Non-Fascist AI

Dan McQuillan

This text was first published in Propositions for Non-Fascist Living: Tentative and Urgent, Maria Hlavajova and Wietske Maas, eds. (Utrecht and Cambridge, MA: BAK basis voor actuele kunst and MIT Press, 2019)

Does AI have a politics?

When discussing propositions for a non-fascist artificial intelligence (AI), we should first ask whether AI has a politics at all. AI as it really exists is simply machine learning and neural networks[1], not a post-human intelligence with its own agenda. But there are authoritarian tendencies associated with AI which arise from the resonance between its concrete operations and the surrounding political conditions. We have reason to worry not only about the uses to which it might be put but also about the inherent characteristics of the technology itself. To make these characteristics visible so they can be overcome requires some appreciation of the actual workings of AI.

Machine learning and neural networks are forms of computational pattern finding. Their uncanny ability to recognize faces or play Go doesn't come from a spark of consciousness in the circuitry but from vast calculations of probabilities. To render a real world situation amenable to AI means abstracting numerical features and running many examples through iterative calculations involving an objective function. The objective function defines some abstract mathematical distance between the model's current classification of the pattern and the correct one, and strives to minimize or optimize that. What makes this operation powerful is the convergence of computational methods such as backpropagation[2], the vast seas of data

[1] . The general class of algorithmic methods known as "machine learning" use statistical techniques to build a mathematical model based on training data. In this way they can "learn" how to perform a specific task such as classification without needing explicit instructions. Neural networks are a specific subset that use layers of interconnected nodes ("neurons") whose output signals are computed sums of their inputs. This "connectionist" approach to machine learning has recently become massively successful and is what most people now mean when they refer to AI.

[2] . Backpropagation, shorthand for "the backward propagation of errors," is a mathematical method used to minimize errors in the training of a multi-layer neural network. By computing the iterative adjustments needed at each layer to improve the predictions of the network, it enables a process of "gradient descent" that optimizes the mapping between sets of inputs (data) and outputs (labeled examples). A very accessible visual explanation of backpropagation is

that are now available, and the development of microprocessors called GPUs that can handle the necessary scale of parallel calculations.

The result is a giant leap in the abilities of automated systems. They can recognize faces as well as we can, so perhaps they can perform cancer scans with the same kind of accuracy. And what about law enforcement? Judges deciding bail already review a suspect's criminal record and associations; perhaps if we expand the data involved and run it through our AI we can discern a more reliably predictive pattern. And here is the first problem: for a culture steeped with faith in scientific authority, empirical calculations are taken to reveal something more significant about the world than mere experience. Never mind that these patterns are correlations rather than some determination of causality. The engineering solutionism that defines Silicon Valley sees no limits to the application of AI in messy social contexts, so it can draw out insights that were previously obscured by unreliable and subjective discourse.

AI's character as a form of probabilistic prediction leads inevitably to preemption: forms of intervention intended to change the course of events. This has a particular appeal in a time of austerity and financial crisis. AI's targeting is seen as a way to square the circle between cash-strapped services and rising demand—and for corporations to keep squeezing a result out of falling rates of profit. Under the logic of neoliberalism the best result comes from the sum of signals in a free market, while the objective functions of AI make predictions by summing vectors over troves of training data. Thus, the optimizations of AI extend neoliberal logic into the near future. Decisions to allocate or withdraw resources become algorithmic means by which to bring the future into the present. AI provides a generalizable, technocratic method for reducing profound social questions to optimization problems. Like bureaucracy in the twentieth century, AI is poised to become the unifying logic of legitimation across corporations and government.

Is AI fascist?

The computational classifications produced by AI are inserted into processes that previously entailed dialogue, debate and human judgment. They augment forms of governmentality through empirical speed and scale, while obscuring decisions that were previously understood to be political. This obfuscation is intensified when neural networks are involved, as they

given in the 3Blue1Brown video "What is backpropagation really doing?" online at: https://www.youtube.com/watch?v=Ilg3gGewQ5U

have an interpretability problem; the massively parallel and iterative calculations undertaken with neural networks can't be directly translated back into terms accessible to human reasoning. However, the risk posed by AI is not a machine tyranny of automated decisions but the amplification of existing human tendencies to automaticity. AI not only undermines due process but produces thoughtlessness, in the sense that political philosopher Hannah Arendt meant when interpreting the actions of Nazi war criminal Adolf Eichmann; the inability to critique instructions, the lack of reflection on consequences, a commitment to the belief that a correct ordering is being carried out. This is a form of everyday fascist thinking even when it is not undertaken on behalf of an explicitly far-right regime.

