

Towards an anti-fascist AI

Dan McQuillan

This talk was given at the launch of the 'All Access AI' network at Goldsmiths, University of London 1st April 2019.

intro

This talk is about some pressing issues with AI that don't usually make the headlines, and why tackling those issues means developing an antifascist AI.

When I talk about AI i'm talking about machine learning and about artificial neural networks, also known as deep learning¹. I'm addressing actual AI not a literary or filmic narrative about post-humanism.

AI is political. Not only because of the question of what is to be done with it, but because of the political tendencies of the technology itself. The possibilities of AI arise from the resonances between its concrete operations and the surrounding political conditions. By influencing our understanding of what is both possible and desirable it acts in the space between what is and what ought to be.

concrete

Computers are essentially just faster collections of vacuum tubes. How can they emulate human activities like recognising faces or assessing criminality?

Think about a least squares fit; you're trying assess the correlation between two variables by fitting a straight line to scattered points, so you calculate the sum of the squares of distances from all points to the line and minimise that². Machine learning does something very similar. It makes your data into vectors in feature space so you can try to find boundaries between classes by minimising sums of distances as defined by your objective function³.

These patterns are taken as revealing something significant about the world. They take on the neoplatonism of the mathematical sciences; a belief in a layer of reality which can be best perceived mathematically⁴. But these are patterns based on correlation not causality; however complex the computation there's no comprehension or even common sense.

Neural networks doing image classification are easily fooled by strange poses of familiar objects. So a school bus on it's side is confidently classified as a snow plough⁵. Yet the hubristic knights of AI are charging into messy social contexts, expecting to be able to draw out insights that were previously the domain of discourse.

Deep learning is already seriously out of its depth.

callousness

Will this slow the adoption of AI while we figure out what it's actually good for? No, it won't; because what we are seeing is 'AI under austerity', the adoption of machinic methods to sort things out after the financial crisis.

The way AI derives its optimisation from calculations based on a vast set of discrete inputs matches exactly the way neoliberalism sees the best outcome coming from a market freed of constraints. AI is seen as a way to square the circle between eviscerated services and rising demand awithout having to challenge the underlying logic⁶.

The pattern finding of AI lends itself to prediction and therefore preemption which can target what's left of public resource to where the trouble will arise, whether that's crime, child abuse or dementia. But there's no obvious way to reverse operations like backpropagation to human reasoning⁷, which not only endangers due process but produces thoughtlessness in the sense that Hannah Arendt meant it⁸; the inability to critique instructions, the lack of reflection on consequences, a commitment to the belief that a correct ordering is being carried out.

The usual objection to algorithmic judgements is outrage at the false positives, especially when they result from biased input data. But the underlying problem is the imposition of an optimisation based on a single idea of what is for the best, with a resultant ranking of the deserving and the undeserving⁹.

What we risk with the uncritical adoption of AI is algorithmic callousness, which won't be saved by having a human-in-the-loop because that human will be subsumed by the self-interested institution-in-the-loop.

By throwing out our common and shared conditions as having no predictive value, the operations of AI targeting strip out any acknowledgement of system-wide causes hiding the politics of the situation.

far right

The algorithmic coupling of vectorial distances and social differences will become the easiest way to administer a hostile environment, such as the one created by Theresa May to target immigrants.

But the overlaps with far right politics don't stop there. The character of 'coming to know through AI' involves reductive simplifications based on data innate to the analysis, and simplifying social problems to matters of exclusion based on innate characteristics is precisely the politics of right wing populism.

We should ask whether the giant AI corporations would balk at putting the levers of mass correlation at the disposal of regimes seeking national rebirth through rationalised ethnocentrism. At the same time that Daniel Guerin was writing his book in 1936 examining the ties between fascism and big business¹⁰, Thomas Watson's IBM and its German subsidiary Dehomag were enthusiastically furnishing the nazis with Hollerith punch card technology¹¹. Now we see the photos from Davos of Jair Bolsonaro seated at lunch between Apple's Tim Cook and Microsoft's Satya Nadella¹².

Meanwhile the algorithmic correlations of genome wide association studies are used to sustain notions of race realism and prop up a narrative of genomic hierarchy¹³. This is already a historical reunification of statistics and white supremacy, as the mathematics of logistic regression and correlation that are so central to machine learning were actually developed by Edwardian eugenicists Francis Galton and Karl Pearson.

antifascist

My proposal here is that we need to develop an antifascist AI.

It needs to be more than debiasing datasets because that leaves the core of AI untouched.

It needs to be more than inclusive participation in the engineering elite because that, while important, won't in itself transform AI. It needs to be more than an ethical AI, because most ethical AI operates as PR to calm public fears while industry gets on with it¹⁴. It needs to be more than ideas of fairness expressed as law, because that imagines society is already an even playing field and obfuscates the structural asymmetries generating the perfectly legal injustices we see deepening every day.

