

Book Review and Discussion:

The Singularity is Near: When Humans Transcend Biology

Ray Kurzweil, *Viking Penguin*, New York, 2005, 602 pages plus index; \$29.95

Joseph Coates

Ray Kurzweil describes the "Singularity" as "a future period during which the pace of technological change will be so rapid, its impact so deep, that human life will be irreversibly transformed."

Harold A. Linstone, the editor of *Technological Forecasting and Social Change*, asked Joe Coates and four other practitioners of technological forecasting to review and discuss Kurzweil's new book, *The Singularity is Near: When Humans Transcend Biology*, in light of what Linstone calls its "great potential significance and controversial nature." The other commentaries may be found in *Technological Forecasting and Social Change* (vol. 73 no. 2, Feb. 2006).

A bit of deconstruction brought to the title of the book makes it clear that this is a secular religious tract. Kurzweil, not above a little humor, pictures himself on page 368 carrying a sandwich board with the title of the book as the message. The religiosity of the book shows in the frequent allusions to mortality. Perhaps consciousness leads to the awareness of our own inevitable death. Kurzweil describes how each of us, or at least our descendents, will be able to achieve immortality. Early in the book, he recounts his episode of type two diabetes. He recovered after rejecting his physician's advice, which he found useless, and doing his own research on what was known about the disorder.

Today, Kurzweil sees himself in tiptop physical shape. He consumes about 250 pills daily and takes frequent injections. He even goes further to say that since this began at age 38 (he is now 56), he so improved by all the medical measures of age, that he now is functionally a man of 40 (p. 211).

Interesting, engaging, and by most standards peculiar behavior. Peculiar behavior is usually driven by some hidden or less visible factors. In this case, I conjecture, by fear of death. Later in the book, he characterizes himself as a patternist, which I presume is a neologism, making the point that the essence of being a person are the patterns that comprise him or her. The patterns are the only thing that counts. Patterns are intelligence; intelligence is pattern. For him information is nothing; it is knowledge, namely the pattern, laid on information that counts. That line of reasoning allows him to see that as we move through a world in which non-biological things (let us call them "robots" for convenience) have the capability of

generating, producing, expanding and analyzing patterns, they will become, first, our aides; second, our peers; third, our masters.

Showing his secular religious components and his fear of death is that if the essence of a person is the patterns that comprise his or her life, then we will be able to transfer those patterns to the robots. As the robots then become increasingly sophisticated, we will survive and will then become immortal. That is the secular religious message of the book. Unfortunately, it has the effect of distorting the analysis in the first part of the book, where he writes about the singularity being near. Its closeness is conveniently fixed to be comforting to an adult 40 years old or younger. Much of what he claims in terms of the pace of information technology appears to be excessively fast. I have rarely read a book in which numbers like 1016 to 1055 were so much in use and presented with certitude and so little thoughtful criticism.

Before turning to what I see to be the critical scientific and social defects in his argument for the singularity, let me make two key points. First, the book is extremely interesting. I urge you as a reader of this journal to read it and, second, to be wary of his seductive central message.

Kurzweil is a master of expository writing. I have rarely seen such breadth in material handled so clearly. The book has a number of engaging features. The introductory quotations at the beginning of each chapter are mostly fresh and insightful. At the end of each chapter, and sometimes at a critical point in the chapter there is a dialogue between himself and Molly (every woman) 2004 or Molly at some future time. Other people, like Freud, get dragged into the discussion

as necessary.

I assume that the other contributors have substantially treated the technological details of the book, therefore I will try to avoid redundancy and assume that you have read or will read those discussions.

With regard to futurists, among whom he rightly counts himself, when he wants to beat on us he uses the label prognosticators. His key point is quite stimulating—that most futurists' work assumes a linear trend from wherever now is into the future, whereas his own analyses show that almost everything of importance is exponential in growth. So one of the problems for the prognosticators is to get over their commitment to linear, usually short term, approaches and recognize that the most important things in the technological world are going at a much, much more rapid pace. Furthermore, as that exponential development increases the exponent itself increases. What it all boils down to is that it is extremely hard to imagine what the future will be like if you stick with linear developments.

His analysis is clear from a quantitative point of view. What is implied, from a human day-to-day - living point of view, is a void except for the inevitable forecasting chestnuts such as elimination of disease. Unfortunately the only technologies he discusses are electronics and genetics, because his single focus of attention is information manipulation as the key to all developments and prospects for the long term future. A further crippling consequence of his narrow focus is that the role of society in the forecasted developments is ignored. Getting from here to wherever there is, is trivialized. Maybe it was just a cloudy day on Mount Olympus.

