© 2012 World Future Society • 7910 Woodmont Ave., Suite 450, Bethesda, MD 20814, U.S.A. • www.wfs.org • All rights reserved.

Toward a Standard Metric of Machine Intelligence

By Richard Yonck

Introduction

Our world is becoming increasingly, inexorably intelligent. For ages, intelligence was deemed the exclusive domain of human beings. Later, as various sciences developed, the notion that animals could exhibit degrees of intelligence became increasingly supported and accepted.

Today, with technology racing ahead at an ever-accelerating pace, the idea of intelligent machines has also gained acceptance among many scientists, philosophers, and theorists. While such intelligence is still very rudimentary, some consider it possible—or even likely—that the exponential growth of these technologies will in time lead to machine intelligence exceeding our own. By how much the machines might exceed us remains unknown, but it could potentially be by a very significant degree.

As our technologies develop, we may even gain the capacity to better our own minds, making ourselves vastly smarter than any unaugmented person. In fact, such improvements may not be limited only to human beings, but could extend to other members of the animal kingdom, as well. Of course, such "uplifting" will create many ethical issues, as well as practical ones.

While we may well stand on the cusp of an intelligence explosion, we are far from ready for it. There are many ways in which we will need to prepare ourselves to properly deal with such fundamental changes. But our single most important tool may be one we haven't yet developed: the ability to accurately measure universal intelligence.

The Coming Intelligence Explosion

The exponential growth of computing power has been well documented ever since Gordon Moore's famous article "Cramming more components onto integrated circuits" was published in *Electronics* magazine (19 April 1965). The trend identified by Moore's Law has continued with astonishing consistency for a half-century. Others have noted similar growth in everything from hard-drive storage density to the price and speed of DNA sequencing. In fact, many of these trends have been shown to extend back many decades and perhaps even farther, demonstrating a pattern in technological advancement that appears almost innate. Technologists and futurists such as Ray Kurzweil and Kevin Kelly have written extensively supporting just this idea.

One key feature of this technological growth that has not been adequately measured is the degree to which technology is becoming more intelligent. While there is an intuitive sense that AI programs today are more capable than those of thirty years ago, and that those were considerably "smarter" than the serial instructions that passed through the first supercomputers, we really can't accurately say how much more intelligent they

Richard Yonck is a foresight analyst with Intelligent Future. His perspective on the future is informed by over 20 years experience as a computer systems programmer-analyst. He is also a contributing editor for *The Futurist* magazine and a blog editor for the World Future Society. He can be reached at ryonck@intelligent-future.com.

are or in what ways. For that matter, we can't even say with absolute assurance whether or not these machines truly exhibit intelligence at all. Projects exist that proclaim they have achieved the processing equivalent of a rat's cortical column or have emulated a cat's brain, but what does this actually mean? We can quantify so many other aspects of our world and our universe. Shouldn't we be developing the means to properly quantify non-human intelligence?

For years, AI scientists have worked to design methods of testing machine intelligence. Gregory Chaitin, David L. Dowe, José Hernández-Orallo, Marcus Hutter, Shane Legg, Ray Solomonoff, and Chris Wallace are but a few of these. According to Hernandez-Orallo and Dowe, "it is clear that there is a lack of measurement devices to assess the progress of artificial (general) intelligence, but the need for intelligence tests and tests for other cognitive abilities will be indispensable in the forthcoming decades for a plethora of bots, robots, avatars, "animats," etc. and any hybrid or collective system of these and biological systems (humans and non-human animals)."

Superintelligence

The idea of a superintelligent computer has been around a long time. But as work in AI progressed, it soon became apparent such a goal would be much more difficult to achieve than early expectations and popular media might have led us to believe. Nevertheless, work toward an artificial general intelligence, or AGI, has been underway in many different projects around the world. These include OpenCog, an open-source AGI framework; the Non-Axiomatic Reasoning System (NARS); and the Learning Intelligent Distribution Agent (LIDA). Assuming there is no inherent reason why human-level intelligence or greater cannot be achieved in a non-biological substrate (e.g., Orch-OR, the Penrose-Hameroff quantum-microtubule theory of consciousness), it seems all but inevitable that we will one day attain it.