I think a good start is to take some guidance from the feminist and decolonial technology studies that have cast doubt on our cast-iron ideas about objectivity and neutrality¹⁵. Standpoint theory suggests that positions of social and political disadvantage can become sites of analytical advantage, and that only partial and situated perspectives can be the source of a strongly objective vision¹⁶. Likewise, a feminist ethics of care takes relationality as fundamental¹⁷, establishing a relationship between the inquirer and their subjects of inquiry would help overcome the onlooker consciousness of AI.

To centre marginal voices and relationality, I suggest that an antifascist AI involves some kinds of people's councils, to put the perspective of marginalised groups at the core of AI practice and to transform machine learning into a form of critical pedagogy¹⁸. This formation of AI would not simply rush into optimising hyperparameters but would question the origin of the problematics, that is, the structural forces that have constructed the problem and prioritised it.

AI is currently at the service of what Bergson called ready-made problems; problems based on unexamined assumptions and institutional agendas, presupposing solutions constructed from the same conceptual asbestos¹⁸. To

have agency is to re-invent the problem, to make something newly real that thereby becomes possible unlike the probable, the possible is something unpredictable, not a rearrangement of existing facts.

Given the corporate capture of AI, any real transformation will require a shift in the relations of production. One thing that marks the last year or so is the sign of internal dissent in Google¹⁹, Amazon²⁰, Microsoft²¹, Salesforce and so on about the social purposes to which their algorithms are being put. In the 1970s workers in a UK arms factory came up with the Lucas plan which proposed the comprehensive restructuring of their workplace for socially useful production²². They not only questioned the purpose of the work but did so by asserting the role of organised workers, which suggests that the current tech worker dissent will become transformative when it sees itself as creating the possibility of a new society in the shell of the old²³.

I'm suggesting that an antifascist AI is one that take sides with the possible against the probable, and does so at the meeting point between organised subjects and organised workers. But it may also require some organised resistance from communities.

A thread is a sequence of programmed instructions executed by microprocessor. On an Nvidia GPU, one of the AI chips, a warp is a set of threads executed in parallel²⁴. How uncanny that the language of weaving looms has followed us from the time of the Luddites to the era of AI. The struggle for self-determination in everyday life may require a new Luddite movement, like the residents and parents in Chandler, Arizona who have blockaded Waymo's self-driving vans 'They didn't ask us if we wanted to be part of their beta test' said a mother who's child was nearly hit by one²⁵ The Luddites, remember, weren't anti-technology but aimed 'to put down all machinery hurtful to the Commonality'²⁶.

The predictive pattern recognition of deep learning is being brought to bear on our lives with the granular resolution of Lidar. Either we will be ordered by it or we will organise. So the question of an antifascist AI is the question of self-organisation, and of the autonomous production of the self that is organising.

Asking 'how can we predict who will do X?' is asking the wrong question. We already know the destructive consequences of on the individual and collective psyche of poverty, racism and systemic neglect. We don't need AI as targeting but as something that helps raise up whole populations²⁷.

Real AI matters not because it heralds machine intelligence but because it confronts us with the unresolved injustices of our current system. An antifascist AI is a project based on solidarity, mutual aid and collective care. We don't need autonomous machines but a technics that is part of a movement for social autonomy.

Notes

¹Nielsen, Michael A. 2015. 'Neural Networks and Deep Learning'. 2015. <http://neuralnetworksanddeeplearning.com>.

²Chauhan, Nagesh Singh. n.d. 'A Beginner's Guide to Linear Regression in Python with Scikit-Learn'. Accessed 1 April 2019. <https://www.kdnuggets.com/2019/03/beginners-guide-linear-regression-python-scikit-learn.html>, <https://www.kdnuggets.com/2019/03/beginners-guide-linear-regression-python-scikit-learn.html>.

³Geitgey, Adam. 2014. 'Machine Learning Is Fun!' Adam Geitgey (blog). 5 May 2014. <https://medium.com/@ageitgey/machine-learning-is-fun-80ea3ec3c471>.

⁴McQuillan, Dan. 2017. 'Data Science as Machinic Neoplatonism'. *Philosophy & Technology*, August, 1–20. <https://doi.org/10.1007/s13347-017-0273-3>.

⁵Alcorn, Michael A., Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, and Anh Nguyen. 2018. 'Strike (with) a Pose: Neural Networks Are Easily Fooled by Strange Poses of Familiar Objects'. *ArXiv:1811.11553 [Cs]*, November. <http://arxiv.org/abs/1811.11553>.

⁶Ghosh, Pallab. 2018. 'AI Could Save Heart and Cancer Patients', 2 January 2018, sec. Health. <https://www.bbc.com/news/health-42357257>.

