

The Robot Within or The Ghost is The Machine

I. I Need My Mobile and My Mobile Needs Me

If one reads the epochally influential paper, "Computing Machinery and Intelligence"(1950), by Alan Turing, the opening sentences come as something of a surprise to those familiar with popular representations of the paper's contents. "I propose to consider the question, 'Can machines think,'" Turing begins. No surprises here; the tests held annually to determine if one or another computer programme has officially crossed the line into consciousness bear Turing's name in honour of this proposal he outlines in the paper. The sentences that follow, however, while less consistent with conventional understandings of Turing's argument, are perhaps even more important for a global culture facing the advance of large-scale AI technologies and the integration of "smart" devices into nearly every sphere of daily activity. Turing writes that before answering the question, the terms "machine" and "think" must be considered in themselves:

The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words 'machine' and 'think' are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question ... is to be sought in a statistical survey such as a Gallup poll. But this is absurd (Turing 1950).

Turing was right to be suspicious, and not just of Gallup polls. Today, theories about the character and nature of artificial intelligence appear at a rate that may outpace Moore's Law itself. Thus, Turing's caveat is an important one, because it highlights a key factor for understanding and assessing the putative cognitive capacities of computers: the inescapable fact that the interpretation of artificial intelligence is predicated on human intelligence; what humans believe about AI will be equally important as the actual structures and capacity of machines dedicated to AI capacities. This is an increasingly pressing matter as human interaction with smart devices increases and these devices are built into the infrastructure of our cultural lives. Beliefs about the capacities of AI are, therefore, as important as AI's technological status itself.

For decades, technology has been integrated into the structures that underpin governments and economies (where these concepts can be understood as distinct). Much of this process has been tacitly experienced with only intermittent peaks of interest; for example, in the advance of surveillance technology throughout London, or the momentary outrage that greeted the revelations of wholesale global surveillance by Edward Snowden. This dynamic is changing as digital devices become more and more intimately connected to our lives. Everyone may have a passive relationship with the National Security Administration which they may choose to acknowledge, resist or ignore, but the relationship persists whatever an individual's attitude toward it may be. Smartphones, however, and the apps that run on them, only really work when you pay attention to them.

In both the case of NSA surveillance and smartphone apps, technologies are paying attention to you constantly. The NSA may benefit from invisibility, but an app needs to respond to you and to become visible to be useful. Like the animals their cartoon spokes-beings often simulate, if you show devotion to your app and feed it regularly (with personal information) your relationship will only deepen and become more rewarding, at least from the app's perspective.

II. One Algorithm to Rule Them All

As an app feasts on the information banquet its user provides, it forms a picture of its human user. The picture is more or less representative of one's life patterns the more frequently one uses an app. The more proficient--perhaps "sensitive" is the word--the app is in understanding and addressing your needs, the more frequently it is likely to be used. The old model of such human machine relations was articulated, perhaps somewhat labouriously, in Marshall McLuhan's and Quentin Fiore's book, *The Medium is the Massage* (1967). McLuhan takes an almost 19th century materialist view of the ways humans and technologies interact: "a wheel is an extension of the foot. A book is an extension of the eye. Clothing is an extension of the skin" (McLuhan and Fiore 1967). A wheel may or may not be "an extension" of the foot, but the relationship McLuhan defines in terms of power hierarchy in the book is clear. Machines are versions of us, and crucially, they are versions of our *bodies*. This line of thinking has long infused thinking about AI, but it has been reformulated to use the mind (not simply the brain) as the fundamental reference point. In the case of Andy Clark and David Chalmers' influential 1998 paper "The Extended Mind", the role of computing functions is explicitly understood in terms of human cognitive functions:

Epistemic action, we suggest demands the spread of epistemic credit. If we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognising as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process (Clark and Chalmers 1998).

One need not agree with McLuhan or Clark and Chalmers to concede the efficacy of keeping such metaphors visible. In explicitly theorising human-object relations in the terms they do, one may observe a dynamic of feedback, but lines dividing clothing and skin and computer and mind, even in the case of Chalmers and Clark, remain readily identifiable. The relationship to apps, however, is much more nebulous--both metaphorically, and literally.