The milestone of a computer of human-level or better intelligence creates the potential for machines of still greater intelligence. As computational power soars and the price-performance ratio improves, such systems will continue to speed up. Moreover, an increasing percentage of processor cycles may become available for use in the self-improvement of the system through such means as reconfigurable circuits, evolvable hardware, and self-improving algorithms. As each of these factors progresses, the relative intelligence of the system increases. Perhaps most importantly, this could potentially be a recursive process, which would result in a positive feedback loop. Should this occur, it could rapidly drive the evolution of a superintelligence. Such an event would be so world-changing that it has come to be referred to by many as *the technological singularity.*

The singularity was popularized by author and mathematician Vernor Vinge in a seminal paper published in 1993. Vinge used the term "singularity" to evoke the image of a gravitational singularity—a black hole—because the information cannot travel beyond its event horizon. Similarly, many think that the future beyond the technological singularity is unknowable, or even unimaginable. Since its conception, the idea has been examined and explored by futurists, as well as technologists and writers. Their conclusions have varied widely. Some believe such an event will never occur, while others think it could be just around the corner. Inventor and futurist Ray Kurzweil has gone so far as to specify an exact year when the singularity will occur: 2045. But whether such an event happens next year or late in the century, it would be most prudent of us to be prepared for the possibility.

Still more important is the ability to influence the outcome of such an event. If, as Kurzweil argues, such a progression is inevitable, we would no longer be the smartest entities on the planet. While it can be argued that an actively malevolent superintelligence is unlikely, even one that is indifferent to humanity could result in existential disaster. For this reason, concepts such as Friendly AI, as put forth by Eliezer Yudkowsky, are very important and need to be carefully explored. Ensuring a superintelligence that will do no harm may be humanity's best chance of surviving this next evolutionary leap.

Human-Equivalent AI vs. Human-Level AI

One of the major flaws in many machine intelligence tests is their exclusive focus on identifying systems that achieve human-equivalence. This anthropocentric approach presumes that nothing short of—or different from—human intelligence is considered valid. Considering our experience with studying nonhuman animal intelligence, this would appear short-sighted.

The need to distinguish human-*equivalent* AI from human-*level* AI seems critical. While perfect emulation of human intelligence *in-silico* may be feasible in theory, in practice it is likely to be extremely difficult to achieve. Potentially, every aspect of the biological processes involved would need to be translated with very high fidelity. Even minor variations could result in an intelligence that does not process "humanly."

Human-*level* AI is another matter. Achieving capabilities that are equivalent to those of the human mind could be very feasible if it is not limited to perfectly mirroring the underlying processes involved. For instance, some pattern-recognition algorithms are already superior to human abilities. This machine ability is not achieved by duplicating the processes our brains use, though some of the methods have been inspired by them. If researchers had been limited to replicating the brain's underlying processes, we would still be waiting for the development of a machine equivalent.

Society of Mind

MIT professor and AI pioneer Marvin Minsky wrote of the mind being built up from many non-intelligent processes or "agents." He suggested that these agents were features of the brain that had developed and been acquired through a billion or more years of evolution. Agents could interact and share information about the world through a number of different mechanisms. Perhaps most importantly for us, at some point some of these processes became self-referential and self-reflecting, which led to human self-awareness and conscious.

Building on this concept, Minsky showed how a series of computer functions could be structured in a similar way. These functions would combine into a "society of mind," not unlike, though also not identical to, our own.

Psychometrics/Human Intelligence Testing

Interest in human intelligence has existed for thousands of years, but only in the last hundred years or so have intelligence tests become formalized and standardized. At the turn of the last century, Alfred Binet, a French psychologist, formulated the first formal intelligence test, which sought to determine the mental age of students in an effort to identify those needing special help. From this, German psychologist William Stern proposed a method for calculating the intelligence quotient, or IQ, of the child being tested. Stern's method calculated a value from the ratio of the subject's mental age and their chronological age: (MA/CA x 100). This method had significant limitations, not least of which was the inability to apply it to adults.

