Divinatory Computation : Artificial Intelligence and the Future of the African continent

Faeeza Ballim and Keith Breckenridge, 25 July 2018

This, apparently, is the fourth industrial revolution. Klaus Schwab, the founder of the World Economic Forum, announced this intriguing fact at Davos in January 2016, and he later explained what he meant in a short book of the same name. His claim was overshadowed by the events of that year, which included – in case our reader has forgotten – the European migration crisis, Brexit, ongoing brutal civil war in Syria, nuclear postering on the Korean peninsula and the Trump election. All of which seemed very much to be extensions of the long term conflicts of the 20th century. These political continuities are interesting because the core of the claim that we are experiencing the fourth industrial revolution is that it marks a break with the epoch that immediately precedes it. Schwab calls this period – the years from the 1960s to 2007 – the third – digital – industrial revolution. The new era is "characterized by a much more ubiquitous and mobile internet, by smaller and more powerful sensors that have become cheaper, and by artificial intelligence and machine learning."(Schwab 2017, 12) Certainly by the end of 2017 – nearly two years after the announcement at Davos – the global celebrity intelligentsia had recovered from the events of 2016 and returned to the fourth industrial revolution. The idea that the global economy had moved into a fundamentally new technological epoch captured (in the title of another best seller) as *The Second Machine Age* and driven by ubiquitous internet and artificial intelligence had entered into common sense. (Google Trends 2018)

Two big, general, political problems hang over this excitement. The first is the claim that the current tools of AI – and the broader field of machine learning from which they are derived –

mark a break with the histories of automation and cybernetics that track back, at least, to the late 19th century. (Wiener 1988; Noble 1977; Beniger 1986; Zuboff 1988; Kline 2015) Evaluating this claim requires careful study of what artificial intelligence can do now, how the technologies actually work, and the ambitions of their developers. This is not a simple task. After decades of decline that began in the 1970s, the field of machine learning has been growing very rapidly over the last ten years. It is now large and boisterous, and shot through with fierce and interesting debates about the philosophical implications and cognitive models of learning machines. In this paper we agree that something very important is happening, but we also insist that it has a long and fraught history that is a more useful guide to its future politics than the optimistic projections of the advocates of AI.

The second momentous claim is that we stand on the edge of the development of the autonomous artificial intelligence. This has triggered a philosophically and politically provocative debate, which can be unpacked into three competing premises. First is the claim that the techniques of AI, and deep neural networks in particular, are modelled on the workings of the human brain and, because they work like the human mind and can train themselves, hold out the possibility of general machine intelligence (Lee 2018; Ng 2017; Hinton 2017). A second is something like the opposite, that deep learning is so mathematically convoluted, demanding of standardised data for algorithm development and, thus, so tightly focused on carefully framed problems that it is intellectually naive and already at the limits of its potential application (Marcus 2018; *Does AI Need More Innate Machinery?* 2017). The third – also driven by the convoluted and illegible structures of deep neural networks and, in particular, their dependence on opaque and provisional techniques for the reduction of errors – is the recent claim by veterans in the field that the popularity of these machines entails the restoration of alchemy (which, as the authors argued at the most recent machine learning conference, often worked well without its practitioners understanding why). (Peng 2017; Matthew Hutson 2018; Sculley et al. 2018; Alister 2017) In each case there are firmly-held and well developed positions often pitting engineers employed to implement AI systems that "just work" against researchers based in the universities. But these conflicts also reflect the long history of existential struggle within the field of machine learning, dating back to the 1950s, pitting five different schools against each other in a high-stakes struggle to determine the future of computer science, and much else besides. (Compellingly explained by Domingos 2015a) This also means that hanging over the current debates is the unease of those who can remember the field, when – in the 1980s or even the 1950s – it last promised a revolution in machine intelligence. (Ng 2017; Lee 2018; Peng 2017; Crevier 1993)

The most popular and powerful forms of machine learning – deep neural networks – are algorithms for iterative matrix calculations aimed at reducing errors in the automation of bottlenecks in the global information processing economy. These software machines are currently deleting spam, selecting advertisements, reading images, determining credit worthiness and replacing taxi drivers. The most revolutionary applications are being set to work on automating facial recognition images from identity databases, cell phones and surveillance video. At their core the new machine learning technologies apply simplified mathematical techniques for determining non-linear probabilities that long predate Galton and Pearson's more exact (and linear) statistics. (S. M Stigler 1986, pt. 1) In short, AI is motivated and resourced by an already successful and very lucrative effort to automate and control work -- in particular the work of placing and selling on-line advertisements. It is important to notice that, to date, the unambiguous commercial success of the software machines has been much more influential than the difficult, and often arbitrary, mathematical techniques that describe them.

Much of this success has to do with the utility and simplicity of software-based machines, notwithstanding some of the very grand claims about general intelligence. The programming required to develop robust, machine-taught algorithms for character recognition, for example, can be developed in less than a hundred lines of code, and, once the algorithm has learned the most effective weights it needs to use in the matrix calculations, it can be published as a simple, but powerful, script on the web. (See Nielsen 2015 for example.) This reliance on the device that works as explanation and motivation in science is old and, now, commonplace (Shapin and Schaffer 1985; Stephen M. Stigler 1989) and it mobilizes the influential history of engineering automation in the United States. (Noble 1984; Zuboff 1988) But these new machines are no longer confined to the workplace. Unlike the earlier technologies of feedback and surveillance that dominated the automation of the factory in the 20th century, the new forms of AI have very broad – indeed potentially global – appetites for training data and similarly wide economic and political capacities. It is these global ambitions that bring us to the main objectives of this paper.

