NYUx Theories of Media and Technology - Philip Hughson - philipfhughson@gmail.com - January 2021

Synthesis 3 - Cybernetic Organisms as Creative Conversationalists

In this essay I explore the idea that, if a cyborg is a human being whose capacities have been extended by the integration of technology into a greater whole, and if artificial intelligence (AI) is the key technology of the current age, then we might wonder how our capacity to create and recreate ideas - our imagination in short - might be enhanced by AI. I argue that Conversation Theory, developed by Gordon Pask and more recently expounded by Paul Pangaro, offers an interesting structure for understanding how this might happen, and suggests creative directions which are interesting to explore as AI technologies such as recursive neural networks are integrated into our technological landscape. Otherwise put, I think that AI can be used to help us explore creative possibilities in an interesting way. I do this with reference to two main creative works: the beatboxing of Harry Yeff, also known as "Reaps One"; and the Google Magenta project, which contains a range of applications of recent AI to creative processes such as drawing and music making.

This idea might seem a little far from common notions of what a cyborg is: we might readily think of characters from science fiction films such as Terminator, Robocop, Blade Runner, or perhaps the "precogs" of Minority Report. But these represent only a limited set of what might qualify as a cyborg. Manfred Clynes and Nathan Kline's definition of cyborg - "the exogenously extended organizational complex functioning as an integrated homeostatic system unconsciously" (1960) - shows its origins in early cybernetic theory, where the notion of feedback systems creating homeostasis is central. The idea idea of exogenous extension, when applied to a human, clearly suggests technology: but which technologies? And which human capacities, among the many our species has?

Given a broad enough perspective, many technologies could be considered part of cybernetic systems. Consider the Roman Legion: a homeostatic arrangement of human flesh, stone, cloth, wood and metal, carefully organised, created the largest empire the world had seen. Or perhaps our capacity to write - on paper, on Neolithic caves walls, on touch sensitive LCD screens - is an example of the cybernetic organism. Or contraception. Or ice skates. Scuba gear. The space suit. Indeed, the notion of post-humanism argues that we have long since moved past the point where our species can be easily considered in isolation, away from the many artificial appendages that surround us and are essential to our homeostasis and expression. Perhaps it is easier to understand the cyborg as an exploration of the possibilities of the technologies of each age for expanding our capabilities as a species.

For me though, the key talent of our species is our imaginative capacity: our ability to explore new ideas and possibilities. Karl Popper pithily summarised this as our ability to have our ideas die in our stead [POPPER]. This talent is not trivial. One means of understanding imagination is that we can self stimulate our sensory organs, either within the brain, by using the "minds eye" to create images in our head, by stimulating the same neurons as are stimulated by sense data from the retina, or similarly with words. Another is through our social capacities and language: we can discuss ideas and improve them before carrying them out. And when something works

well for someone else, we are readily adept at imitating their behaviour, even when it is novel for us. And so I think it is interesting to ask how our imaginative or creative abilities can be enhanced cybernetically.

The obvious answer from the perspective of the current age is to use artificial intelligence, which after long decades of doubt has in the part ten years flourished, producing triumphs such as Deep Blue, Watson and AlphaGo. Much of this has been based on artificial neural networks, a suitably cyborgian idea: networks of abstracted signal processing units which excel at pattern recognition much as our brains do. Indeed, technologies such as Recursive Neural Networks are behind things familiar technologies such as predictive text, offering likely next words or sentences based on what we have already written. In domains such as translation, image processing and speech recognition, recent years have seen vast improvements based on artificial neural networks, such that they are now effective and commonplace.

Google Magenta is an example of a recent project which has sought to apply neural networks to creative work.[1]. It covers creative domains such as music and drawing, and showcases projects which include the automatic completion of sketches based upon a few strokes; or the creation of variations on a given melody or rhythm; or the creation of rhythmic or harmonic accompaniment for a melody play by the user. Projects created with Magenta are frequently interactive to the point that they could certainly be described as cyborgian[2]. Their output is often something that we would not have created by ourselves, and also something that would not have been created solely by the artificial neural network either. In short, Magenta's output is often the product of both human and artificial neurons, and this represents something of a new frontier in creative expression.