In response to some of these criticisms, subsequent tests such as the Wechsler Adult Intelligence Scale (WAIS) were made up of sub-tests that more thoroughly gauged different abilities. But perhaps most important was the way these tests changed how their results were calculated. Instead of dividing the subject's mental age by their chronological age, tests such as Wechsler's compared the subject's performance against that of all others in the same age group. By fixing the median score at 100, with approximately 68% of the group falling between 85 and 115 (standard deviation 15), two important things were achieved. First, IQ results could be valid across all ages. Second, it was now possible to establish equivalence between different types of tests by standardizing the results across the entire population. This method of using the statistical tool of standard deviation has become an established technique applied to such tests.

Non-Human Intelligence Testing

Animal cognition is the study of animal mental abilities. It uses test methods from comparative psychology, including challenges such as mazes and operant conditioning chambers. While these yield useful information for comparative purposes, they fall far short of providing rigorously defined values that can be applied across species. Considerable differences in intelligence exist between different species and even between members within the same species, yet there is no unified standard for quantifying this. Existing behavioral tests only yield comparative results and human intelligence tests are completely inappropriate for these nonhuman intelligences.

Machine Intelligence Tests

The origin of testing machine intelligence is generally attributed to Alan Turing, who, among his many achievements, conceived of the Turing Test. According to Turing, a machine that could engage one or more people in a text-based dialogue and convince them it was human passed the test and was considered to have achieved human-level intelligence. While a ground-breaking concept, the Turing Test met with numerous criticisms over the years and isn't considered a serious goal by most working in the field of artificial intelligence. Besides objections such as Searle's Chinese Room argument, the Turing Test has an extremely anthropocentric bias. Essentially, it tests for humanity rather than intelligence.

Over the years, a number of tests have been proposed, although few have been broadly implemented. Linguistic Complexity, Psychometric AI, and several variations on the Turing test itself are among these. Many of the tests are impractical to implement, and their results are of questionable value, particularly because their expectation of testing for human-equivalent intelligence introduces an inherent bias.

A major problem when applying the strategies used in human psychometrics to testing machine intelligence is the inability to standardize values across a uniform population. Not only do computers not have a well-defined population, but unlike human intelligence, the upper bounds of their abilities are continually shifting to a significant degree.

Tests that would seem to have a reasonable chance of successfully testing nonhuman intelligence are those that use mathematics to rigorously define the value of a given challenge. Such an approach would help eliminate cultural (or species) bias, while providing well-defined values for interpretation. Machine intelligence tests such as the C-test and text compression meet this requirement by using algorithmic information theory, a subfield of information theory that studies the complexity of data structures.

The C-test has some potential to generate meaningful data about machine intelligence. Based on work done by Gregory Chaitin and others in complexity theory, it presents a series of abduction and prediction problems, similar to those in standard IQ tests. But unlike the questions in these tests, questions in the C-test are based on Kolmogorov complexity, a formal measure of complexity. As with many IQ tests, these are also formatted and presented as a series of increasing complexity. When administered to human subjects, the C-test showed a high correlation with classical psychometric tests. But when applied to computers, it was determined that the programs could be modified to better their scores for that specific task, thereby gaming the test and skewing the score.

In algorithmic information theory, the Kolmogorov complexity of an object is a measure of the computational resources needed to specify it. For example, the string 0101010101 is much more compressible than XMJAQBVILF and so has a lower Kolmogorov complexity. Such compression improves with a better ability to recognize patterns, identifying, learning, and reusing them. Using this approach, the degree to which an intelligent agent can compress and summarize information can be said to be one mathematically definable measure of its intelligence. However, Kolmogorov complexity has been shown to be incomputable, so an alternate method, Levin's Kt complexity, has been proposed in order to calculate approximate but practical values.

