Response to Kurzweil

In *Exponential Growth an Illusion?: Response to Ilkka Tuomi*, Ray Kurzweil comments two of my papers that discuss the development of semiconductor and computing technology.[1] Kurzweil used his response in his Accelerating Change Conference plenary keynote 14 September 2003, where we also had a debate on Kurzweil's hypothesis. My latter paper was written as a background paper for the conference. Kurzweil's comments now give me an opportunity to clarify some perhaps confusing points in my previous papers.

My first paper *The Lives and Death of Moore's Law* was a historical study on the origins and validity of common claims concerning Moore's Law. To interpret its statements correctly, it is helpful to note that in the paper I did not make any specific claims concerning the future of computing. The paper was simply a search for the "historically correct" version of Moore's Law, and an empirical study on the validity of its popular versions. Using the data that the paper describes, I tested the various versions of Moore's Law, and didn't find scientifically valid support for them.

The paper made a number of observations. First, Moore's Law has clearly escaped its original domain, and most existing descriptions of Moore's Law are historically wrong. It appears that Moore's Law has often been extended without understanding what kind of evidence would be needed to support the extensions.

A second observation was that claims about exponential trends have typically been based on selecting only those data points that support the hypothesis of exponential development. This is a rather elementary methodological error. If we only accept exponential trends, all data points that do not align with the trend are easily regarded as irrelevant and exceptional. A more methodologically sound approach is to start from the data. If there are "outliers" or "exceptions," it is useful to explain why they are special. We do bad science if we drop data points simply because they do not fit our theory.

For example, Kurzweil argues in his response that outliers are irrelevant because we are interested in the most "cost-efficient" chips. This allows us to ask whether in fact the mentioned outlier points represent "not cost-efficient" chips. I don't find any evidence on this from Kurzweil's texts. In fact, based on normal economic rationality one would expect that all existing chips are economically cost-efficient. Otherwise they would not be introduced, produced, and sold.

In general, cost-efficiency of course depends on the use of chips. The "outliers" in MIPS graphs typically represent slow processors that are optimized either for low chip cost or low power consumption. These points are easily dropped from the "real trend" when a specific use of microprocessors is assumed to be the relevant one. From

---

[1] http://www.kurzweilai.net/meme/frame.html?main=/articles/art0593.html. My first paper is Tuomi, I. (2002) The Lives and Death of Moore's Law. First Monday 7(11), November 2002, http://firstmonday.org/issues/issue7_11/tuomi/index.html; and the second paper is: Tuomi, I. (2003) Kurzweil, Moore, and Accelerating Change. Working paper prepared for the 2003 Accelerating Change Conference, September 13-15, Stanford, available at http://www.jrc.es/~tuomiil/moreinfo.html.

Kurzweil's discussion it is clear that he means "cost-efficient" within a class of applications where chip cost or for example power consumption are not constraints. One problem with this approach is that only a very small minority of semiconductor chips belongs to this class.

For a similar reason, it is, for example, empirically untrue that "semiconductor component counts are doubling regularly." Different semiconductor product classes scale at very different rates. Just to give an example included in my second paper: if we take the worst performer, communication chips in the first half of the 1990s, the doubling time for performance per constant dollars would be 22 years. For this reason, I discuss the different types of chips in my second paper, but try and not mix them into a single abstract bundle of "chips." Instead, I stick to empirically relevant chip categories, such as Intel 80x86 microprocessors.

This hopefully clarifies one important difference in the ways Kurzweil and I read our graphs. Whereas Kurzweil selects points that can approximately be aligned with an exponential trend curve, and drops data points that do not fit his theory, I simply use all chips within a reasonably coherent empirical chip category. Where Kurzweil sees a line, I see a scatter plot full of dots that can be connected in many different ways. This difference is clear, for example, in the way Kurzweil comments my graph that shows Intel 80x86 MIPS ratings. Kurzweil cannot understand how I am able to hold a view that the MIPS ratings don't support the claim about exponential growth. He sees the exponential segments that I fit to the underlying data, and becomes confused as I still claim that the data does not support claims of exponential growth.

