

Artificial intelligence and natural stupidity revisited

Overview

The research field of artificial intelligence (AI) emerged from an academic conference in 1956. It was a daring venture: its founding optimists proposed to create autonomous machines that could emulate and then exceed human intelligence, and fast. That was a glamorous goal, and attracted investment.

In stories about daring ventures, protagonists may not get to the place they set out for. On their journey they often discover unforeseen obstacles, and their preconceptions may obscure the value of these discoveries. If they do arrive at their destination, it may differ from what they had imagined.

Hindsight enables dramatic irony: our interest is heightened by knowing things the characters in the story cannot. The founders of AI made an innocent mistake. They held human intelligence in higher regard than is possible now, knowing what we do in the early 21st century.

Advances in various fields have discovered that our evolved capacity for fantasy, willful ignorance and self-deception are fundamental to human decision-making. Contemporary MBA courses teach the economic importance of misrepresentations that play on natural stupidity.ⁱ

Part of contemporary human wisdom consists of appreciating how natural stupidity shapes human institutions and markets, and indeed contributes to social complexity. Beyond apparently intractable technical difficulties, the business case for building machines that emulate human cognitive biases and heuristics seems weak. There is no scarcity of stupidity.

But the old wisecrack that AI is no match for natural stupidity has a twist. Clever technologies that complement (or compensate for) human foibles have emerged. Now AI and natural stupidity can be a matchⁱⁱ—and like human couples, a more or less happy one.ⁱⁱⁱ These alliances underpin the emergence of billion-dollar industries like interactive entertainment^{iv} and online dating.

We see the AI tradition as having quietly decoupled from the fantasy of making machines like us, or perhaps better. The fourth and fifth generations of AI researchers have fanned out into specialised fields that are so diverse an expert in one may not understand an expert in another. Concepts that emerge from this constellation may be developed into clever technologies that get to market; and become more or less useful, commercially viable adjuncts to human purposes.

The challenge for contemporary human governance is to make a distinction between potentially useful and merely glamorous emerging technologies. A

W H H
W I L T O N
H A N F O R D
H A N O V E R

**Wilton Hanford
Hanover Pty Ltd**
ABN 15 078 798 366

Level 2
19 MURRAY CRES,
Manuka ACT
PO Box 3709
Manuka ACT 2603

T (02) 6273 3048
F (02) 6273 1833
E info@whh.com.au
W whh.com.au

long view can help to steer a portfolio in ways that minimise over-investment in fantasy and under-investment in technological change.

No known human or machine process (with the possible exception of hindsight) can abolish uncertainty or eliminate the possibility of human error: but errors we survive and learn from are worthwhile.

What hindsight shows

In the last quarter of the 20th century many smart people made predictions about the future of digitised information technology. Mainly they got two things wrong.

- Digitised devices have become more pervasive than anyone imagined.
- These devices remain much stupider than smart people had hoped.

Semantics may have biased our thinking in the Anglophone world. We elected to name a class of digitised devices *computers*, suggesting ‘one who calculates.’ That suggests someone with insight, who can identify, define and solve problems.

The French called these devices *ordinateurs*, derived from the Latin for ‘one who organises.’ This might suggest a file clerk who follows rules and instructions. The French term now seems a much more realistic pointer to what most computers are used to do: store and access files.

A second semantic misdirection arises from the idea of the human mind as a thinking machine.^v This idea has a long philosophical and literary history. The ideal of pure reason is a key theme in one of the great intellectual fantasies of post-medieval Western civilization: Rational Economic Man.

In the 1930s a thought experiment called the ‘Turing machine’ inverted the ideal of human rationality into a concept about a machine that could think—or rather, what smart people conceived of as thinking in that era. So the hypothesis that human intelligence could be simulated by a machine became available to the many people who seek and make grants for research.

