Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Author's personal copy

Available online at www.sciencedirect.com

SciVerse ScienceDirect

Int. J. Human-Computer Studies 71 (2013) 135–156

International Journal of Human-Computer Studies

www.elsevier.com/locate/ijhcs

Knowledge acquisition: Past, present and future ☆

Brian R. Gaines[a,b,*]

[a] University of Calgary, Canada
[b] University of Victoria, Canada

Accepted 25 September 2012
Available online 17 October 2012

Abstract

As we celebrate the 50th knowledge acquisition conference this year it is appropriate to review progress in knowledge acquisition techniques not only over the quarter century since the conference series began but backwards through the millennia to the beginnings of knowledge capture and forwards through the foreseeable future to speculate on reasonable expectations, appropriate targets and potential surprises in the next quarter century.

"Time present and time past are both perhaps present in time future, and time future contained in time past. What might have been is an abstraction remaining a perpetual possibility only in world of speculation." (T.S. Eliot, Four Quartets).
© 2012 Elsevier Ltd. All rights reserved.

Keywords: Knowledge acquisition; Human evolution; Information technology evolution; Propagation scholarly errors; Information overload; World Wide Web; Semantic web; Technological forecasting

1. Introduction

Human–computer interaction research is inevitably forward-looking because of the extraordinary pace of improvements in the technology and the expectations that this generates. We know that we will be able to do tomorrow what we barely dreamed of yesterday, and that successful innovations will be assimilated rapidly into the texture of our society to become part of our everyday routines and expectations. Our focus is to transform possibilities in our worlds of speculation from abstractions to effective technologies assimilated in, and enhancing, our lifeworlds.

Occasionally, however, it is useful to look back, assess how far we have come, recognize how much was unexpected, and use this to refine our projections, noting that the future is amazingly open. We need to be able to recognize and utilize what we did not expect as much as we also need to strive particular objectives by particular means.

The uncertainties of the technologies and the openness of the future should not, however, blind us to the underlying human needs and intrinsic human capabilities that have remained constant for millennia. Our research serves a society where individuals, social groups, organizations, nations and civilizations attempt to survive and prosper. Their specific aspirations reflect the ethos of each era but their fundamental needs and unenhanced individual capacities have remained the same throughout recorded history.

The accumulation of techniques for addressing these needs and enhancing these capabilities we term human "knowledge." As a species we are "doomed to think in order to live" (Castañeda, 1990, p. 3). Knowledge acquisition and utilization to support our thinking has been a major imperative of human civilization throughout the millennia. In our era, computer technology and human–computer interaction have come to play a major role in knowledge processes, facilitating a level of knowledge generation, dissemination, access and utilization beyond that we have ever known.

This journal has played a significant role in supporting the development of improved human–computer interaction technologies and has had strong engagement with

☆ This paper is a substantially enhanced version of a keynote address presented at KCAP'2011 (Gaines, 2011).

*Correspondence address: University of Calgary, Canada.
E-mail addresses: gaines@ucalgary.ca, gaines@uvic.ca.

1071-5819/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ijhcs.2012.10.010

Author's personal copy

research in such interaction at the knowledge level. In particular it has been a major forum for research on the support of human knowledge acquisition processes and as we celebrate 25 years of the knowledge acquisition conferences world-wide it is appropriate to reflect not only on the progress over that period but also on how that progress is situated in the growth of human knowledge over much longer time scales.

This paper traces human knowledge acquisition processes, and techniques for their support, from time long past through time recent and time present to speculations on time future.

2. Knowledge in human evolution

Where does one start in situating knowledge acquisition processes? We characterize our era as *the information age*, but there are many different perspectives on when that era began. Castells (1996) in his monumental studies of the economic, social and cultural aspects of the information age dates it to the rise of the *network society* based on computer technologies. Some date it later from the invention of the transistor (Riordan and Hoddeson, 1997), whilst others date it back to 1800 and the spread of industrialization (Beniger, 1986), to 1700 and the dissemination and organization of the knowledge arising from of the scientific revolution (Headrick, 2000), or trace it through two and half millennia from the shift from oral to written communication through the development of printing to the growth of computer networks (Hobart and Schiffman, 1998).

Along that path there are major milestones such as the Greek enlightenment, its continuing impact through various 'renaissances' as Greek philosophy, mathematics and science was rediscovered in various civilizations at different times, the growth of monasteries and universities as locations for the preservation and dissemination of knowledge, the scientific revolution, the age of enlightenment and the industrial revolution that culminated in our present age when the massive accumulation of knowledge acquired along the way is being digitized to make it accessible in computational form on a scale never previously possible.

The grandest sweep of all is to date the growth of information and knowledge in the universe back to the big bang and see it as part of a continuous evolutionary process of structure formation in the physical universe. Ayres (1994), a well-respected technological forecaster, wrote a remarkable book, *Information, Entropy and Progress: A New Evolutionary Paradigm* that provides a coherent systemic model of physical, geological, biological, social, cultural, psychological and economic evolution, and models skilled activity such as manufacturing as an information process that, for example, creates an automobile by imposing information on matter.

If we conceive of knowledge abstractly as the information we impute to a system to account for its behavior (Clancey, 1990) then Ayres' framework shows knowledge processes playing a far wider role than any we normally envision. If we characterize living systems abstractly as *autopoietic* (Maturana, 1975) in actively creating conditions for their own persistence, then Ayres' informational formulation allows us to model the fundamental processes of life as being those of knowledge creation, capture and transmission.

Cybernetic/systemic models of such broad scope are fascinating and inspiring but perhaps too remote to have a direct impact on the diverse disciplines they encompass. However, in the past 20 years advances in molecular biology have made DNA sequencing technologies available to archeologists and anthropologists, and enabled information-flow models to be used to expose not just the systemic commonalities but also the mutual constraints coupling genetic, cultural and behavioral processes in living systems. Oyama's (1985) *Ontogeny of Information* is arguably the first such analysis to become widely influential through the *developmental systems theory* community. Jablonka and Lamb's (2005) *Evolution in Four Dimensions* provides a unified model of the transmission of variation between living systems encompassing genetic, epigenetic and behavioral sub-systems and their interactions.

From a knowledge acquisition perspective, we can see such unified models as providing a detailed account of how:

- genomes adapt to the environment through random mutation, encoding and propagating information that may enhance the fitness of future generations (Altenberg, 1995);
- epigenetic processes manage the expression of particular capabilities encoded in the genome 'library' to more rapidly propagate adaptations to major environmental change (Harper, 2005);
- behavioral adaptations are propagated through reinforcement and mimicry, both intrinsically and through pedagogy (Thornton and Raihani, 2010);
- symbolic representations of the information involved in all these processes may be used to facilitate them, amplify their effect, and enable them to be widely diffused through both space and time (Noble and Davidson, 1996).

The exchange of information between all levels and partitions of living systems provides a common framework for biological symbiosis, psychological foundations of socio-cultural systems and, through the symbolic signaling system of 'money,' for economic models of those systems.

Physicists have set a not unrealistic target of a unified *theory of everything* in the physical sciences, but those facing the complexities of the biological and human sciences have felt it foolish even to dream of such for their disciplines. However, quite suddenly, as an outcome of advances in molecular biology and the human genome project, such unification is occurring without it ever having been an envisioned target.

3. Evolution of knowledge acquisition, dissemination and utilization

Ayres evolutionary model leads through the processes of formation of stars and planets, to those of complex molecules

Author's personal copy

and living systems, and multidimensional biological models characterize the evolutionary processes of various forms of life including the human beings. Archaeological science has made major advances through the deployment of modern technologies such as remote sensing (Wiseman and El-Baz, 2007), computer modeling (Lock, 2003), and biogenetics (Renfrew and Boyle, 2000), and this has resulted in much greater understanding of the evolution of the brain, tools, language and mind (Gibson and Ingold, 1993; Noble and Davidson, 1996; Smail, 2008).

Our species, *Homo sapiens sapiens*, diverged from *Homo erectus* some 500,000 years ago, from *Homo sapiens neanderthalis* some 300,000 years ago, developed some form of language some 50,000 years ago, was reduced by environmental catastrophe to a population of some 3000 in Africa some 50,000 years ago, and through migration commencing in the Levant expanded worldwide, developing community infrastructures and agriculture some 10,000 years ago, commencing the Neolithic era of modern humanity. The details are contested in a massive research literature, but the overall framework is widely accepted (McBrearty and Brooks, 2000; Stringer, 2002; Liu et al., 2006; Endicott et al., 2010).

For most of our history, genetic, epigenetic and behavioral processes dominated our evolution as they do in other animal species, but at some time in the past 100,000 years information came to be communicated and captured symbolically to an extent that gradually came to differentiate us from other species—"humans became behaviourally modern when they could reliably transmit accumulated informational capital to the next generation, and transmit it with sufficient precision for innovations to be preserved and accumulated." (Sterelny, 2011, p. 809).

The capability to capture and transmit the knowledge created by individuals and communities is generally taken in the archeological and anthropological literatures to be the major factor in the explosion of the human population. Whereas the rate of unconstrained population growth in other species is proportional to the population size, and hence exponential, for the human species it is proportional to the square of the population, and hence hyper-exponential—until the 1960s when the population growth rate dramatically declined (Korotayev, 2005). The additional multiplier is attributed to the generation and diffusion of knowledge being proportional to the size of the population (Korotayev, 2005). It is not clear how the decline in population growth in the past 50 years will affect knowledge growth. For example, whether knowledge growth is being sustained by the substitution of computer-based knowledge processing for the equivalent human activity as suggested by the human–computer symbiosis models of Martin (2000), Kurzweil (2005) and others.

Human population growth does not show a smooth progression over recorded history. There have been major die-offs due to climatic factors such as the ice ages, and diseases such as the Black Death, but the overall trend has been hyper-exponential. One can discern a pattern of trends encouraging the generation and diffusion of knowledge, such as the development of communities around population centres, which also increase the risk to life, for example, by facilitating the development and spread of disease (McNeill, 1989; Cantor, 2001) requiring the further development of knowledge.

Language and knowledge are not intrinsically 'survival traits.' Bickerton (1990) notes that one possible outcome of the power of intelligence is species destruction. Wojciechowski (2001) models the growth of knowledge as process whereby more knowledge must be continuously created to combat the adverse side-effects of the application of prior knowledge. He develops 25 *laws of knowledge*, for example:

Law 1: The number and variety of causes of stress are proportional to the amount of knowledge.
Law 2: The perception of the complexity of the consequences of knowledge is proportional to the development of knowledge.
Law 3: The knowledge of knowledge is a function of the general development of knowledge.
Law 4: The size and complexity of the problematic of knowledge is proportional to the general development of knowledge.
Law 5: Thought induces change.
Law 6: Humans' ability to determine the development of humanity is proportional to their knowledge.

The 'laws' provide a general phenomenology of the knowledge construct as an autopoietic, self-perpetuating, system from a perspective similar to that of the "selfish meme" (Distin, 2005) and situates knowledge together with bacteria such as e-coli as essential symbionts parasitic upon the human species. They provide substance for Musil's aphorism in *The Man Without Qualities*: "the compulsion to know is a mania, just like dipsomania, erotomania, homicidal mania: it produces a character out of balance. It is not at all true that a scientist goes after truth. It goes after him. It is something he suffers from" (Musil, 1930, p. 254), and for Dostoyevsky's *Underground Man*: "Leave us to ourselves, without our books, and at once we get into a muddle and lose our way—we don't know whose side to be on or where to give our allegiance, what to love and what to hate, what to respect and what to despise" (Dostoyevsky, 1864, p. 123).

A complementary socio-economic perspective is provided in Snooks' (1996, 1998) monumental series of books on the laws of global economic history which identify the major strategies through which societies acquire resources, notably through *population growth*, *conquest*, *commerce* and *technology*. All four strategies are essentially knowledge intensive, and Snooks' models of the cycles of strategies adopted in ancient and modern civilizations enable one to trace through the ages the knowledge processes involved that support medicine, warfare, commerce and science/technology. He also proposes a series of

Author's personal copy

laws characterizing the dynamics of civilizations, for example, "the law of cumulative technological change that the relationship between a series of technological paradigm shifts is geometric owing to the accelerator effect when the output of one paradigm becomes the input of the next." Section 7 explores this phenomenon in the evolution of information technologies.

4. Early knowledge capture

When did processes of knowledge capture and transmission originate? The major problem with answering this question is that the media used for knowledge capture have limited lifetimes, and often do not survive decades let alone millennia (Diringer, 1982). Archaeologists are left with a highly biased sample of the few originals that survived, and historians with the residues of the transcription and copying processes that have attempted to preserve the content as the medium decays. That situation continues in our era as our computer media all have short life expectancies and rely on continuing backup processes for the preservation of their content. However, effective digitization procedures can now guard against transcription errors and ensure exact copying (Gladney, 2007).

Evidence of the transmission of accumulated informational capital in behaviorally modern humans used to be associated with the post-glacial Neolithic era commencing some 10,000 years ago. However, advances in archaeology are continually pushing back that date (Mellars et al., 2007) and, while our current civilization may have its roots in Neolithic Mesopotamia, there is now evidence of many earlier civilizations that developed but died out leaving a small residue of cultural artefacts (Stringer, 2006). Recent studies have dated some cave paintings as being over 40,000 years old (Pike et al., 2012). An outstanding example of early knowledge capture is that of the detailed paintings of animals in the Chauvet Caves illustrated in Fig. 1 that may possibly be dated to Aurignacian culture of the Upper Paleolithic some 30,000 years ago (Pettitt, 2008).

While such images are commonly analyzed as 'art' (Lewis-Williams, 2002), they may equally well be construed as knowledge representation for purposes such as education, and their location suggests they were intended to be preserved, as they have been, very successfully. Mellars (2004, p. 461) takes these and other representational materials as indicating that "the Aurignacian period shows an apparently sudden flowering of all the most distinctive features of fully 'modern' cultural behavior," and that "communication and expression at this level of complexity would be almost inconceivable in the absence of complex language systems and in the absence of brains structured very similarly, if not identically, to our own."