The reason is simple. The exponential segments are generated as least-squares exponential fit to the data (I don't fit curves by hand, as Kurzweil for some reason assumes). By definition, the fitted curves are exponential and on a logarithmic scale they look straight lines. I have segmented the underlying data roughly according to main processor architectures to see whether there is a stable trend. The answer is no. First, the actual evolution of data points is not captured by exponential trends. Second, if you connect segments from two time periods, the exponential "doubling time" changes. By looking the data instead of the line segments that I have generated, one can see that there is no good exponential trend that would capture the characteristics of the underlying data.[2]

I do not think that we disagree with Kurzweil on the underlying data in this specific case. The difference arises because he assumes that my data points include all kinds of irrelevant points that confuse the grand evolutionary picture of technical progress. He projects an imagined exponential curve on top of my data points, and sees order where it actually may not exist. Kurzweil says that his way of connecting the dots is the right one and other possibilities are irrelevant. For me, exclusion of data points without good justification represents bad methodology.

---

[2] The segmentation that I use is of course somewhat arbitrary. I tried to bundle together microprocessor chips in sets that represent developments over roughly a decade, but the architectural changes makes it more reasonable to check the more advanced processor families also separately. I used MIPS ratings that were available at Intel's processor data set. The newest processors did not have MIPS ratings. I assume that the reason is that Intel knows that architectural changes have made MIPS comparisons meaningless and that usable MIPS data does not exist.

Even if a good justification were to be found, generic claims about semiconductor development would usually be incorrect, as many developments in this technology category would be left unexplained. This becomes particularly relevant when Kurzweil links his law of accelerating returns with Darwinistic models of evolution, arguing that technical change can be understood as the selection of most successful technologies. Most semiconductor chips in any well-defined category (such as 80x86 Intel processors) do not represent maximal "performance." Most semiconductors are not manufactured using best available manufacturing technologies and chips that are do not usually represent maximum complexity, as I note in my second paper. My point is here that Kurzweil talks about evolution of information technology, but he leaves almost all examples of this technology out of his pictures. When exceptions become the majority, they are not exceptions anymore.

As Maslow and others have noted, if we only have a hammer, the whole world looks like a nail. Kurzweil states in his response that the empirical results that I review in my papers "defy common sense and clear observation." Yet, simply browsing through the pictures that Kurzweil presents as evidence for exponential growth, one cannot but wonder whether Kurzweil has fallen into the Maslowian cognitive trap. For example, Kurzweil's graphs of "Total Bits Shipped," "Microprocessor Clock Speed," and "Internet Hosts" show deceleration instead of accelerating returns.[3] Kurzweil also argues that my graph that shows computer and software investment growth in the US actually looks exponential. To me, it consists of three relatively linear segments, a segment of about 3.4 percent growth in the 1960-69 period, almost flat segment in 1970-79, and an almost linear faster growing segment for 1980-2002. Between the first two periods, there is clear slowing down, and during the last period there is no exponential growth. Kurzweil also apparently sees exponential growth in curves that mathematicians would not call exponential. An example of such a non-mathematical exponential trend with doubling time of 12 months is Kurzweil's figure on Internet Service Providers, shown below. I get a bit confused when I see a graph that has nothing to do with exponential growth, labeled with a constant doubling time. The confusion increases when I try to figure out what are Kurzweil's ISPs over a decade before ISPs were allowed to provide services, or at times when the internet protocols were not yet in use. Furthermore, Kurzweil is perhaps too US-centric in this case if he really wants to prove a point about technical evolution. The ISP dollar costs and their rates of change have varied greatly across different countries.

---

[3] Of course, anyone who has studied Internet host counts knows that these numbers do not represent anything real. Here I am not arguing that Kurzweil uses bad data; instead, I'm simply pointing out that he apparently sees exponential curves where other people probably would not see them.