We are interested here to map the likely political consequences for people and societies on the African continent of the turn to machine learning as the basis of the global economy. What does it mean for Africans that the richest companies, and the most powerful governments, are investing heavily in technologies, programmes and infrastructures of artificial intelligence?

There is, already, a long list of well-articulated dangers. The most obvious are problems of bias that derive from an excess of information – from the dense, hidden and ingrained structures of racism that infect the autonomous development of algorithms, especially in the United States. (Mohamed 2017; Angwin et al. 2016) There are also problems of bias that are the result of the absence of high-quality training datasets -- for example of African names or facial images or words (Marwala 2018; Raji 2018; Lohr 2018). These are actually problems of over-fitting, in which machines will learn algorithms around training data that excludes, or under-estimates, African examples. A third risk is that AI will exaggerate the already existing brutal deficits of infrastructure – of high-speed network connections, reliable power supplies, data processing centres and, especially, of human expertise. For centuries the products that African people, firms and societies have produced have been monopolised and discounted by metropolitan corporations with the energetic assistance of local elites. Will the growing power of the centres of artificial intelligence in the United States and China – and the global monopoly power of a small number of firms secured by AI -- produce a new era of data-driven extraversion and dependency? (Ake 1981, chap. 4 and 5; Bayart 2000) And finally there is the problem of work itself. Economists worry that Africa's historically unprecedented labour market growth -- which sees 12 million new young workers every year -- can only be met by labour-intensive industrial investments, and that those are specifically endangered by the new forms of automation. (Rodrik 2016; Chutel 2017; African Development Bank 2018, 48)

These are all serious risks and, if the long duree history is any guide, it is likely that they will all be realised in some meaningful way. But, as we show here, there are also interesting and compelling ambiguities and ambivalences in the specific histories of artificial intelligence on the African continent. This makes the long-term politics of AI -- like the other Internet technologies -- difficult to assess. None of these questions can be answered categorically and, in each case, there are already vocal and persuasive advocates of the contending answers. In each case -- in the development of massive new sources of training data, the fostering of science institutions and expertise, the rollout of robust infrastructures, and a general interest in leapfrogging and disruption across the continent -- the expansion of machine learning is having ambiguous effects. Much of this ambiguity is a consequence of the basic workings of the technology.

The Geopolitics of Machine Learning

It is a measure of how seriously governments take the problems and possibilities of artificial intelligence that China, the United Kingdom and the United States have all published major plans for industrial policy aimed at fostering the infrastructures, skills and industries that will compete for global ascendancy in this field. Wired magazine has described the Chinese plan, which focuses on the teaching of artificial intelligence "as a discipline" in schools and universities, as the "secret weapon in the global AI arms race." (Executive Office of the President and National Science and Technology Council 2016; Executive Office of the President 2016; Royal Society 2017; State Council 2017; Beard 2018) While some of this is Cold War rhetoric renewed, there is also unambiguous evidence of an existential struggle for dominance in AI amongst the world's largest Internet-based firms (which are, also, the only entities capable of generating the scale and quality of training data that is required to develop effective machine learning algorithms.) Google, Apple, Facebook and Amazon have all made enormous investments in firms specialising in machine learning over the last four years. In 2017 alone over 100 AI startups were acquired by the largest companies, and it is clearly partly the speculative opportunity that emerges with these investments that is fostering the renewed global obsession with machine learning.(CB Insights Research 2018) But this is not a purely speculative frenzy, not least because two of the firms with the largest demonstrated capacity in artificial intelligence – Baidu and Alibaba – are based, primarily, in China. Their dominance in revolutionary fields like universal facial recognition and citizenship scoring derives in large part from unrestricted support from the Chinese government.(*The Economist* 2018; Captain 2016; Hvistendahl 2017)

Something very serious is underway that goes well beyond the tricks of voice recognition and anticipatory on-line content. When a figure like Kai-fu Lee, who combines decades of experience with the largest US companies with contemporary dominance in China, warns that African countries with large poor populations face a frightening future dependency we should take him seriously. "The countries that are not in good shape are the countries that have perhaps a large population, but no AI, no technologies, no Google, no Tencent, no Baidu, no Alibaba, no Facebook, no Amazon," Lee warns. "These people will basically be data points to countries whose software is dominant in their country. If a country in Africa uses largely Facebook and Google, they will be providing their data to help Facebook and Google make more money, but their jobs will still be replaced nevertheless."(Lee 2018) But evaluating this claim – especially in relation to the resources of information and institutions that are available on this continent – requires some understanding of how AI actually works.