The earliest examples of symbolic knowledge capture where we have a substantial body of material is Babylonian cuneiform impressed on clay tablets from some 2000 to 5000 years ago (Robson, 2008). Modern scholarship has decoded many tablets which originated to keep track of trade transactions and inventories (Nissen et al., 1993) and

Fig. 1. Multimedia knowledge capture at Grotto Chauvet some 30,000 years ago.

were repurposed to capture mathematical and military procedures (Neugebauer et al., 1945; Melville and Melville, 2008). We can also see the beginnings of scientific data collection and modeling in the Babylonian materials where astronomical and weather phenomena are tracked and used to predict political and economic events (Swerdlow, 1998), possibly with some partial success in both cases since the weather affects harvests and prosperity which in turn affects the popularity of rulers.

There was probably some diffusion of Babylonian knowledge into later Greek astronomy but overall the outcome appears to be what Burnet (1920) in his comments on early Greek science terms one of the *periodical bankruptcies of science*. In this respect knowledge evolution parallels biological evolution in that most innovations end in failure and only a few propagate to become assimilated into the 'memome' of science.

Fig. 2 shows BM85194, a well-preserved tablet that has been studied in depth. It gives solutions to several mathematical problems such as finding the length of a chord of a circle from its diameter of the circle and the distance of the middle of the chord from the circumference (Høyrup, 1999), and military problems such as the construction of earthwork ramps to enter a city under siege (Melville and Melville, 2008).

There are strong parallels between the Babylonian development of cuneiform writing and later developments of knowledge capture technologies, including that of computers. What is common is the addressing of timeless human needs with the best available technology of each era:

- The environmental stress of warfare was addressed with cuneiform tablets detailing siege techniques—the first digital computers were developed under the stress of the Second World War for the purposes of code breaking and ballistics calculations.
- The cuneiform tablets supported administrative record keeping—IBM adopted computer technology postwar to enhance its existing card-based census and business record-keeping systems.

Author's personal copy

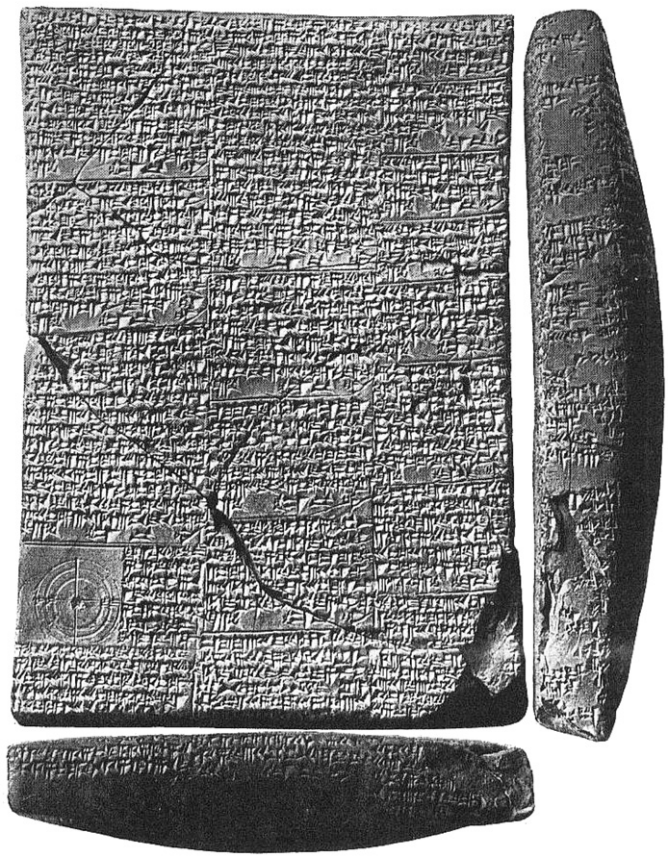

Fig. 2. Mathematical and military knowledge capture.

- Cuneiform tablets captured the surprisingly sophisticated mathematical algorithms of that era—computers make operational those of our era.

And so on—the most powerful approach to technological forecasting is to identify the primary social needs of an era and assume that major social resources are being applied to develop and apply effective technologies to address them (Gilfillan, 1937), and the four major strategies for acquiring resources that Snooks (1996, 1998) has analyzed characterize the needs that have been common to every human civilization.

5. The Greek enlightenment

In the era immediately after Babylonian innovations in knowledge capture we find civilizations in India and China making major advances in mathematical, scientific, medical and legal knowledge (Katz and Imhausen, 2007) and capturing them in a variety of scripts on a range of media such as animal hides (Diringer, 1982). The developments that had most impact on western civilization were those of the Greek enlightenment some 2500 years ago when Euclidean geometry, Socratic dialectic, Platonic philosophy and Aristotelian logic, metaphysics, science and ethics provided the foundations of modern logical, mathematical, scientific, medical, ethical and legal systems (Russo, 2004).

Early Greek civilization captured knowledge that had been created primarily in the brains of people and propagated it through an oral tradition that probably extends back at least 100 millennia but cannot be tracked because it left no record other than brief historical accounts in the later written record. However, by the time of Plato knowledge was being captured in written form using an alphabet deriving from an earlier Phoenician script (Powell, 1991) that continued in a variety of forms thereafter, including an Etruscan variant in Rome that constitutes our current Latin alphabet.

The first major library for which we have detailed accounts are those of Aristotle some 2400 years ago, collected despite the sarcastic comments of his peers because he regarded it as important to understand the ideas of others in developing his own (Blum, 1991, Section 2.6). A succession of national leaders also saw the importance of collecting the world's knowledge of their era, forming major libraries such as that of Ptolemy at Alexandria some 2400 years ago where Kallimachos developed techniques of cataloguing and indexing library materials that are similar to those in use today (Blum, 1991).

The preservation of written knowledge was erratic until the invention of printing facilitated the wide dissemination of many copies of major works making it probable that some copies would survive local catastrophes (Eisenstein, 1979). Aristotle's library was passed to three generations of successors but then stored under conditions where much material was severely damaged (Laughlin, 1995). The library at Alexandria was completely destroyed. The Greek knowledge base that provided the intellectual foundations of modern science only survived in substantial part because several later societies attempted to collect and capture it for their own use, notably the Arabic translation movement in Baghdad some 1300 years ago that both captured the material in Arabic and stimulated an industry of making additional copies of the Greek originals for translation purposes (Gutas, 1998).

One can continue the story of knowledge capture and translation, but not within the scope of these few pages—the relevant literature constitutes a substantial part of national library collections. The account above is sufficient to show how major roles now being played by the World Wide Web (web) have their parallels through the ages:

- The web provides a compendium of human knowledge fulfilling the role of the library at Alexander and its later formulations such Diderot's Encyclopédie (Collison, 1966) and Wells′ (1938) World Brain—both of which were seen by their proponents as socially egalitarian and liberating, much as the web is seen today.
- Discussion in the Athenian agora is emulated by mailing lists and interactive blogs where questions and issues may be raised and discussed in a community—some participants also exemplify Sextus Empiricus′ (1933) critical skepticism that provides counter-examples to any established position.

Author's personal copy

- Aristotle's codification of the abstract patterns of knowledge representation and inference underlies the description logic foundations (Baader et al., 2003) of the semantic web (Shadbolt et al., 2006).

6. The growth of knowledge, information overload, authenticity, suppression

The growth of recorded knowledge was already explosive at the time of the Library at Alexandria which was reported to have over 500,000 volumes, of which Kallimachos had catalogued some 133,000 (Parsons, 1952, p. 112). The number was already sufficient to generate notions of information overload, as noted in an admonition from the same era, "of making many books there is no end; and much study is a weariness of the flesh" (Ecclesiastes 12:12).

The copying of books by scribes who translated and changed the text to suit their audiences and purposes also led to problems of authenticity and quality of knowledge and, just as cataloguing had been invented, so was literary scholarship to compare different versions of what was nominally the same text. In the third century Origen is reputed to have written some 6000 books many of which were analyses of Christian doctrine based on many different versions of the same text (Crouzel, 1989). Issues of information overload were raised again such as St. Jerome's remark "Which of us can read all he has written?"

Origen also provides an early example of the suppression of literature. While his tracing of Christian doctrines to their origins to the early Christian literature was welcome, his continuation of his literary analyses to trace them back further to pagan literature was not. He was put to the rack and tortured and, three centuries after his death, his works were declared anathema by Second Ecumenical Council of Constantinople in 553. What survives is still a significant part of modern religious scholarship and also played a significant role in the methodology of the scientific revolution; Newton drew heavily on Origen's textual interpretation methodology in his own biblical scholarship, and Newton's rules for scientific method in the *Principia* of 1687 are a subset of those he developed for biblical interpretation in his *Treatise of Revelation* from the 1670s (Manuel, 1974; Force, 1990).

Knowledge production continued at an increasing rate over the centuries and continued to lead to complaints of information overload in every era. Four centuries ago Barnabe Rych (aka Rich or Riche) whose writings had a major influence on Shakespeare remarked "One of the diseases of this age is the multiplicity of books; they doth so overcharge the worlde, that it is not able to digest the abundance of idle matter that is every day hatched and brought into the worlde that are as divers in their formes, as their Authors be in their faces" (Rich, 1610). The exchange of information through written correspondence showed a similar massive growth; between 1364 and 1410 the Italian merchant, Datini, exchanged some 125 thousand letters with his factors or agents (Jardine, 1996, p. 111). The infrastructure developed for such correspondence supported the international dissemination of knowledge in the scientific revolution (Gingras, 2010), eventually leading to scholarly journals as a means of open correspondence.

In 1665 the first two scholarly journals came into being: the Journal des Sçavans in January 1665 in France, and the Philosophical Transactions of the Royal Society in England. The founding editor of the Transactions, Henry Oldenburg, came to Britain in 1653 as Bremen's ambassador to Oliver Cromwell and remained after the monarchy was restored. He supported himself through "merchandizing knowledge" (Hunter, 1989, p. 249) by selling his correspondence with European scientists to other scientists and then using his links through Hartlib and Boyle with the newly formed Royal Society to formalize this as the *Proceedings of the Royal Society*. Scholarly journal publication grew slowly at first to 10 in 1750 and 100 in 1800, but by 1830 the number of papers being published was such that abstract journals were created to provide a more accessible overview of the growing literature (Price, 1963, p. 9).

In 1827, Faraday apologized to an Italian scientist for not quoting his prior work on the fluidity of sulphur, remarking that: "It is certainly impossible for any person who wishes to devote a portion of his time to chemical experiment, to read all the books and papers that are published in connexion with his pursuit; their number is immense, and the labour of winnowing out the few experimental and theoretical truths which in many of them are embarrassed by a very large proportion of uninteresting matter, of imagination, and of error, is such that most persons who try the experiment are quickly induced to make a selection in their reading, and thus inadvertently, at times, pass by what is really good" (Faraday, 1859, p. 215).

A century later, Bernal, Churchill's wartime scientific advisor, emphasized the growing problem, "The present mode of scientific publication is predominantly through the 33,000 odd scientific journals. It is, as we have already shown, incredibly cumbersome and wasteful and is in danger of breaking down on account of expense" (Bernal, 1939, p. 292). Bush, Roosevelt's scientific advisor, remarked "The difficulty seems to be not so much that we publish unduly in view of the extent and variety of present-day interests, but rather that publication has been extended far beyond our present ability to make real use of the record" (Bush, 1945, p.102).

In our era the Barnabe Rych quotation provides a nice example of the propagation of incorrect information through scholars copying the errors of others rather than checking primary sources—'scholarly virus propagation.' Price (1963, p. 63) quotes Rych, giving no citation but misdating the quotation to 1613, and other authors using the quote have copied this date, sometimes citing Price, but not checking Rych's books for the proper citation.

Author's personal copy

Fortunately, in recent years Rych's books have been made available in digital form and I was able to trace the source and date of the citation with a few hours of research (Fig. 3).

There is a similar situation with the "the compulsion to know is a mania" quotation of Section 7. I first came across it in Judson's (1979, p. 79) *Eighth Day of Creation*, a history of the discovery of the structure of DNA, where he attributes it to Wilkins who attributed it to Kierkegaard. I used it in a paper (Gaines, 1984a), citing this chain of attributions, but later came across it when re-reading Musil's *Man Without Qualities* and have not been able to trace it in Kierkegaard's works. Out of curiosity I searched for it on the web whilst writing this paper and found it used many times with a simple attribution to Kierkegaard, but no citation, presumably all originating from Judson. However, on the *Science Surf* blog in 2007 there is a question about the source, noting that Delbrück had failed to trace it and had offered $50 to anyone who could, and McGrath (2007) responded with the correct attribution to Musil.

The Internet is making scholarship substantially easier, but not effortless; one still has to sift information and trace it back to its primary source, but both information and sources are now much more readily accessible. It is also changing the nature of scholarly publication as the ease of providing additional versions of an article on the web encourages authors to issue new versions with corrections, comments and other enhancements. However, those critiquing or quoting an article need a static entity to reference, and a plethora of versions can create its own confusion. Web material published outside the quality control processes of peer review and critical commentary, sometimes including material intended to mislead, presents even greater problems of evaluation but, again, the availability of multiple sources, commentaries, and other related material, makes it primarily a problem of being able to filter and assess a large volume of information from a critical perspective.

Thus, issues that we currently associate with the Internet and web such as information overload, authenticity of information and suppression of material considered undesirable in particular cultures, have been with us for over two millennia, probably rather longer. Cuneiform material so far decoded does not contain complaints of excessive numbers of tablets, incorrect versions or inappropriate content, but one suspects they may well have existed.

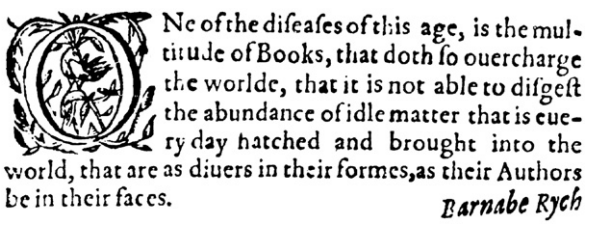

Fig. 3. Quotation whose incorrect dating has propagated in the literature as a scholarly virus.

However, as an example of Wojciechowski's notion that more knowledge may provide solutions to the problems created by knowledge, the digitization of the scientific literature in recent years and its indexing by content make it easier today to trace relevant knowledge and check the veracity of assertions than it was a decade ago.