Within the project, there are a number of different applications of artificial neural networks explored. One that I find interesting is "sketch-rnn"[3], which, given input of a few squiggles, draws nine variations of a chosen doodle, such as a bird or a cat. These are generated by Recurrent Neural Networks (RNNs) which have been trained on lots of human drawn examples of sketches of the given subjects. These RNNs can then generate more examples of the sketches. If this begins to look like progress on the Lovelace test for machine intelligence (roughly, the ability to produce something novel, or surprising[4]), it is also interesting for its ability to suggest multiple alternatives in a given situation, something generally essential to creative thinking processes. This is perhaps not dissimilar to the multiple words offered by predictive text applications on smartphones - or, indeed, to the notion of the "minority report" on likely future events expressed in the film Minority Report (Spielberg 2002).

Another interesting example is the project "AI Duet"[5], described as "A piano that responds to you.". In this example, a human can play a melody on a piano keyboard, and the "computer" artificial intelligence responds with a melody which is based on the melody played by the human user, but varied. For me, it is a little like doing question-and-answer jazz improvisation with a talented but slightly ill-disciplined child: the responses clearly relate to the initial input, but are a little erratic, and overall little thematic direction, or musical development, can be noted. This is cybernetic conversation, then, even if of a relatively stunted sort: there's a clear boundary to the system, feedback systems and the beginning of a reflexivity, where we might not know whose ideas are whose....

"AI Duet" works on the level of MIDI, or digitisation on the level of notes being played, with start times, loudness and duration values. But other projects, beyond Google Magenta, work on the

level of audio (and so timbre), generating waveforms rather than notes. "SampleRNN" (Mehri 2017) is one example of a project which does this. It uses an RNN to generate novel audio in the style of audio on which it has been trained. Applications of this have included the generation of endless death metal music [6], or infinite bass solos[7]. In the latter case, the generated music is integrated into a composition by a human composer in the form of Adam Neely, suggesting a new form of cyborg music: music created on top of selectively curated artificial creations of a neural network.

For me though, the most astonishing example of this kind of approach is exemplified an interaction between SampleRNN and the beatboxer Harry Yeff, also known as "Reeps One". Yeff is a musician, or vocal artist, who specialises in beatboxing, producing sounds using his vocal tract which sound a lot like electronic beats and dance music. Yeff's performance "Move" (2013) exemplifies this[8]: it contains an astonishing range of sounds, far beyond what I had thought the human voice was capable of.

On a basic level, we might consider this an example of cyborg music, and indeed an interesting one, because it uses sounds which are clearly inspired by electronic music technology, but in this case the technology has been removed entirely from the performance (excepting the microphone, perhaps). In other words, the "artifical" element of the cyborg music hear persists only in the sounds produced, or the ideas, which otherwise would not have been conceivable. The artform of beatboxing may refuse technology, and insist on the human, but its creative output is unquestionably cyborgian.

The "Second Self" (Yeff 2019) performance is even more interesting. It consists of a Yeff performing in alternation - duet, if you like - with an RNN training on Yeff's own vocalisations. This is similar to the "AI Duet" described above from Google Magenta, but it works on the level of audio to produce novel responses based on Yeff's input. Yeff describes his experience with this duet in a later TEDx talk (Yeff 2020). He suggests that AI can be used in this way to make us explore the power and range of our voices in a way that we would not otherwise do. For him, we can "use technology as a mirror", and "we could use machine learning to push us in any way we want... I could play phrase chess with myself".

Perhaps this is the cyborgian notion of creativity or imaginative exploration which I was considering at the beginning of this essay? The idea of using technology as a mirror, or perhaps as a kind of sketch pad, is an interesting one: it contrasts sharply with the notion of technology replacing our capacities, typical tropes of which might include improvements in transportation reducing our physical prowess, producing the "couch potato"; or the familiar news item about somebody losing their ability to make basic navigation decisions because of over-reliance on satnavs. This instead is an idea of dialogue with artificial intelligence producing creative and aesthetic outcomes which we would not otherwise have arrived at. We might think of the AI as a teacher, or as a muse, or perhaps as something we play with in order to develop our capacities.... this latter idea reminds me of the ideas of the psychologist Piaget, for whom the capacity to invent new games is one of the highest forms of human expression.

