

Transhumanism: a Complex Adaptive Systems perspective Neil Davison

A friend recently asked me to write an article on transhumanism.

His exact words were *I am after your take on TH from a systems perspective, that's all – however you see it, de novo*. I guess it is because he knows a broader systems' perspective would be interesting to apply to such a complicated topic.

Thx Neil please consider writing as if you were a whole systems theorist giving some advice to TH fan (outlining) about the sorts of systems implications involved in TH. Possibly you could then include a note as to whether you think this advice is being heeded and if no what would you suggest please.

Where to start...?

To give me a hint he had tacked on the end "P.S. Transhuman is the humanising of a robot"; but, hang on a minute... I thought I'd also heard somewhere that a human that had body parts replaced by machined parts could also be regarded as 'transhuman'. And that augmentation of natural human capabilities might also count...

How to define the topic and clarify the task? What sources should I look at? Who should I believe? What if there was more than one way to interpret the information I found? How long should I look – just until I had a range of perspectives, or until I had reached a definitive conclusion, you know, one I was absolutely certain about?

I was still bemused... that despite being a relatively intelligent and informed sentient being I was still uncertain, and aware that it would require effort to identify the sources of knowledge and how to interpret them, before I was able to synthesise the information in order to provide my best answer!

Imagine how much easier, more efficient and effective I would be if I had instantaneous access to a better memory, or to all the information ever written! Imagine if I had enhanced analytical skills that enabled me to flawlessly use reductionist approaches to drill through to the right answers, every time! Imagine if I automatically reached all the right conclusions, faster, and with greater certainty!

Wow, wouldn't that be great!!

Then I realised, this was actually one of my biggest misgivings about TH.

Maybe it is our uncertainty that makes us human? Maybe it is the ability to see the plurality of perspectives that forces us to use considered reason to get to truth? If we knew the answers, then purely instinctive actions would suffice, and we would be programmed automatons...wouldn't we?

But things are not simple black and white decisions. We live in complex adaptive systems, made more complex by human reflexivity – making us, humanity, complex adaptive reflexive systems – and how we respond depends on multiple factors. Our worldviews, our values, and the environmental conditions at the time, our level of understanding of the linked social-ecological systems we live in, our willingness to engage in analysis trends or the synthesis of alternative design options, our degree of prescience about likely futures – all of these make the choices on behalf of humanity complex, not easily automated with black and white decisions.

You see, I devote much time to trying to understand the complexity of humanity, the various levels of intelligence, expertise, maturity and capability, disciplinary knowledge and decision-making abilities, and the values and worldviews that flavour their decisions in this messy thing we call 'living'.

I do this because, frankly, I think despite our apparent evolutionary success to date (if measured by our domination of natural systems through the overwhelming weight of populations of human flesh, and presence in most habitats on Earth) we seem incapable of recognising and responding to the feedback from the complex adaptive systems that we live in. We certainly have major compounding problems (e.g. climate, population, ...) and an inability to respond in ways that are mutually beneficial to the systems on which we depend.

It sometimes appears that no amount of individual knowledge or understanding can be translated into a collective response in our collective interests. Why would a sentient being with foresight compete for finite resources to exploit, exploit them until they are depleted, poison the air, water and soil and decimate the ecosystem services and biodiversity we depend on for survival?

Is it this inability to agree across difference that makes us human? Or is it the assumption that we are different from the rest of nature? Or the recognition that we are interdependent with nature?

My own questions still ringing in my own ears, I went searching for more information on TH. I came across the more detailed definition of TH in Humanity+ <http://humanityplus.org>:

What is a transhuman?

In its contemporary usage, "transhuman" refers to an intermediary form between the human and the posthuman?"].

And

"Transhumanism takes a multidisciplinary approach in analyzing the dynamic interplay between humanity and the acceleration of technology. In this sphere, much of our focus is on the development and ethical use of biotechnology, nanotechnology and artificial general intelligence."

I was pleased to see they were taking a multidisciplinary approach, but concerned by the apparent assumptions that an accelerated technological focus could improve humanity.

I wondered, given our inability as a species to yet fully understand ourselves, and our patchy collective recognition of our role in the complex adaptive natural systems we form part of, how automating any of these less-than-perfect functions would be of any benefit, when it might have unanticipated consequences and potentially accelerate the rate at which we do the wrong things righter?

I would expect this risk to be exacerbated by the choices made by those with a technocratic worldview; for, if we can assume that human frailties and limits – and even the bounds of our fragile and finitely resourced planet – are no longer 'limits', then what responsibility do we need to have for those people and things we will leave behind (physically, metaphorically and technologically)?

I have long recognised that both nature and nurture are critical elements of human development, and what is 'uploaded' and experienced through safe-to-fail trial and error during the nurture phase is a significant determinant of the ideology and future decision-making of the individual/s involved.

What if a set of self-serving values and operating instructions (logic base) based on untested or incomplete assumptions were uploaded to a future sentient being capable of self-replication... and isn't that what current western society already has problems with, despite individual humans not having been technically augmented to the fullest extent possible? What if the wrong operating system is installed?