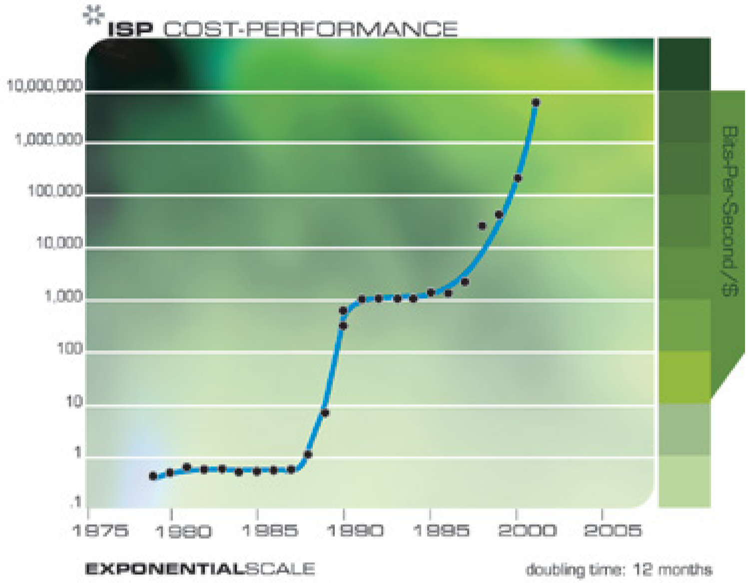

Figure 1. Source: Kurzweil, The Law of Accelerating Returns.[4]

One way to simplify the descriptions of the dynamics of semiconductor technology is to interpret exponential trends as envelope curves. Implicitly, Kurzweil seems to rely on such envelopes, and in some cases this approach makes sense. This approach is closely related to Moore's complexity limit law. In 1979, Moore drew an exponential envelope curve, under which all existing chips could be located. According to Moore, this graph represented the maximum theoretically attainable complexity. In Moore's graph of Intel's chips, most semiconductor chips were located one or two orders of magnitude below the curve. This approach, however, requires that we draw envelope curves using dimensions that are continuous and well defined.[5]

Moore's "limit law" was theoretically coherent as it was based on the well-defined dimension of number of transistors on a chip. The position of data points is unambiguous as soon as we know how many components they have.[6] In contrast, Kurzweil's limit law of most powerful processors and his limit law of most cost-efficient chips are not well defined. We cannot locate data points on these axes without making further assumptions.

One should also note that historical envelopes have no real predictive power. With envelope curves we can study historical development and, for example, its regularity, but the underlying dynamic remains a black box in this approach. Good predictive models require that we can know when they break down and when the black boxes need to be opened. Reliance on historical trends to predict future is a simple example of technological determinism. Moore's Law is sometimes used in a way that comes

---

[4] http://www.kurzweilai.net/meme/frame.html?main=/articles/art0134.html

[5] Futuristic extensions of "Moore's Law" that, for example, see quantum computing as the next paradigm in computing, have to address the fact that quantum logic is not a simple extension of two-valued logic. Quantum computing machines are not Turing machines.

[6] Here I am of course simplifying a bit. The number of components is not exactly unambiguous and the timing of the chips requires some care. Looking the original patent applications for the first integrated chips, one can see that they consisted of several components and transistors, instead of a single transistor.

close to naïve technological determinism, but I would not expect anyone with real interest in technological change to ride such a dead horse.[7]

The problem of choosing relevant data points relates to the third observation made in my first paper. Comparisons of "processing power" ratings are really very tricky. Probably most professionals in this area would say that MIPS stands for "meaningless instructions per second." Processing power depends on the configuration of processor architecture, hardwired instruction sets, the problem at hand, the problem definition as represented in software, and the compiler that is used to translate the problem definition into machine instructions. Kurzweil probably would not intentionally chart computers on a trend line using MIPS for some computers and MFLOPS for others.[8] Yet, similar comparisons between qualitatively different "MIPS" are necessary if we want to draw long-term trends of "MIPS" increase. MIPS, and other processing power measures, make sense only when we compare similar processors and similar programs. "Processing power" measures are widely known to be empirically misleading and they are mainly used for marketing purposes.[9] This is one complaint I make. If we draw trends, the axes should have some empirical validity.

Kurzweil seems to be immune to this problem. He draws, for example, graphs where mammals, sailing boats, television sets, and the World Wide Web happily co-align in an exponential chain of being. Of course, Kurzweil could argue that the Internet is an evolutionary successor of TV and the sailing boat, which in turn are offsprings of iron and humanoids. If that is the case, I think we need to add Barbie dolls, toasters, surfboards, Coca-cola cans, toads and humu-humu-nuku-nuku-apua'a[10] to get a fuller picture of the underlying logic. More fundamentally, my point is simply that if we create surprising associations, we have to explicate their justification and underlying logic. Only then we can discuss whether the underlying logic makes sense.