Clearly, while they obviously involve feedback and a measure of homeostasis, these ideas have moved beyond early ideas of cybernetics. The ideas of reflexivity and virtuality explored in the second and third wave of cybernetics perhaps have some relevance, but so far the most convincing means of understanding them for me has come from the "Conversation Theory" developed by British cyberneticist Gordon Park in the latter half of the 20th century. As yet I have

not had time to develop a thorough understanding of this theory, but I will sketch some brief notions of what I have understood so far and why I think it is an interesting perspective.

Conversation theory is similar to constructivist ideas of knowledge. It posits that shared understanding is arrived at through conversation, which involves an exchange expressed in a language. Each of the participants in the conversation has a goal, perhaps to arrive at some knowledge, for a student, or to impart information, for a teacher. There is typically a shared conceptual background, but not everything need be the same. During the conversation, the participants can share and interrogate perspectives on the concepts they are using: for example, if one person has a concept of a "cup", the other might ask what it is used for, or what the cultural contexts might be (having tea in Japan is different to in the UK....). By this dialogue, we arrive at an agreement, a shared understanding or consciousness, one which transcends our selves, and which can help us move forward in understanding. Conversation need not be with natural language: it could be through physical interaction with objects or others, as in children's play; or it could be in musical form, as with Yeff's musical interaction, or the "Duet AI" project. In all cases, there is a notion of interface, which might be as literal as the interaction between two faces, or might be a digital or computer interface. The following diagram illustrates the basic structure of the exchange:

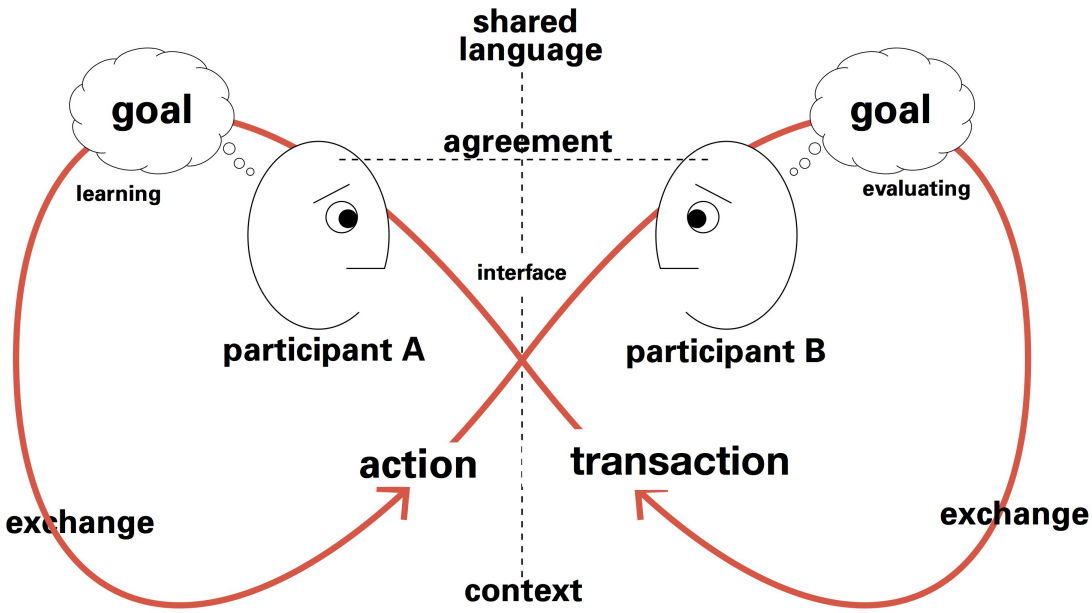

after Dubberly Design Office 2008

One of the key ideas in Conversation Theory is the notion of progress towards a goal, or something new: in short, it helps us to become. For designers, the sketch book has long been an essential tool for such progress: with it, designers alternate between conversational roles, between drawing and interpreting. But what of contemporary AI? Paul Pangaro (2019) offers this critique of it:

> "[...] recommendation engines just moor us to the past, to what we used to like and do — to who we used to be. They rob us of “living in the now” and deciding for ourselves, thus pre-defining our selves. It can be great but this is the hidden stain of AI. Let’s design a "Question Engine” that poses novel questions as possibilities — valid choices we might

> really want to take — to let us decide where we want to go and who we want to become."