7. Evolution of information technology

I took up the Killam Research Chair at the University of Calgary in 1985 in order to form a *Knowledge Science Institute* (KSI) with the aim of "facilitating the worldwide process of change to knowledge-based economies" though "multi-disciplinary studies of knowledge science and technology" and "system development and application projects involving knowledge-based systems" (Gaines, 1985). A major activity of the KSI distinct from those reported at the knowledge acquisition conferences was the modeling and forecasting of information technology. We tracked the structural changes in the infrastructure of information technology and developed a model of them as a tiered structure of the learning curves of sub-disciplines built upon the layers below, and applied this to model the past impact of information technology on many economic sectors and industries and project the likely future impact (Gaines and Shaw, 1986; Gaines, 1991a, 1998).

The major impact on knowledge processes in our own era has been through the invention of digital computers in the 1940s and integrated circuits in the 1960s enabling computers to become low-cost, ubiquitous and of increasing capability at a sustained exponential growth rate beyond that of any other technology (Gaines, 1998). As illustrated in Fig. 4, the number of devices on a chip has undergone a continuing exponential growth since 1959 with a doubling period of some 1.5 years (Mollick, 2006; Rupp and Selberherr, 2011), and this rapid sustained quantitative growth over five decades has triggered qualitative structural changes in the nature of the information sciences and their applications.

Fig. 5 shows the overall structure of the model from (Gaines, 1998) extended to 2012. The underlying learning curve for each tier may be characterized by six phases:

1 The era before the learning curve takes off, when too little is known for planned progress, is that of the *inventor* having very little chance of success but continuing a search based on intuition and faith.
2 Sooner or later some inventor makes a *breakthrough* and very rapidly his or her work is *replicated* at research institutions worldwide.
3 The experience gained in this way leads to *empirical* design rules with very little foundation except previous successes and failures.
4 As enough empirical experience is gained it becomes possible to inductively model the basis of success and

Author's personal copy

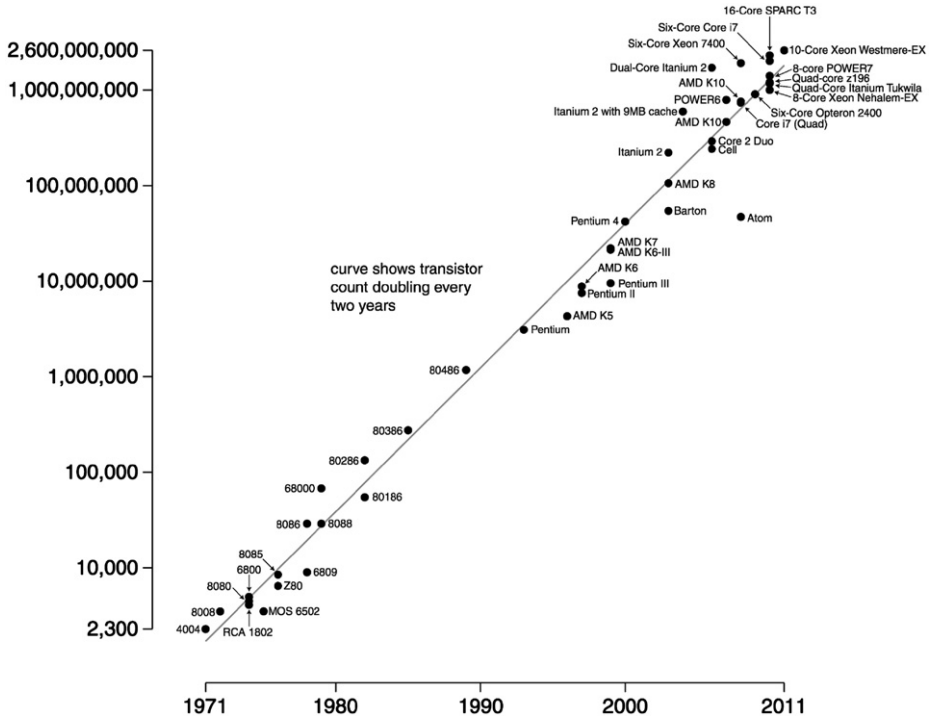

Fig. 4. Growth in umber of transistors on a chip (Wgsimon, 2011).

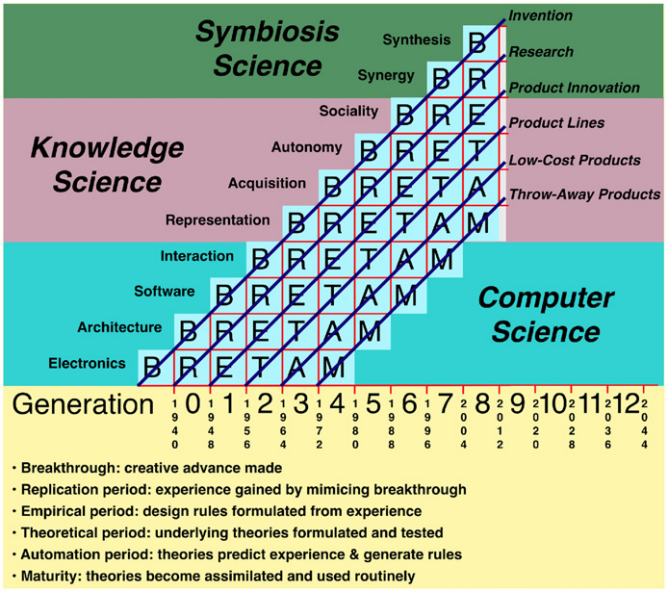

Fig. 5. Infrastructure of information technology.

failure and develop *theories*. This transition from empiricism to theory corresponds to the maximum slope of the underlying logistic learning curve (Gaines, 1998).

5 The theoretical models make it possible to *automate* the scientific data gathering and analysis and associated manufacturing processes.

6 Once automation has been put in place effort can focus on cost reduction and quality improvements in what has become a *mature* technology.

We identified the different tiers as being:

- the underlying digital *electronics*;
- its application in computer *architectures*;
- the programming of general-purpose computers through *software*;
- the development of computer–people and computer–computer *interactivity*;
- the *representation* of human knowledge;
- the *acquisition* of additional knowledge from interaction with the world, people and stored knowledge;
- the development of goal-directed *autonomous* knowledge creating processes;
- the increasing coupling of knowledge processing entities in *social* networks;
- the development of techniques to facilitate the *synergy* between human and computer knowledge processes;
- the *synthesis* of both into a unified system.

We characterized the theory of the lowest four tiers as constituting *computer science*; of the next four as *knowledge science*—the focus of the KSI; and project those above to form a currently developing *symbiosis science*.

Empirically, from an analysis of several thousand events in the history of the computing and the information sciences, we identified the time-scale of each phase, and hence each computer generation, as 8 years. We also found that after the initial phase of invention, research was associated with the replication phase, product innovation with the empirical phase, product lines with the theory

Author's personal copy

phase, low-cost products with the automation phase, and ultra-low cost, throw-away products with the maturity phase. This generates the diagonal trajectories shown where each tier of the hierarchy has its own research, development and product cycle.

The throw-away product trajectory is particularly interesting because it corresponds to technologies becoming so low in cost that they are essentially free, either literally so, or consumables purchased routinely with little budget impact. Web browsers have long been freeware as has access to much information through the Internet. Even major hardware items with significant manufacturing costs can follow this trajectory. For example, while writing this paper, I needed a new toner and toner drum for my high-speed, high-resolution, duplex laser printer and found that the cost of these was $280 but the cost of a new printer that included them was $150; the cost of printer itself had essentially become negative. New manufacturing technologies such as 3D printers (Dimitrov et al., 2006) are accelerating this trend, and it has also generated a major secondary industry of recycling electronic systems (Lounsbury et al., 2003) and new engineering issues of *design for recycling* (Kriwet et al., 1995); examples again of new knowledge being developed to cope with problems created by advances in knowledge.

There is also a strong positive feedback between the levels in the tiers below any given tier support its processes, and its innovations support those of the tiers below—computers are used in computer-aided design of electronic devices, Internet communities collaborate in the design of software and the sharing of knowledge, and so on.

When the KSI was formed in 1985 we saw the knowledge acquisition tier as having achieved major breakthroughs and to be in a research phase on the verge of the transition to the empirical phase where product innovation occurs. The development of knowledge acquisition techniques and tools was established as a second major project area, and contacts were made internationally with others having related interests, in particular John Boose at Boeing Computer Services, the co-founder of the knowledge acquisition conference series.

Twenty-five years later the knowledge acquisition learning curve is on the verge of the transition from automation to maturity. What characterizes that stage is the weight of knowledge required to make a meaningful contribution, and a small number of major products which are well-established and difficult to displace, for example, the Protégé tool for logical knowledge modeling (Gennari et al., 2003), the WebGrid tool for conceptual modeling (Gaines and Shaw, 1997) and socio-cognitive inquiry (Gaines and Shaw, 2012b), and the Google search engine for textual knowledge indexing (Langville and Meyer, 2006).

Note that the maturity of the learning curve does not indicate that research and innovation ceases, only that it becomes increasingly difficult. Innovative research continues on syllogistics that matured two millennia ago and on matrix algebra that matured over a century occur, but it is very rare and requires an immense depth of knowledge of the existing literature. From this perspective, one would expect the knowledge acquisition research community to have become increasingly specialized, managing a repository of expertise that is significant to research and development in the tiers that build upon it, and responsible for incorporating advances in the fields upon which it builds in order to provide state-of-the-art techniques and tools.

8. Artificial intelligence research

The tier below knowledge acquisition in Fig. 5 is that of knowledge representation and artificial intelligence where advances in expert systems in the 1970s provided the stimulus for knowledge acquisition research in the 1980s. Artificial intelligence research began in the late 1950s with the *Dartmouth Summer Project* (McCarthy et al., 1955) in the USA and the *Mechanization of Thought Processes Symposium* (NPL, 1959) in the UK. These occurred as computers came into their second generation, before the advent of computer science as an academic discipline, and when the artificial intelligence metaphor might well have become the core of such a discipline. There was a crisis in the 1970s as computers came into their fourth generation and embryonic computer science departments had to vie with nascent artificial intelligent departments in their requests for major funding to purchase the next generation of computers such as the DEC PDP10.

In Britain the conflict led the Science Research Council to commission a distinguished applied mathematician not associated with the contending applicants to report on the state of the art and future potential of artificial intelligence research. The infamous *Lighthill Report* damned both in sarcastic terms:

> "It has sometimes been argued that part of the stimulus to laborious male activity in 'creative' fields of work, including pure science, is the urge to compensate for lack of the female capability of giving birth to children. If this were true, then Building Robots might indeed be seen as the ideal compensation! One piece of evidence supporting that highly uncertain hypothesis: most robots are designed from the outset to operate in a world as like as possible to the conventional child's world as seen by a man; they play games, they do puzzles, they build towers of bricks, they recognise pictures in drawing-books ('bear on rug with ball'); although the rich emotional character of the child's world is totally absent. Builders of Robots can justly reply that while robots are still in their infancy they can mimic only pre-adult functions and a limited range of those at most, and that these will lead on to higher things. Nevertheless, the view to which this author has tentatively but perhaps quite wrongly come is that a relationship which may be called pseudo-maternal rather than Pygmalion-like comes into play between a Robot and its Builder." (Lighthill, 1973)

The report undermined the funding of artificial intelligence research in the UK and USA for seven years—the first so-called *AI winter* (McDermott et al., 1985).

Author's personal copy

It is ironic that the report only briefly mentions the already recognized achievements of DENDRAL (Buchanan et al., 1967), overlapped the publication of Winograd's (1972) doctoral thesis on SHRDLU which marked a major advance in natural language understanding, and was shortly followed by Shortliffe's (1976) doctoral thesis on MYCIN that provided the foundations for expert systems research and application. Whilst the immediate impact of the Lighthill report was highly negative for artificial intelligence research funding, the effect was alleviated in 1980 when MITI in Japan announced its national program for the development of a fifth generation of computers targeted on *knowledge-processing systems* (Moto-oka, 1982; Feigenbaum and McCorduck, 1983; Gaines, 1984b).

The Japanese program triggered a competitive response in governments around the world: in the USA DARPA's *Strategic Computing Program* (SCC); in the UK the *Alvey Program*; and in the EEC the *Esprit Program*. Major corporations funded their own internal programs and through the 1980s artificial intelligence boomed as never before, particularly that associated with expert system development. The AAAI annual conference had continuing attendances above 5000 with very large accompanying commercial exhibitions—the call for sites for the 1988 conference estimates attendance at 6500 and exhibition space at 80,000 square feet.

However, the failure of the Japanese program to produce any meaningful outcomes led to the SCC, Alvey and Esprit programs being discontinued in the late 1980s, and the MITI sixth generation program (STA, 1985; Gaines, 1986) attracting little interest. A second *AI winter* (Hendler, 2008) commenced in 1987, primarily in the USA, as the associated technology bubble burst and commercial interest in expert systems waned, but had surprisingly little effect on the knowledge acquisition and other specialist artificial intelligence communities world-wide.

The reduction in the hype and excessive expectations provided breathing space to reflect on fundamentals, continue to develop and refine techniques and tools, and integrate frameworks from the diverse fields relevant to knowledge acquisition. It had been obvious all along that available information technology provided the means to support and amplify human intelligence, rather than replace it, and we needed to understand the requirements in greater depth and match the technology to them—a slow process of incremental improvement rather than dramatic breakthrough.

9. Knowledge acquisition research

Artificial intelligence and knowledge-based system entered the product innovation phase in the 1980s as new companies were formed to commercialize university research in 'expert systems' largely based on production rules, and the products being the software frameworks used in the initial systems, stripped of their content, and termed expert system 'shells.' The explosion of industrial and academic interest in the potential of expert systems is evidenced by the attendance of over 7000 at the 1985 joint IJCAI/AAAI Conference at UCLA. The three largest tracks were: *Learning and Acquisition* with 31 papers; *Expert Systems* with 28; *Natural Language* with 28. Some 36% of the 245 papers presented were on these knowledge acquisition themes, but the unexpectedness of this is illustrated by the conference planning where none of the 4 keynotes and only one of the 12 panels addressed these themes.

