1984) by systems providing e.g. prespecified control strategy, production rules formalism and by tools for sensitivity analysis, execution tracing and explanation.

4. PROBLEM SOLVING: THE QUALITY OF KNOWLEDGE PROCESSING

The static definitions of intelligence, from Homer (who called it a gift of grace (The Odyssey, Book 8)) to the modern psychologist's concept of intelligence being "the sum of the attributes of a prototypically intelligent person" or "what-an-intelligence-test- measures", are consensus-dependent and offer no help in practical Knowledge Engineering. Earl Hunt (Hunt 1983) identifies intelligence with the individual's mental ability and performance in

- choosing an internal represenation for a problem,
- strategies for manipulating the representation, and
- executing elementary information processing steps.

Hunt's concept is more in line with the needs of AI which is interested in solutions to complex information processing problems (Marr 1977). It takes into account the abstract formulation of "what" is being computed. The "why" and "how" specify the set of particular algorithms for implementing given computations.

In our review of quality issues, the "what" is correlated with knowledge, its representation and its storage in knowledge bases. The "why" (being domain defined) is deliberately excluded here because of the unmanageable complexity the sum of unique applications would present. The set of knowledge processing algorithms (the "how") that are generally useful in AI and Knowledge System Engineering will be addressed in a simple empirical sequence of dominant methods. This is because the categorical structure of thinking process models did not stabilize adequately. This will also avoid unresolved differences in viewpoints that these algorithms should model reasoning according to well worked out mathematical logic languages or imitate the way the human mind works, which is almost certainly not the way of mathematical logic (Kolata 1982).

4.1 Search Processes

Independent of their basic importance, search algorithms in Artificial Intelligence and Knowledge Systems Engineering (Barr and Feigenbaum 1981) evidently do not present any unique quality problem. They are not different from any other algorithm implemented in computer programs. The basic question here relates to the implementation correctness and its demonstrated verification. Care must be taken for strategies that assign numerical values to tree nodes using an evaluation function to allow ordering and assuring the admissability of heuristic techniques (Williams 1981).

Some new challenges may be found in implementing associative searches, but even there, the main quality issue will be correctness.

4.2 Deductions

The use of predicate logic to represent declarative knowledge and the use of resolution refutation to mechanize deductions (Nilsson 1980) is now firmly established and, in a similar way to search processes, does not present quality problems beyond assuring their correctness. Again, more design care must be taken while working with nonstandard logics, during transitions from first-order to higher order logic, or in logics that allow intensional operators (Cohen and Feigenbaum 1982, Weber and Nilsson 1981).

Deduction is basically a chain of statements, each of which is either a premise or something which follows from a statement occurring earlier in the chain, leading to the problem's solution. The knowledge system quality requirements lead to implementation which must allow for deduction's visibility, ease of inspection, backward traceability, etc., to verify its correctness and improve solution credibility.

4.3 Inductive Inferences And Learning

The most important factor affecting the effectiveness of inductive inferences and learning is the level and quality of information available. The quality of its representational forms in terms of fitness for use is describable in four important dimensions: expressiveness, ease of inference, modifiabililty, and extendability (Cohen and Feigenbaum 1982). Quality aspects of the learning system's performance element can be associated with its ability to diagnose incorrect rules and to avoid their integration into the system.

The learning system's overall performance can be measured by the rate of knowledge compilation, especially by the rate of composition of new productions and the rate of "proceduralization" (Marr 1977). Overall system quality is a function of the application level of meta-rules:

A. Content referencing (Davis 1980), as a strategy of applying meta-rules to reasoning about rules, will:

- avoid saturation, in which so many sources are potentially useful that it is unrealistic to consider unguided, exhaustive invocation cycles.
- improve the degree of system flexibility in responding to changes, especially in programs with large, constantly evolving knowledge bases. Flexibility will also allow changes in the knowledge base between performance runs, a feature called "Compile Time Flexibility."
- improve system credibility, which is a measure of its ability to accurately reflect a knowledge source's content.

B. Control and reasoning about invocations (Davis 1980) which brings advantages to traditional problem solving paradigms, including means-ends analysis, resolution, heuristic search, problem reduction, etc.

The quality of the learning system is also reflected in its

- generality (the ability to perform succesfully in novel situations) which is especially important in language acquisition.
- ability to discriminate between critical and peripheral information.
- ability to restrict the range of applications of productions to only appropriate circumstances.
- ability to control "interestingness" of inferences (Schank 1979) to comply with principles of significance and goal satisfaction.

