

Introduction

What is AI

This book is about how and why we should resist the introduction of artificial intelligence, or AI. It hopes to persuade the reader that resistance is what is needed, by showing how AI represents a technological shift in the framework of society that will amplify austerity while enabling authoritarian politics. However, despite the presentation of the varieties of AI harmfulness in the first part of the book, it is intended as an ultimately optimistic text, one that holds out the possibility of a radically transformative approach to AI that aligns with wider values of care and the common good. But before we get into discussing these developments, let alone what part we can play in them, we need to clarify what we mean by AI itself.

The book is concerned with actual AI as it operates in the world, not with the grandiose rhetoric or sci-fi storylines that obscure it. AI is, on a basic level, a set of specific computational operations, and Chapter 1 sets out to demystify these operations by bringing them out from behind the veil of technical obfuscation. However, AI is always more than a set of machine learning methods. When we're thinking about the actuality of AI, we can't separate the calculations in the code from the social context of its application. AI is never separate from the assembly of institutional arrangements that need to be in place for it to make an impact in society. Likewise, these institutions are immersed in wider frameworks of understanding that carry implicit and explicit assumptions about how the world is to be differentiated and valued. AI, as it is talked about in this book, is this layered and interdependent arrangement of technology, institutions and ideology. The general term we will use for this arrangement is 'apparatus'.

Most of this book uses deep learning as its technical reference point because deep learning is the dominant form of AI at the time of writing. It's important to refer to the actual technology because one of the themes of this text is that political impacts arise from resonances between concrete technical characteristics and the surrounding social and political conditions. Understanding AI means understanding its specific computational operations and everything that is being carried along by them; the history that AI has absorbed, the world in which it is emerging, and the futures that it calls forth. Some of what may seem, at the start, like nerdish technical detail will turn out to have significant political implications.

Having said that, the analysis presented here is not limited to deep learning. On the one hand, as the intent of the text is to interrupt the most dangerous tendencies incipient in AI before they come to pass, some of the case studies are not applications of AI as such but of precursor algorithmic systems; that is, algorithms that play some role in automated decision making but which are not themselves forms of machine learning. On the other hand, the broader thrust of the argument addresses not only deep learning, and its close cousins like reinforcement learning, but any subsequent computational system that offers a form of statistical optimization as a solution to social problems. As we'll see in more detail as we go through the book, any AI-like system will act as a condenser for existing forms of structural and cultural violence.

AI, as we know it, is a kind of computing, but it's also a form of knowledge production, a paradigm for social organization and a political project. While it might be interesting in another context to ask philosophical questions about the meaning of intelligence and whether it can ever be artificial, that's not the concern of this book, which instead sets out to ask what part AI plays in history as we are living it. Whatever else AI is, it is not neutral, and neither can we be. AI is political because it acts in the world in ways that affect the distribution of power, and its political tendencies are revealed in the ways that it sets up boundaries and separations. The apparatus of AI forms feedback loops with the rest of society: it's "a structured structure that becomes a structuring structure" (Bourdieu, 1980, cited in

Castelle, 2018). The focus here is on the ways that AI will alter the landscapes of our lives.

Resisting AI

The public narrative around AI has created high expectations. In the last few years AI seems to have accelerated from movie trope to material reality, with our cities about to be filled with self-driving cars and our health conditions diagnosed earlier and more accurately by apps. AI is being heralded as a potential solution to societal ills from child protection to climate change. On the other hand, this very acceleration has stirred up apocalyptic fears, from predictions by business pundits that AI will take all our jobs to the vision of AI as a dystopian superior intelligence. The superintelligent AI apocalypse is taken sufficiently seriously to occupy the full attention of both philosophers (Bostrom, 2014) and leading computer scientists in the field (Russell, 2020).

This book agrees that AI is important but not for any of the reasons given above. The theme explored throughout the text is that AI is a political technology in its material existence and in its effects. The concrete operations of AI are completely entangled with the social matrix around them, and the book argues that the consequences are politically reactionary. The net effect of applied AI, it is claimed, is to amplify existing inequalities and injustices, deepening existing divisions on the way to full-on algorithmic authoritarianism. In the light of these consequences, which are justified more fully in the following chapters, the book is titled after the stance it hopes to encourage, namely that of ‘resisting AI’.

Rather than focusing on what might happen if AI developed superintelligence, we look in Chapter 1 at the narrower reality of what AI technologies actually do; how their algorithms work, where the data comes from, and what social patterns feed in and out of these computational operations. The chapter digs into deep learning to reveal both its clever statistical manipulations and the gulf between this and anything we’d acknowledge as human-like intelligence. More importantly, it traces how the specific data transformations of deep learning shape its likely social effects. The chapter also looks at the hidden labour

relations without which deep learning would not exist, and at the substrate of circuits and servers that require vast systems of cooling and energy supply.