The fourth observation in my first paper was that claims about "cost of computing" are economic arguments. They require some economic sophistication. Just to give an example, the GDP and "constant dollar" numbers used by Kurzweil already discount improvements in computing technology. The U.S. National Accounts and CPI are based on hedonic estimates of price changes in computing. Although Kurzweil claims, "the hedonic model has little validity," in fact he completely relies on it in his cost

---

[7] Studies on technical change and history of technology have repeatedly shown that technologies change in constant interaction with social and economic processes. Technological "improvements" do not drive change, as technological determinism assumed. There is no abstract "selection" of "best technologies" because the difference between good and less good technologies depends on the way they are adopted and applied in social life. Economic determinism is a variation of technical determinism, as it assumes that new technologies are adopted simply because they are cheap. This type of economic determinism would require, for example, that people breathe at increasing speeds simply because air is free. It is however, in theory, possible to argue that in specific circumstances technological determinism is a good approximation. I have not seen such arguments, although my first paper comes close to arguing that. I note that semiconductor technology has developed in a way that has been surprisingly independent of the rest of reality, and that this can economically be viewed as effectively infinite demand. The infinite demand is an oversimplification, as Kurzweil notes. I borrowed the concept from Moore.

[8] He may, in fact, do this for some older computers, as their processing power is sometimes given as floating point operations per second (FLOPS) instead of MIPS.

[9] I don't believe that my shirts are getting exponentially whiter just because I'm year after year told that there is now an even better detergent on the market.

[10] It is a fish. Rhinecanthus rectangulus, to be exact.

estimates. It is also useful to note that most of the measured increases in economic growth rates in the U.S. come from these hedonic corrections in the latter half of the 1990s. Before that, the famous “productivity paradox” stated that computers and information technology do not have any visible impact on productivity. In fact, if the quality corrections in semiconductors are not taken into account, the U.S. productivity growth disappears or becomes negative after 1974. This is clearly against the intuitions of people who work with computers. On the aggregate level, however, computers seem to be irrelevant. The name of the “productivity paradox” comes from this: we think it is obvious that computers increase productivity but empirical studies do not support the claim.[11]

Engineers easily react to such economic challenges by noting that they don’t believe in economics. If that is the case, they should, however, drop all arguments about “constant dollar prices of computing,” as well as arguments that use economic growth numbers. You cannot have it both ways; either you need to stick to engineering, or you have to study economics enough to be able to say when exactly it breaks down.

The problem of measuring prices in “constant dollars” becomes exceptionally interesting when we study long-term trends and technological change. With some exaggeration, one could ask what the price of computation was before there were dollars. Where should we put the dot for the abacus? This is a theoretically interesting and open question. We already know, however, that consumer price indices cannot be used to measure historical costs of computing in the way Kurzweil uses them. Here is a fertile area for Kurzweil’s team for further research.

Whereas my first paper does not mention Kurzweil, in my second paper I explicitly discussed Kurzweil’s claims. I understood that one of his main claims is that evolutionary selection of the best technologies leads to increase in the resources that are used to develop them. Kurzweil calls this the “law of accelerating returns.” It is closely related to arguments that have been earlier put forward by Kuznets and Schumpeter, and versions of it have become popular in economics as the idea of positive returns, path dependency, and network effects. Kuznets, for example, argued that investments flow to those sectors of economy that show promise for exceptional returns. Schumpeter described this as “swarming” of entrepreneurs into new technology areas that creatively destruct old industries and products.

Kurzweil, therefore, is in a good company with his “law of accelerating returns.” I think he misses, however, the main underlying argument. Schumpeter’s core idea was that old technologies and technical “paradigms” repeatedly become superseded by new innovations. For example, Carlota Perez has argued in the Schumpeterian framework that changes in techno-economic paradigms and key technologies define long economic waves of economic growth and decline. According to Perez, the current economic growth wave is driven by microelectronics but, by definition, it will be followed by other technologies. There is no obvious reason why computing or information processing would be related to the next key technologies. According to

---

[11] The literature on the impact of information technology on economic growth is vast, and there is no general agreement on the matter. There is some evidence of productivity increase at the firm level, but the link between aggregate productivity increase and firm-level task productivity measures is unknown. I’m currently writing an article on the topic, which summarizes some of the discussions and includes references to the relevant literature.