How AI Works

Machine learning means that computers are able to draw conclusions about their environment without being explicitly programmed to do so. Since 2015, machine learning practitioners have espoused the virtues of an algorithmic system known as Deep Learning. This marks the return to dominance, in a long history of contestation, of the "empiricists" who emphasise learning through trial and error as against the "rationalists" who utilise predetermined frameworks of logic.(Domingos 2015b, 59) At the operational core of Deep Learning lies the neural network, which as its name suggests, aims to simulate the neural structure of the human brain, particularly the way it receives and interprets sensory stimuli. The neural network has deep roots in behavioural psychology. It gained traction among adherents of the cybernetic movement in the 1950s, who combined mathematics, psychology and engineering to develop self-governing control systems (Minsky and Papert, 1987: ix). In 1957, a New York-based psychologist, Frank Rosenblatt, published his findings about a neural network model he called the Perceptron (Rosenblatt, 1957). Rosenblatt had applied the reinforcement learning technique on an earlier model of a neuron, known as the McCulloch-Pitts model, developed in 1947 (Crevier, 1993: 104). Reinforcement learning is used by psychologists to encourage behavioural change by rewarding certain activities the subject should repeat and punishing those it should cease performing. Rosenblatt introduced a system of reward and punishment to the neural network through the adjustment of its weights (also known as parameters). During training, a machine would correctly identify an object if a light labelled with the correct identifier flashed. If an incorrect light flashed, he decreased the weights that connected the object to this light and increased the weights leading to the light that was supposed to have flashed (Crevier, 1993: 104). While the Perceptron is rudimentary by contemporary standards, it is Rosenblatt's technique of adjusting weights by trial and error from some random initial that persists.

By the 1970s, the Perceptron had fallen out of favour in the academic field of artificial intelligence. One of the reasons for its downfall was that it could not compute the XOR function – that is, being able to tell when two inputs differ. Inputs are the stimuli a machine receives that require interpretation and they can take the form of images, sounds or a numerical dataset. The neural network's presumed impotence led to the rise of knowledge engineers -- or rationalists -- who insisted that the foundations artificial intelligence required specific coded instructions and legible, mathematical proofs. (Domingos, 2015: 101). Neural networks were rescued from likely obscurity in 1982 when the physicist John Hopfield introduced intermediary computing layers – which are famously called hidden layers now – between the input data and the final solution. The hidden layers allowed the network to target constituent features of the input and better focus on assembling the correct solution. In other words the overarching problem to be solved is broken up into a series of fractional questions that can be answered in sequence.

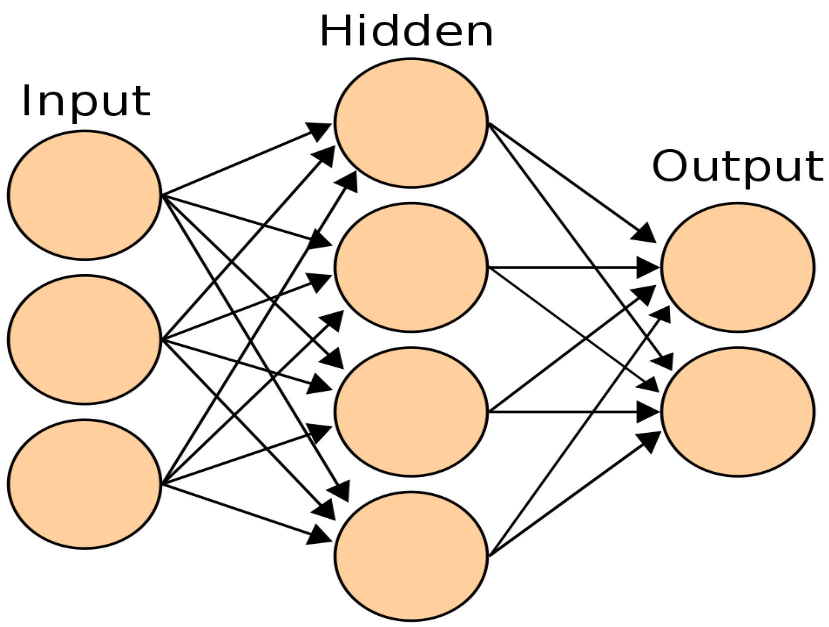

Caption: Neural network with one hidden layer. The connecting lines between neurons are the weights or parameters.
(https://commons.wikimedia.org/wiki/File:Artificial_neural_network.svg)

The architects of neural networks utilize what is called an activation function to simulate the transmission of signals between the neurons of the human brain. The sigmoid function, also known as the S-curve, which Domingos describes as the “most important curve in the world” enables the activation of neurons on a non-binary basis (2015: 104). While in the Perceptron, the output neuron was either fully activated or entirely inactive (the light could only be on or off), the sigmoid function allows a neuron to be activated in degrees based on the scale of stimulation it receives. The sigmoid function, depicted in the graph below, defines the threshold that, as in the human nervous system, signals have to break through to allow the neuron to fire to the next neural layer. The sigmoid function also converts the input it receives from the neurons connected to it, into a value that can only lie between 0 and 1. The neural network's eventual classification (or its final solution) is one hundred percent accurate when the probability is one and one hundred percent false when zero. These values correspond to the zero and one y-values of the sigmoid function in the graph below.