The initial success of expert system developments (Michie, 1979) and the development of a number of reasonably domain-independent software support systems for the encoding and application of knowledge (Hayes-Roth et al., 1983) had opened up the possibility of widespread usage of expert systems. However, what Feigenbaum (1980) had termed 'knowledge engineering', the reduction of a large body of knowledge to a precise set of facts and rules, has already become a major bottleneck impeding the application of expert systems in new domains. It was recognized that we needed to understand more about the nature of expertise in itself (Hawkins, 1983) and to able to apply this knowledge to the elicitation of expertise in specific domains:

> "Knowledge acquisition is a bottleneck in the construction of expert systems. The knowledge engineer's job is to act as a go-between to help an expert build a system. Since the knowledge engineer has far less knowledge of the domain than the expert, however, communication problems impede the process of transferring expertise into a program. The vocabulary initially used by the expert to talk about the domain with a novice is often inadequate for problem-solving; thus the knowledge engineer and expert must work together to extend and refine it. One of the most difficult aspects of the knowledge engineer's task is helping the expert to structure the domain knowledge, to identify and formalize the domain concepts" (Hayes-Roth et al., 1983, p. 129).

This was the ethos in which John Boose and I launched the first Knowledge Acquisition Workshop (KAW) at Banff in 1986 amidst in a wave of academic and commercial interest in research and development of knowledge acquisition techniques and tools. We took advantage of the enthusiasm to set high entry standards, a journal-quality paper of about 20 pages each refereed by some seven reviewers, and enforce a workshop format where more time was allocated to discussion than to presentations. One hundred and twenty papers were submitted, 500 applications to attend from 30 countries, 42 papers were accepted and 60 researchers invited to attend. A week in the isolated and awe-inspiring environment of the Canadian Rockies in a conference centr where the rooms did not have telephones or television sets promoted a very high degree of engagement. While we had originally intended that the meeting circulate in North America, each year the

Author's personal copy

participants voted to return to Banff, and KAW continued there as a one-week intensive workshop until its 12th meeting in 1999, the year that I retired and moved out to the British Columbia coast on Vancouver Island.

Twenty-seven of the papers from the 1986 KAW were developed into full journal papers and published in a succession of special issues in this journal in early 1987. This pattern of publication was followed in subsequent years as the reputation of the meeting attracted continuing high levels of submission and the acceptance that a refereed journal-level paper was the only 'ticket' for attendance. There was also a strong culture that the systems presented must be operational, and it was expected that a tool described would be made available for interaction in the demonstrations track. We also followed the lead of the machine translation and machine learning communities in developing over the years a series of knowledge acquisition challenge problems, named *Sisyphus* tasks by Marc Linster that presented to the community in order to exhibit how their knowledge acquisition techniques might address them. The first problem was one of room allocation, the second of elevator design, the third of igneous rock classification and the fourth of porting knowledge acquisition tools to the web. They each resulted in many papers at the conferences together with journal special issues that enabled different knowledge acquisition and modeling techniques, and different tools, to be compared.

The published papers in the early years were also reissued in book form at low cost to make the material widely available to those unable to attend. In 1988 John and I persuaded Academic Press that a sister journal specializing in knowledge acquisition research would be viable, and the journal *Knowledge Acquisition* carried much of the KAW material as well as other submitted papers from 1989 to 1994. In 1995 we merged it with this journal because it had substantially more material to publish than its page limits allowed, and the publisher was prepared to expand the size of the well-established journal but not the new one. Thus this journal has had a major track of knowledge acquisition research for over 25 years.

At the 1986 meeting there were knowledge acquisition researchers from Europe and Australia and it was decided to establish a second series of workshops in Europe and a third in South East Asia. John and I organized the first European KAW in London in 1987 co-chaired by Tom Addis, and the first in South East Asia in Japan in 1990 co-chaired by Hiroshi Motoda and Riichiro Mizoguchi. We followed the same format as at Banff except, since we were located in major cities, the first day was open with a presentation format to provide a larger community with access to the state-of-the-art in knowledge acquisition. Fig. 6 shows the dates and locations of the 50 meeting worldwide to the end of 2012.

In 1992 Mark Musen joined as co-chair of the North American KAW in preparation for John stepping down as he had moved to another division of Boeing Computer Services when their knowledge-based systems research program was completed. After I retired in 1999, a new series of meetings was commenced under the title, *Knowledge Capture*, and K-CAP 1 was held in Victoria, BC, in 2001.

10. Porting knowledge acquisition tools to the web

Although it is difficult to recollect nowadays, in 1986 when KAW commenced the web had not been envisioned. The Internet and email were in place, we set up a workstation with a modem connection, and many of those attending brought their own portable computers and modems. However, the web of documents and interactive services was still some eight years away, taking off in 1994/1995 as the Internet was commercialized and the HTTP protocol became its greatest source of traffic (Gaines, 1998).

The web came into being through a similar series of serendipitous processes commencing in 1989 when Berners-Lee (1989) presented CERN with a proposal for

1986	KAW 1	Canada	**1994**	KAW 8	Canada	**2002**	EKAW 13	Spain
1987	KAW 2	Canada		EKAW 8	Belgium		PKAW 9	Japan
	EKAW 1	UK		PKAW 4	Japan	**2003**	KCAP 2	USA
1988	KAW 3	Canada	**1995**	KAW 9	Canada	**2004**	EKAW 14	UK
	EKAW 2	Germany	**1996**	KAW 10	Canada		PKAW 10	New Zealand
1989	KAW 4	Canada		EKAW 9	UK	**2005**	KCAP 3	Canada
	EKAW 3	France		PKAW 5	Australia	**2006**	EKAW 15	Czech Republic
1990	KAW 5	Canada	**1997**	EKAW 10	Spain		PKAW 11	China
	EKAW 4	Holland	**1998**	KAW 11	Canada	**2007**	KCAP 4	Canada
	PKAW 1	Japan		PKAW 6	Singapore	**2008**	EKAW 16	Italy
1991	KAW 6	Canada	**1999**	KAW 12	Canada		PKAW 12	Vietnam
	EKAW 5	UK		EKAW 11	Germany	**2009**	KCAP 5	USA
	PKAW 2	Australia		PKAW 7	Australia	**2010**	EKAW 17	Portugal
1992	KAW 7	Canada	**2000**	EKAW 12	France		PKAW 13	Korea
	EKAW 6	Germany		PKAW 8	Australia	**2011**	KCAP 6	Canada
	PKAW 3	Japan	**2001**	KCAP 1	Canada	**2012**	PKAW 14	Malyasia
1993	EKAW 7	France					EKAW 18	Ireland

Fig. 6. Fifty Knowledge Acquisition Workshops 1986–2012.

Author's personal copy

managing its documents effectively: *"We should work toward a universal linked information system, in which generality and portability are more important than fancy graphics techniques and complex extra facilities."* The major impetus behind the web implementation we know today came in 1992 when Marc Andreessen joined the *www-talk* list server and started the development of what became the *Mosaic* browser and eventually *Netscape* and *Internet Explorer*. Marc took the full affordances of printed documents as his target model, and added the fancy graphics techniques and complex extra facilities that Berners-Lee had disclaimed. He and Eric Bina radically changed the nature of the web when they enabled web documents to become interactive through the embedding of HCI widgets such as text edit boxes, check boxes, radio buttons, selection lists and menus. This allowed web documents to be used as graphic user interfaces, web browsers as universal interfaces, and HTML as a generic interface programming language (Gaines and Shaw, 2001).

At the 8th KAW in 1994 Tom Gruber drew attention to the potential of Netscape's development of interactive protocols for the web, and set the knowledge acquisition community the challenge of porting its knowledge acquisition tools to the web. We ported the repertory grid and semantic network conceptual modeling tools that Boeing and we were using for knowledge elicitation and, in 1995, presented papers on them at the Computer-supported Collaborative Learning Conference (Gaines and Shaw, 1995b; Shaw and Gaines, 1995), and gave a paper on them and a tutorial on the porting techniques developed at the World Wide Web Conference (Gaines, 1995; Gaines and Shaw, 1995a). We added the rule induction and inference algorithms we had developed (Gaines, 1989), and at the 10th KAW in 1986 demonstrated the complete knowledge engineering cycle from elicitation through to testing and application as a web service (Gaines and Shaw, 1996).

This rapid response to challenges, prototyping, demonstration, dissemination and publication cycle was by no means peculiar to the KSI but was endemic to all the research groups attending the KAW meetings worldwide. The friendly competition between groups, exchange of ideas, suggestions for enhancements, collaborations in joint projects to integrate tools and share datasets, pressure to produce original papers for two or three meetings a year (and for the many other workshops and conference organized by those within the KAW community), provided an extremely productive international research environment.

11. Porting knowledge in books and journals to the web

However, while we were operating within a computational knowledge representation paradigm that eventually became the basis of the *semantic web* (Shadbolt et al., 2006), a much older but at least equally significant paradigm for supporting human knowledge processes that of textual knowledge capture in books and journals was also making major advances through the use of web technologies. Journal publishers' digitization of the bulk of the scientific literature, the growth of mutual help mailing list archives, blogs and wikis encompassing all aspects of human life, and the ongoing project of digitizing all the world's literature, led to an explosion of digitally encoded knowledge in textual and other media becoming not only almost universally available but also a widely accepted and utilized resource in a very short span of time.

It also created a massive information overload that was eventually tamed through Google's indexing technologies (Langville and Meyer, 2006), and similar developments in the content-based indexing of all materials on corporate and personal computers. In parallel with these new technologies many of the tasks that had been targets in the early days of artificial intelligence, such as text and speech recognition and machine translation, were incrementally addressed to the extent of becoming routinely useful, as much through the continuing exponential growth in computing power as through improvements in understanding the nature of the problems and solution techniques.

Web technology is providing access to a high proportion of human knowledge available through contextual indexing in a manner that seems to match well the processes of human memory (Griffiths et al., 2007), and provides much of the support for human knowledge processes that had been expected from expert systems. It does so largely through document-retrieval techniques based on content and human-generated linkages with little use of the logical methods and knowledge structures that have dominated artificial intelligence and knowledge acquisition research.

It 'gives advice' and 'answers questions' by finding relevant pre-existent material that provides the information to enable the human users to infer possible solutions to their problems. DeepQA/Watson (Ferrucci et al., 2010) does use logical ontologies for some inferences but its answers are based in large part on massive information retrieval and probabilistic evidence combination from many diverse sources.

The use of information indexing and retrieval techniques to produce powerful knowledge support tools that provide valuable services to people whilst having minimal semantic foundations is a powerful technique, but the lack of 'understanding' involved can have curious side effects. For example, Google's automatic translation services based on statistical modeling of a large corpus of existing translations have great practical value. However, the approach can lead to major errors. For example, I noted in a draft of this paper that

> "In 1964 bar-Hillel provided convincing arguments that machine translation was impossible, and for many years his prognosis appeared correct but nowadays machine translation of usable quality is available and we have come to take it for granted. How did this happen?"

and decided to use this text to illustrate automatic translation. Google generated a very reasonable German version:

> "Im Jahr 1964 bar-Hillel vorausgesetzt überzeugende Argumente, dass die maschinelle Übersetzung nicht

Author's personal copy

möglich war, und seit vielen Jahren erschien seine Prognose richtig, aber heutzutage maschinelle Übersetzung von brauchbarer Qualität zur Verfügung steht und wir sind gekommen, um es für selbstverständlich halten. Wie konnte das passieren?"

I then had it translate the German back into English:

"In 1964, Bar-Hillel provided convincing arguments that the machine translation was possible, and for many years his prediction came true, but nowadays machine translation of reasonable quality is available and we have come to hold it for granted. How could this happen?"

Compare the text in bold to see that the "nicht" in the German has been lost, and Bar-Hillel is represented as making the opposite statement to that with which we started.

The point of this is not to poke fun at machine translation as we could so easily do in its early years. What is produced now is often amazingly good in its accuracy and colloquialism. However, the lack of knowledge-level representation can lead to egregious errors. What has been achieved is not machine understanding but rather a very effective, and useful, demonstration of Searle's (1980) 'Chinese room' in action.

This raises questions about the integration of the techniques developed in 25 years of knowledge acquisition research, most of which have logicist foundations and many of which have been successfully deployed, with techniques of information retrieval and natural language analysis and generation. Will the *semantic web*, based on description logics and formal ontologies to support machine understanding of the information stored on the web, actually become a significant component of the evolution of the web, and, if so, what is the most effective research agenda for its development?

12. Changes in knowledge-based system research in the last 25 years

Computer technology changed dramatically over the quarter century of knowledge acquisition conferences. At the first KAW, in order to write his paper during the meeting, Jeff Bradshaw brought a backpack with a state-of-the-art Macintosh Plus computer—7 MHz processor, 1 MByte main memory and 20 MByte external hard drive. Our elicitation and conceptual modeling tools ran on such a machine with the program occupying 384 KByte of the main memory. Nowadays a reasonably priced workstation might have 16 or more 3 GHz processors, 64 GByte or more main memory, 10 TByte or more of backing store that is increasingly solid state, plus high speed Internet connections that provide access to a world of information and services.

The number of transistors on a chip has grown by factor of 100,000, computers are three generations on and some one million times more powerful. To a large extent we no longer have to be concerned about optimizations to minimize memory usage, maintain effective interactivity, and so on, certainly not for the scale of problems we tackled in the early days, such as extracting rules from 'large' datasets such as the 47,000 Garvan thyroid cases (Gaines and Compton, 1993). However, increased computing power has also led us to address very much more demanding real-world requirements, and such considerations still arise as we tackle some of the very much larger datasets of the present day.

For example, reducing the time to classify the several hundred thousand concepts in the SNOMED ontology from a day to a few seconds to enable interactive identification and corrections of errors was presented as a major challenge by Rector in his keynote address at K-CAP 2011, and optimizations of knowledge representation and inference engines to achieve this is a widespread research activity (Cornet and Spackman, 2008). Classifying a very large number of individuals in real-world applications of description logics also requires optimizations such as modularizing the A-Box so that sub-sets of instances can be checked within main memory, and the techniques involved raise issues of problem-description at the ontological level that go beyond merely computational optimization (Wandelt and Möller, 2012).