However, implementing these quality attributes is not easy. Design and Quality Assurance challenges grow by an order of magnitude when we depart from categorical (or deterministic) reasoning and turn to probabilistic reasoning (Szolovits and Pauker 1978) with all its trappings of controversies about psychological probabilities (Cohen 1973) and documented failures of man as an intuitive statistician. Or, when we turn to inferences based on social judgement (Nisbett and Ross 1980), which are plagued with errors of insufficient evidence, misleading vividness of information, illusory correlations, misguided parsimony, circumstantial promptings, belief perseverance in the face of evidence, self-serving biases, prejudice, etc. In all these situations, unrecognized side effects of applied algorithms may also adversely impact the actual and perceived quality of knowledge processing.

4.4 Programming Languages

Programming languages have always played a central role in AI and have served two important purposes (Barr and Feigenbaum 1982):

- to provide convenient implementation of programs to test and demonstrate ideas.
- to act as vehicles of thought; higher level languages allow communication of higher level concepts.

The Quality Assurance approach to development or selection of programming languages is traditionally very pragmatic. Quality requirements usually concentrate on

- ease of program design, coding, debugging, verification, understanding, learning, and maintenance.
- uniformity of error handling, automatic error correction, entry into debugging facilities, handling of special error conditions by user functions, etc.
- automatic backtracking facilities that allow reasoning control.
- minimizing differences in styles between conventional and knowledge systems programming (Sammet 1969).
- machine and compiler independence.
- level of standardization.

This traditional approach is not powerful enough to contribute to the resolution of debates e.g. LISP vs PROLOG (Nii 1983).

A similar situation exists in evaluating database languages (Stamen and Costello 1981) where quality attributes usually are

- the level of nonprocedurality,
- generality,
- completeness,
- consistency and
- human factors.

Applying natural languages will enrich the scope of quality problems by demonstrating the need to work with Fuzzy Information and to facilitate certain types of reasoning with Fuzzy Sets (LaFaivre 1977).

5. QUALITY OF THE HUMAN INTERFACE

As the user community expands beyond Data Processing professionals to include novices with very little technical training, the quality of the human interface and the level of human factors influencing its design will play a larger and larger role (Schneiderman 1982). Design strategies must go beyond traditional hardware human engineering, ergonomic regulations, and intuitive design of system friendliness. They must address the physical/physiological (e.g. work posture, character size, noise, etc.) as well as the job's intellectual aspects (the worker's role, habit patterns, job frustration levels, human memory requirements, etc. (Fried 1982)). Solutions must be based not only on anthropometry but also developed in the context of the system's lexical, syntactic and semantic requirements (Foley 1982). Both the designer and the psychologist see an opportunity to improve current approaches to complex human problem solving and human-computer-human interaction and communication (Card, Morgan and Newell 1983). There is a growing recognition of the need for serious scientific experimentation (Grimes 1984) and the number of system usability labs is increasing. The documented frustration level of designers who attempt to improve a computer systems' human engineering aspects by altering its surface structure after all the fundamental architectural decisions have been made (Branscomb and Thomas 1983) often leads to architectural guidelines for

- separation of the user interface,
- layered interfaces,
- input/output media translatability,
- hooks for behavior observation and monitoring,
- synonymity,
- adaptability to changing user needs and
- consideration of the limits of basic human capabilities.

In assessing the current state-of-the-art by looking e.g. at the results of the March 1982 National Bureau of Standards Conference in Gaithersburg, MD (900 attendees, 200 papers presented) or other human factors society conferences, it is evident that further research in cognitive science, computer science, and computational linquistics is needed.

6. HARDWARE QUALITY

Even after considering that the new generation of computers will be much faster, more complex, and will have much more computing power, we assume that the traditional hardware quality issues (reliability, safety, environmental ruggedness, ergonomics, electromagnetic compatibility, availability, and serviceability) will remain substantially unchanged.

6.1 Parallelism

New strategies will be required to deal with the expected scope of parallelism (Wallich 1983) and depth of pipelining. The answers will come from the Fault-tolerant Design which addresses

- execution tracking and monitoring,
- recovery, restart, and reconfiguration,
- processor switching control reliability, and
- error detection and correction.

System testability and diagnosis will be a very complicated issue because in highly parallel processing environments even a small level of asynchronism might make accurate repetitions of a computational task impossible.

6.2 Very Large Scale Integration (VLSI)

Expected levels of VLSI complexity and density will reemphasize difficulties already encountered in assuring the chip design quality and verification, and the quality of processing materials and manufacturing processes. The development of high capacity and heat-dissipation packaging technologies will play a fundamental role in assuring basic reliability and serviceability.