Perez, the techno-economic paradigm that preceded the microelectronic one was based on petroleum.[12]

At times Kurzweil seems to believe that the exponential trend is actually created by a sequence of logistic curves. A related view underlies influential studies on disruptive innovation.[13] The difference between long-wave theories and sequential models of disruptive innovation is that the latter focus on specific product categories and relatively well-defined industries (e.g. hard disk industry) instead of overall growth. Kurzweil, in effect, assumes that there is one fundamental product category that can be traced through millennia (i.e. “technology”) or through centuries (i.e. “computation”), but on the other hand he promotes the universal validity of exponential trends to the extent that he does not say much about the situations where exponential trends disappear. Of course, a simpler argument would be that technical trajectories are created by underlying s-curves, and that overall improvements occur when we jump from old product categories to new ones. This, however, would mean that the Singularity is not necessarily near. We get to singularity only if the envelope of s-curves is exponential, and if there is quantitative continuity across sequences of technologies. Someone may assume that common sense and everyday experience proves that exponentials are everywhere. Others may think that s-curves are closer to common sense and that the “overall” trajectory is simply abstract fiction that reflects some taken-for-granted assumptions about the nature of progress. For one direction of another, we need some argumentation and justification.

In my second paper I therefore tried to see whether there, in fact, has been accelerating returns in semiconductor and computing industries. The answer is a qualified yes if we measure the overall growth of the computer and software sectors in the U.S. This is simply because the investment stock of computers and software was close to zero in the 1960s. Today there are more computers and we buy them more than before. This is what we would expect for any new product category, for example, Barbie dolls. I have nothing against the argument that information processing technologies are now more important than forty years ago. On the other hand, I am not completely sure how important computers are and to what their importance should be compared. Mattel has sold over billion Barbies since the birth of the first chip. I guess a generation or two of American girls have well understood the point that the future is in plastics. It is, however, not completely clear whether plastics are more important than computers in modern life. It is difficult to make computers without plastics and it is difficult to make plastics without computers.[14]

If we look the rate of growth in the semiconductor industry, we see deceleration since the early 1960s. Here Kurzweil misunderstands the logic of calculations that I make. To check the validity of the “accelerating returns” hypothesis, we have to check

---

[12] More recently, Perez has characterized the previous techno-economic paradigm as based on oil, automobile, and mass production (cf. Perez, 2003 Technological Revolutions and Financial Capital, Edward Elgar.)

[13] E.g. Tushman & Anderson, Utterback & Abernathy, Abernathy & Clark.

[14] Of course, Barbies invite through the back door discussions on mind and body, and abstract spirit and concrete materia. Computers perhaps look important because they promise one day to fulfill the dream of abstract intelligence. In fact, however, toys seem to be quite important for the evolution of technology. Much of the recent developments in computers have been driven by computer games. Perhaps at some deeper cosmic level Barbie and Max Payne are connected and express the forces of evolution, and the structure and order of nature.

whether the resources in this “successful technology niche” have grown across time faster than other less successful niches. In such a calculation we have to compare the overall economic growth rate with the growth rate of the specific industry. If the industry grows as fast the rest of economy, we don’t have accelerating returns in Kurzweil’s sense. I do this analysis by subtracting global semiconductor growth rates from the annual growth rates of the US economy. This is an imperfect procedure, but gives a reasonable first estimate of the validity of the accelerating returns hypothesis. As the results show deceleration in returns, I state that the history semiconductor industry does not seem to support Kurzweil’s hypothesis.

Kurzweil believes that I quote a World Semiconductor Trade Statistics report that shows annual industry growth of about 18 percent in the 1958-2002 period. There is no such report. I use the original nominal sales data and it would of course normally be corrected for general inflation. I do the additional adjustment for overall growth, as noted above. It should also be noted that much of this growth comes from population increase. The fact that we can calculate an average growth rate does not mean that the growth would have been exponential. Indeed, looking my Figure 4 on semiconductor shipments in Kurzweil’s response, one can see that the short term growth has been cyclical and the long-term trend has been declining. In the 1990s the global sales were relatively flat, except for the bubble years, as can be seen from the graph in my second paper.