As we move up from the computer level in Fig. 5 to the knowledge level, other very significant advances have occurred over the last quarter century. The expert systems breakthroughs that triggered the explosion of activity in the 1980s were based on rule-based knowledge representation (Wandelt and Möller, 2012), and modular production rules (Buchanan and Shortliffe, 1984) became the underlying representation and inference technology of expert systems shells. Whatever the natural representation of an expert's conceptual model in a knowledge acquisition tool, it ultimately had to be compiled into the forms of rules that the shell supported together with the control structures that the shell required to respond to interactive inquiries.

We collaborated with colleagues to integrate our knowledge acquisition tools with their shell technologies (Rappaport and Gaines, 1988; Gaines and Linster, 1990) in such a way as to support the complete knowledge engineering life cycle from elicitation, through testing, to application and continuing development as issues were discovered in use. From a user perspective it was surprisingly easy to create systems integrating heterogeneous tools never designed for such integration, but the internal 'knowledge structures' that one passed around were messy and opaque, a communication protocol between tools rather than a knowledge presentation medium for people (Gaines and Shaw, 1993).

Euphoria about the power of production rules was probably what led the Japanese Fifth Generation project down the fruitless path of designing special-purpose Prolog machines in competition with the equally fruitless path in the USA of designing special-purpose Lisp machines.

Author's personal copy

There has been a long history in computing of special-purpose computers, such as digital differential analyzers (Gaines, 1968), being designed to address specific problems but rapidly becoming obsolete because general purpose computers achieve lower cost through high volume, software algorithms are optimized to achieve the speed required, and general purpose computer technology has exponentially increased in power with time. Nowadays, the only significant special-purpose chips are those at the interface supporting graphic interaction and communication protocols.

General-purpose inference engine technology in the form of SAT algorithms has had major advances over the years commencing with the initial logical constraint propagation algorithms half a century ago (Davis and Putnam, 1960; Davis et al., 1962), and continuing today as more effective algorithms that model the problem structure and are parallelizable continue to be developed (Gu et al., 1996; Babic et al., 2006) and used as a foundational module of theorem provers for a range of logics (Armando et al., 2000; Giunchiglia et al., 2002; Pipatsrisawat and Darwiche, 2011) including description logics (Vescovi, 2011). One can see these developments as part of the maturity phase of the representation tier in Fig. 5, the foundational, 'machine code' level, of knowledge representation and inference analogous in their role to the continuing optimization of processor and compiler architectures at the lower levels.

The major advance in the representation tier at the knowledge level has been the development of *description logics* as decidable fragments of first order predicate logic with known complexity bounds, and algorithms that approach optimality within those bounds (Baader et al., 2003). This is another development that commenced well before the first KAW with Quillian's (1967) development of semantic networks, Minsky's (1975) of frames, Brachman (1977) logical analysis, Fahlman's (1979) implementation in NETL, Woods (1975), Hayes (1977, 1979) and Brachman's (1983, 1985) critiques of the weaknesses in the logical foundations, culminating in its first stage in Brachman and Levesque's (1984) discovery that inference in knowledge representation schema with even very basic primitives could be intractable.

As the knowledge acquisition conferences commenced, knowledge representation research entered its second stage with Borgida et al. (1989) implementation of CLASSIC as a logically sound and tractable knowledge representation and inference tool. Such knowledge representation schemes were initially called *terminological logics*, then *term subsumption* languages (Patel-Schneider et al., 1990) and eventually came to be called by their current name, *description logics* (Baader et al., 2003). It was realized that the representation primitives available in CLASSIC were too weak for many purposes and research focused on the computational complexity of algorithms for inference with a variety of combinations of more comprehensive representation primitives, and on the implementation of algorithms to provide tools that were as computationally effective as was possible (Haarslev and Moller, 2001; Tsarkov et al., 2007).

Skuce et al. (1989) presented and demonstrated a term subsumption workbench written in Smalltalk at the 4th KAW. It had an elegant graphical interface and powerful representation and inference capabilities and was clearly a more principled and attractive system than the available expert system shells. The paper describing CLASSIC (Borgida et al., 1989) that also became available that year was presented at a database conference in a highly tutorial form and, perhaps because of this, it was immediately apparent how to implement it. Based on that paper the KSI developed a knowledge representation server (KRS) in C++ that readily integrated with our elicitation and machine learning tools, described it at JKAW '90 in Japan (Gaines, 1990a) and demonstrated it at the 5th KAW (Gaines, 1990b).

It also proved simple to add rules, rules with exceptions and ripple-down rules (Compton and Jansen, 1990) to KRS in a principled fashion (Gaines, 1991c), provide an interactive graphical editor for knowledge structures represented as semantic networks (Gaines, 1991b), modularize KRS to accept new data types such as numbers and dates through plugins providing well-defined services (Gaines, 1993), solve the Sisyphus room allocation problem (Gaines, 1994b), extend our semantic network tools to support advances in description logics (Gaines, 2009), and express the logical foundations in terms of cognitive psychology (Gaines and Shaw, 2012a).

Again, I have discussed how advances elsewhere were incorporated in the KSI research because I know this first hand, but what I have described is representative of all the research groups participating in the conferences. In particular, the focus on knowledge capture for description logics and the ontological structures that resulted made it natural for knowledge acquisition researchers to play a major role in the development of the *semantic web* and for that to become a central topic at recent conferences (Musen and Corcho, 2011). There is a sense in which the semantic web technology has become a delivery vehicle for knowledge acquisition studies replacing the expert system shells with which we commenced and the more principled stand-alone knowledge representation servers which supplanted them.

The knowledge acquisition research community drew upon advances in the knowledge representation community for its preferred representation and inference technologies, and upon the machine learning and text analysis communities for some of its knowledge acquisition techniques. Colleagues from those communities came to the knowledge acquisition conferences to present the state-of-the-art in their areas and to investigate how techniques in their areas were being applied in ours. The conferences were very eclectic and, to the extent that there was a unique focus, it was on elicitation and modeling of expert knowledge that domain experts could understand and validate the models.

The issues of human sociology, culture, psychology and rationality that were involved were also apparent from the start and were a major theme at every meeting. We were

Author's personal copy

acutely aware of the arguments that human expertise could not be emulated by computers (Dreyfus, 1979) and that knowledge was not a substance (Clancey, 1989) but that did not prevent us from taking a pragmatic approach, trying to model and emulate, imputing knowledge as the basis of expertise and assessing the outcomes.

13. Future directions

Where do we go next? Forecasting is a dangerous game in which I have been involved throughout my career. Surveys of the state of the art in various areas inevitably lead to extrapolations and very often, not many years later, they are ridiculously incorrect. In 1978 I wrote a state-of-the-art report on minicomputers in business applications in the next decade and predicted that bubble memory technology might play a major role in providing high-speed backing stores (Gaines, 1978). In 1981 IBM offered its *Winchester Disk* in the OEM market and it rapidly became the backing store of choice with bubble memories finding no significant applications. As is often the case, the social need had been identified correctly but the technology that would actually come to satisfy it had not.

One widely agreed projection is that as technology at the interaction level of Fig. 5 advances our current connection to the Internet through keyboard, mouse and display screen will become replaced by direct neurological connection (Cochrane, 2012). Neurological prostheses have become a major focus of research and the primitive 'chip in the head' is already with us and becoming a valuable aid to the severely disabled (Ohl and Scheich, 2007). The social need driving such development more generally is our increasing dependence on the Internet as an essential extension of our selves, providing multimedia memory expansion, instant access to a dynamic corpus of knowledge well beyond our capabilities to assimilate and store, and instant communication with our social networks. More contentious projections at the symbiosis level of Fig. 5 are that information technology is moving towards *alien intelligence* (Martin, 2000) and a *singularity cusp* (Kurzweil, 2005) where human and computer intelligences merge.

The knowledge acquisition research community has generally pursued projects between these two extremes, at the knowledge levels in Fig. 5, based on formal knowledge representation supporting deductive, inductive and abductive inference. A difficult problem area for such an approach is that of 'symbol grounding' and 'cognitive commonality,' of trying to ensure that the primitive concepts used in our knowledge structures have the same meaning throughout the community of practice managing and using those structures. This raises issues of *collective cognition* (Resnick et al., 1991; Gaines, 1994a), *collective rationality* (Gaines, 2010; Goldberg, 2010), *organizational knowledge* (Weick, 1995; Gaines, 2003) and the extent to which we do actually use what we regard as *shared concepts* in the same way (Shaw and Gaines, 1989).

These are not new issues. Eighteen centuries ago Sextus Empiricus criticized Greek philosophers' focus on exact definitions, noting that "we must allow ordinary speech to use inexact terms" (Sextus, 1933, Adv. Math., II, 129). More recently, Hattiangadi (1987, p. 15) noted that "our understanding of language is approximate—I do not believe that we ever do understand the *same* language, but only *largely similar* ones." A nice example of the issues involved is to be found in research on merging what appear to be well-defined biological taxonomies of the same species from different sources (Thau and Ludäscher, 2007), or of merging them across species to create an overall *tree of life* (Bininda-Emonds, 2004). Pioneers in this area have become disenchanted by the lack of consistency in what initially appeared to be a logically and conceptually well-founded discipline and have documented in some detail the ways in which in practice the taxonomists have muddled through in a way that currently undermines attempts to create computational ontologies (Franz and Thau, 2010).

A miracle of human social existence is that we manage to 'muddle through' in most areas of knowledge despite major lack of cognitive commonality (Fortun and Bernstein, 1998). Computer tools have the same issues as those of a human learner coming to calibrate their cognition against the norms of their communities, and can only develop approximate models with which we can, hopefully, muddle through in an improved fashion. We need to develop a science of such *muddling through* that models human use of open concepts (Waismann, 1945; Weitz, 1977) to capture and transmit knowledge, and to come to comprehend the value of what might appear to be a logical defect as a necessary capability for coping with a complex and incompletely knowable world.

The one area of human knowledge where issues of cognitive commonality can be completely overcome is that of mathematics where the number of ungrounded primitive concepts is minimized and major conceptual structures are built axiomatically upon these few primitives. As Hilbert emphasized, the interpretation of primitive concepts is outside the realm of mathematics and any interpretation is valid provided it is consistent with the axioms. In discussing his axiomatization of Euclidean geometry with Hilbert remarks:

> "If in speaking of my points I think of some system of things, e.g. the system: love, law chimney-sweep…and then assume all my axioms as relations between these things, then my propositions, e.g. Pythagoras' theorem, are also valid for these things." .

However, the mathematical sciences are unique in being characterized in terms of formal structures specified arbitrarily by the human mind and requiring no interpretation, referents or application to be meaningful—quintessential abstract myths or fictions (Tharp, 1989). Many issues that arise in applications of mathematical logic to domains other than mathematics are not problematic in the domain of mathematics. For example, as Hilbert notes the issues of reference, individuals and existence are simple within that

Author's personal copy

domain: "if arbitrarily given axioms do not contradict one another, then they are true, and the things defined by the axioms exist" (Frege, 1980, p. 42). Similarly, issues of defeasibility and paraconsistency do not arise in the mathematical domain, but are of major significance to the way in which we muddle through in our rational discourse in more naturalistic domains.

Our current semantic web tools are based on algorithms that implement fragments of mathematical logic, largely because at the time we realized we needed to formalize the heuristic procedures of the early semantic networks the best understood and implemented deductive processes were those of mathematical logic. It was appropriate in the 1980s to address the well-founded criticisms of Woods (1975), Hayes (1977) and Brachman (1977, 1983, 1985) by adopting the most well-developed logical systems of that era, and this resulted in the OWL/Datalog framework for the semantic web.

The semantic web framework has replaced expert systems shells as the target representation and inference framework for knowledge acquisition. It is a rather more principled Procrustean bed than that of the shells, but it does place very significant limitations on what can be represented and the ease of representing such essential notions as change and defeasible assumptions. However, it is also a rapidly developing technology improving in the light of user experience (Rector and Brandt, 2008) while retaining sound formal foundations and reasonable tractability for real-world knowledge bases, and it is interesting to speculate on its future development.

The two major areas of criticism of the early knowledge representation were: first, the failure to distinguish between structural and assertional links (Woods, 1975); and, second, the attempt to treat any link as defeasible (Brachman, 1985). The first issue is ancient and much studied in the logical literature, the distinction between particulars and universals. The second also goes back at least to Aristotle, but has been less studied until recent years when non-monotonic reasoning has become a major research area. There are also other relevant areas where major advances have been made in recent years such as paraconsistent logics (Béziau et al., 2007), causal reasoning (Pearl, 2000), abductive reasoning (Aliseda, 2006), and so on, but I will focus on the first two topics.

One may analyze some major problems of current description logics as stemming from representing the distinction between universals and particulars as that between concepts and individuals. This was a natural way to represent the difference between structural and assertional links as the ontological concept of an individual had been hard-wired into mathematical logic by Frege and Russell in order to provide extensional, set-theoretic foundations for first-order logic, as formalized later through Tarski's model-theoretic semantics (Feferman, 1999). The treatment of the concept of an individual as a fundamentally different from that of other concepts was a powerful and fruitful constraint but created representational issues that plague us today. Ontological assumptions should be represented in logical terms, not postulated as part of the foundations of a logical framework. Ramsey (1925) criticized the distinction as being artificial and without logical foundation; Russell noted that he was uncomfortable with Ramsey's arguments but could not refute them, "to me, therefore, the question remains open" (Russell, 1931, p. 479).

The question remained open for many years but Ramsey's position was gradually formalized by a number of logicians, notably Kripke and Marcus, culminating in what is now termed the "new theory of reference" (Humphreys and Fetzer, 1998), although the core notion goes back to Mill (1875), that an individuating proper name has only denotation, not connotation. Proper names are intended to be unique identifiers used normatively to label what are asserted to be phenomena associated with the 'same entity,' where what 'same' means and how we track it are extra-logical concerns. That is proper names are properties of concepts representing phenomena.