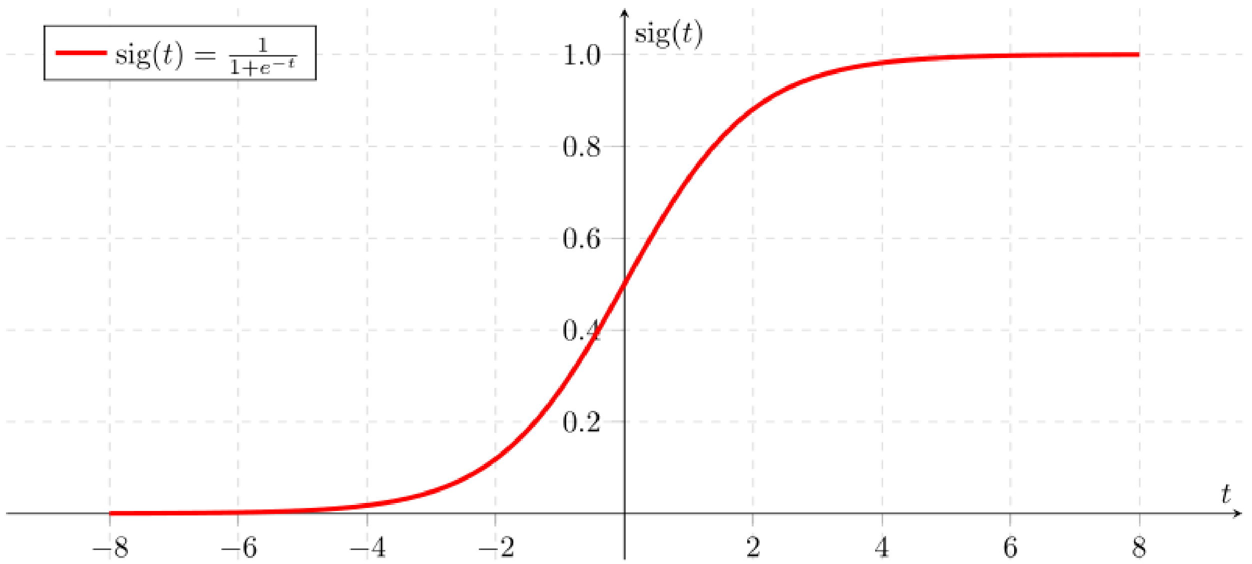

Caption: Sigmoid Function (https://commons.wikimedia.org/wiki/File:Sigmoid-function-2.svg)

Hidden layers raised the problem of measuring the extent to which each contributing parameter or weight was at fault when the network arrived at a wrong answer during training. The solution was a technique called backpropagation, which is driven by a long-standing optimisation technique in calculus called gradient descent.[1] In the modern neural

[1] As an example of the application of gradient descent, assume that data has been collected on

networks gradient descent works in a similar way but the calculation of the cost function differs. The new techniques use gradient descent to find the minimum of the cost function – a multi-dimensional equivalent of Pearson's least squares test of fit. This technique is used by the learning software to optimise the weights of each of the activations performed by the nodes of the network. The key to this process is called backpropagation. In the training set, where the solution that the network should be providing is known, backpropagation begins at the very last output unit. The neural network has already computed the function with, again, randomly chosen parameters. At the last unit, it calculates the difference between the output that the neural network computes and the known value. This difference is known as the error value. This error value is then projected backwards in sequence over each hidden layer. The assumption is that the error value of any given unit (or neuron) in the network is the sum of the error values multiplied by the respective parameter (θ) of each unit in the succeeding layer that it is connected to. The cost function is a summation of all these error values and gradient descent operates to find the minimum of the cost function for the parameter (θ). The parameters are then adjusted accordingly so that the final output the network computes comes closest to the known value.

the relationship between a person's age and monthly income. The goal is to draw a pattern through this data that would allow the machine (often called the learner) to predict a person's income level based only on their age (where age is the x-value). This means drawing a line through the disparate data of the training set, composed of the historical data where both the x and y-values are known, which can then be used for prediction. Gradient descent is used to find the parameters of this hypothesis line, where the parameters are denoted by the symbol θ (theta) and the equation of the hypothesis is: $\theta(0) + \theta(1)x$. In this case the parameter θ (0) is the constant and $\theta(1)$ is the gradient of the equation of the hypothesis line.

While it is the most important component of neural networks, gradient descent suffers from important limitations. It demands that the machine learning practitioner surrender control over the direction of its mathematical operations, in a manner akin to stepping out into the dark. This can be frustrating for those who seek an overarching theory to explain and thus replicate success. Over hundreds of step-wise iterations, the error value (theta) gradually descends to what is hoped to be the absolute minimum of the function. But the cost function often has many troughs (called local minima), some lower than others, and there is no guarantee that one is descending towards the lowest. This is a dilemma that practitioners have learned to live with because the machines have demonstrated that they work. As Domingos writes, “Backprop, with its incremental weight changes, doesn't know how to find the global error minimum, and local ones can be arbitrarily bad, like mistaking your grandmother for a hat. The global minimum is hidden somewhere in the unfathomable vastness of hyperspace—and good luck finding it … But what we've come to realize is that most of the time a local minimum is fine. The error surface often looks like the quills of a porcupine, with many steep peaks and troughs, but it doesn't really matter if we find the absolute lowest trough; any one will do.” (2015:134)

Notwithstanding the public (and some researchers’) claims to the contrary, the operations of the neural network differ fundamentally from that of the human brain because they are entirely governed by a mathematical system of matrix multiplication. In this system, all inputs – images, sounds, words, or actions – are converted to columns and rows of numbers, evocative of the large computer screens of running numbers in the 2001 film *The Matrix.* The most common function that the neural network has to perform is classification and, when doing this, it also assesses the likelihood that its classification is accurate – often in relation to other possibilities. This can take the form of “binary classification”, consisting of ones and zeroes where one denotes true and zero false. For example, a network might be tasked with identifying a particular image as one of only three objects. Each object that it could possibly be is assigned particular vector notation (a vector is a single line or column of numerical values in a matrix) and these then become the possible outputs or solutions of the network. The first is identified as [1,0,0]; the second as [0,1,0] and the third as [0,0,1]. Each of the hidden layers is programmed to ask a single, distinct question of the image it is fed in sequence. The answer each neuron provides is a product of matrix multiplication: it is a sum of the input vector multiplied by its connecting weights or parameters.