In my paper I only checked the industry inflows from global sales. During the Accelerating Change Conference, Steve Jurvetson tried to convince me that I should instead look the investment inflows. This I have not done. Kurzweil’s research team may find something interesting by combining all resource flows to semiconductor industry. Capital investments, however, are supposed to reflect future income flows and profits, so that revenue-based calculations should not lead to qualitatively different results, at least if we assume competitive capital markets.

A couple of specific points in Kurzweil’s response also deserve comment.

Kurzweil quotes my second paper as stating that also the number of scientific papers, important innovations and engineers have been growing exponentially for centuries. In my paper I quoted the founder of scientometrics, Derek de Solla Price, who noted the exponential trends half a century ago, and made them popular in discussions on technology and science. I did not, however, say that Price’s exponential trends would be empirically valid. In fact, more recent research has shown that there are no exponential trends.

Kurzweil also repeatedly quotes the average rates of change that I give in my papers. Kurzweil sees them as an evidence of exponential growth. There are three points to note here. First, as noted above the fact that we can calculate average rates of change does not imply that the rate of change during the specific time periods would equal the average rate. Linear, logistic, exponential and sinusoidal curves all have average aggregate growth rates over a given period of time. As I note in my papers, there have been big differences in the actual growth rates in the different time periods. Most importantly, the rates often seem to slow down. The main reason why I discuss the different rates is that they are different. This is also the reason why I note that the change has not been exponential. One could argue that year-to-year changes are simply noise in the long-term exponential trend, and that we need to smooth out the short-term variations. The different time periods, however, typically have different averages, which is the basis for my claim that the data do not support the claim that there has been a constant exponential trend. Of course, when I focus on short time trends, Kurzweil may say that I miss the real long time evolution. But as Gilda Radner used to say on Saturday Night Live, "it's always something." In long-term trends we easily start to compare apples with orangutans.

Second, Kurzweil argues that I miscalculate doubling times and give systematically too long doubling times. I tried to clarify the issue in a footnote in my second paper, apparently without success. When I talk about doubling, I talk about doubling. When Kurzweil talks about doubling, he sometimes talks about doubling and sometimes about halving. These are different things. As I note in my paper, 0.7x0.7 is about 0.5. Annual decrease of about 30 percent, therefore, leads in two years to a halving of the original amount. As mathematics tells us, square root of two, however, is about 1.41. To double the original quantity, we need to add 41 percent for the accumulated quantity for two years. Theoretically correct way of using rates of decline or rates of growth depends on the process that we are talking about. If we talk about growth, we have to count upwards, if we count decrease, we have to count downwards. Kurzweil's arguments typically deal with doubling and increase, and the correct approach therefore typically is to count upwards. If we would simply make the argument that computing cost is declining, we could count downwards. If we try to argue that processing power doubles we have to count upwards. In any case, we have to be explicit what we are doing.

Third, as I also note in the same footnote, people often mix geometric growth and exponential growth. They are different. Kurzweil uses geometric growth and interprets my numbers apparently assuming that I do the same. This is incorrect. A careful reading of my papers shows when I'm using geometric rates and when I'm referring to research that uses exponential rates. Most "doubling times" and annual aggregate growth rates (AAGR) that I give in my papers are based on exponential growth rates. As I note, we need 100 percent annual growth to double the original amount in a year if the underlying growth process is geometric with annual increase, whereas we need about 69 percent growth rate if it is exponential.

Kurzweil also claims that I mischaracterize Moore when I note that in his 1975 paper he pointed out that the speed of change in integrating circuits was slowing down. I cannot but conclude that Kurzweil has not read Moore's paper. Moore has repeatedly made the point that there was a slowing down from the original one year rate of change. Most recently he made this point in his 2003 International Solid-State Conference keynote. The core point in his 1975 paper is that the component increases that came from "circuit cleverness" would be lost in the future. Moore shows that the increase in the 1959-1975 period had come from shrinking feature size, increasing die size and more clever use of the area on the die. He argues that developments in feature size reduction and die size increase could continue at their historical rates at least until 1980 but that the main source of growth, architectural improvements in the use of die space, which had created about 100 fold increase, could probably deliver only a four-fold increase before hitting a wall. He was wrong, but this error was compensated by the fact that the die area did not grow as fast as he expected. One reason for limited die area growth, of course, has been the fact that modern chips have multiple layers.