Understanding what exactly the relationship an individual has to an app is increasingly difficult and the roots of this question touch on both aspects of the philosophy of the development of AI, but also fundamental aspects of human cognition. In many ways, contemporary apps are the inverse of the material relations implied in McLuhan's variation of the "extension" model; rather than being "extensions" of obviously material parts of the body like hands or eyes or even brains. Conceptually, apps are much more like Clark and Chalmers' notion of an extension of mind, but they are not exactly like this model either (what is the difference between the mind and the brain? Such boundaries are notoriously porous, but for the purposes of this argument, it is most useful to understand "mind" as the cognitive structures and faculties either instantiated in the human brain or the ends toward which human cognitive activity is directed). Apps are not just memory stores or calculating capacities housed in circuits rather than neurons, rather they can be thought of as an extension of the intentions of the mind as realised by the brain. If they are an extension of any basic human property, they are an extension of need.

Thus, the present argument does not see the mind as being simply "extended" but, instead, it sees the behaviours and outputs of the mind being *simulated*. The importance of this point can be seen in returning to Turing's original proposal. The machine in Turing's experiment is not, necessarily, "thinking" in the way a human being might be thinking; it is merely effectively simulating human

thought in such a way as to be accepted by human minds as appearing to do so. Strong AI theorists frequently argue that there can be no meaningful difference between the output of a state and its internal characteristics. This is not the place for an argument about ontology, the present argument is merely interested in the implications of such simulation in relation to Turing's crucial distinction between the appearance of identity and the actuality of identity.

If a machine can appear to simulate thought effectively, a human being can make a choice in relation to that simulation: the human can admire the fidelity of the simulation or a human can treat the simulation as identical with thought as realised in a human mind. In the latter case, humans may then chose to treat the "thoughts" that emerge as the same thing--qualitatively as well as empirically--as human thoughts. In the world of interactive apps which establish behavioural feedback loops with their users, this is an important distinction to highlight. Where cognitive faculties like memory, pattern recognition, navigation, even aesthetic taste and social connection, are increasingly facilitated by, or even produced by, apps, the question of what input the app has in cognitive processes becomes one of considerable significance.

III. PanAppticon: Nudged Gently By Machines of Loving Grace™

The economists and writers, Cass Sunstein and Richard Thaler, in their 2008 book, *Nudge*, argue for a philosophy of government to which they give the name "Libertarian Paternalism". Sunstein and Thaler, perhaps obviously, do not work in areas of economics which consider the theory of branding, but their ideas have been quite influential in governments on both sides of the Atlantic. In its essence, Libertarian Paternalism seeks to present choices to populations in such a way that the most socially beneficial choices are the easiest to access. An easy to understand example of this would be if a hardware store only displayed energy efficient lightbulbs on its shelves, forcing customers to enquire about less energy efficient bulbs. You are still free to waste your money, but the shop won't make it easy for you. Such ideas, rooted in evidence from the field of behavioural economics, make the fair point that humans fairly rarely if ever "fi[t] within the textbook picture of human beings offered by economists" (Sunstein and Thaler 2008). Critical to the Libertarian Paternalist model is the notion of the nudge, from which the book takes its name. Sunstein and Thaler define a nudge thus:

A nudge ... is any aspect of the choice architecture that alters people's behaviour in a predictable way without forbidding any options or significantly changing their economic incentives (Sunstein and Thaler 2008).

The "choice architecture" of which they speak is key to understanding the social role and potential risks of an increasingly app driven world. Sunstein and Thaler refer to the choice architecture as being "the context in which people make decisions" which is, of course, arranged by a "choice architect". In some cases, this architect may be a politician, for example, the Obama administration became interested in Sunstein's and Thaler's ideas and attempted to integrate "nudge logic" in areas of policy including health, education and energy. Sunstein's and Thaler's model is, in their own presentation of its ends, clearly directed to facilitating socially optimal modelling, and they are perfectly frank about the aspects of social engineering that the philosophy entails. Despite such purity of intention, nudge logic must be understood as, in essence, morally neutral. The virtue and value of the nudge is only as virtuous as the nudger.