Powered by hundreds of millions of transistors on the GPUs that were developed for the matrix calculations required for high resolution computer games, by cloud-based linkages between networks of processors, and by truly vast standardised datasets of web-derived attributes, programmers have been able to set computers tasks that assemble networks that defy mathematical explanation. The algorithms combine millions of attributes, of activation nodes, the weights of each of them, and of the relationships between them in combinations -- all generated by the learning algorithm. It is this complexity, and opacity, that has prompted the criticism that neural networks -- and the pragmatic reliance on automated gradient descent -- amount to a form of computational divination. In a recent influential critique, researchers Ali Rahimi and Ben Recht have described the current machine learning models as alchemy, so contrasting them with the reproducibility that is supposed to characterise principles of scientific enquiry. They have taken issue with the fact that machine learning lacks a core set of explanatory principles that could be applied to all scenarios. Gradient descent, they suggested, is an unreliable and inefficient way of arriving at the predictor function (Rahimi and Recht, 2017). The two veterans have argued that the inability of practitioners properly to understand the workings of deep learning algorithms means that the technology works as a black box (Hutson, 2018).   These criticisms are provocative (not least because they suggest links to some of the most common forms of knowledge and prediction on the African continent (Turner 1975)), but they done little to temper the machine learning frenzy.

African Training Data

Neural networks only function well after they have been trained on enormous datasets.   The example often used in contemporary debates is the dataset of 200 million identification photographs that were provided by the Chinese government to Baidu and Alibaba to allow them to train the neural net for facial recognition. The machine tuning of weights and biases in the development enormous networks of activations means that data has become the power broker among corporations and researchers in the artificial intelligence world. This is also one of the reasons that machine learning algorithms are not jealously guarded and are, often, freely available for public perusal on open source platforms like arXiv (Andrew Ng). But it also means that the machine learning systems are “data hungry” -- especially for noiseless training data -- and that innovation depends on their ability to learn from as large a base of training examples as possible.

The importance of collections running into hundreds of millions of carefully assembled examples for training machine learning algorithms raises the obvious question of whether the relative absence of data places the African continent beyond the operations of artificial intelligence. It is true that many of the largest datasets that are abundant outside of the continent – land titles, census returns, civil registrations, credit histories – are simply not present in most African countries. But many of the core sources of data that have fostered recognition algorithms are available, as well as some unusual ones. Many African countries have national identity databases that include photographs and fingerprint biometrics as images or templates for tens of millions of their citizens.

China looms large in the area of the identification and exploitation of training data for machine learning on the continent. Firms like Huawei and ZTE already are responsible for the bulk of the new high speed networks, and there is a well established concern that the Ethiopian state, which has a distinctive history of using centralised internet technologies to bolster central control, is using these network infrastructures, and machine learning analysis, to monitor and control information and participation on the Internet in general. (Gagliardone 2016; Gagliardone and Golooba-Mutebi 2016; Kuo 2016; World Wide Web Foundation 2017) Recently, the Chinese AI facial recognition specialists, CloudWalk Technology, has turned to African governments to address the obvious, race-based, limits of their existing algorithms.

In April, 2018, the Zimbabwean government announced that it had selected CloudWalk to provide an AI-based facial recognition system based on the extensive government national identity database.(Reporter 2018) Only later did the fact emerge that the state, which is notoriously short of foreign currency and in the middle of a contested military seizure of power, gave the Chinese company millions of records from the national population register -- including photographs and real name and identity numbers -- in exchange for the new surveillance tools. "The Zimbabwe deal is unique," Amy Hawkins observed, "in that as part of the agreement—the value of which CloudWalk declined to share—Harare will send data on millions of black faces to the Chinese company to help train the technology toward darker skin tones. The currency here is data as well as dollars." (Hawkins 2018) Part of the self-interest for the companies involved is that this allows a wide expansion of the scope of the training data by including millions of African faces and names, which may correct the racist bias that has long plagued facial recognition systems. There is wide concern that Chinese companies have found easy accommodation in those African countries -- including Ethiopia, Tanzania, Uganda and Rwanda -- that share a common vision of bureaucratic control and surveillance and weak privacy laws. These relationships are important and easily mobilised for sources of standardised data, but they are not the most extensive or valuable.

The Mobile Network Operators – MTN, Orange, Vodacom – maintain databases of hundreds of millions of subscribers, which include identifying information like names and photographs, and in some countries (famously Nigeria) biometrics as well. The licensing requirements for cellular providers also impose an intriguing paradox on the companies: they are required to maintain databases that can assist the state in allowing "authorised interception of communications" while, at the same time, they are specifically constrained to gather only the information necessary for "business purposes" and prohibited from sharing it with other firms.(Federal Republic of Nigeria 2007) The result is that the companies store very large databases, including airtime transactions and the meta-data of calls, data connections and texts, but – with the exception of Safaricom in East Africa and isolated experiments by Orange in Senegal and Cote D'ivoire – they do not process the information or sell it to other companies. Given the explosive intimacy of internet transactions it is easy to see why these firms have been reluctant to convert these data into commercial assets. Phone use is also unusual, as many people maintain multiple devices and multiple sim cards in order to maximize their access to unevenly distributed and priced networks. (Bezuidenhoudt and Breckenridge 2018; Mann 2017) Yet, as Safaricom's success in building an automated scoring system for credit allocation suggests, the commercialisation of the meta-data of the 600 million phone users is entirely possible.(Cook and McKay 2015) The Grameen bank has several projects currently underway in Uganda that mine clients'