Castañeda (1975) analyses such individuation in detail and notes that the individuating name is not simply a property but one that is used an indexical *operator* retrieving the concepts representing the phenomena associated with the individual, what has come to be termed the "mental file" model of singular thought (Azzouni, 2011; Crane, 2011). Castañeda (1983) also emphasizes the need for a non-reductive, pluralist theory of reference, that we reference what we deem to be individuals in many ways, corresponding to the complexity of the world and our relations to it and one another.

These developments suggest that the A-Box component of description logic implementations should be conceptualized as a database of singular concepts and their properties and relations, with the name/identifier of an individual providing a retrieval path to concept representing instances of that individual. This conceptualization and implementation would allow for multiple instances of the 'same' individual that would need to be disambiguated by additional indexical properties beyond names, such as *contexts*, *time slice*s and *worlds*, complicating the A-Box but also subsuming temporal, contextual and modal logics within a unified framework. Such an approach could, for example, provide better support for reasoning about changing individuals than is possible with existing logical frameworks for the semantic web (Artale and Franconi, 2000; Krieger, 2008). This will require more complex reasoning in the A-Box but it is consistent with the restructuring and optimization of that technology that is already underway (Wandelt and Möller, 2012).

Another area where major advances have been made is that of non-monotonic reasoning, of making default assumptions when information is not available and retracting the assumptions if necessary as further information is acquired. This includes formalization of inheritance systems in semantic networks (Gabbay and Schlechta, 2009) and also the development of general principles of rational non-monotonic reasoning (Gabbay, 1985; Lehmann, 1995;

Author's personal copy

Bochman, 2001). Kraus et al. (1990) showed that the principles could be represented as preferences over the epistemic states consistent with the available information, and Lehmann (2001) showed that these results could also be interpreted in terms of Arrow's (1959) *rational choice functions* which are the foundation of *social choice theory* in economics (Arrow et al., 2002).

The choice over epistemic states can be implemented in a description logic as a disjunction over those states together with a preference relation over the states. When a disjunction over states that is asserted true is evaluated, if a state that is consistent with the data is linked to a preferred consistent state then the non-preferred state is set false; that is, it is not considered to be the reason why the disjunction is true. This simple extension of the normal processing of a disjunction can be used to implement all aspects of default reasoning in a principled fashion. It enables defaults to be specified in the T-Box as part of the generic knowledge structure that supports rational default reasoning. It may be seen as a principled implementation of *case based reasoning* through *abduction to the only consistent epistemic state* (Bird, 2007) or, more generally, to the common aspects of all those possible epistemic states that are consistent with the data.

I predict that the implementation of modern theories of reference and non-monotonic reasoning in description logics will be a major activity in the next decade and lead to enhanced foundations for the semantic web that address in a principled fashion some of the significant knowledge representation and inference needs perceived by pioneers of semantic networks that were not within the scope of the technology of their time. This will significantly ease problems of encoding the results of knowledge acquisition projects in a natural form comprehensible to domain experts and users since temporal, contextual and non-monotonic inference is natural to people.

The implementation will probably involve the need to factor the representation and inference systems more effectively than is possible currently, drawing on work on generic *logical frameworks* and *environments* (Huet and Plotkin, 1991, 1993), *universal logics* (Béziau, 2007; Brady, 2007) that provide a *sub-structural* kernel for most logical systems (Schroeder-Heister and Došen, 1993) with the structural operators defined parametrically (Koslow, 1992), and implemented through back-ends that focus on massive dynamic optimizations using the characteristics of particular problems to maximize speed in a way that is transparent to users (Babic et al., 2006; Pipatsrisawat and Darwiche, 2011).

Such implementation requires a major effort by teams with a wide-ranging skill set and will only be possible if the semantic web becomes sufficiently commercially and socially significant. However, the role of logic in our life processes is still very much an open question with too many issues and debates to be covered here. One concise and balanced perspective is provided by Benthem (1999, p.22) who characterizes logic as a formal system representing "ubiquitous patterns in information representation and reasoning, which arise in any subject matter whatsoever." He also notes logic's role as a cognitive science, and this aspect has been nicely characterized by Brandom (2002, p.10), whilst avoiding both psychologism (Woleński, 2003) and normativity (Vidal, 2004), by presenting logic as supporting *meta-rationality* enabling us to reason and communicate about our pre-formal rational behavior, "logic does not define rationality in the most basic sense, but by making it possible for us to express explicitly the already rational relations articulating the contents of all our thoughts, it ushers in a higher level of rationality."

However, such considerations do not address the balance between the need for basic rationality and that meta-rationality in our everyday lives and the processes of our civilizations. How many of our needs can be satisfied by access to the accumulated rationality of humankind as increasingly provided by knowledge retrieval through the web, and how many require the meta-rational codification of that knowledge as envisioned for the semantic web? There is rivalry between two paradigms that have dominated human knowledge processes through the millennia, the accumulation and dissemination of relatively unstructured knowledge and skills through experience, education, and access to literature, as supported by the web in general, and the principled encoding of that material in a logico-mathematical inferential framework as supported by the semantic web. These two paradigms both have much to offer and I can safely predict that bridging between them will become an increasingly significant area of research and development, and one with major surprises beyond any that we might reasonably expect at this stage.

14. Conclusions

The historian, White (1972), has emphasized the plasticity and metaphorical power of historical accounts—we construct histories for ourselves that both empower and constrain our futures. Setting the history of the knowledge community within the ethos of artificial intelligence and expert systems studies alone may not provide an appropriate framework for our current and future research in the much wider context of the semantic web and its role in human society.

This paper and the others in this special issue present a number of perspectives on knowledge acquisition that enable us to construct histories for our community that empower and constrain our futures in interesting ways. Considering the choices available seems an appropriate agenda for our 25th anniversary. They are neither mutually exclusive nor exhaustive—just food for thought—what will be our themes, targets and agendas for the next 25 years?

In general there is a continuing need to consolidate and extend all that we know of knowledge acquisition processes and techniques, drawing upon all literatures and disciplines to support our stewardship of the state-of-the-art in knowledge acquisition. That includes the need to

Author's personal copy

continue to enhance the tools we make available to take advantages of developments in knowledge representation and computer technologies.

We also need to track user requirements for knowledge acquisition technologies to support both the needs of those applying them and the innovations in knowledge acquisition that may be outside of, or substantially extend, our current frameworks. In particular, the original logicist framework that has dominated artificial intelligence and expert systems research, may need substantial extension to support knowledge acquisition systems that incorporate the information indexing techniques of the web.

Beyond all that we need to appreciate the foundational role of knowledge in the evolution of our species and civilizations, a role so fundamental to the essence of what we are that is subject to social dynamics on a scale that we do not yet fully comprehend. The utopian dream of those who saw universal access to knowledge as improving the human condition has to be balanced by the reality that powerful new resources like the Internet can be used by any sector of society to support their activities regardless of how society at large might evaluate them. The technology is neutral to its usage, supporting scientific collaboration (Olson et al., 2008) and discovery (Nielsen, 2012), political and special-interest groups (Ferdinand, 2000), the impoverished in under-developed countries (Burrell, 2012), and criminal activities (Jaishankar, 2011) without discrimination; that is what it means for a technology to become an integral part of our society.

The future will continue to surprise us not only in the technologies we shall be using in another quarter century but also in the ways in which our fellow humans are using them to create a world that, perhaps, even the most speculative science fiction authors have not yet envisioned. We do live in very interesting times. May the knowledge acquisition research community live long and prosper.

Additional materials

Much knowledge acquisition material from the Calgary and Boeing research groups and early knowledge acquisition conferences is archived at http://cpsc.ucalgary.ca/~gaines/reports.

Acknowledgments

I am grateful to Paul Compton, Enrico Motta and Mildred Shaw for critical comments that greatly improved the first draft of this paper. This special issue celebrating 25 years of the knowledge acquisition conferences was Enrico's vision and I am pleased to have an opportunity to participate, and to work with him in his editorial role. The diversity of perspectives on both the past and the future suggests further research, and we hope that stimulation will be of value to others, particularly graduate students, whose support has been one of the major functions of the knowledge acquisition meetings. For me, participating in the worldwide knowledge acquisition community for so many years has been a pleasure and an inspiration, and I thank the many colleagues who have become friends, co-workers and partners in debate. I am also grateful to Enrico and Susan Wiedenbeck for taking on the task of editing this journal; IJHCS has thrived under their stewardship.

References

Aliseda, A., 2006. Abductive Reasoning: Logical Investigations into Discovery and Explanation. Springer, Dordrecht.

Altenberg, L., 1995. Genome growth and the evolution of the genotype-phenotype map. In: Banzhaf, W., Eeckman, F. (Eds.), Evolution and Biocomputation. Springer, Berlin, pp. 205–259.

Armando, A., Castellini, C., Giunchiglia, E., 2000. SAT-based procedures for temporal reasoning. In: Biundo, S., Fox, M. (Eds.), Recent Advances in AI Planning. Springer, New York, pp. 97–108.

Arrow, K.J., 1959. Rational choice functions and orderings. Economica 26, 121–127.

Arrow, K.J., Sen, A., Suzumura, K., 2002. Handbook of Social Choice and Welfare. Elsevier, Amsterdam.

Artale, A., Franconi, E., 2000. A survey of temporal extensions of description logics. Annals of Mathematics and Artificial Intelligence 30, 171–210.

Ayres, R.U., 1994. Information Entropy and Progress: A New Evolutionary Paradigm. AIP Press, Woodbury, NY.

Azzouni, J., 2011. Singular thoughts (objects-directed thoughts). Aristotelian Society Supplementary 85, 45–61.

Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P. (Eds.), 2003. The Description Logic Handbook. Cambridge University Press, Cambridge.

Babic, D., Bingham, J., Hu, A.J., 2006. B-cubing: new possibilities for efficient SAT-solving. IEEE Transactions Computers 55, 1315–1324.

Beniger, J.R., 1986. The Control Revolution: Technological and Economic Origins of the Information Society, Cambridge, MA. Harvard University Press.

Benthem, J. v., 1999. Wider still yet wider… resetting the bounds of logic. In: Achille, C.,.V. (Ed.), European Review of Philosophy: The Nature of Logic, vol.4. CSLI Publications, Stanford, CA, pp. 21–44.

Bernal, J.D., 1939. The Social Function of Science. Routledge, London.

Berners-Lee, T., 1989. Information Management: A Proposal.CERN, Geneva ⟨http://www.w3.org/History/1989/proposal.html⟩.

Béziau, J.-Y. (Ed.), Birkhäuser, Basel.

Béziau, J.-Y., Carnielli, W.A., Gabbay, D.M. (Eds.), 2007. College Publications.

Bickerton, D., 1990. Language and Species. University of Chicago Press, Chicago.

Bininda-Emonds, O.R.P. (Ed.), Kluwer, Dordrecht.

Bird, A., 2007. Inference to the only explanation. Philosophy and Phenomenological Research 74, 424–432.

Blum, R., 1991. Kallimachos: The Alexandrian Library and the Origins of Bibliography. University of Wisconsin Press, Madison, WI.

Bochman, A., 2001. A Logical Theory of Nonmonotonic Inference and Belief Change. Springer, Berlin.

Borgida, A., Brachman, R.J., McGuinness, D.L., Resnick, L.A., 1989. CLASSIC: a structural data model for objects. ACM SIGMOD Record 18, 58–67.

Brachman, R.J., 1977. What's in a concept: structural foundations for semantic networks. International Journal of Man-Machine Studies 9, 127–152.

Brachman, R.J., 1983. What is and isn't: an analysis of taxonomic links in semantic networks. Computer 16, 30–36.

Brachman, R.J. Levesque, H.J., 1984. The tractability of subsumption in frame-based description languages. In: Proceedings of the 4th National Conference on Artificial Intelligence (AAAI), vol. 84, pp. 34–37.

Brachman, R.J., 1985. "I lied about the trees," or, defaults and definitions in knowledge representation. AI Magazine 6, 80–93.

Author's personal copy

Brady, R.T., 2007. Entailment logic—a blueprint. In: Béziau, J.-Y., Carnielli, W.A., Gabbay, D.M. (Eds.), Handbook of Paraconsistency. College publications, London, pp. 128–151.

Brandom, R., 2002. Introduction: five concepts of rationality. In: Tales of the Mighty Dead: Historical Essays in the Metaphysics of Intentionality. Harvard University Press, Cambridge, MA, pp. 1–17.

Buchanan, B., Sutherland, G., Feigenbaum, E.A., 1967. Heuristic-Dendral: a program for generating explanatory hypotheses in organic chemistry. In: Meltzer, B., Michie, D. (Eds.), Machine Intelligence, vol. 4. Oliver & Boyd, Edinburgh, pp. 209–254.

Buchanan, B.G., Shortliffe, E.H., 1984. Rule-based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project. Addison-Wesley, Reading, MA.

Burnet, J., 1920. Early Greek Philosophy. Black, London.

Burrell, J., 2012. Invisible Users: Youth in the Internet Cafes of Urban Ghana, Cambridge, MA. MIT Press.

Bush, V., 1945. As we may think. Atlantic Monthly 176, 101–108.

Cantor, N.F., 2001. In the Wake of the Plague. Free Press, New York.

Castañeda, H.-N., 1975. Individuation and non-identity: a new look. American Philosophical Quarterly 12, 131–140.

Castañeda, H.-N., 1983. Reply to Burge. In: Tomberlin, J.E. (Ed.), Agent, Language and the Structure of the World: Essays Presented to Hector-Neri Castañeda. Indianapolis, Hackett, pp. 355–372.

Castañeda, H.-N., 1990. Introduction: experience and the semantics of thinking. In: Jacobi, K., Pape, H. (Eds.), Thinking and the Structure of the World: Hector-Neri Castañeda's Epistemic Ontology Presented and Criticized. de Gruyter, Berlin, pp. 3–10.

Castells, M., 1996. The Rise of the Network Society, Cambridge, MA. Blackwell Publishers.

Clancey, W.J., 1989. The knowledge level reinterpreted: modeling how systems interact. Machine Learning 4, 285–291.

Clancey, W.J., 1990. The frame of reference problem in the design of intelligent machines. In Architectures for Intelligence. In: van Lehn, K., Newell, A. (Eds.), The Twenty-Second Carnegie Symposium on Cognition. LEA, Hillsdale, New Jersey, pp. 357–424.