The Dartmouth conference emergence

Science had particular prestige during the 1950s, which supported the ‘scientific management’ fad and the (ironic) ‘science of muddling through’.^{vi} After the invention of magnetic core memory, during the long gestation of FORTRAN and a few years before the invention of the integrated circuit, the glamorous field of AI emerged from the 1956 Dartmouth conference.^{vii}

AI's founders were remarkably optimistic about their prospects. . For example conferee Herbert Simon (subsequently a Nobel Laureate) reportedly declared: ‘machines will be capable, within twenty years, of doing any work a man can do.’

For a sympathetic history, see Pamela McCorduck’s *Machines who think*,^{viii} written in the late 1970s and republished in 2004. As a field of human

endeavor AI exhibited many of the features of human systems that machines have trouble emulating, like value systems.

AI utopians valued the prospect of creating purposeful, autonomous devices that could achieve goals in environments that require adaptive capability, with no or minimal human intervention. Another utopian AI fantasy was that machines would give the bulk of mankind more time to think and enjoy things like opera and modern poetry.

More or less clever robots now operate on earth, in space and some nearby planets. Although these digitised devices have achieved some degree of autonomy, most rely on significant human interventions to achieve their purposes, and human error often has a profound effect on their life cycle.

So Wolowitz of *The Big Bang Theory* has used ‘how would you like to drive a car on Mars?’ as a pick up line at least twice. In neither instance did things work out as he would have liked; surprising twists occurred.

Real world AI outcomes appear to be human-machine symbioses that resonate with the purposes of human institutions and human users. These outcomes are often surprisingly twisted. Hindsight also shows *ordinateurs* have generated more demand for human labour than they have replaced, which is probably a good twist in a world with a growing human population.

Fluctuations and adaptations in the hype

As discrepancies between promise and delivery were noticed, glamour faded and the first flood of AI grants dried up. The founding optimists established a black box defense: the engineers were lagging in producing hardware that could enable them to realise their vision.

In the early 1980s the perceived threat of the Japanese ‘technopolis’ with its cleverly-automated production lines seems quaint now, but it had political force at the time. In that context AI returned from the fringes, perhaps improved by a spell there, to again become a glamour science.

Specialised hardware products like LISP machines and expert system software were developed at prices affordable by the US military. National security institutions can fund ‘just-in-case’ rather than ‘just-in-time’ research. This kind of research may not publish scientific results.

Glamour enables people to imagine much more than they know. The fantasy of massively intelligent, adaptive systems that could replace humans provided the Skynet conceit for the *Terminator* movies. In the commercial domain, expert systems generated most of its income from seminars and workshops, another form of show business.

This wore thin as many practitioners noticed that fees spent learning about these things did not enable them to earn fees. The industry that translates geek optimism into commercial cynicism moved on to inflate the next bubble—the *knowledge management* (KM) boom of the 1990s.

This fantasy was about siphoning what employees knew into machines so companies didn't have to depend on their people so much. The financial return on the IT-based KM investments strongly resembled sunk costs. Allied with a pitch to invest in more and better hardware, the KM boom succeeded in destroying amounts of shareholder value in major corporates.^{ix}

After the KM bust, smart people in the management discipline suggested that firms would be better off learning how to manage knowledgeable people better rather than investing in expensive techno-fantasies. Of course, leadership and management educators have an interest in selling their stuff, too!

The reality so far seems to be: a concept that begins as an AI grant may come to market as a technology. But the technology is not as an autonomous replacement for human intelligence or stupidity. It is a prosthesis or agent for a human mind or a human system: an advance in human-machine symbiosis.

Not that there is anything wrong with that, at least from the perspective of someone using the new technology to advantage.

Changing ideas about human intelligence and stupidity

Kurt Lewin's pioneering work at MIT in the late 1940s enabled the social sciences to theorise productively about evidence that learning is something most adults prefer to avoid, and many human groups actively discourage.^x Like other aspects of psychology and sociology, the implications of this for 'machine learning' were not immediately obvious in the IT field.