A similar situation has developed around Facebook and the other popular social media platforms. After initially offering platform-subsidized free access using the smart card applications most of the operators have now moved to subsidised data packages that allow parsimonious access to specific platforms. In South Africa MTN offers a 30 day 1gb single-purpose tariff for WhatsApp or Twitter at a cost of R10 (less than a dollar). (McKane 2018) In Nigeria they offer the same product – access using the app only to WhatsApp, Instagram or several others on an unmetered basis -- for about much less than a dollar a month. These walled-garden data packages, and, especially communications on WhatsApp, are producing the volume of data that machine learning algorithms require. Nearly 200 million people on the continent are currently active users of the personal conversation networks supported by WhatsApp -- which typically involve users in overlapping but discrete conversations structured by home, family, lineage, friends, work and politics.(Note the historical potency of these networks: Geschiere and Nyamnjoh 2000) In some countries the platform now consumes more than half of all the available traffic. This data would lend itself to very powerful algorithmic classifications. At least as far as the company will acknowledge publicly, only the commitment to end-to-end encryption has to date prevented the development of very powerful sorting and diagnostic algorithms. That is apparently now changing.(Dahir 2018; Solon 2018) Assembling big data, to use the name of an older ICT obsession, is not the problem for artificial intelligence on the continent.

Systematic racial, gender and class bias is a much more serious problem. Researchers have already demonstrated convincingly that the scoring algorithms widely used by US courts to assess the likely risks of an accused committing crime in the future discriminate methodically against black people. The bias is imported into these human-coded algorithms by relying, amongst other secret indicators, on survey questions like “How many of your friends are taking drugs illegally?” There is no doubt amongst machine-learning developers and researchers that self-taught algorithms – trained on data from the existing criminal justice system – will adopt similar, or even more direct, proxy variables for race in developing predictive scoring systems. (Angwin et al. 2016; Winston 2018) It is this obvious bias – and the new political danger that it generates -- that has encouraged the development of activist groups like Black in AI. (“Black in AI - Home” n.d.) It is also important to notice how quickly the field of machine learning has responded to these problems of bias, with a significant portion of the key papers at the NIPS conferences addressing problems questions of fairness and bias, and searching for algorithmic tools of diagnosis and remedy. (Pleiss et al. 2017; Kusner et al. 2017; Adler et al. 2018; Thomas et al. 2017)

In comparison with the entrenched, dispersed and authoritarian forms of bureaucratic racism that exist everywhere, artificial intelligence, as a field, is strongly inclined toward a radical egalitarianism. “Machine learning scientists come from every technical background imaginable: probability theory and statistics, physics and economics, neuroscience and psychology, operations research and signal processing, and many beyond,” as the organisers of the recent IndabaX in Johannesburg have observed: “These scientists are also men and women, are Africans and Asians, are black and white, are lesbian and gay and transgendered. They come from every part of the world, speak in different languages and accents, hold different beliefs, and are shaped by different histories. Amidst all this diversity, they are singularly committed to discovering the principles of learning in brains and machines, and to the beneficial applications of this knowledge."(Mohamed 2017, 17) It seems, at this stage of the rebirth of AI, that the prospects for careful self-awareness about the indirect importation of racialised models of classification into machine learning are actually very good.

References:

Adler, Philip, Casey Falk, Sorelle A. Friedler, Tionney Nix, Gabriel Rybeck, Carlos Scheidegger, Brandon Smith, and Suresh Venkatasubramanian. 2018. “Auditing Black-Box Models for Indirect Influence.” *Knowledge and Information Systems* 54 (1): 95–122.

African Development Bank. 2018. “African Economic Outlook.” Abidjan, Ivory Coast: African Development Bank Group.

Ake, Claude. 1981. *A Political Economy of Africa*. Harlow, Essex: Longman.

Alister. 2017. *Ali Rahimi’s Talk at NIPS(NIPS 2017 Test-of-Time Award Presentation)*. https://www.youtube.com/watch?v=Qi1Yry33TQE&feature=youtu.be.

Bayart, Jean-Francois. 2000. “Africa in the World: A History of Extraversion.” *African Affairs* 99 (395): 217–67. https://doi.org/10.1093/afraf/99.395.217.

Beard, Alex. 2018. “China’s Children Are Its Secret Weapon in the Global AI Arms Race.” WIRED UK. April 19, 2018. http://www.wired.co.uk/article/china-artificial-intelligence-education-superpower.

Beniger, James R. 1986. *The Control Revolution: Technological and Economic Origins of the Information Society*. Cambridge MA: Harvard University Press.

Bezuidenhoudt, Johann, and Keith Breckenridge. 2018. “African Meta-Data and ML,” May 6, 2018.

“Black in AI - Home.” n.d. Accessed March 30, 2018. https://www.facebook.com/blackinai/.

Breckenridge, Keith. 2014. *Biometric State: The Global Politics of Identification and Surveillance in South Africa, 1850 to the Present*. Cambridge: Cambridge University Press.

“Baidu offers facial recognition technology to help Beijing airport streamline boarding, traffic”, *South China Morning Post*, 24 August 2017. https://www.scmp.com/tech/china-tech/article/2108163/baidu-offers-facial-recognition-technology-help-beijing-airport.