Fabricating devices with submicron geometries on a wafer-scale integration level will stimulate expansion (and sometimes breakthroughs) in a host of supporting assurance technologies, such as dimensional controls, failure analysis, material analyses, automated inspections, and Statistical Quality Control methods.

6.3 New Architectures

Quality issues associated with new, non-von Neumann architectures remain mostly undefined.

7. KNOWLEDGE SYSTEMS QUALITY

The technical aspects of quality attributes on all system levels, knowledge levels (Newell 1982) or paradigms (Nii 1983) must reflect the users' requirements and expectations that are associated with the complete knowledge system. Here, the domain specialist and user will use the traditional language of fitness for use, correctness, system reliability and maintainability.

7.1 Fitness For Use

This classical view of quality, established clearly by Dr. Juran (Juran, Gryna and Bingham Jr. 1974), helps distinguish this quality attribute from other objectively based notions of performance, reliability, etc.

and subjectively based customer satisfaction. It also helps to relate system quality to its value and competence.

The value of a knowledge system, especially its social value, reflects the user's level of understanding and the structure of the particular information processing problem. It does not reflect the mechanism through which it is implemented (Marr 1977). We might think about systems which unenlightenedly mimic some aspects of human mental performance as having low social value. We might also think about systems which advance our understanding and reflect results of a genuine exploration as having high value. Weizenbaum (Weizenbaum 1965) now judges his ELIZA program to belong to the low value category.

D. Marr (Marr 1977) also reminds us of Chomsky's (Chomsky 1965) useful notion of system "competence." Competence of an AI system consists of its ability to

- isolate a particular information-processing problem,
- formulate a computational theory for it,
- construct an algorithm that implements it, and
- practically demonstrate that the algorithm is successful.

When competence has been established, it never requires reestablishment. This notion of system competence, especially the judgements applied for deciding if competence is adequate, can be used to determine the system's fitness for use. Adequacy of the system's computational, reasoning and information storage resources can be a part of this assessment.

7.2 Correctness

Correctness of the information processing problem's solution (its freedom from errors) is a quality attribute that is as fundamental to information processing as notions of accuracy and precision are for physical measurement quality.

Automatic deduction (mechanical theorem proving), which has been a major concern of AI since its earliest days (Cohen and Feigenbaum 1982), will remain as an ideal for tools that assure quality of the knowledge system, its implementation, and results. But the road to an ideal is neither short nor easy. In many situations, we must be satisfied with plain credibility of inferences or even assessed credibility of final results. This is especially true in situations where new, previously unknown results have been reached, or where situation complexity or its real-time aspects prohibits repetition. The most difficult challenges are presented by problems modelled by a complicated set of feedback loops or interactions among many stochastic processes that generate counter-intuitive results. The known verification difficulties in current expert systems (Duda and Shortliffe 1983) with discoveries of incomplete knowledge, problems with rules, surprising reasoning sequences, etc. give us an early indication of things to come.

7.3 Reliability And Maintainability

The known problems of hardware reliability, availability, and serviceability will have new solutions based on economically affordable, VLSI-implemented, fault-tolerance. Automatic reconfigurations in highly parallel systems will be a necessity.

A new set of knowledge system maintenance issues is associated with the programming paradigm as it relates to the application domain. There is a need for knowledge-base maintenance which is especially difficult if real-time upgrades are necessary and frequent. Maintenance of the production system might be simplified by the presence of meta-knowledge or in learning systems. But many basic issues of software maintenance (which are making our life interesting now) will remain (Martin and McClure 1982).

7.4 System Support

Similar to the current experience, system support requirements are expected to be functions of

- the level of user involvement in system definition and design,
- the competence level of knowledge engineering and human engineering applied during the design,
- the volume and quality of the "help" information residing in the system and available in user documentation, and
- the quality and scope of the initial user training.

8. SUMMARY

This introductory and elementary review of quality issues generated by the advancement of Artificial Intelligence and its commercialization via knowledge systems partially indicates the size and complexity of the task the design and quality assurance communities in the computer industry face. Successful development of new quality strategies addressing these new issues plus a creative application of current assurance expertise, technology, and experience will help the quality specialist participate in, and contribute to, the advancement of computer technology. It will also improve the probability of the success of individual products.

The current conceptual framework of Quality Assurance seems generally adequate to address the needs of knowledge systems design, manufacture and use. New technologies, new architectures, new system parts and functions will, however, require intensive research and applied engineering work. Some breakthroughs are expected in scientific instrumentation that will support progress in VLSI manufacturing, software quality assurance, fault-tolerance of highly parallel systems and in the Quality Assurance of knowledge .