At the same time, the universal probability of a particular string or data structure can be defined through the mathematical theory of inductive inference, which was developed by AI pioneer, Ray Solomonoff. A string whose shortest description is small is given a higher probability than one that is longer. Solomonoff prediction theory weighs all of the possibilities to arrive at the most likely one. This essentially formalizes Occam's razor—the principle that, given the choice between two explanations, the one that is simpler is more likely correct. An agent that performs better using this method should be better able to predict, anticipating subsequent information based on prior patterns. Since the ability to identify and utilize patterns is a form of predictive learning, these combine to give a measurable aspect of intelligence.

In 2006, the Human Knowledge Compression prize, also known as the Hutter prize, was established to promote research supporting such compression tests. The challenge involves the compression of a 100-megabyte extract of Wikipedia. It's theorized that the best result will come from an agent with vast language and real-world knowledge. But whether such a system would translate into an artificial general intelligence remains debatable.

Other limitations for these tests include problems applying them to systems of low intelligence—for instance, those unable to understand natural language. In order to be useful across a wide range of machine intelligences, tests will need to be able to adapt to the level of the agent being tested.

Universal Intelligence Test

The Universal Intelligence test is another type of formalized complexity test. In it, the agent is challenged in a fully interactive environment using reinforcement learning methods, its actions generating observations and rewards. This agent works to maximize its reward, not unlike the player of a computer video game. As with the other tests that apply Kolmogorov complexity and Levin's Kt complexity, these challenges are designed so their complexity can be formally defined.

The difficulty of the environments and challenges increases over time in order to assess the intelligence of the agent. Because the environments can be tailored to any level, it is possible, at least in theory, to adapt the test to any intelligence, be it a simple adaptive agent or a superintelligence. The Universal Intelligence test is referred to as an "anytime" test because it can generate a useful value even over a short period, while converging on a more accurate value as more time is allotted.

The test is called "universal" because in theory it can assess the intelligence of any human or nonhuman, be it animal, machine or, if need be, alien. But in practice, the trick is to create tests and means of interfacing which can be successfully used with the test subject. While humans and even machines can be interfaced with through any of several methods, doing this with an animal or theoretical alien will be much more challenging. Nevertheless, a fully interactive environment would seem to offer a much better opportunity for successfully testing all forms of intelligence. Currently, work is underway to devise a series of environments and challenges capable of delivering an assortment of complexities, which would therefore be suitable for a range of intelligences.

Is It Possible?

There are those who remain unconvinced that useful values can ever be obtained from machine intelligence tests. AI researcher and Novemente founder Ben Goertzel states that he is "a bit skeptical of the practical value of mathematically defined tests for measuring real-world general intelligence." Given the view that minds develop and adapt to a world environment, it's easy to see why applying a numerical value would fall well short of defining any given mind.

However, much the same argument has often been made regarding psychometrics and tests of human intelligence. After all, how can a number capture the depth and intricacies of the mind? Nevertheless, results from these tests do have some value and have been shown to correlate strongly with a number of social outcomes, including education level attained, average income, and even health. In this case, even delivering a single value representing the g factor—general intelligence—has both merit and application.

Spanning the Gamut

As nonhuman intelligences develop and grow, there will be a significant need to rate, rank, and categorize them. Due to the nature of technological intelligence, these could potentially cover a much broader range than the baseline human population. Additionally, unlike the human population, the range and extremes of these intelligences will probably be developing and changing rapidly, relative to the timescale of biological evolution.

Human intelligence tests use a linear scale, typically ranging from zero to a few hundred. But because machine intelligences have an extremely basic lower bound and the potential for a very open-ended upper range, the use of such fixed scales will not be feasible in the long run. Given the potential for exponential increases in AI, it may make sense to use tests that can adapt to the intelligence level being tested and rely on logarithmic scales to gauge and track them. Examples of such exponential scales that are used in other fields include the Richter scale, used for determining the energy of earthquakes; and the Kardashev scale, which gauges the technological advancement of a civilization. Such an approach would allow for an open-ended upper scale while providing the means for useful analysis.

A Vector-Based System

Human psychometrics test a person for a handful of intelligence factors relative to the rest of the population. Values for these intelligence factors can be extracted from the raw data and analyzed, providing useful information about human intelligence.