Captain, Sean. 2016. “Baidu Says Its New Face Recognition Tech Is Better Than Humans At Checking IDs.” Fast Company. November 17, 2016. https://www.fastcompany.com/3065778/baidu-says-new-face-recognition-can-replace-checking-ids-or-tickets.

Chutel, Lynsey. 2017. “Robots Are Set to Take Africa’s Manufacturing Jobs Even before It Has Enough.” *Quartz* (blog). July 26, 2017. https://qz.com/1037225/robots-are-set-to-take-africas-manufacturing-jobs-even-before-it-has-enough/.

Cook, Tamara, and Claudia McKay. 2015. “How M-Shwari Works: The Story so Far.” Forum 10. Washington, DC: Consultative Group to Assist the Poor (CGAP) and Financial Sector Deepening (FSD). http://www.cgap.org/sites/default/files/publications/multimedia/CGAP_Annual_Report_2015/files/forum/forum-how-m-shwari-works-apr-2015.pdf.

Crevier, Daniel. 1993. *AI : The Tumultuous History of the Search for Artificial Intelligence*. USA: Basic Books. http://espace2.etsmtl.ca/id/eprint/3778/.

Dahir, Abdi Latif. 2018. “WhatsApp Is the Most Popular Messaging App in Africa.” *Quartz* (blog). February 14, 2018. https://qz.com/1206935/whatsapp-is-the-most-popular-messaging-app-in-africa/.

Does AI Need More Innate Machinery? 2017. https://www.youtube.com/watch?v=aCCotxqxFsk.

Domingos, Pedro. 2015a. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York: Basic Books, a member of the Perseus Books Group.

Executive Office of the President. 2016. “Artificial Intelligence, Automation and the Economy.” Washington DC: United States Government.

Executive Office of the President, and National Science and Technology Council. 2016. "Preparing for the Future of Artificial Intelligence." Washington DC: United States Government.
Federal Republic of Nigeria. 2007. *Nigerian Communications Act : Consumer Code of Regulations*. Vol. FGP 122/82007/300 (OL 80).
Hinton, Geoffrey. 2017. Heroes of Deep Learning: Andrew Ng interviews Geoffrey Hinton Interview by Andrew Ng. https://www.youtube.com/watch?v=-eyhCTvrEtE.
Hutson, Matthew. 2018. "AI Researchers Allege That Machine Learning Is Alchemy." *Science*, May 2, 2018. http://www.sciencemag.org/news/2018/05/ai-researchers-allege-machine-learning-alchemy.
Gagliardone, Iginio. 2016. *The Politics of of Technology in Africa: Communication, Development, and Nation-Building in Ethiopia*. New York: Cambridge University Press.
Gagliardone, Iginio, and Frederick Golooba-Mutebi. 2016. "The Evolution of the Internet in Ethiopia and Rwanda: Towards a 'Developmental' Model?" *Stability: International Journal of Security and Development* 5 (1). https://doi.org/10.5334/sta.344.
Geschiere, P., and F. Nyamnjoh. 2000. "Capitalism and Autochthony: The Seesaw of Mobility and Belonging." *Public Culture* 12 (2): 423–52. https://doi.org/10.1215/08992363-12-2-423.
Google Trends. 2018. "Global Search Interest for the Fourth Industrial Revolution, Jan 2010 - May 2018." Google Trends. May 8, 2018. /trends/explore.
Hawkins, Amy. 2018. "Beijing's Big Brother Tech Needs African Faces." *Foreign Policy* (blog). July 24, 2018. https://foreignpolicy.com/2018/07/24/beijings-big-brother-tech-needs-african-faces/.
Hvistendahl, Mara. 2017. "In China, a Three-Digit Score Could Dictate Your Place in Society." WIRED. December 14, 2017. https://www.wired.com/story/age-of-social-credit/.
Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. "Machine Bias." Text/html. ProPublica. May 23, 2016. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.
Lee, Kai-Fu. 2018. We Are Here To CreateEdge.org. https://www.edge.org/conversation/kai_fu_lee-we-are-here-to-create.
Kline, Ronald R. 2015. *The Cybernetics Moment: Or Why We Call Our Age the Information Age*. JHU Press.
Kuo, Lily. 2016. "Only China, Syria, and Iran Rank Worse in Internet Freedom than Ethiopia." *Quartz* (blog). November 16, 2016. https://qz.com/838908/internet-freedom-in-ethiopia-is-the-fourth-worst-in-the-world-after-iran-syria-and-china/.
Kusner, Matt J, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. "Counterfactual Fairness." In *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 4066–4076. Curran Associates, Inc. http://papers.nips.cc/paper/6995-counterfactual-fairness.pdf.
Lohr, Steve. 2018. "Facial Recognition Is Accurate, If You're a White Guy." *The New York Times*, February 9, 2018, sec. Technology. https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html.
Mail Reporter. 2018. "Zimbabwe Imports Facial ID Technology from China." *The Zimbabwe Mail* (blog). April 12, 2018. http://www.thezimbabwemail.com/technology-science/zimbabwe-imports-facial-id-technology-from-china/.
Mann, Laura. 2017. "Left to Other Peoples' Devices? A Political Economy Perspective on the Big Data Revolution in Development." *Development and Change*, n/a-n/a. https://doi.org/10.1111/dech.12347.

Marcus, Gary. 2018. “Deep Learning: A Critical Appraisal.” *ArXiv:1801.00631 [Cs, Stat]*, January. http://arxiv.org/abs/1801.00631.