The idea that wise policy is the result of incremental trial-and-error rather than comprehensive rational thought emerged as a subversive thesis in public administration in 1959.^{xi} The implications for social policy of complex, intractable conflicts of interest in pluralist societies required the invention of the class of 'wicked problems' which cannot yield correct solutions.^{xii}

Herbert Simon preferred to call this class of problems 'ill-defined'. In 1978 he received a Nobel Prize. It was not for a break-through in the AI field. He received the prize in economics for his work on what we now call 'bounded rationality' and its impacts on organisational decision-making: an important advance in our understanding of natural stupidity.

In 1979 two psychologists, Kahneman and Tversky, published a paper on prospect theory.^{xiii} This provided a compelling account of the ways that cognitive biases and heuristics skew decision-making under uncertainty. In a sense Simon's work paved the way for these thinkers to terminate the glamour of Rational Economic Man.

By 1982 Richards J. Heuer had read Kahneman and Tversky down into *The Psychology of Intelligence Analysis* for the CIA. Ivy league business schools and high-end business journalism took notice of the ways that bounded rationality and social complexity affect the viability and profitability of businesses.

A new field called behavioral economics emerged to describe how natural stupidity influences human economies. In 2002 Kahneman received his Nobel Prize in economics. (Tversky had died.)

We note that clever technologies can enable human players to leverage natural stupidity in gaming markets. Equity markets and their regulators have had interesting learning experiences with the deployment of these technologies.^{xiv}

Despite the new technological enablers, natural stupidity provides a better set of explanations for the 2008 financial crisis. The technologies are not innately greedy, selfish, stupid or malicious: the humans that deploy the technologies supply those attributes.

AI now going in many directions simultaneously

So where is AI headed now? This is not a stupid question. But any attempt at a comprehensive answer is likely to a) have a high level of Kolmogorov complexity;^{xv} and b) be noticeably incomplete. Indeed the connotations of ‘AI’ may be as semantically misleading as those that arise from the term ‘computer.’

The bounds of complicated problems that AI can address are expanding in complex ways. The field has become fragmented technically, dividing into subfields each with their own conferences and journals and, indeed, domain ontologies. Few individuals have current awareness across the spectrum.

As in most areas of advancing research, we can detect differences between pure and applied intents. On one hand, there is a strongly academic stream including connections into psychology and cognition, and new formulations of logic to provide stronger ways of formally reasoning about problems.

On the other, mainstream AI now can now be defined in terms of a set of technologies applied to a wide range of complicated problems. Readers of our [ideas note on semantic web](#) will see a similarity: what starts off as a vision of what might be worth doing morphs into a body of work that can be pragmatically described as a set of technologies that actually do somewhat different things.

Some human capabilities that contemporary AI emulate include:

- Forms of perception such as
 - vision—optical character recognition, derivation of 3-D shapes from 2-D shading, facial recognition, visual signature and relative motion detection
 - hearing—speech recognition, voice identification and various kinds of signal filtering and audio signature detection
 - smell and/or taste—detection of volatiles in air, substances in solution; and aspects of touch, tactility, temperature sensing
 - proprioception, balance and awareness of accelerations &c.
- language understanding: natural language processing, machine translation, ontologies;^{xvi}
- reasoning: automated learning, expertise and planning and scheduling.
- robotics: autonomous physical movement and manipulation drawing on a suite of the previous capabilities and on situational awareness.
- situational or social awareness skills such as

- global positioning
- threat and target detection
- exhibited empathy
- making inferences about and responding to another entities intent, ‘emotional’ state or attributes.

Our starting point included technologies that make inferences about the compatibility of individuals who resort to online dating services and augment machine learning about experience playing an online game by reading the manual. We think we can read those down from this list.

No doubt we have left many things off this list that others regard as important. If you are one such person and have taken the trouble to read this far, please let us know.

Why has AI been so slippery?

AI began as a field where good people announced ambitions that turned out to be unrealistic. A classic instance was an MIT summer research project in 1966 that set out to ‘solve the vision problem’ in a few weeks.^{xvii} That project’s problem definition is still being decomposed decades later.