Machine intelligences could potentially cover a much wider gamut of types and implementations of intelligence factors. In fact, it's very likely that many intelligences and intelligence factors will arise that we can't even anticipate and for which we have no conceptual precedent. Therefore we need a system that is flexible in the range that it can cover, and can also adapt to the emergence of novel types of intelligence. In theory, the concept of testing within a fully interactive environment could be used to evaluate novel intelligences. But given the potential range of minds and intelligence-factors, how might these be recorded, evaluated and studied?

One possible strategy may be a multidimensional, vector-based system. Such a system would allow for tracking, comparison and analysis of any number of intelligence factors by assigning a dimension to each. By implementing a vector system, each intelligence factor could be easily assessed and, where appropriate, recombined for analysis. Such a structure could also lend itself to use within the developing field of 3-D visual analytics, improving our ability to spot trends and patterns within the data.

Taxonomy

Such a universal intelligence test could allow us to identify and track the development of any number of nonhuman, transhuman, and posthuman intelligences. Many types of intelligences could develop over time to fill various niches in a vast intelligence ecology. Technologist Kevin Kelly has identified a range of possible types of mind:

- Super-fast human mind.
- Mind with operational access to its source code.
- Any mind capable of general intelligence and self-awareness.
- General intelligence without self-awareness.
- Self-awareness without general intelligence.
- Super-logic machine without emotion.
- Mind capable of imagining greater mind.
- Mind capable of creating greater mind. (M2)
- Self-aware mind incapable of creating a greater mind.
- Mind capable of creating greater mind, which creates greater mind, etc. (M3, Mn)
- A mind incapable of designing a greater mind, but capable of creating a platform upon which greater mind emerges.
- Mind requiring protector while it develops.
- Very slow "invisible" mind over large physical distance.
- Mind capable of cloning itself and remaining in unity with clones.
- Mind capable of immortality.
- Rapid dynamic mind able to change its mind-space-type sectors (think different).
- Global mind—large supercritical mind of subcritical brains.
- Hive mind—large super-critical mind made of smaller minds, each of which is supercritical.
- Low-count hive mind with few critical minds making it up.
- Borg—supercritical mind of smaller minds supercritical but not self-aware.
- Nano mind—smallest (size and energy profile) possible supercritical mind.
- Storebit—Mind based primarily on vast storage and memory.
- Anticipators—Minds specializing in scenario and prediction making.
- Guardian angels—Minds trained and dedicated to enhancing your mind, useless to anyone else.
- Mind with communication access to all known "facts." (F1)
- Mind that retains all known "facts," never erasing. (F2)
- Symbiont, half-machine/half-animal mind.
- Cyborg, half-human/half-machine mind.
- Q-mind, using quantum computing.
- Vast mind employing faster-than-light communications.

As Kelly himself says, this list is far from exhaustive. Also, while it defines general types of mind, a vastly larger taxonomy becomes possible if the overall balance of agents or intelligence factors is used to specify a unique type or species of mind—for example, a mind with only intuitive abilities and a total absence of analytical ones; or one that can readily grasp spatial relationships in a six-dimensional Calabi-Yau manifold, has rudimentary social skills, and no natural language ability whatsoever. The gamut of such an ecology of minds is truly vast, and human intelligence would be only an extremely small fraction of it. The range of possible minds could be visualized as in Figure 1.

Alternately, a series of qualities or intelligence factors could be identified and graphed in the following way, perhaps with base-line human intelligence set at the origin (Figure 2).

Figure 1.

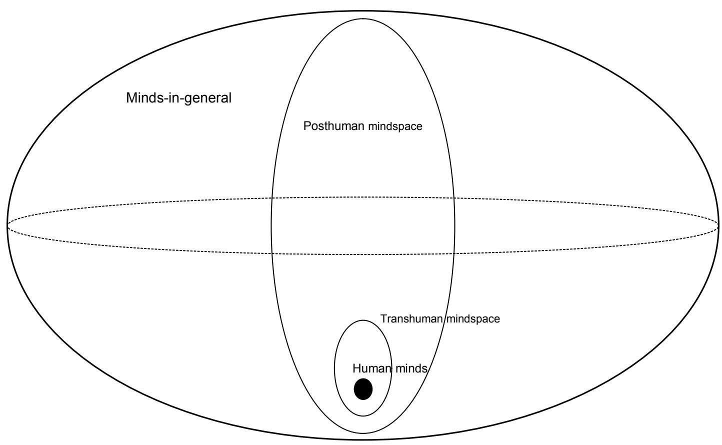

Figure 2.