Cochrane, P., 2012. 100 years in the past and 100 years in the future. Proceedings of the IEEE 100, 1278–1284.

Collison, R.L., 1966. Encyclopaedias: Their History Throughout the Ages. Hafner, New York.

Compton, P., Jansen, R., 1990. A philosophical basis for knowledge acquisition. Knowledge Acquisition 2, 241–258.

Cornet, R., Spackman, K.A., (Eds.), 2008. Representing and Sharing Knowledge Using SNOMED: KR-MED 2008 Third International Conference on Knowledge Representation in Medicine 〈http://kr-med.org/2008/proceedings/proceedings-kr-med2008.pdf〉.

Crane, T., 2011. The singularity of singular thought. Aristotelian Society Supplementary 85, 21–43.

Crouzel, H., 1989. Origen. Harper & Row, London.

Davis, M., Putnam, H., 1960. A computing procedure for quantification theory. Journal of ACM 7, 201–215.

Davis, M., Logemann, G., Loveland, D., 1962. A machine program for theorem-proving. Communications of ACM 5, 394–397.

Dimitrov, D., Schreve, K., Beer, N.d., 2006. Advances in three dimensional printing – state of the art and future perspectives. Rapid Prototyping Journal 12, 136–147.

Diringer, D., 1982. The Book Before Printing: Ancient, Medieval, and Oriental. Dover, New York.

Distin, K., 2005. The Selfish Meme: A Critical Reassessment. Cambridge University Press, Cambridge.

Dostoyevsky, F., 1864. Notes from the Underground. Penguin, Middlesex, UK (trans. J Coulson).

Dreyfus, H.L., 1979. What Computers Can't Do: The Limits of Artificial Intelligence. Harper & Row, New York.

Eisenstein, E.L., 1979. The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early Modern Europe. Cambridge University Press, Cambridge.

Endicott, P., Ho, S.Y.W., Stringer, C., 2010. Using genetic evidence to evaluate four palaeoanthropological hypotheses for the timing of Neanderthal and modern human origins. Journal of Human Evolution 59, 87–95.

Fahlman, S.E., 1979. NETL, a System for Representing and Using Real-world Knowledge. MIT Press, Cambridge, MA.

Faraday, M., 1859. Experimental Researches in Chemistry and Physics: Reprinted from the Philosophical Transactions of 1821–1857, the Journal of the Royal Institution, the Philosophical Magazine and other Publications. Taylor & Francis, London.

Feferman, S., 1999. Logic, logics, and logicism. Notre Dame Journal of Formal Logic 40, 31–54.

Feigenbaum, E., 1980. Knowledge Engineering: the Applied Side of Artificial Intelligence (Report Number STAN-CS-80-812 Department of Computer Science, Stanford University).

Feigenbaum, E.A., McCorduck, P., 1983. The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World. Addison-Wesley, Reading, MA.

Ferdinand, P., 2000. The internet, Democracy and Democratization. Frank Cass, London.

Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A.A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J., Schlaefer, N., Welty, C., 2010. Building Watson: an overview of the DeepQA project. AI Magazine 31, 59–79.

Force, J.E., 1990. Newton's God of dominion: the unity of Newton's theological, scientific and political thought. In: Force, J.E., Popkin, R.H. (Eds.), Essays on the Context, Nature, and Influence of Isaac Newton's Theology. Kluwer, Dordrecht, pp. 75–102.

Fortun, M., Bernstein, H.J., 1998. Muddling Through: Pursuing Science and Truths in the 21st Century. Counterpoint, Washington.

Franz, N.M., Thau, D., 2010. Biological taxonomy and ontology development: scope and limitations. Biodiversity Informatics 7, 45–66.

Frege, G., 1980. Philosophical and Mathematical Correspondence. Blackwell, Oxford.

Gabbay, D., 1985. Theoretical foundations for non-monotonic reasoning in expert systems. In: Apt, K.R. (Ed.), Logics and Models of Concurrent Systems. Springer, Berlin, pp. 439–457.

Gabbay, D.M., Schlechta, K., 2009. Defeasible inheritance systems and reactive diagrams. Logic Journal IGPL 17, 1–54.

Gaines, B.R., 1968. Varieties of computer—their applications and interrelationships. In: Proceedings of IFAC Symposium on Pulse Rate and Pulse Number Signals in Automatic Control, IFAC, Budapest, pp. 1–16.

Gaines, B.R., 1978. Minicomputers in business applications in the next decade. In: Infotech State of Art Report on "Minis Versus Mainframes", Infotech International, Berkshire, UK, pp. 51–81..

Gaines, B.R., 1984a. Methodology in the large: modeling all there is. Systems Research 1, 91–103.

Gaines, B.R., 1984b. Perspectives on fifth generation computing. Oxford Surveys in Information Technology 1, 1–53.

Gaines, B.R., 1985. The Formation of a Knowledge Science Institute in Canada. Knowledge Science Institute, University of Calgary, Calgary 〈http://www.cpsc.ucalgary.ca/~gaines/reports/KM/KSI85/〉.

Gaines, B.R., 1986. Sixth generation computing: a conspectus of the Japanese proposals. ACM SIGART Newsletter 95, 39–44.

Gaines, B.R., Shaw, M.L.G., 1986. A learning model for forecasting the future of information technology. Future Computing Systems 1, 31–69.

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. In: Proceedings of the Sixth International Workshop on Machine Learning, Morgan Kaufmann, San Mateo, California, pp. 156–159.

Gaines, B.R., 1990a. Knowledge representation servers: a generic technology for knowledge acquisition and knowledge-based systems. In: Motoda, H., Mizoguchi, R., Boose, J., Gaines, B. (Eds.), Knowledge Acquisition for Knowledge-Based Systems. Ohmsha, Tokyo, pp. 413–430.

Gaines, B.R., 1990b. An architecture for integrated knowledge acquisition systems. In: Boose, J. H., Gaines, B. R. (Eds.), Proceedings of the Fifth AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop, University of Calgary, Calgary, pp. 8-1–8-22.

Author's personal copy

Gaines, B.R., Linster, M., 1990. Integrating a knowledge acquisition tool, an expert system shell and a hypermedia system. International Journal of Expert Systems Research and Applications 3, 105–129.

Gaines, B.R., 1991a. Modeling and forecasting the information sciences. Information Sciences 57–58, 3–22.

Gaines, B.R., 1991b. An interactive visual language for term subsumption visual languages. In: IJCAI'91: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, Morgan Kaufmann, San Mateo, California, pp. 817–823.

Gaines, B.R., 1991c. Integrating rules in term subsumption knowledge representation servers. In: AAAI'91: Proceedings of the Ninth National Conference on Artificial Intelligence. AAAI Press/MIT Press, Menlo Park, California, pp. 458–463.

Gaines, B.R., 1993. A class library implementation of a principled open architecture knowledge representation server with plug-in data types. In: IJCAI'93: Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, Morgan Kaufmann, San Mateo, CA, pp. 504–509.

Gaines, B.R., Compton, P., 1993. Induction of meta-knowledge about knowledge discovery. IEEE Transactions on Knowledge and Data Engineering 5, 990–992.

Gaines, B.R., Shaw, M.L.G., 1993. Eliciting knowledge and transferring it effectively to a knowledge-based systems. IEEE Transactions on Knowledge and Data Engineering 5, 4–14.

Gaines, B.R., 1994a. The collective stance in modeling expertise in individuals and organizations. International Journal of Expert Systems 7, 21–51.

Gaines, B.R., 1994b. A situated classification solution of a resource allocation task represented in a visual language. International Journal of Human–Computer Studies 40, 243–271.

Gaines, B.R., 1995. Porting interactive applications to the web. In: 4th International World Wide Web Conference Tutorial Notes, O'Reilly, Sebastopol, CA, pp. 199–217.

Gaines, B.R., Shaw, M.L.G., 1995a. WebMap: concept mapping on the web. World Wide Web Journal 1, 171–183.

Gaines, B.R., Shaw, M.L.G., 1995b. Collaboration through concept maps. In: Schnase, J.L., Cunnius, E.L. (Eds.), Proceedings of CSCL95: Computer Support for Collaborative Learning. Lawrence Erlbaum, Mahwah, New Jersey, pp. 135–138.

Gaines, B.R., Shaw, M.L.G., 1996. WebGrid: knowledge modeling and inference through the World Wide Web. In: Gaines, B. R. Musen, M. A., (Eds.), Proceedings of Tenth Knowledge Acquisition Workshop, pp. 65-1–65-14.

Gaines, B.R., Shaw, M.L.G., 1997. Knowledge acquisition, modeling and inference through the World Wide Web. International Journal of Human–Computer Studies 46, 729–759.

Gaines, B.R., 1998. The learning curves underlying convergence. Technological Forecasting and Social Change 57, 7–34.

Gaines, B.R., Shaw, M.L.G., 2001. Human–computer interaction in online communities. Journal of Research and Practice in Information Technology 33, 3–15.

Gaines, B.R., 2003. Organizational knowledge acquisition. In: Holsapple, C.W. (Ed.), Handbook on Knowledge Management: 1. Springer, Berlin, pp. 317–347.

Gaines, B.R., 2009. Designing visual languages for description logics. Journal of Logic, Language and Information 18, 217–250.

Gaines, B.R., 2010. Human rationality challenges universal logic. Logica Universalis 4, 163–205.

Gaines, B.R., 2011. Knowledge capture through the millennia: from cuneiform to the semantic web. In: Musen, M. A., Corcho, O. (Eds.), Proceedings K-CAP '11: Sixth International Conference on Knowledge Capture ACM, New York, pp. 1–8.

Gaines, B.R., Shaw, M.L.G., 2012a. Computer aided constructivism. In: Caputi, P., Viney, L.L., Walker, B.M., Crittenden, N. (Eds.), Constructivist Methods. Wiley, New York, pp. 183–222.

Gaines, B.R., Shaw, M.L.G., 2012b. Sociocognitive inquiry. In: Ting, I.-H., Hong, T.-P., Liang, L.S. (Eds.), Social Network Mining, Analysis and Research Trends: Techniques and Applications. IGI Global, PA, USA, pp. 35–55.

Gennari, J.H., Musen, M.A., Fergerson, R.W., Grosso, W.E., Crubèzy, M., Eriksson, H., Noy, N.F., Tu, S.W., 2003. The evolution of Protégé: an environment for knowledge-based systems development. International Journal Human–Computer Studies 58, 89–123.

Gibson, K.R., Ingold, T., 1993. Tools, Language, and Cognition in Human Evolution. Cambridge University Press, Cambridge.

Gilfillan, S.C., 1937. The prediction of inventions. In: Ogburn, W.F. (Ed.), Technological Trends and National Policy, Including the Social Implications of New Inventions. U.S. Government Printing Office, Washington.

Gingras, Y., 2010. Mapping the structure of the intellectual field using citation and co-citation analysis of correspondences. History of European Ideas 36, 330–339.

Giunchiglia, E., Tacchella, A., Giunchiglia, F., 2002. SAT-based decision procedures for classical modal logics. Journal of Automated Reasoning 28, 143–171.

Gladney, H.M., 2007. Preserving Digital Information. Springer, Berlin.

Goldberg, S., 2010. Relying on Others: An Essay in Epistemology. Oxford University Press, Oxford.

Griffiths, T.L., Steyvers, M., Firl, A., 2007. Google and the mind. Psychological Science 18, 1069–1076.

Gu, J., Purdom, P.W., Franco, J. Wah, B.W., 1996. Algorithms for the Satisfiability (SAT) Problem: A Survey. Defense Technical Information Center.

Gutas, D., 1998. Greek Thought, Arabic culture: The Graeco-Arabic Translation Movement in Baghdad and Early Abbasid Society. Routledge, London.

Haarslev, V., Moller, R., 2001. RACER system description. In: Proceedings of the International Joint Conference on Automated Reasoning (IJCAR 2001), vol. 2083, pp. 701–705.

Harper, L.V., 2005. Epigenetic inheritance and the intergenerational transfer of experience. Psychological Bulletin 131, 340–360.

Hattiangadi, J.N., 1987. How is Language Possible?: Philosophical Reflections on the Evolution of Language and Knowledge. Open Court, La Salle.

Hawkins, D., 1983. An analysis of expert thinking. International Journal of Man–Machine Studies 18, 1–47.

Hayes, P.J., 1977. In defence of logic. In: Proceedings International Joint Conference Artificial Intelligence, pp. 559–565.

Hayes, P.J., 1979. The logic of frames. In: Metzing, D. (Ed.), Frame Conceptions and Text Understanding. de Gruyter, Berlin, pp. 46–61.

Hayes-Roth, F., Waterman, D.A., Lenat, D.B. (Eds.), 1983. Addison-Wesley, Reading, Massachusetts.

Headrick, D.R., 2000. When Information Came of Age: Technologies of Knowledge in the Age of Reason and Revolution. Oxford University Press, Oxford, pp. 1700–1850.

Hendler, J., 2008. Avoiding another AI winter. IEEE Intelligent Systems 23, 2–4.

Hobart, M.E., Schiffman, Z.S., 1998. Information Ages: Literacy, Numeracy, and the Computer Revolution. Johns Hopkins University Press, Baltimore.

Høyrup, J., 1999. Pythagorean "Rule" and "Theorem" – Mirror of the Relation Between Babylonian and Greek Mathematics.

Huet, G., Plotkin, G. (Eds.), 1991. Cambridge University Press, Cambridge.

Huet, G., Plotkin, G. (Eds.), 1993. Cambridge University Press, Cambridge.

Humphreys, P., Fetzer, J.H., 1998. The New Theory of Reference: Kripke, Marcus, and its Origins. Kluwer, Dordrecht.

Hunter, M.C.W., 1989. Establishing the New Science: The Experience of the Early Royal Society. Boydell Press, Woodbridge, Suffolk.

Jablonka, E., Lamb, M.J., 2005. Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life. MIT Press, Cambridge, MA.

Jaishankar, K., 2011. Cyber Criminology: Exploring Internet Crimes and Criminal Behavior. Taylor & Francis, Boca Raton.