I looked further and found much more information had been written on the topics, for example the *Transhumanist Declaration* <http://humanityplus.org/philosophy/transhumanist-declaration>. I don't wish to denigrate the good work, and have not got the space to go into the implications here, but I think their Declaration gives me a few anchors for my (naive?) concerns. I'll explore a few points:

1. *Humanity stands to be profoundly affected by science and technology in the future. We envision the possibility of broadening human potential by overcoming aging, cognitive shortcomings, involuntary suffering, and our confinement to planet Earth.*

I agree, humanity will be profoundly affected by science and technology in the future – we already are! It is how we make decisions about how we allow ourselves to be affected and what we do about it that will make the difference. And we must do this in the context of a world with massive inequity, concentrations of wealth in the hands of a few, diminishing resources, and impending catastrophic compounding systemic shocks. Taking responsibility for these things, as the humans that created them, should be a human endeavor, and one requiring many of the resources that could otherwise spent on augmenting those able to avoid the consequences (at least in the short-term).

2. *We believe that humanity's potential is still mostly unrealized. There are possible scenarios that lead to wonderful and exceedingly worthwhile enhanced human conditions.*

I agree, but perhaps humanity's potential, on a planet with multiple systemic crises and diminishing resources, is to look further inside itself, not outside itself, for the collaborative solutions that will be needed as we learn to live together more equitably and necessarily scale back from the conditions we take so much for granted in a cheap fossil-fueled, world of relative plenty with a stable climate?

3. *We recognize that humanity faces serious risks, especially from the misuse of new technologies. There are possible realistic scenarios that lead to the loss of most, or even all, of what we hold valuable. Some of these scenarios are drastic, others are subtle. Although all progress is change, not all change is progress.*

I agree, it is what I have been saying, but it is unclear whether this refers to misuse of technologies for, on, or in, humans – we have multiple examples of all three.

4. *Research effort needs to be invested into understanding these prospects. We need to carefully deliberate how best to reduce risks and expedite beneficial applications. We also need forums where people can constructively discuss what should be done, and a social order where responsible decisions can be implemented.*

I agree, yet again, it is unclear whether this is focused on the prospects humanity must address for itself on a planet with multiple systemic crises and diminishing resources, what sort of 'beneficial applications' are envisaged, for whom, with what means and to what ends. As mentioned elsewhere, we already have a crisis of faith in a 'social order where responsible decisions can be implemented'.

5. *Reduction of existential risks, and development of means for the preservation of life and health, the alleviation of grave suffering, and the improvement of human foresight and wisdom should be pursued as urgent priorities, and heavily funded.*

No question about the need for reduction of existential risks; I would ask whether these are seen as individual or collective (Universal) risks. If the ethical choice is to alleviate grave suffering, then don't we have enough foresight already to know that the existing economic growth paradigm, fuelled by technologies deployed to increase productive efficiency for those they work for, but often without mature and capable whole systems wisdom oversight to diminish suffering, is already causing grave suffering (and will lead to more as natural systems collapse)? This process will not be assisted by more/ better technology. It requires the capability and maturity to stop doing harmful things and a change in the system ethics that direct priorities.

6. *Policy making ought to be guided by responsible and inclusive moral vision, taking seriously both opportunities and risks, respecting autonomy and individual rights, and showing solidarity with and concern for the interests and dignity of all people around the globe. We must also consider our moral responsibilities towards generations that will exist in the future.*

How could one argue with this statement?

In terms of inclusivity, as a friend said to me recently "...engagement *must not* depersonalize and alienate relevant individuals from meaningful action. Above all things, this may be my biggest problem with fairly standard "policy solutions." They are top-down, impersonal, and inherently *political/bureaucratic* rather than *democratic*. They advance inputs and outputs through incentives and mandates, rather than focusing on maximising the efficacy of underlying social and human processes."

The questions here are: what Institutions will develop and how will they be constituted to administer the Policies i.e. will they be a cookie cut of the political/bureaucratic ones under the guise of democracy like we have now? (I think even moreso) What systems' understandings and System Ethics will inform the logic base of these Institutions, and, to be inclusive, how will representative groups come together to have meaningful discourse (within meaningful discourse ethics) to enable deliberative democracy?

7. *We advocate the well-being of all sentience, including humans, non-human animals, and any future artificial intellects, modified life forms, or other intelligences to which technological and scientific advance may give rise.*

Again, how could one argue with the intent, except that we cannot guarantee the 'character' of life forms we have not yet experienced, so blanket advocacy for well-being could be a little premature?

8. *We favour allowing individuals wide personal choice over how they enable their lives. This includes use of techniques that may be developed to assist memory, concentration, and mental energy; life extension therapies; reproductive choice technologies; cryonics procedures; and many other possible human modification and enhancement technologies.*

By inference, the individuals with 'wide personal choice' right now are those with the power and resources (i.e. money) to enable them to choose how they enable their lives. If all else was equal/

equitable perhaps this would be a fair principle. However, when the existing situation is so imbalanced and inequitable, how could one justify the augmentation of one life over another?

Review

So, after my brief foray into trying to get a grip on Transhuman issues, am I any wiser? I don't think so... but I'm open to persuasion... would an artificial intelligence be so unsure, or open to challenge?

In short, I see the sorts of existential threats already being posed by humanity on itself far outweigh the likely benefit of any technologies we might implant, augment or throw at the problems we have created – unless we augment our systems' understanding and systems' ethics first. My current gut-feeling (where's my artificial brain when I need one;-) is that until we learn how to reach collective agreements about how to collaboratively act together in society's Universal Interests, it will be dangerous to empower some individuals over others – the playing field is already unlevel, and many already don't play by the rules.

Neil Davidson, April 2013