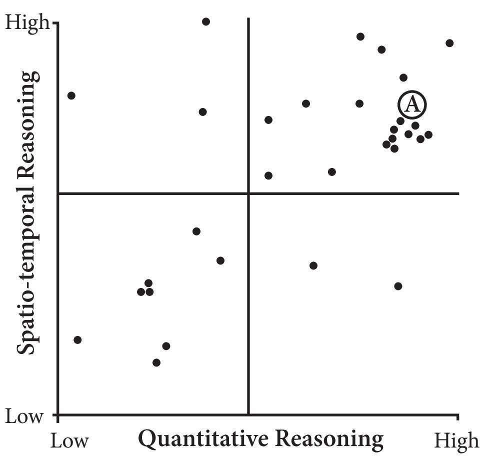

As well as taking an inventory of different types of nonhuman intelligences, analysis could identify under-developed niches, as well as optimal conditions. For instance, clustering (A) might indicate a grouping of intelligences competitively suited to specific conditions or environments.

One anthropocentric argument against such a plethora of minds is that we would never design and develop most of them. But this presumes that such control remains in our hands. Assuming this is not the case, and intelligences begin to generate and evolve according to needs and pressures that are not our own, it is easier to see the many possible intelligence niches being filled.

Futures Tools

The ability to accurately measure at least certain aspects of non-human intelligence provides information that could be used in the course of futures work. With sufficient data points, one can perform trend monitoring, projection, and analysis on a range of qualities and progressions. For instance, it might be possible to assess the likelihood of a particular type of intelligence coming into existence. Or we might be able to improve estimates of when, if ever, the technological singularity will occur.

A better understanding of possible developments in the intelligence ecosphere could also aid in the building of scenarios. Knowing that certain types of machine intelligence correlate to certain actions and behavior could provide better backgrounds against which possible futures can be set.

Conclusions

The nature of technological progress is to continually build upon and improve prior innovations. Nowhere is this probably truer than in the field of computing. AI computer technology can be said to exhibit a degree of intelligence, in that it is able to learn from its environment. While this machine intelligence is not yet on a par with human intelligence, it is improving at a rate many orders of magnitude faster than would be possible through biological evolution.

As a result, barring limiting conditions, machine intelligence will reach, and probably surpass, human intelligence at some point in the future. While estimates of when that will occur range from years to never, a significant number of experts consider it to be likely during this century.

In order to anticipate this event and hopefully be better prepared for it, a means of accurately measuring any type of intelligence is needed. Various approaches to this problem have been made, though many are less than up to the task due to factors such as lack of rigor in defining intelligence and anthropocentricity. Current tests that may be more adequate are those that use complexity theory and algorithmic information theory to judge the sophistication of the subject being tested. The use of a fully interactive environment offers the potential to test any level of intelligence, be it a basic adaptive agent or a superintelligence.

While such tests are still at a very early stage, there are indications that they represent a rigorous method of determining intelligence levels in both humans and nonhumans. If a fully operational test is devised, it will provide information that will be useful for everything from making technology-related policy decisions to futures work. Such a tool will go a long way in helping us influence the course of the future, a future in which humans will no longer be the most sophisticated minds on the planet.

References

Bostrom, N. "When Machines Outsmart Humans." *Futures*. Vol. 35:7, pp. 759–764. 2000.

Dowe, D.L., Hajek, A.R. "A Computational Extension to the Turing Test." Proceedings of the 4th Conference of the Australasian Cognitive Science Society. 1997.