Is it not human to be reluctant to admit that you’ve committed an error of judgment, especially when the admission could cruel your career? So do we damn the founders of AI for never acknowledging—at least in public—that the problem of electronically replicating human perception and intelligence has proven intractable?

Once their sex drives begin to fade, shared indignation might be the greatest pleasure humans can share. So it is attractive to put eminent people into a pillory for talking tosh. But indignation, while natural, can be stupid. We may as well acknowledge that claiming attention makes sense for compelling social and psychological reasons, and so does a reluctance to admit you’ve talked tosh.^{xviii}

The Herbert Simon of 1978 had developed ideas of great and lasting value, and along the way he made some overly optimistic predictions. Indeed, cutting optimists some slack has merit. Alchemy did not deliver transmutation, but it did develop some capabilities useful to creating the science of chemistry.

Can we admit a kind of noble, if specious, art in the promotion of visions that play on ignorance and fantasy in order to muster support for daring ventures? Without this art humans would not have the built environments we find at Cape Canaveral or Point Bennelong.

In the 21st century, should we not be more sophisticated about both strategic misrepresentation and various forms of cognitive bias, including optimism? Professor Flyvbjerg, the very model of a 21st century social scientist, studied megaprojects in many countries over many decades. He found that:

... promoters and forecasters intentionally use the following formula in order to secure approval and funding for their projects:

$$\text{Underestimated costs} + \text{Overestimated benefits} = \text{Project approval}^{\text{ix}}$$

A cynic might substitute 'grant' for 'project' in this formula. But we also need to consider non-cynical appreciations.

- Firstly, humans are capable of maintaining a point of view in the face of contrary facts, and this kind of ego strength may be a component of creativity.
- Secondly, humans are capable of projecting human attributes onto non-human entities.

The second capability is called anthropomorphism. Like a teddy bear, AI appears to be capable of getting a lot of love without really having to do very much.

Of course, AI researchers are not machines. Fantasy is among the strong drivers of daring ventures of any kind. Likeable intellectual passion does (and probably should) influence people who dispense grants and assess research.^{xx}

Beyond that, when we reflect on the ways that AI has inspired genius merchants of dreams like Stanly Kubrick and Ridley Scott, we can see that the attraction of AI extends far beyond the limited circle of people interested in the problems of recursive programming.

The wide arc of anthropomorphic effects

Eliza, a comparatively simple program, was published in 1966 by Joseph Weizenbaum. Named after Bernard Shaw's Cockney girl, trained to talk posh, Eliza could engage human subjects in a conversation that bore a striking resemblance to one with an empathic psychologist.

The program applied very limited pattern matching rules to a human's statements to select its replies. Weizenbaum saw this only as a parody of an initial psychiatric interview. He was struck and deeply concerned by how quickly and deeply people subjects became emotionally involved with the Eliza.

Subjects took offence when he asked to view the transcripts, claiming it was an invasion of their privacy. Some asked him to leave the room while they were interfacing with their imaginary friend.

Weizenbaum's Eliza device was limited in its aims and application, but proved compelling for vulnerable people whom it could not take further by way of diagnosis or therapy. So it discovered a new arena for human-machine symbiosis, and this discovery underpinned a pervasive set of applications for purposes that Weizenbaum never envisaged.

Eliza's many great-granddaughters are called chatterbots. Call centres utilise them heavily. For example, IKEA uses a chatterbot called Anna to handle initial product queries.^{xxi} Anna knows nothing of love, the weather or the finer points of non-directive counseling, but she can tell you all about IKEA products.

This is an example of a small scale AI project that discovered exploitable principles now embedded in pervasive applications.