Kurzweil describes right off, on page 7, what the singularity is:

It's a future period during which the pace of a technological change will be so rapid, its impact so deep, that human life will be irreversibly transformed. Although neither utopian nor dystopian, this epoch will transform the concepts that we rely on to give meaning to our lives, from our business models to the cycles of human life, including death itself.

Notice how contentless the definition is. He regards people who subscribe to this as

singularitarians, and has a charming discussion of the diverse range of beliefs of singularitarians. It sounds very much like reading about the Christian church fathers before the year 600 (pp. 370-376).

Kurzweil divides the present and the emerging future into six epochs. First is the epoch of physics and chemistry, where information is on atomic structures. The second epoch is the biological-DNA epoch. Third is the epoch of neuro patterns. The fourth is the epoch of information in hardware and software design. The fifth epoch is the merger of technology with human intelligence, and in the sixth epoch the universe wakes, and patterns of matter and energy in the universe become saturated with intelligent processes and knowledge.

Notice that all six of these epochs are statements about information; the first information in atomic structure; the second, information in DNA; the third, information in neuro patterns, and so on. He says, on page 21, that the sixth epoch "...is the ultimate destiny of the Singularity and of the universe."

Kurzweil's second chapter is devoted to a theory of technology evolution and is instructive from the point of view of forecasting. He proposes his own law, "the law of accelerating returns," meaning that with more and more development of technologies, the payoff becomes greater and greater. Each epoch in the evolution of technology is the preparation for the next epoch, and he specifically says that in the sixth epoch there will be no distinction between humans and technology; "this is not because humans will have become what we think of as machines today, but machines will have progressed to be like humans and beyond."

Chapter 3, dealing with achieving the computational capacity of the human brain, is where the serious weaknesses in the book begin to show up. He never defines or indicates what intelligence is. He rather sees whatever the brain does as a computational exercise, and since the machines that we have do computation very well, we soon will be able to imitate in machines what the brain does. Kurzweil definitely disallows the notion of duplicating the brain. He rather sees that the problem for the computer is duplicating the brain's functionality. That is a big blind step because what exactly is that functionality he wants us to duplicate? He primarily writes about cognitive things and he talks about them in terms of computer capacity. But he nowhere describes, in a

way that you or I, an ordinary person, would think about intelligence, much less the other functions in the brain—the emotional, the memorial, problem solving, and so on.

By treating the brain as a number machine, he throws away most of the interesting parts of the problem as he would “duplicate its functionality.” However, later in the book, he refers to uploading the brain into the computer—and what exactly is he going to upload? He talks about this being accomplished with the use of “nanobots,” which will be circulating in our blood by the hundreds of millions and presumably will know where to stop at the brain and go from place to place to upload. But if the structure of the brain is, as he admits, not as good as a modern computer, what exactly is going to be uploaded—just the structure of the brain? He clearly does not mean that, but it is a blur as to what he does mean and hence, how it would be accomplished.

Chapter 3, as with the earlier and later ones, is an interesting digest of various kinds of biological and technical information, but it does not all hang together in the way one would hope. He writes extensively about the computer capacity and the brain’s computational capacities, but he fails to inform this reader and, I think, every other reader, about how exactly, by 2030, a thousand-fold increase in capability and a thousand-fold increase in speed of machines will relate to the human being. On the reasonable assumption that our brains are not designed to deal with that load of information, how will we deal with it?

The common notion that we only use 10 percent of our brain is just fiddle-faddle, cooked up by somebody probably as a joke or an off-the-cuff remark. There is no evidence about what percentage of the brain we use, but common sense suggests that we would not have a brain in its size, shape and complexity if we were not using all of it. Therefore, what good is a thousand-fold increase in memory and a thousand-fold increase in speed, if it cannot be assimilated. If it is only going to be exogenous, it is nothing more than a Library of Congress at my fingertips or perhaps some kind of didactic aide to problem solving, like getting the cube root of a 27-digit number. Kurzweil consistently fails to deal with the question of the linkage and the consequences and the assimilability of this material into the human being in the short run and with the human being in the long run.

The author goes farther off the deep end in Chapter 4, “Achieving the Software of Human Intelligence,” where he writes about reverse engineering the human brain. “Reverse engineering” is a well-known concept in which one takes a device, manufactured presumably, one of some simple or complex sophistication, and then figures out from analyzing it how it was put together and, in many cases, how it operates. Reverse engineering is the great tool of the counterfeiter.