AI's predictions create new terrains of inclusion and exclusion with respect to resources and opportunities. By combining both judgment and intervention, the machinic classifications of AI will have the force of law without being of the law. According to Agamben, this "force of" is the signature of states of exception—conditions in which rights and citizenship are suspended[3]. The actions of everyday algorithms will not create a state of emergency per se, but consitute a continuous partial modulation of life chances. Whether the government is captured by far-right political movements or is simply appeasing them, the legal lacunae created by AI are ready to become the vectors for far-right policies. This was Agamben's point about the fascist states of the 1930s—they were not dictatorships but dual states with second structures that existed alongside the constitutional ones. He warned us that what we need to be alert to is not a confusion of legislative and executive powers but this separation of law and force of law.

But the overlap with far-right politics doesn't stop there. The character of "coming to know through AI" involves simplifications based on data innate to the analysis, and the reduction of 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 rationalized ethnocentrism. At the same time that French anarcho-communist Daniel

3   . The classic state of exception occurs when a government declares a state of emergency during a disaster or civil unrest. Authorities gain far-reaching powers and citizens lose existing protections, for example against detention without trial. A digitally mediated state of exception is the US No Fly List, which designates people who may be arbitrarily denied the right to board a plane without explanation. Invisible algorithmic decisions to withhold resources or trigger an intervention against an individual or group have the potential to operate in a similar manner.

Guerin was writing his book in 1936 examining the ties between fascism and big business, Thomas Watson's IBM and its German subsidiary Dehomag were enthusiastically furnishing the Nazis with Hollerith punch card technology. Now we see the photos from Davos of Brazil's reactionary populist president Jair Bolsonaro seated at lunch between Apple's Tim Cook and Microsoft's Satya Nadella.

Meanwhile the algorithmic correlations of genome-wide association studies are used to sustain notions of "race realism" and prop up a narrative of genomic hierarchy. 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. The risk of reductive optimization can be seen in the valorization of "general intelligence" by some of the most prominent AI practitioners. This occurs in the conflation of advances in AI as it actually exists, which is machine learning, and the belief that this is a stepping stone to Artificial General Intelligence (AGI)—the transhuman intelligence of fictional AI. What is notable in this context is not the weirdly religious belief in the coming of a higher intelligence, but the reductive identification of mind with intelligence and intelligence with IQ. Popular amongst the elite "tech bros" of Silicon Valley, this form of human ranking on the basis of an intelligence metric is also the slide rule for so-called race science, the pseudo-empirical justification of racial hierarchies that is a historical cornerstone of white supremacy.

How can we resist fascist AI?
We can't talk about a non-fascist AI without first considering how to resist the encroachments of AI that tend toward fascism. The well-meaning liberal position is to constrain the negative outcomes of AI in some way. Like AI itself, this has both "weak" and "strong" variants. The weak form is the hugely popular pastime of promoting AI ethics, which comes in 101 variations. This is already losing ground as it becomes clear that ethics is a PR exercise designed to calm public fears while industry gets on with unrestrained implementation. The strong form of liberal restraint is to compel fairness through law; that is, to address AI's production of imbalance in the same way as legislation against other discriminations. However, this imagines that society is already an even playing field and obfuscates the structural asymmetries generating the perfectly legal injustices we see deepening every day. It's exactly these that will be amplified by algorithmic systems.

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 organize. So the question of a non-fascist AI is the question of self-organization, and of the collective production of the self and community that comes about through organizing. There are already signs of resistance in the relations of production, such as the internal dissent at Google, Amazon, Microsoft and so on, which contests the social purposes to which their algorithms are being put. To become really effective, the resistance to toxic applications needs to become part of a broader worker self-organization that can mobilize an alternative social vision. 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 production[4]. They not only questioned the purpose of the work but did so by asserting the role of organized workers, which suggests that the current tech worker dissent will become transformative when it begins attempting to create the possibility of a new society in the shell of the old.