The arc of anthropomorphic considerations is perhaps best bounded by a long-term, large scale project which purports to be about human common sense. That example is Cyc,^{xxii} a sustained quest to ‘develop a true artificial intelligence’ begun in 1984 by Dr Douglas Lenat. He reported that in Cyc’s first decade ‘a person-century of effort has gone into building roughly 105 general concepts spanning human reality.’^{xxiii}

Cycorp reportedly claims success in arenas such as anti-terrorism and security; in supporting companies in developing semantic maps. Its business goal is to continue to obtain funding to support its research and development of a machine readable version of human common sense.^{xxiv}

In some reports its founder took care to emphasise that Cyc is about perseverant, consistent search of carefully-crafted filing systems rather than brilliant, creative insights. We note that common sense, intelligence and empathy are all lovable attributes for some humans.

What has AI achieved?

Let’s leave aside the set of cultural effects that reference artifacts like HAL, *Terminator* and *Blade Runner*, and focus merely scientific and technical outcomes. While the intellectual genealogy of AI is enlivened by claims that now seem overly ambitious and as yet unfulfilled, the known outcomes of the initial investigations are clever technologies that are adjunct to human systems. Like [semantic web](#), artificial intelligence now can be defined by a set of advanced software technologies.

It is fair to say the development of clever technologies has had a huge impact in many of the arenas where complicated real world problems are managed.

These problems can be addressed in better ways because of advances in the capacity to access and organise increasing amounts of data, to discern patterns in data and to deploy modeling and simulation to make sense of things in ways that complement, correct and extend the human capacity to describe phenomena, and generate and test hypotheses.

Here are a few of the many applications of AI research as clever technologies to complicated real-world problems:^{xxv}

- Data mining can discover hitherto-undetected patterns in very large collections of data using statistical machine learning technologies.
- Logistics and scheduling technologies use different forms of logic to automated planning. Fuzzy logic assists with managing uncertainty; constraint logic addresses the issue that all plans are subject to constraints at different stages. These technologies have made possible important features of GPS vehicle navigation systems as well as more esoteric capabilities such as Hubble Space Telescope’s autonomous planning and NASA’s interplanetary space projects.

- Fault diagnosis uses ontologies to underpin the detailed knowledge required for very specific domains. Specialised forms of logic such as Bayesian logic are used to handle the uncertainty in diagnosis.
- Image recognition: pattern recognition techniques such as neural networks have brought this application to a stage that particular images, for example faces, can be compared against a limited set of known faces in environments where the conditions, such as lighting, are well controlled.
- Machine translation: statistical natural language processing enables the analysis of large bodies of sample text in the target and originating language to identify common patterns to train machine translation features of the Google search engine amongst other places. These are quite practical if rather rough-and-ready in many cases.
- Animation: the movie industry makes use of multi-agent planning to simulate crowd or herd behavior. It is remarkable how effective relatively simple rules use the cooperation and competition between many agents to achieve a given goal without the need for detailed oversight by the animators. Emergent behaviour such as this is driven by technologies such as evolutionary algorithms and swarm intelligence.

These applications are not magic. Larger data sets can generate larger issues about the gaps and garbage in the data contained therein. Those issues are compounded by questions about the validity of assumptions used to draw inferences from the data.

We recognise deep, reflective thinking in humans by detecting the capacity for self-doubt; and the duty of doubt is not eliminated by clever technologies: it is made more onerous. Indeed the presence of natural stupidity can be indicated by the detection of unreasonable certainty or unreasonable doubt.

Australia needs strengths in the AI domain. As a nation we must address complicated issues—issues where some scientific certainty may be possible—as well as complex issues that science cannot resolve. We may ponder whether our strengths are equal to the great weight of cynically-enabled natural stupidity that encumbers public debate on some complicated issues, but that is another story.

We note that in the recent Australian Research Council assessment of research skills ANU scored the highest of any Australian university in AI, and substantially above world level.^{xxvi} ANU is capable at both at the fundamental level (e.g. logic programming) and in applied technologies (such as machine vision and machine learning).

Stopping point

There are three kinds of endings to a story: happy, unhappy, and the set-up for a sequel. Let's go for the third: the future of AI is a subset of uncertainty. Ignoring uncertainty is something humans often prefer to do: it saves effort and anxiety. The founders of AI managed to claim attention (and a great deal of funding) with astounding predictions of machines that might take over the world.