The central flaw in this key chapter is that Kurzweil overlooks, neglects, or is unaware of, the fact that there are central problems in neuroscience today for which there is virtually nothing known of any significant fundamental value, that is, knowledge linking to computers functionally imitating the subjects of those central problems. Among these are:

- Why do we sleep?
- Creativity in all its forms, is not touched on.
- Kurzweil has a long discussion on consciousness that is sophomoric, uninformed, and largely a waste of time. He cannot simply deal with the question of consciousness and self-consciousness as a conundrum in the brain, but sets it aside as an issue for philosophers to chew over. That strategy is shockingly unscientific.
- How is memory stored and recalled? He takes a pass at the latter by saying that memory is holographic, but that is nothing more than a weak metaphor. He sheds absolutely no light on how the structure of the brain either creates a memory or reconstructs a memory when you want it to. If you think about last night’s dinner and try to reconstruct it, it is not stored in any one place. There is no neuron or clutch of neurons that are last night’s dinner. The great mystery is how is that memory stored in a half dozen different locations by the size, shape, color, etc.? Then, how are the components drawn back together? The obvious homunculus model of a brain within the brain that pulls it together does not make any sense. Kurzweil chooses to just gloss over these fundamental problems.

Chapter 4 is, however, full of good and interesting material. There is a rash forecast on page 145, “as we saw earlier, computational capacity needed to

emulate human intelligence will be available in less than two decades.” It is that tricky word, “emulate.” This retreat into the fog of vagueness is not only of no use; it subverts his arguments.

As the book goes deeper and deeper into human intelligence in relation to the machine, the author puts greater and greater dependence on nanobots, that is, nano-sized machines that will be pumped into the body, used universally, and will have many, many functions. He writes, earlier in the book, on how those nanobots will purge us of all the traditional ills of the human, replacing parts as needed, tracking down viruses and bacteria and disposing of them, and so on. When it gets to the brain, he says (p. 197),

Once the nanobot era arrives in the 2020s, we will be able to observe all of the relevant features of the neuro performance with very high resolution from inside the brain itself, sending billions of nanobots through the capillaries will enable us to non-invasively scan our entire working brain in real time.

I wonder what curious meaning Kurzweil gives to “non-invasively?”

Later on, he writes about the robots uploading the brain into a robot, and yet he completely overlooks the question of what exactly will be uploaded. If the neurons do not, one by one, comprise the basis of human memory, knowledge, or intelligence, then what is the point of just reproducing it in a robot?

Chapter 5, “Three Overlapping Revolutions,” is about GNR-genetics, nano-technology, and robotics. They are what Kurzweil characterized earlier as epoch 5, the beginning of the singularity. He sees that we are now at the early stages of the “G” revolution. The “N” revolution will enable us to redesign and rebuild—molecule by molecule—our bodies and brains and the world with which we interact, going far beyond the limitations of biology. He writes (p. 206);

The most powerful impending revolution is “R”: human-level robots with their intelligence derived from our own, but redesigned to far exceed human capabilities. R represents the most significant transformation, because intelligence is the most powerful “force” in the universe.

Intelligence, if sufficiently advanced, is, well, smart enough to anticipate and overcome any obstacle that stands in its path.

On page 260, Kurzweil makes the point that the most important revolution involving robotics is really artificial intelligence that exceeds human intelligence and does not necessarily have to be associated with what we think of as a robot. On page 261, he observes that “...once a computer is able to match the subtlety and range of human intelligence, it will necessarily soar past it and then continue its double-exponential ascent.”

On page 262, he makes the extreme claim “...but super intelligence innately cannot be controlled.” One of the ironies of the book is that the author is a man of incredibly broad knowledge and inventiveness. But dealing with the technology he promises will completely transform the world, he seems oblivious of people. He takes no account of the six billion who now live, or the X billion who will live, in the 20, 30, or 40 year future that he sees so singularly important. He writes as if everyone everywhere were from the upper quarter of income earners in Silicon Valley. It is a pity that Kurzweil has given no attention to the real consequences along the way of what he is proposing. Hence, he throws away all subtlety connected with their significance for us ordinary people. He can only speak in abstract generalizations which, to some people, would be alarming; to others, be incredible, but with no meaning for most of us.