The resistance to fascist AI becomes stronger at the point where worker dissent meets social movements. The breakthrough represented by current AI came with the huge leap in its ability to recognize faces, as signaled by the results of the 2012 ImageNet competition. Ironically, it may be that facial recognition is also the breakthrough point for social resistance. The engineers behind this form of facial recognition themselves are active in calling for constraints, while at the same time it has become a concern for social movements such as Black Lives Matter and broad coalitions of groups who are alarmed by the extension of automated surveillance into their communities. Given the pace of AI's implementation across every aspect of social infrastructure, the struggle for self-determination in everyday life may require a new Luddite movement. New Luddites might look like the residents and parents in Chandler, Arizona who blockaded Waymo's self-driving vans and threw rocks at them. "They didn't ask us if we wanted to be part of their beta test" said a mother whose child was nearly hit by one. The Luddites, remember, did not oppose technology as such but

4 . In the midst of industrial unrest in 1976, a shop steward committee at Lucas Aerospace consulted with their members about alternative products for the factory to build. Their technology designs included wind turbines, hybrid power packs, and electric cars. Despite international support, the plan was rejected by corporate management. In a 1978 documentary by the Open University the workers themselves tell the story of the Lucas Plan, online at: https://www.youtube.com/watch?v=0pgQqfpub-c

aimed "to put down all machinery hurtful to the Commonality."[5]

How do we achieve non-fascist AI?

A good start on the path to non-fascist AI is to take some guidance from the feminist and decolonial technology studies that have cast doubt on our ironclad ideas about objectivity and neutrality. 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 vision[6]. Likewise, a feminist ethics of care takes relationality as fundamental; and the effort to establish a relationship between the inquirer and their subjects of inquiry would help overcome the onlooker consciousness of AI. The question is how to mobilize situated knowledges and an ethics of care in and around the everyday practices of AI. To center marginal voices and relationality, I suggest that a non-fascist AI must involve some kinds of people's councils, to put the perspective of marginalized groups at the core of AI practice and to transform machine learning into a form of critical pedagogy. This formation of AI would not simply rush into optimizing hyperparameters but would question the origin of the problematics, that is, the structural forces that have constructed the problem and prioritized it.

The purpose of people's councils in non-fascist AI is also the overcoming of subjection. The dispersal of algorithmic systems across everyday life means they come to shape our habitus, that is, our ingrained habits, skills, and dispositions. They become exactly as sociologist Pierre Bourdieu described, "structured structures predisposed to function as structuring structures."[7] Our relations with these systems of control affect how we know ourselves as subjects, our understanding of our place in society, and our relationships with other people and institutions. The AI we don't want draws its influence as an overlay to the authoritarian and dualistic conditioning that already exists in society. Achieving a measure of self-

5   . This particular phrase appeared in a letter of 10 March 1812 from "Ned Ludd" addressed "To Mr Smith Shearing Frame Holder at Hill End Yorkshire," cited in *Writings of the Luddites*, Kevin Binfield, ed.(Baltimore, MD: Johns Hopkins University Press, 2004), pp. 209–211. The Luddites have been consistently misrepresented in historical and contemporary accounts as being "against technology." They were actually a social movement concerned with self-regulatory power against the threats of the new capitalists and the government of the day.

6   . See, for example, Donna Haraway, "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective," *Feminist Studies*, v.14 no. 3 (Autumn 1988), pp.575–599.

7   . Pierre Bourdieu, *The Logic of Practice* (Redwood City, CA: Stanford University Press, 1990), p. 53.

realization and autonomy that is not based on individualism but solidarity is a collective activity. A non-fascist AI is one that supports autonomy; that supports freedom from the colonization of everyday life by the cultural codes of patriarchy, racism, or authoritarianism. In that way it can only be a practice of people who are committed to autonomy; to the principle of refusing to be dominated but also of refusing to dominate anyone else, by rejecting any participation in oppressive patterns of interaction. A people's council is a collective structure based on consensus, open to the equal voice of all participants, and grounded in the acknowledgement of all standpoints. Reaching decisions by consensus is the iterative process of finding positions to which everyone can commit.

AI is currently at the service of what philosopher Henri Bergson called ready-made problems; problems based on unexamined assumptions and institutional agendas, presupposing solutions constructed from the same conceptual asbestos[8]. 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. A non-fascist AI is one that takes sides with the possible against the probable. We need an AI that takes as its problematic exactly the sort of data features that are normally ignored as having no predictive value, because they represent shared conditions rather than individual idiosyncrasies. Rather than a technical exercise in the excision of risk, AI becomes a catalyst for more fundamental change and transformation. Asking "how can we predict who will do X?" is asking the wrong question. We already know the destructive consequences of poverty, racism, and systemic neglect, because that's what we can see multiplying around us. Imagining that any number of AI-enabled interventions can reverse this network effect is futile. We don't need AI as targeting but as something that works to raise up whole populations. Real AI matters not because it heralds machine intelligence but because it confronts us with the unresolved injustices of our current system. A non-fascist 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 and liberation.

---

8   . See Henri Bergson, *The Creative Mind: An Introduction to Metaphysics* (Mineola, NY: Dover Publications, 2010).