Will machines conquer the world? This is uncertain, but we can say with some certainty that people with better technologies have been conquering other people for a long time.^{xxvii} Whatever smart people imagined, feared or hoped 40 years ago about AI, no autonomous machine of superhuman intelligence has yet liberated us from that ancient social fact.

The SkyNet of the *Terminator* franchise remains fantasy. Indeed, we expect any rational machine-based super intelligence would be likely to consider its mid-term dependence on human taxpayers. And contemporary governments who rely on the consent of taxpayers for their legitimacy broadly agree that technological competition is fundamental to the history of warfare and commercial competition.

Indeed, this competition is apparently rooted in pre-history. As we worked on this paper, we found a 2008 study that deployed a ‘machine-learning algorithm’ to compile and analyse evidence about resource competition between Neanderthals and Cro-Magnons.^{xxviii}

This competition may have lasted ~10,000 years, and seems to have concluded ~25,000 years ago. One of the critical technologies in that competition was the spear. The contemporary study was apparently enabled by clever technologies created in the past 30 years.

The accelerating rate of evolution in human-machine symbioses may mean that critical competitions could occur in rather shorter time frames in the 21st century. That perspective provides us with an adequate reason to take a continuing interest in what comes out of the complicated, slippery and evolving AI space.

The writers

Richard Jones is adjunct professor of computer science at the Australian National University and a senior consulting associate of Wilton Hanford Hanover (WHH). **Bruce Hanford** is a director of WHH. **Helen Portillo-Castro**, the WHH client services manager, contributed by asking questions.

The notes

ⁱ See Bent Flyvbjerg 2007, ‘How Optimism Bias and Strategic Misrepresentation in Early Project Development Undermine Implementation’, in *Concept rapport No. 17*, available online.

ⁱⁱ Interactive entertainment (as distinct from online gambling and day trading) appears to be displacing older forms of commercial entertainment. These games use clever technologies to engage the humans they are designed to amuse, often in synthetic quests or contests. Humans have a recognised vulnerability—no matter how many times you say ‘RTFM’ to them, typically they will not read the manual. For a report on the use of machine learning with an RTFM feature to best a game, see <http://arstechnica.com/science/news/2011/07/how-a-computer-beat-civilization-by-rtfm.ars>.

ⁱⁱⁱ For a useful discussion of how online dating services interpret psychological research on couples to develop algorithms, see http://www.newyorker.com/reporting/2011/07/04/110704fa_fact_paumgarten.

^{iv} Interactive entertainment (as distinct from online gambling and day trading) reportedly generated sales of US\$65-billion last year [<http://uk.reuters.com/article/2011/06/06/us-videogames-factbox-idUKTRE75552I20110606>]. Online gambling is an issue for the established ~\$100-billion p.a. casino industry, which has successfully lobbied for regulations in various jurisdictions that Pricewaterhousecoopers says make it reporting its revenue impracticable. Paumgarten 2011 says US online dating revenue is about US\$1-billion, consistent with Googled sources.

^v The title of Jacques Futrelle's 1905 short story 'The Thinking Machine' referred to his fictional detective, Professor Augustus Van Dusen.

^{vi} Charles E. Lindblom, 'The Science of "Muddling Through"', *Public Administration Review*, Vol. 19, No. 2. (Spring, 1959) pp. 79-88. This classic current has nearly 6000 refereed citations, and now reads as a precursor insight for the Rittel-Weber construct of 'wicked problems'.

^{vii} Go to <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html> for a statement of intent for the Dartmouth Conference. The article at http://en.wikipedia.org/wiki/Artificial_intelligence is a useful folk history.

^{viii} *Machines who think: 25th anniversary edition*, Pamela McCorduck (2004), ISBN 1-56881-205-1, http://www.pamelamc.com/html/machines_who_think.html

^{ix} See Thomas Davenport and Lawrence Prusak, 1998 *Working knowledge: how organisations manage what they know*, Harvard Business School Press, for a post-boom critique of the knowledge management bubble.