Marwala, Tshilidzi. 2018. “The Fourth Industrial Revolution Is Discriminating against Africans.” *Sunday Independent*, April 15, 2018.

McKane, Jamie. 2018. “Social Data Bundles Compared – MTN vs Cell C vs Telkom vs Vodacom.” May 5, 2018. https://mybroadband.co.za/news/cellular/258761-social-data-bundles-compared-mtn-vs-cell-c-vs-telkom-vs-vodacom.html.

Minsky, Marvin, and Seymour A. Papert. 1987. *Perceptrons: An Introduction to Computational Geometry, Expanded Edition*. Expanded edition edition. Cambridge, Mass.: The MIT Press.

Mohamed, Shakir. 2017. “Strengthening African Machine Learning : Outcomes of the First Deep Learning Indaba.” The Deep Learning Indaba.

Ng, Andrew. 2017. “Artificial Intelligence Is the New Electricity.” February 2. https://www.youtube.com/watch?v=21EiKfQYZXc&Z=630s.

Nielsen, Michael A. 2015. “Neural Networks and Deep Learning.” http://neuralnetworksanddeeplearning.com.

Noble, David F. 1977. *America By Design: Science, Technology and the Rise of Corporate Capitalism*. Oxford: Oxford University Press Oxford: Oxford University Press.

Pleiss, Geoff, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. 2017. “On Fairness and Calibration.” In *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 5680–5689. Curran Associates, Inc. http://papers.nips.cc/paper/7151-on-fairness-and-calibration.pdf.

Raji, Inioluwa. 2018. “Black Panther Face Scorecard: Wakandans Under the Coded Gaze of AI.” *Medium* (blog). March 13, 2018. https://medium.com/@inioraji/black-panther-face-scorecard-wakandans-under-the-coded-gaze-of-ai-588e3d711dab.

Rodrik, Dani. 2016. “Premature Deindustrialization.” *Journal of Economic Growth* 21 (1): 1–33. https://doi.org/10.1007/s10887-015-9122-3.

Royal Society. 2017. “Machine Learning : The Power and Promise of Computers That Learn by Example.” https://royalsociety.org/topics-policy/projects/machine-learning/.’

Rosenblatt, F. 1957. *The Perceptron, a Perceiving and Recognizing Automaton Project Para*. Cornell Aeronautical Laboratory.

Schwab, Klaus. 2017. *The Fourth Industrial Revolution*. First U.S. edition. New York: Crown Business.

Sculley, D., Jasper Snoek, Alex Wiltschko, and Ali Rahimi. 2018. “Winner’s Curse? On Pace, Progress, and Empirical Rigor,” February. https://openreview.net/forum?id=rJWF0Fywf.

Shapin, Steven, and Simon Schaffer. 1985. *Leviathan and the Air-Pump*. Princeton: Princeton University Press.

Solon, Olivia. 2018. “WhatsApp CEO Jan Koum Quits over Privacy Disagreements with Facebook.” *The Guardian*, April 30, 2018, sec. Technology. http://www.theguardian.com/technology/2018/apr/30/jan-koum-whatsapp-co-founder-quits-facebook.

State Council. 2017. “State Council Notice on the Issuance of the Next Generation Artificial Intelligence Development Plan.” New America. July 20, 2017. https://www.newamerica.org/cybersecurity-initiative/blog/chinas-plan-lead-ai-purpose-prospects-and-problems/.

Stigler, S. M. 1986. *The History of Statistics: The Measurement of Uncertainty before 1900*. Belknap Press.

Stigler, Stephen M. 1989. “Francis Galton’s Account of the Invention of Correlation.” *Statistical Science* 4 (2): 73–79.
Peng, Tony. 2017. “LeCun vs Rahimi: Has Machine Learning Become Alchemy?” *Medium* (blog). December 12, 2017. https://medium.com/syncedreview/lecun-vs-rahimi-has-machine-learning-become-alchemy-21cb1557920d.
The Economist. 2018. “Non-Tech Businesses Are Beginning to Use Artificial Intelligence at Scale,” March 31, 2018. https://www.economist.com/printedition/2018-03-31.
CB Insights Research. 2018. “The Race For AI: Google, Intel, Apple In A Rush To Grab Artificial Intelligence Startups.” CB Insights Research. February 27, 2018. /research/top-acquirers-ai-startups-ma-timeline/.
Thomas, Philip S., Bruno Castro da Silva, Andrew G. Barto, and Emma Brunskill. 2017. “On Ensuring That Intelligent Machines Are Well-Behaved.” *ArXiv Preprint ArXiv:1708.05448*.
Wiener, Norbert. 1988. *The Human Use of Human Beings: Cybernetics and Society*. Cambridge MA: Da Capo Press.
Winston, Ali. 2018. “Palantir Has Secretly Been Using New Orleans to Test Its Predictive Policing Technology.” The Verge. February 27, 2018. https://www.theverge.com/2018/2/27/17054740/palantir-predictive-policing-tool-new-orleans-nopd.
World Wide Web Foundation. 2017. “Artificial Intelligence : The Road Ahead in Low and Middle-Income Countries.” Washington DC: WWW Foundation. https://webfoundation.org/docs/2017/07/AI_Report_WF.pdf.
Zuboff, Shoshana. 1988. *In the Age of the Smart Machine: The Future of Work and Power*. Boston: Basic Books.