The expected partnership between knowledge system design teams and Quality Assurance teams should be beneficial. Just as it has proven beneficial in other high-technology environments (Nilsson 1980, Politakis and Weiss 1984, Kohoutek 1983 and 1984).

ACKNOWLEDGEMENTS

This work benefitted highly from the suggestions of reviewers and from the editorial expertise of Doug Newlin of Hewlett-Packard's Fort Collins Systems Division. To him, our thanks for his significant improvements in the readability and writing style of this paper.

REFERENCES

Anderson, J.R.
The Architecture of Cognition
Harvard Univ. Press, Cambridge, MA 1983

Barr, A. and Feigenbaum, E. (ed.)
Handbook of Artificial Intelligence
Vol. I, Chapt. III.
Heuristech Press, Stanford, CA 1981

Barr, A. and Fiegenbaum, E. (ed.)
Handbook of Artificial Intelligence
Vol. II, Chapt. VI.
Heuristech Press, Stanford, CA 1982

Bernstein, P., Goodman, N. and Hadzilacos, V.
Recovery Algorithms for Data Base Systems
Proceedings of 1983 IFIP Congress
Elsevier Science Publishers B.V. (N. Holland)
Amsterdam, 1983 pp.799-807

Brachman, R.J. and Smith, B.C.
Special Issue on Knowledge Representation
SIGART Newsletter, No. 70. (1980), pp.1-138

Brady, E.L. and Wallenstein, M.B.
National Standard Reference Data Series, No. 1.
National Bureau of Standards, 1964

Branscomb, L. M. and Thomas, J. C.
Ease of Use: A System Design Challenge
Proceedings of the 1983 IFIP Congress
Elsevier Science Publishers B.V. (N. Holland) Amsterdam, 1983 pp.431-438

Card, S., Morgan, T. and Newell, A.
The Psychology of Human-Computer Interaction
Lea Publishers, Hillsdale, NJ 1983

Chomsky, A.N.
Aspects of the Theory of Syntax
MIT Press, Cambridge, MA (1965)

Cohen, J.
Psychological Probabilities
Schenkman Publishing, Cambridge, MA 1973

Cohen, P. and Feigenbaum, E. (ed.)
Handbook of Artificial Intelligence
Vol. III, Chapt. XII
Heuristech Press, Stanford, CA, 1982

Cohen, P. and Feigenbaum, E. (ed.)
Handbook of Artificial Intelligence
Vol. III, Chapt. XIV
Heuristech Press, Stanford, CA, 1982

Davis, R.
Content Reference: Reasoning About Rules
Artificial Intelligence, Vol. 15 No. 3 (1980), pp.223-239

Davis, R.
Meta-Rules: Reasoning About Control
Artificial Intelligence, Vol. 15 No. 3 (1980), pp.179-222

Duda, R. O. and Shortliffe, E. H.
Expert Systems Research
Science, Vol. 220, No. 4594, 15 April 1983, pp.261-268

Foley, J. and vanDam, A.
Fundamentals of Interactive Computer Graphics
Addison-Wesley Pub. Co., Reading, MA 1982

Fried, L.
Nine Principles for Ergonomic Software
Datamation, Nov. 1982, pp.163-166

Gray, J.N.
Notes on Database Operating Systems: An Advanced Course
Springer-Verlag, New York, NY 1979

Gregory, R.L.
Mind in Science
Cambridge Univ. Press, Cambridge, U.K. 1981

Grimes, B.
The Difficult Path to the Easy-to-Use Computer
Business Week, Feb. 27, 1984 pp.93-94

Hayes-Roth, F.
An Industry Emerges: Knowledge Systems
Presentation at TTI Seminar
Wash. DC, Nov. 7, 1983

Hintikka, J.
Knowledge and Belief
Cornell Univ. Press, Ithaca, NY 1962

Homer
The Odyssey, Book 8, Lines 167-168

Hunt, E.
On the Nature of Intelligence
Science, Vol. 219, 14 Jan. 1983, pp.141-146

Juran, J., Gryna, F. and Bingham, R.
Quality Control Handbook, 3rd Ed.
McGraw-Hill Book Co., NY 1974

Kassarjian, H. and Robertson, T.
Perspectives in Consumer Behavior, 3rd Ed.
Scott, Foresman & Company, Glenview, IL 1981

Kohoutek, H. J.
R&D and Quality Assurance Partnership
IEEE Transactions on Engineering Management
Vol. EM-30, No.4, Nov. 1983, pp.200-204