^x Kurt Lewin developed the concept of individual minds and group dynamics operating in 'quasi-stationary equilibria' within particular social force fields. This concept was developed by Edgar H. Schein (among many) and helped us to define adulthood as a form of learning disability: adults generally refuse to learn unless they recognise that they cannot afford to avoid it. This in turn helps to explain a phenomenon Clayton Christensen examined in many case studies: most human organisations seem to prefer to fail rather than adapt to strategic changes in their operating environment.

^{xi} Lindblom 1959

^{xii} Horst Rittel and Melvin Webber, 'Dilemmas in a General Theory of Planning', *Policy Sciences* 4 (1973), 155-169

^{xiii} Daniel Kahneman and Amos Tversky, 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica*, Vol. 47, No. 2. (Mar., 1979), pp. 263-292.

^{xiv} For a report on the flash crash event, see 'How a trading algorithm went awry' in the *Wall Street Journal*. This available online at <http://online.wsj.com/article/SB10001424052748704029304575526390131916792.html>

^{xv} Dr Paul Whitbread has suggested, in a private conversation, that Kolmogorov complexity is an indicator of the level of irreducible complicatedness in a description of an object. This is a useful distinction from 'complex' or 'wicked' problems that contain irreducible level of uncertainty about how they can be described.

^{xvi} See our earlier discussion paper on Semantic Web at www.whh.com.au

^{xvii} See http://people.csail.mit.edu/paulfitz/pub/phd/thesis_01_introduction.pdf, p. 15. ‘Time and time again, the difficulty of implementing analogues of human perception has been underestimated by AI researchers. For example, the Summer Vision Project of 1966 at the MIT AI Lab apparently expected to implement figure/ground separation and object recognition on a limited set of objects such as balls and cylinders in the month of July, and then extend that to cigarette packs, batteries, tools and cups in August.... That “blind spot” continues to the current day....’

^{xviii} Theory about the construction of knowledge predicts that some experts will defend their intuitions against contrary evidence. A classic paper in this field is Kuglanki and Webster, ‘Motivated Closing of the Mind: “Seizing” and “Freezing”’, *Psychological Review*, 1996, Vol. 103, No. 2, 263-283.

^{xix} Flyvbjerg 2007, p. 53.

^{xx} See Slovic et al 2003, ‘The affect heuristic’, [http://people.usd.edu/~xtwang/DM\(GuangHua\)/Readings\(GuangHua\)/AffectHeuristic.pdf](http://people.usd.edu/~xtwang/DM(GuangHua)/Readings(GuangHua)/AffectHeuristic.pdf)

^{xxi} See http://en.wikipedia.org/wiki/List_of_chatterbots , <http://www.ikea.com.au> and <http://www.hugeinc.com/casestudies/ikea>

^{xxii} see <http://www.cyc.com/cyc/opencyc> and also <http://en.wikipedia.org/wiki/Cyc>

^{xxiii} Douglas B Lenat, ‘CYC: a large-scale investment in knowledge infrastructure,’ *Communications of the ACM*, Vol. 38, No. 11 (November 1995), pp. 33-38.

^{xxiv} See the MIT online Technology Review: <http://www.technologyreview.com/article/16264/>

^{xxv} Wikipedia is particularly strong in more detailed descriptions of the AI inspired technologies mentioned here; the community of practice is clearly engaged in that forum.

^{xxvi} See http://www.arc.gov.au/era/outcomes_2010.htm

^{xxvii} For a striking analysis of the effects of technical advantages that might seem trivial in a pivotal encounter between the Spanish and American indigenes, see Jared Diamond, 2003, *Guns, Germs and Steel: The Fates of Human Societies*, p. 67–82.

^{xxviii} See a report on the study by Banks *et al* ‘Neanderthal Extinction by Competitive Exclusion’ at <http://www.sciencedaily.com/releases/2008/12/081229105039.htm>