Chapter 6 on “The Impact...” is where he further weakens his grip on reality. There is a long list of things that will be new or newly done or newly conceived or newly operated, but beyond whatever the noun is following “new,” he has nothing significant to say. The chapter is full of *avant garde* speculations which are fun to read, but difficult to assimilate. He quotes, on page 320, Richard Feynman, that

...it is one of the most remarkable things that in all of the biological sciences, there is no clue as to the necessity of death. If you say we want to make perpetual motion, we have discovered enough laws, as we studied physics, to see that that is either absolutely impossible or else the laws are wrong, but there is nothing in biology yet found that indicates the

inevitability of death.

It is hard to believe that Feynman had never heard of hunting and fishing, had never heard of predators and prey, had never heard that these are inevitabilities, certainly in the whole world of chordate animals. We also know, from all studies of all animals, that they age and that aging leads to disabilities and incapacities and, finally, the inability to function which is then succeeded by death. We also know that for all species, including the human, as death approaches (assuming it is not traumatic or it does not involve disease causing excruciatingly pain) our mind prepares us for death.

Chapter 7, “Ich bin ein Singularitarian,” ropes us into juvenility, an amateurish and unsatisfying discussion of such things as consciousness. Having been so precise, bold and extensive in his thinking, Kurzweil says about consciousness, on page 380, “precisely because we cannot resolve issues of consciousness entirely through objective measurement and analysis [science], a critical role exists for philosophy.” There is nothing to substantiate the point that he is making here. There are always relatively gray or blank areas in science, but there is nothing that I know of that suggests that we cannot deal with the problems of consciousness with more scientific research. Effectively, he is trying to back off from one of the limitations of his computational *uber alles* model by declaring that consciousness is somebody else’s game. Then he gets into an inane discussion of why, when he wakes up in the morning, is he him and not someone else. He apparently is totally ignorant of the possibilities that have been discussed scientifically and are subjects of research as well as of clinical treatment: such as amnesia, multiple personalities, the problem of trans-gender people who feel they belong to one sex but have the apparatus of the other, those people who feel they are possessed by another spirit, the anorexics who have a grossly distorted body image, and finally, those people who have some questions about who they really are. It is a pity he could not bring his dogged research skills to explore those areas and turn his intelligence onto saying something that was more than Philosophy 101.

Chapter 8, “The Deeply Entwined Promise and Perils of GNR,” is unsatisfactory. It amounts, more than anything else, to some casual discussions of the precautionary principle, simulations, strategies for holding back or limiting or relinquishing

capabilities, but really makes no contribution to either understanding current, short-term, or longer term risks. He gives some attention to fundamentalism. Under the threat of fundamentalism, he says (p. 414):

The world is struggling with an especially pernicious form of religious fundamentalism in the form of radical Islamic terrorism. Although it may appear that these terrorists have no program other than destruction, they do have an agenda that goes beyond literal interpretation of ancient scriptures: essentially to turn the clock back on such modern ideas as democracy, women’s rights, and education.

It is hard to believe that a man of his curiosity, resources, and scope of friendships could write such a simpleminded inanity about the problems of fundamentalism and the threats that it does or does not present to the future.

To summarize: two questions about the book are the essence of the criticism. Will the singularity be upon us in 40 years? Can a singularity occur? To the first question, the answer is almost certainly not. While 40 years is a good number for Kurzweil (his biological age as opposed to his chronological age is 40), the next 20 to 40 years gives him psychological comfort that he will benefit in the ways he describes. But it is hardly likely that it will occur, since three of the great mysteries of the brain—memory, sleeping, consciousness and its derivative, self-consciousness—are far from adequate understanding. Secondly, the notion that it will sweep the world is virtually impossible in terms of any even casual techno-economic analysis of the X billion people and the implication of “sweeping the world.” No doubt, there will be tremendous progress on all the fronts that he discusses, but that progress is hardly likely to be as exponential or as productive and beneficial to most people as he makes it out to be, simply because the solution of problems is not merely the development of computational competence.

On the question of whether singularity will ever occur, as he presents it with his epoch 6 machines basically taking over being a gazillion times smarter than we are and the universe becoming smart, he has moved into science fiction. The flaw in science fiction that makes it near worthless as serious

forecasting illustrates a principle of formal logic: that a falsehood makes anything else acceptable. There is little doubt that the capability to understand all aspects of natural, animate, inanimate, human and nonhuman existence more fully is the unfolding program and agenda for the future, whether we plan it or not. But the likelihood of it going through the six phases he mentions is simply not clear as even possible, much less as in the cards, because he ignores personal and institutional consequences and responses as his dream unfolds.

So, all in all, this is a fun book to read. Enjoy it and learn from it, but do not take the central thesis or claim seriously by joining Kurzweil's secular sect.