In contrast, he offers a suggestion for the characteristics of "great conversation" - it should "tell you things you enjoy learning", "delight you", "[be] surprising", "go places you didn't expect to go", "[be] generative" and "evolve in ways you couldn't evolve on your own". These ideas are hopefully familiar to us all from inspiring conversations we have had, and yet I doubt many of us have had such interactions with Siri, Alexa or Cortana, the familiar speech-enabled personal assistants that presently haunt our zeitgeist.

For me, Pangaro a critique of contemporary AI practice which I have long felt, but struggled to find at theoretical basis for: AI seems to explore the space of the known effectively, and it offers many useful labour saving devices which I would not be without - as well as dystopian features such as omnipresent facial recognition. But many of the interactions I have with AI are lacking in depth, absorption or sensitivity: I become bored quickly at move on. I become frustrated at AI that thinks it understands me without bothering to check, such as the Strava bot which told me to go running, not realising that I had simply got bored of it and left it behind the last time I did.

In conversation theory, though, there are suggestions of directions for AI which might make me want to become more of a cyborg. Conversation, as expressed in the Harry Yeff's "Second Self" dialogue, or perhaps in Google Magenta's project "AI Duet", is a rich metaphor which can be exploited to produce more meaningful and creative AI applications: perhaps we can create artificial minds we want to interact with in order to develop our creative capacities? A sketch book with a neural network that helps us towards our creative goals - or, better still, helps as to perceive new and previously unimagined possibilities? The goal for me, for the next week at least, will be to see if I can produce a project which develops some of these ideas.

REFERENCES

Manfred CLYNES and Nathon KLINE (1960) "Cyborgs and space". Astronautics, September 1960. Available at http://web.mit.edu/digitalapollo/Documents/Chapter1/cyborgs.pdf

Soroush MEHRI et al. (2017) “SampleRNN: An Unconditional End-to-End Neural Audio Generation Model.” Available at https://arxiv.org/pdf/1612.07837.pdf

Paul PANGARO (2019) "Less interference/More Dance". Available at https://pangaro.com/lasg2019/index.html and https://www.youtube.com/watch?v=GfP6CiAcDLs

Aaron PARKHURST (2012) Becoming Cyborgian, The New Bioethics, 18:1, 68-80, DOI: 10.1179/2050287713Z.0000000006. Available at https://www.tandfonline.com/doi/abs/10.1179/2050287713Z.0000000006?journalCode=ynbi20

Gordon PARK (1976) Conversation Theory: Applications in Education and Epistemology. Elsevier. Amsterdam, Oxford, New York.

Karl POPPER (1977). "Natural Selection and the Emergence of Mind" (lecture). Available at https://www.informationphilosopher.com/solutions/philosophers/popper/natural_selection_and_the_emergence_of_mind.html

Stephen SPIELBERG (2002). Minority Report (film). 20th Century Fox.

Harry YEFF (2013) “Move” (beatbox performance) Available at https://www.youtube.com/watch?v=YH5ty3Kucz4

Harry YEFF (2019) “Reeps One ft. A.I. ‘Second Self’”. Available at https://www.youtube.com/watch?v=q981cTdL0_Y

Harry YEFF (2020) “How Battling A.I. Unlocked The Power of My Voice”. Available at https://www.youtube.com/watch?v=wTMMopLYJn4

---

1. See https://magenta.tensorflow.org/demos for some examples. ↩
2. See Parkhurst (2012) "becoming cyborgian" for this neologism, descibed as "new frameworks with which we can engage novel intimate relationships that develop between humans and emerging technology". ↩
3. See https://magenta.tensorflow.org/assets/sketch_rnn_demo/index.html ↩
4. see https://hssmi.org/the-lovelace-test-or-can-machines-take-us-by-surprise/ for background on the Lovelace test. ↩
5. See https://experiments.withgoogle.com/ai/ai-duet/view/ ↩
6. See https://www.youtube.com/watch?v=MwtVkPKx3RA. ↩
7. See https://www.youtube.com/watch?v=2xMhRwxXJTc&ab_channel=AdamNeely ↩
8. Please listen at https://www.youtube.com/watch?v=YH5ty3Kucz4 ↩