The advance of "life-managing" apps represents a manifestation of nudge logic. From the fitbit, to calendars and clocks that nudge their users toward optimal behaviour patterns, the model of Libertarian Paternalism is increasingly inscribed in the products which are produced by paternalistic libertarians in Silicon Valley. Whether such apps "work" is one of the most frequently encountered questions on the internet, but a question far less regularly asked is the following: even if such apps perform in the way they are supposed to, for whom do they really work? A diet app, for example, holds data about your eating choices, the frequency of your meals, perhaps even the "mistakes" you make in holding to your chosen diet. Such data may be valuable to a user trying to improve their eating habits, but it is even more valuable to someone interested in knowing what a user is eating. Sleep apps "reminding" their users when it's time for bed, and even offering advice on what to do before going to sleep would seem modeled on an almost literal form of paternalism.

Data is immortal. Data has the capacity to be infinitely replicated at low cost. As data piles grow, the apps and the customers of the information the apps compile become "smarter". They know us more intimately; they learn what kinds of nudges work and what kinds do not. Such nudging may be benign, even beneficial in the first instance, but over time, nudging may come to exclude more choices than it presents. The "choice architecture" that Sunstein and Thaler speak of in their book begins to look more like a prison cell. Indeed the question must be asked: what does the architecture that lies behind the choices actually look like, and who is the architect? Is it the algorithm? Is it the people who design and update the algorithm? Is it the customers of the company who buy user data-pools from the app? Perhaps it is all of the above.

If choices are restricted and nudges become more aggressive, life patterns are almost certain to change. Where Sunstein and Thaler would like to put the nudge at the service of building a better society, a company developing apps or using an app to nudge an app-user need not be so public-spirited. For these organisations and individuals, the nudge that matters is the nudge that pushes a user in the direction that makes them richer, more powerful and more informed about the choice patterns of a person for future exploitation. Apps may be digital, but they represent reifications of ideologies. These ideologies are, by the very nature of the apps they employ, invisible, and the more obscure they appear, the better for the company involved. What is true at app level is also true of the larger economy in which apps function. As paternalistic sleep apps learn which bedtime stories are most soothing to their users, the wider economy is inscribing itself in the life patterns of individuals. One is not only encouraged to go to sleep to become more efficient for work, there are also suboptimal and optimal ways to prepare to go to sleep to become more efficient. Over time, the nudge may wear away the nerve-endings it connects with and the "choice architecture" may simply seem like freedom itself. He loved the way Big Brother nudged him.

This returns the argument to its originary point and to Turing's 1950 paper. Turing may have been uninterested in whether machines could think in the same way as human beings, but a data-driven app economy may very well come to produce human thought and behaviour patterns that are much more like those of machines. The possibility of this happening is far from absurd. As AI advances, so will the belief in the power of AI and the logic and rationality of ostensibly "impartial" processes of thought--whether or not such processes are remotely "impartial" or, indeed "thought". However attractive it may seem to purge irrationality from human thought with algorithms of loving grace, particularly in this Year of Our Lord 2016, human ideologies will never be fully exorcised from the devices we use. Such ideologies may come to feel like our own, but they will not be, they will merely be chambers of the choice architecture with which we are presented. Turing himself, after

saving countless allied lives with his heroic code-breaking work in the Second World War was driven to suicide by a government that, instead of being grateful, chose instead to violently nudge Turing aside for failing to choose a sexuality within the choice architecture provided by national security imperatives. Outsourcing the mind in an essentially mindless fashion may come to have similarly pernicious consequences for global populations. Gilbert Ryle's dismissive term coined in relation to Cartesian understandings of the interaction between mind and body--"the ghost in the machine"--takes on a new dimension in the age of apps: the ghost is the machine, and the machine is the ghost, and it is a ghost that is eternally hungry.

Works Cited

Clark, Andy; Chalmers, David. (1998). "The Extended Mind". *Analysis*, 58: 10-32.

McLuhan, Marshall; Fiore, Quentin. *The Medium is the Message*. London: Allen Lane, 1967.

Sunstein, Cass; Thaler, Richard. *Nudge: Improving Decisions about Health Wealth and Happiness*. New Haven: Yale University Press, 2008.

Turing, Alan. "Computing Machinery and Intelligence". *Mind*, 59; 433-460.