Jardine, L., 1996. Worldly Goods: A New History of the Renaissance. Talese, New York.

Author's personal copy

Judson, H.F., 1979. The Eighth Day of Creation. Simon & Schuster, New York.
Katz, V.J., Imhausen, A., 2007. The Mathematics of Egypt, Mesopotamia, China, India, and Islam: A Sourcebook. Princeton University Press, Princeton, NJ.
Korotayev, A., 2005. A compact macromodel of world system evolution. Journal of World-Systems Research 11, 79–93.
Koslow, A., 1992. A Structuralist Theory of Logic. Cambridge University Press, Cambridge.
Kraus, S., Lehmann, D., Magidor, M., 1990. Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence 44, 167–207.
Krieger, H.-U., 2008. Where temporal description logics fail: representing temporally-changing relationships. In: Dengel, A., Berns, K., Breuel, T., Bomarius, F., Roth-Berghofer, T. (Eds.), KI 2008: Advances in Artificial Intelligence: Lecture Notes in Computer Science, vol. 5243. Springer, Berlin, pp. 249–257.
Kriwet, A., Zussman, E., Seliger, G., 1995. Systematic integration of design-for-recycling into product design. International Journal of Production Economics 38, 15–22.
Kurzweil, R., 2005. The Singularity is Near: When Humans Transcend Biology. Viking, New York.
Langville, A.N., Meyer, C.D., 2006. Google's PageRank and Beyond: The Science of Search Engine Rankings, Princeton, N.J. Princeton University Press.
Laughlin, B., 1995. The Aristotle Adventure: A Guide to the Greek, Arabic, and Latin Scholars who Transmitted Aristotle's Logic to the Renaissance. Albert Hale, Flagstaff, AZ.
Lehmann, D., 1995. Another perspective on default reasoning. Annals of Mathematics and Artificial Intelligence 15, 61–82.
Lehmann, D., 2001. Nonmonotonic logics and semantics. Journal Logic Computation 11, 229–256.
Lewis-Williams, J.D., 2002. The Mind in the Cave: Consciousness and the Origins of Art. Thames & Hudson, London.
Lighthill, J., 1973. Artificial intelligence: a general survey. In: Artificial Intelligence: A Paper Symposium. Science Research Council, UK.
Liu, H., Prugnolle, F., Manica, A., Balloux, F., 2006. A geographically explicit genetic model of worldwide human-settlement history. American Journal Human Genetics 79, 230–237.
Lock, G., 2003. Using Computers in Archaeology: Towards Virtual Pasts. Routledge, London.
Lounsbury, M., Ventresca, M., Hirsch, P.M., 2003. Social movements, field frames and industry emergence: a cultural Àìpolitical perspective on US recycling. Socio-Economic Review 1, 71–104.
Manuel, F.E., 1974. The Religion of Isaac Newton. Clarendon Press, Oxford.
Martin, J., 2000. After the Internet: Alien Intelligence. Capital Press, Washington, DC.
Maturana, H.R., 1975. The organization of the living: a theory of the living organization. International Journal of Man-Machine Studies 7, 313–332.
McBrearty, S., Brooks, A.S., 2000. The revolution that wasn't: a new interpretation of the origin of modern human behavior. Journal of Human Evolution 39, 453–563.
McCarthy, J., Minsky, M.L., Rochester, N., Shannon, C.E., 1955. A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. Dartmouth College, Dartmouth, MA.
McDermott, D., Waldrop, M.M., Schank, R., Chandrasekaran, B., McDermott, J., 1985. The dark ages of AI: a panel discussion at AAAI-84. AI Magazine 6, 122–134.
McGrath, J., 2007. Science Surf: Enter without knocking if you can ⟨http://www.wjst.de/blog/blog/2007/01/29/enter-without-knocking-if-you-can/⟩.
McNeill, W.H., 1989. Plagues and Peoples. Anchor, New York.
Mellars, P., 2004. Neanderthals and the modern human colonization of Europe. Nature 432, 461–465.
Mellars, P., Boyle, K., Bar-Yosef, O., Stringer, C. (Eds.), 2007. McDonald Institute, Cambridge.
Melville, S.C., Melville, D.J., 2008. Observations on the diffusion of military technology: siege warfare in the Near East and Greece. In: Slotsky Ross, M. (Ed.), From the Banks of the Euphrates: Studies in Honor of Alice Louise. Eisenbrauns, Winona Lake, IN, pp. 145–168.
Michie, D. (Ed.), Edinburgh University Press, Edinburgh.
Mill, J.S., 1875. A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation. Longmans Green, London.
Minsky, M., 1975. A framework for representing knowledge. In: Winston, P.H. (Ed.), The Psychology of Computer Vision. Simon and Schuster, New York, pp. 211–280.
Mollick, E., 2006. Establishing Moore's law. Annals of the History of Computing 28, 62–75.
Moto-oka, T., (Ed.), North-Holland, Amsterdam.
Musen, M.A., Corcho, O., (Eds.), 2011. Proceedings K-CAP '11: Sixth International Conference on Knowledge Capture. ACM, New York.
Musil, R., 1930. The Man Without Qualities, vol. 1, Secker & Warburg, London (translated by E. Wilkins & E. Kaiser).
Neugebauer, O., Sachs, A. Götze, A., 1945. Mathematical Cuneiform Texts. American Oriental Society and the American Schools of Oriental Research, New Haven, CN.
Nielsen, M.A., 2012. Reinventing Discovery: The New Era of Networked Science, Princeton, N.J. Princeton University Press.
Nissen, H.J., Damerow, P., Englund, R.K., 1993. Archaic Bookkeeping: Early Writing and Techniques of Economic Administration in the Ancient Near East. University of Chicago Press, Chicago, IL.
Noble, W., Davidson, I., 1996. Human Evolution, Language, and Mind: A Psychological and Archaeological Inquiry. Cambridge University Press, Cambridge.
NPL, 1959. Mechanisation of Thought Processes: Symposium at the National Physical Laboratory. H.M.S.O., London.
Ohl, F.W., Scheich, H., 2007. Chips in your head. Scientific American Mind 18, 64–69.
Olson, G.M., Zimmerman, A., Bos, N., 2008. Scientific Collaboration on the Internet. MIT Press, Cambridge, MA.
Oyama, S., 1985. The Ontogeny of Information: Developmental Systems and Evolution. Cambridge University Press, Cambridge.
Pike, A.W.G., Hoffmann, D.L., García-Diez, M., Pettitt, P.B., Alcolea, J., Balbín, R.D., González-Sainz, C., Heras, C. d. l., Lasheras, J.A., Montes, R., Zilhão, J., 2012. U-series dating of Paleolithic art in 11 caves in Spain. Science 336, 1409–1413.
Parsons, E.A., 1952. The Alexandrian Library, Glory of the Hellenic World: Its Rise, Antiquities, and Destructions. Cleaver-Hume, London.
Patel-Schneider, P.F., Owsnicki-Klewe, B., Kobsa, A., Guarino, N., MacGregor, R., Mark, W.S., McGuinness, D.L., Nebel, B., Schmiedel, A., Yen, J., 1990. Term subsumption languages in knowledge representation. AI Magazine 11, 16.
Pearl, J., 2000. Causality: Models, Reasoning, and Inference. Cambridge University Press, Cambridge.
Pettitt, P., 2008. Art and the middle-to-upper paleolithic transition in Europe: comments on the archaeological arguments for an early upper paleolithic antiquity of the Grotte Chauvet art. Journal of Human Evolution 55, 908–917.
Pipatsrisawat, K., Darwiche, A., 2011. On the power of clause-learning SAT solvers as resolution engines. Artificial Intelligence 175, 512–525.
Powell, B.B., 1991. Homer and the Origin of the Greek Alphabet. Cambridge University Press, Cambridge.
Price, D.J. d.S., 1963. Little Science, Big Science. Columbia University Press, New York.
Quillian, M.R., 1967. Word concepts: a theory and simulation of some basic semantic capabilities. Behavioral Science 12, 410–430.
Ramsey, F.P., 1925. Universals. Mind 34, 401–417.
Rappaport, A., Gaines, B.R., 1988. Integrating knowledge acquisition and performance systems. In: Barter, C. J., Brooks, M. J. (Eds.), AI'88: 2nd Australian Joint Artificial Intelligence Conference, Springer, Berlin, pp. 307–326.

Author's personal copy

Rector, A.L., Brandt, S., 2008. Why do It the hard way? The case for an expressive description logic for SNOMED. Journal of the American Medical Informatics Association 15, 744–751.

Renfrew, C., Boyle, K.V., 2000. Archaeogenetics: DNA and the Population Prehistory of Europe. McDonald Institute for Archaeological Research, Cambridge.

Resnick, L.B., Levine, J.M., Teasley, S.D., 1991. Perspectives on Socially Shared Cognition. American Psychological Association, Washington.

Rich, B., 1610. A New Description of Ireland, Wherein is Described the Disposition of the Irish Whereunto They are Inclined, etc.. T. Adams, London.

Riordan, M., Hoddeson, L., 1997. Crystal Fire: The Birth of the Information Age. Norton, New York.

Robson, E., 2008. Mathematics in Ancient Iraq: A Social History. Princeton University Press, Princeton, NJ.

Rupp, K., Selberherr, S., 2011. The economic limit to Moore's law. IEEE Transaction Semiconductor Manufacturing 24, 1–4.

Russell, B., 1931. Review of the foundations of mathematics and other logical essays by Frank Plumpton Ramsey. Mind 40, 476–482.

Russo, L., 2004. The Forgotten Revolution: How Science was Born in 300 BC and Why it Had to be Reborn. Springer, Berlin.

Schroeder-Heister, P.J., Došen, K., 1993. Substructural Logics. Clarendon Press, Oxford.

Searle, J., 1980. Minds, brains, and programs. Behavioral Brain Sciences 3, 417–424.

Sextus. 1933. Sextus Empiricus, with an English translation by R.G. Bury. Heinemann, London.

Shadbolt, N., Hall, W., Berners-Lee, T., 2006. The semantic web revisited. IEEE Intelligent Systems 21, 96–101.

Shaw, M.L.G., Gaines, B.R., 1989. Comparing conceptual structures: consensus, conflict, correspondence and contrast. Knowledge Acquisition 1, 341–363.

Shaw, M.L.G., Gaines, B.R., 1995. Comparing constructions through the web. In: Schnase, J.L., Cunnius, E.L. (Eds.), Proceedings of CSCL95: Computer Support for Collaborative Learning. Lawrence Erlbaum, Mahwah, New Jersey, pp. 300–307.

Shortliffe, E.H., 1976. Computer-Based Medical Consultations: MYCIN. Elsevier, New York.

Skuce, D., Shenkang, W. Beauville, Y., 1989. A generic knowledge acquisition environment for conceptual and ontological analysis. In: Boose, J. H., Gaines, B. R. (Eds.), Proceedings of the 4th Knowledge Acquisition for Knowledge-based Systems Workshop, University of Calgary, Calgary, pp. 31-1–31-20.

Smail, D.L., 2008. On Deep History and the Brain. University of California Press, Berkeley.

Snooks, G.D., 1996. The Dynamic Society: Exploring the Sources of Global Change. Routledge, London.

Snooks, G.D., 1998. The Laws of History. Routledge, London.

STA, 1985. Promotion of R&D on Electronics and Information Systems That May Complement or Substitute for Human Intelligence. Science and Technology Agency, Tokyo.

Sterelny, K., 2011. From hominins to humans: how sapiens became behaviourally modern. Philosophical Transactions of the Royal Society B: Biological Sciences 366, 809–822.

Stringer, C., 2002. Modern human origins: progress and prospects. Philosophical Transactions of the Royal Society B: Biological Sciences 357, 563–579.

Stringer, C., 2006. Homo Britannicus: The Incredible Story of Human Life in Britain. Allen Lane, London.

Swerdlow, N.M., 1998. The Babylonian Theory of the Planets. Princeton University Press, Princeton, NJ.

Tharp, L., 1989. Myth and mathematics: a conceptualistic philosophy of mathematics I. Synthese 81, 167–201.

Thau, D., Ludäscher, B., 2007. Reasoning about taxonomies in first-order logic. Ecological Informatics 2, 195–209.

Thornton, A., Raihani, N., 2010. Identifying teaching in wild animals. Learning Behavior 38, 297–309.

Tsarkov, D., Horrocks, I., Patel-Schneider, P., 2007. Optimizing terminological reasoning for expressive description logics. Journal of Automated Reasoning 39, 277–316.

Vescovi, M., 2011. Exploiting SAT and SMT techniques for automated reasoning and ontology manipulation in description logics. University of Trento, Trento, Italy.

Vidal, C.M., 2004. Normativity and its vindication: the case of logic. Theoria 19, 191–206.

Waismann, F., 1945. Verifiability. In: Proceedings of the Aristotelian Society, Supplementary vols. 19, 119–150.

Wandelt, S., Möller, R., 2012. Towards A Box modularization of semi-expressive description logics. Applied Ontology 133–167, 211–231.

Weick, K.E., 1995. Sense making in Organizations. Sage, Thousand Oaks, CA.

Weitz, M., 1977. The Opening Mind: A Philosophical Study of Humanistic Concepts. University of Chicago Press, Chicago.

Wells, H.G., 1938. World Brain. Doubleday, Garden City, NY.

Wgsimon. 2011. Transistor Count and Moore's Law – 2011 〈http://commons.wikimedia.org/wiki/File:Transistor_Count_and_Moore%27s_Law_-_2011.svg〉.

White, H.V., 1972. What is a historical system? Biology. In: Breck, A.D., Yourgrau, W. (Eds.), , History, and Natural Philosophy. Plenum Press, New York.

Winograd, T., 1972. Understanding Natural Language. Academic Press, New York.

Wiseman, J., El-Baz, F., 2007. Remote Sensing in Archaeology. Springer, New York.

Wojciechowski, J.A., 2001. Ecology of Knowledge. Council for Research in Values and Philosophy, Washington.

Woleński, J., 2003. Psychologism and metalogic. Synthese 137, 179–193.

Woods, W.A., 1975. What's in a link. In: Bobrow, D.G., Collins, A. (Eds.), Representation and Understanding. Academic Press, New York, pp. 35–82.