In this sense, they already use the third dimension which Kurzweil sees as a potential source of future developments.

Kurzweil also seems to have problems in understanding what I mean when I talk about analog computing and the point that many mathematical problems require an infinite number of algorithmic computations. According to Kurzweil, my statements are illogical and without basis. In fact, I have discussed this issue in some detail in a paper that is available on the net.[15] Although the paper is an old one, from 1988, it gives some useful references to relevant research on computational complexity and the limits to digital computer architectures.

As Kurzweil focuses on the grand plan of evolution, stretching millennia, he also thinks that I am nit-picking and don't see the "forest of exponential trends" from trees. In a way, Kurzweil probably is right. I try to check the validity of claims about exponential progress by checking the details. When I find that Kurzweil moves Babbage's Analytical Engine about half a century in time, or that available literature gives different values for the cost of IBM 7090, one could say that such details don't really matter. In a sense I agree, and that is the reason why I put these observations in a footnote. On the other hand, such details do matter as they also test the reliability of Kurzweil's data. I used these two data points simply because I had data available for them. I was unable to verify Kurzweil's data more generally, as the data linked to Kurzweil's book does not make any obvious sense. Despite several attempts to decipher the dataset and reverse engineer its bugs, I could not figure out what the data was supposed to tell.[16]

Kurzweil agrees that there is an enormous difference in implications between a one-year and a three-year doubling time. He, however, also claims that there is no such variability in the data, unless one is trying to create confusion. To try and avoid adding to the confusion, one may simplify the issue and forget all sophisticated analyses about the evolution of semiconductors. Instead, one can simply focus on three numbers. According to Moore, in the 1960s the component counts doubled roughly every 12 months. According to Moore, after 1975 they doubled roughly every 24 months. According to the most recent ITRS roadmap, the number of transistors on microprocessors is expected to double roughly every 36 months. If this creates confusion, hopefully it is creative confusion. In any case, these numbers do not in any obvious sense represent accelerating speed of change. Time may be relative to the observer, as Kurzweil notes, but one of us is probably moving too fast.

When I note in my paper that the current dynamic in computer industry may change, for example, because the historically important link between software and personal computer hardware may break as open source becomes more common, Kurzweil states that this is a big leap, with not support to land. There is more support to this point that I have explicitly discussed in the paper. I have discussed quite extensively the dynamics of open source and other internet related innovations in my recent

---

[15] Tuomi, I. (1988) Neural networks as measurement type computers: some theoretical reasons for non-algorithmic information processing, available at http://www.jrc.es/~tuomiil/moreinfo.html.

[16] The table at http://www.penguinputnam.com/static/packages/us/kurzweil/excerpts/exmain.htm tells, for example, that the calculations per second per thousand dollars with a 1953 IBM 701 was 1.188E101 and about five times higher with a 1966 IBM 360. According to the same data set a Pentium PC in 1993 handled 1.00E107 calculations per second. The data is obviously corrupted.

book.[17] It would be interesting to expand this discussion, and more clearly articulate where we actually have different views concerning technological development. I hope I will be able to return to this and other interesting points in Kurzweil’s arguments at a later time. I also hope the above discussion clarifies some of the main sources of apparent confusion.

We both know that innovation requires creative and radical ideas, commitment to positions that one knows could be improved, but also a process of editing and quality control, where exciting new ideas are tested and redefined. Illusions may be interesting and useful even when they are not, strictly speaking, true. Even when illusions do not tell us how the world is, they tell us how we see the world, and what we want to see there. Kurzweil’s theories are particularly interesting, as they combine a fundamentally deterministic and conservative belief in predictable technical progress with the idea of completely unpredictable world beyond the rapidly approaching singularity. The underlying logic may be paradoxical, but paradoxical logic may be what we count in the future.

---

[17] Tuomi, I (2002) Networks of Innovation: Change and Meaning in the Age of the Internet. Oxford, Oxford University Press.