Dowe, D.L., Hajek, A.R. "A Non-Behavioural, Computational Extension to the Turing Test." Proceedings of the International Conference on Computational Intelligence & Multimedia Applications (ICCIMA'98), pp. 101–106. 1998.

Goertzel, B. "Human–level Artificial General Intelligence and the Possibility of a Technological Singularity." *Artificial Intelligence*, pp. 171–18. 2007.

Goertzel, B., Wang, P., editors. "Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms," volume 157 of *Frontiers in Artificial Intelligence and Applications*, pp. 17–24. 2007.

Hart, D. and Goertzel, B. "OpenCog: A Software Framework for Integrative Artificial General Intelligence." Proceedings of the First AGI Conference, Ed. Wang, et al, IOS Press. 2008.

Hernández-Orallo, J. "Beyond the Turing Test." *Journal of Logic, Language and Information* 9(4): pp. 447–466. 2000.

Hernández-Orallo, J., Dowe, D.L. "Measuring Universal Intelligence: Towards an Anytime Intelligence Test, Artificial Intelligence." 2010. doi:10.1016/j.artint.2010.09.006

Hernández-Orallo, J., Minaya–Collado, N. "A Formal Definition of Intelligence Based on an Intensional Variant of Kolmogorov Complexity." In Proceedings of the International Symposium of Engineering of Intelligent Systems (EIS'98), pp. 146–163. ICSC Press. 1998.

Hutter, M. The Human Knowledge Compression Prize. 2006. http://prize.hutter1.net.

Kelly, K. "The Landscape of Possible Intelligences." The *Technium*. September 2008. http://www.kk.org/thetechnium/archives/2008/09/the_landscape_o.php.

Kelly, K. "A Taxonomy of Minds." The *Technium*. February 2007. http://www.kk.org/thetechnium/archives/2007/02/a_taxonomy_of_m.php.

Kurzweil, R. *The Age of Spiritual Machines: When*

Computers Exceed Human Intelligence. New York: Penguin. 2007.

Kurzweil, R. *The Singularity Is Near*. New York: Viking Penguin. 2005.

Legg, S., Hutter, M. "A Formal Measure of Machine Intelligence." In Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn'06). 2006.

Legg, S., Hutter, M. "Tests of Machine Intelligence." *50 Years of Artificial Intelligence*. Lungarella, Max, et al. (eds.). (4) pp. 232–242. New York: Springer. 2007.

Legg, S., Hutter, M. "Universal Intelligence: A Definition of Machine Intelligence." *Minds and Machines* 17 (4) 391–444. 2007.

Mahoney, M. "Text Compression as a Test for Artificial Intelligence." Proceedings of the National Conference on Artificial Intelligence, AAAI/John Wiley & Sons Ltd., pp. 970. 1999.

Minsky, M. *The Society of Mind*. New York: Simon and Schuster. 1985.

Searle, J. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3 (3), pp. 417–457. 1980.

Solomonoff, R.J. *A Formal Theory of Inductive Inference. Part I, Information and Control* 7 (1), pp. 1–22. 1964.

Turing, A.M. "Computing Machinery and Intelligence." *Mind* 59, pp. 433–460. 1950.

Vinge, V. "The Coming Technological Singularity, Vision-21: Interdisciplinary Science & Engineering in the Era of CyberSpace." Proceedings of a Symposium held at NASA Lewis Research Center (NASA Conference Publication CP-10129). March 1993.

Wallace, C.S., Dowe, D.L. "Minimum message length and Kolmogorov complexity." *Computer Journal* 42 (4) pp. 270–283. Special issue on Kolmogorov complexity. 1999.

Yonck, R. "An Argument for the Singularity." World Future Society. July 2011. http://www.wfs.org/content/argument-for-singularity.

Yonck, R. "Can Machines Be Intelligent?" Psychology Today. March 2011. http://www.psychologytoday.com/blog/the–intelligence–report/201103/can-machines-be-intelligent.

Yudkowsky, E. "Creating Friendly AI." Singularity Institute. 2002. http://singinst.org/ourresearch/publications/CFAI/index.html.