Kohoutek, H. J.
Effective Quality Management in a High Technology Program Environment
World Quality Congress 1984, Proceedings Vol. 1 pp.517-523, Institute of Quality Assurance, London

Kolata, G.
How Can Computers Get Common Sense?
Science, Vol. 217, 24 Sept. 1982, pp.1237-1238

Korf, R.E.
Toward a Model of Representation Changes
Artificial Intelligence Vol. 14 No. 1 (1980), pp.41-78

La Faivre, R.A.
FUZZY: Reference Manual
Computer Science Dept., Rutgers Univ. 1977

Levi, I.
The Enterprise of Knowledge
MIT Press, Cambridge, MA 1980

Lide, D. Jr.
Critical Data for Critical Needs
Science, Vol. 212, No. 4501, 19 June 1981
pp.1343-1349

Marr, D.
Artificial Intelligence - A Personal View
Artificial Intelligence Vol. 9 No. 1 (1977), pp.37-48

Martin, J. and McClure, C.
Maintenance of Computer Programming
Savant Research Studies, Carnforth, U.K., 1982

McCarthy, J. and Hayes, P.J.
Some Philosophical Problems from the Standpoint of Artificial Intelligence
Readings in Artificial Intelligence,
Tioga Publishing Co., Palo Alto, CA, 1981
pp.431-450

Monroe, K. and Petroshius, S.
Buyers' Perception of Price: An Update Of The Evidence
Journal of Marketing Research 10
(Feb. 1973) pp.70-80

Newell, A.
The Knowledge Level
Artificial Intelligence, Vol. 18 No. 1 (1982), pp.87-127

Nii, H. P.
Other Options
Presentation on TTI Seminar, Wash. DC, Nov. 8, 1983

Nilsson, N. J.
Principles of Artificial Intelligence, Chapt. 5.
Tioga Publishing Co., Palo Alto, CA, 1980

Nisbett, R. and Ross, L.
Human Inference : Strategies and Shortcomings of Social Judgement
Prentice-Hall, Inc., Englewood Cliffs, NJ 1980

Page, A.
Utility Theory: A Book of Readings
John Wiley & Sons, Inc., NY 1968

Pease, M., Shostak, R. and Lamport, L.
Reaching Agreement in the Presence of Faults
JACM, Vol. 27, No. 2, 1980, pp.228-234

Politakis, P. and Weiss, S. M.
Using Empirical Analysis to Refine Expert System Knowledge Bases
Artificial Intelligence, Vol. 22 No. 1 (1984), pp.23-48

Sammet, J. E.
Programming Languages : History and Fundamentals
Prentice-Hall, Inc., Englewood Cliffs, NJ 1969

Schank, R.
Interestingness: Controlling Inferences
Artificial Intelligence, Vol. 12 No. 3 (1979), pp.273-297

Shneiderman, B.
How to design with the user in mind.
Datamation, Apr. 1982, pp.125-126

Stamen, J. and Costello, W.
Evaluating Database Languages
Datamation Vol. 27 (May 1981), pp.116-122

Szolovits, P. and Pauker, S.
Categorical and Probabilistic Reasoning in Medical Diagnosis
Artificial Intelligence, Vol. 11 No. 1,2 (1978) pp.115-144

Verity, J. W.
PROLOG vs. LISP
Datamation, Jan. 1984, pp.50-54

Wallich, P.
Designing the next generation
IEEE Spectrum, Nov. 1983, pp.73-109

Warfield, J.
Some Principles of Knowledge Organization
IEEE Transactions on Systems, Man and Cybernetics,
Vol. SMC-9, No. 6, June 1979, pp.317-325

Weber, B. and Nilsson, N. (ed.)
Readings in Artificial Intelligence, Chapt. 2.
Tioga Publishing Co., Palo Alto, CA., 1981

Weizenbaum, J.
ELIZA–A computer Program for the Study of Natural Language Communication Between Man and Machine
Communications of ACM, Vol. 9 No. 1 (1965) pp.36-45

Wiener P., (ed.)
Dictionary of the History of Ideas
Vol. IV pp.450-458
Charles Scribner's Sons, NY 1973

Wilks, Y.A.
Methodological Questions About Artificial Intelligence: Approaches to Understanding Natural Language
Journal of Pragmatics, Vol. 1 (1977), pp.69-84

Williams, G.
Tree Searching, Part 2: Heuristic Techniques
Byte, Oct. 1981, pp.195-210

Ziman, J.
Reliable Knowledge
Cambridge Univ. Press, Cambridge, U.K. 1978