⁷Knight, Will. n.d. 'There's a Big Problem with AI: Even Its Creators Can't Explain How It Works'. *MIT Technology Review*. Accessed 1 April 2019. <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>.

⁸Arendt, Hannah. 2006. *Eichmann in Jerusalem: A Report on the Banality of Evil*. 1 edition. New York, N.Y: Penguin Classics.

⁹Big Brother Watch. 2018. 'The UK's Poorest Are Put at Risk by Automated Welfare Decisions'. 7 November 2018. <https://bigbrotherwatch.org.uk/2018/11/the-uks-poorest-are-put-at-risk-by-automated-welfare-decisions/>.

¹⁰Guerin, Daniel. 2000. *Fascism and Big Business*. Translated by Francis Merrill and Mason Merrill. 2nd edition. Pathfinder Press.

¹¹Black, Edwin. 2012. *IBM and the Holocaust: The Strategic Alliance Between Nazi Germany and America's Most Powerful Corporation-Expanded Edition*. Expanded edition. Washington, DC: Dialog Press.

¹²Slobodian, Quinn. 2019. 'The Rise of the Right-Wing Globalists'. *The New Statesman*. 31 January 2019. <https://www.newstatesman.com/politics/economy/2019/01/the-rise-of-the-right-wing-globalists>.

¹³Comfort, Nathaniel. 2018. 'Sociogenomics Is Opening a New Door to Eugenics'. *MIT Technology Review*. 23 October 2018. <https://www.technologyreview.com/s/612275/sociogenomics-is-opening-a-new-door-to-eugenics/>.

¹⁴Metz, Cade. 2019. 'Is Ethical A.I. Even Possible?' *The New York Times*, 5 March 2019, sec. Business. <https://www.nytimes.com/2019/03/01/business/artificial-intelligence.html>.

¹⁵Harding, Sandra. 1998. *Is Science Multicultural?: Postcolonialisms, Feminisms, and Epistemologies*. 1st edition. Bloomington, Ind: Indiana University Press.

¹⁶Haraway, Donna. 1988. 'Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective'. *Feminist Studies* 14 (3): 575–99. <https://doi.org/10.2307/3178066>.

¹⁷'Carol Gilligan Interview.' 2011. *Ethics of Care (blog)*. 16 July 2011. <https://ethicsofcare.org/carol-gilligan/>.

¹⁸Solhdju, Katrin. 2015. 'Taking Sides with the Possible against Probabilities or: How to Inherit the Past'. In . *ICI Berlin*. https://www.academia.edu/19861751/Taking_Sides_with_the_Possible_against_Probabilities_or_How_to_Inherit_the_Past.

¹⁹Shane, Scott, Cade Metz, and Daisuke Wakabayashi. 2018. 'How a Pentagon Contract Became an Identity Crisis for Google'. *The New York Times*, 26 November 2018, sec. Technology. <https://www.nytimes.com/2018/05/30/technology/google-project-maven-pentagon.html>.

²⁰Conger, Kate. n.d. 'Amazon Workers Demand Jeff Bezos Cancel Face Recognition Contracts With Law Enforcement'. *Gizmodo*. Accessed 1 April 2019. <https://gizmodo.com/amazon-workers-demand-jeff-bezos-cancel-face-recognition-1827037509>.

²¹Lee, Dave. 2019. 'Microsoft Staff: Do Not Use HoloLens for War', 22 February 2019, sec. Technology. <https://www.bbc.com/news/technology-47339774>.

²²Open University. 1978. Lucas Plan Documentary. <https://www.youtube.com/watch?v=0pgQqfpub-c>.

²³'Preamble to the IWW Constitution | Industrial Workers of the World'. 1905. 1905. <http://www.iww.org/culture/official/preamble.shtml>.

²⁴Lin, Yuan, and Vinod Grover. 2018. 'Using CUDA Warp-Level Primitives'. NVIDIA Developer Blog. 16 January 2018. <https://devblogs.nvidia.com/using-cuda-warp-level-primitives/>.

²⁵Romero, Simon. 2019. 'Wielding Rocks and Knives, Arizonans Attack Self-Driving Cars'. The New York Times, 2 January 2019, sec. U.S. <https://www.nytimes.com/2018/12/31/us/waymo-self-driving-cars-arizona-attacks.html>.

²⁶Binfield, Kevin, ed. 2004. Writings of the Luddites. Baltimore: Johns Hopkins University Press.

²⁷Keddell, Emily. 2015. 'Predictive Risk Modelling: On Rights, Data and Politics.' Re-Imagining Social Work in Aotearoa New Zealand (blog). 4 June 2015. <http://www.reimaginingocialwork.nz/2015/06/predictive-risk-modelling-on-rights-data-and-politics/>. [^