Chapter 2 makes it clear that AI, as it actually exists, is a fragile technology, which should face fundamental questions about its unexpected failure modes, its lack of explainability and its amplification of unwelcome cultural patterns. It explores the way AI's brittleness overwhelmingly causes harm to people who are already marginalized, and sets out the reasons why current remedies, from ethical principles to legal regulation, and from technical fixes to the human-in-the-loop, have little traction on constraining these harms. It highlights the way AI is sold as a solution to social problems, when what it is really doing is applying algorithmic morality judgements to target groups while obscuring the structural drivers of the very problems it is supposedly solving.

It would be troubling enough if AI was a technology being tested in the lab or applied in a few pioneering startups, but it already has huge institutional and cultural momentum. As we see in Chapter 3, AI derives a lot of its authority from its association with methods of scientific analysis, especially abstraction and reduction, an association which also fuels the hubris of some of its practitioners. The roll out of AI across swathes of industry doesn't so much lead to a loss of jobs as to an amplification of casualized and precarious work. Rather than being an apocalyptic technology, AI is more aptly characterized as a form of supercharged bureaucracy that ramps up everyday cruelties, such as those in our systems of welfare. In general, according to Chapter 3, AI doesn't lead to a new dystopia ruled over by machines but an intensification of existing misery through speculative tendencies that echo those of finance capital. These tendencies are given a particular cutting edge by the way AI operates with and through race. AI is a form of computation that inherits concepts developed under colonialism and reproduces them as a form of race science. This is the payload of real AI under the status quo.

What we should also be examining, given the current state of global financial, epidemiological and ecological conditions, are the tendencies enabled by AI in times of crisis, and this

is the focus of Chapter 4. The latest wave of AI has come to prominence in the period following the 2008 financial crash, and its ability to optimize rationing at scale readily fits in with austerity policies based on scarcity. Chapter 4 focuses on the way AI enables the kinds of exclusions that appeal all too easily to carceral states and security regimes. The polarization of outcomes under COVID-19, with their echoes of eugenics, flags up the way a crisis can rationalize the disposability of some for the good of the remainder, and we should be attentive to the ways algorithmic ranking can play a part in that.

Chapter 4 is a call to action regarding the potential of AI under crisis and the way the pseudo-rational ideology of artificial intelligence, with its racist and supremacist undertones, makes it an attractive prospect for the already existing authoritarian and fascist tendencies in political movements around the world. Given this, a shift to resisting AI is not only necessary but urgent. As we look forward with trepidation to the consequences of the climate crisis, with the likelihood that privilege will be defended, responsibility deflected and the vulnerable sacrificed, our priority for advanced technologies like AI should be to ask not only how they can be prevented from intensifying harm but how we can reassert the primacy of the common good.

Anti-fascist approach

At this point, we need to clarify why we're also talking about an anti-fascist approach to AI. In part, it's because fascism never really went away, something that's clearer every day with the rise of fascist-influenced political parties in so many countries. Given the real existing threat of fascist and authoritarian politics, we should be especially wary of any emerging technology of control that might end up being deployed by such regimes. But the main reasons for having an anti-fascist approach to AI run deeper into the nature of the technology itself and its approach to the world. It's not just about the possibility of AI being used by authoritarian regimes but about the resonances between AI's operations and the underlying conditions that give rise to those regimes. In particular, it's about the resonances between AI and the emergence of fascistic solutions to social problems.

To be clear, this book doesn't claim some deterministic link between AI and fascism: it's not saying that AI is fascist. However, what brings an instance of fascism into play as a historical force is a confluence of various factors, and it's in relation to these precursor currents that the character of AI becomes especially relevant. The conditions that need to be present for fascism to become a serious force are both ideological and opportunistic; the ideas have to be present but so do the particular kinds of crises that cause those ideas to look like a solution (Malm and The Zetkin Collective, 2021). AI's potential contribution is as a vector for normalizing specific kinds of responses to social instabilities.

Being alert to this possibility means having some idea about fascist ideology and the conditions under which it tends to thrive. In terms of ideology, we can refer to a widely used, if somewhat condensed, summary of fascism that describes it as 'palingenetic ultranationalism' (Griffin, 1993). These two words distill the ideology into features that are constant over time, and helps us to avoid getting diverted into looking for exact repeats of fascist rhetoric from the 1930s. The palingenetic bit simply means national rebirth; that the nation needs to be reborn from some kind of current decadence and reclaim its glorious past, a process which will inevitably be violent. The term ultranationalism indicates that we're not talking about a nation defined by citizenship but by organic membership of an ethnic community. Hence, with AI, we should be watchful for functionality that contributes to violent separations of 'us and them', especially those that seem to essentialize differences.

In terms of the political and social conditions, what is required to trigger a turn to fascism is a deep social crisis of some kind. The extremist ideas of fascism only start to have mass appeal when there's a sense of existential risk. For a crisis to be 'fascism-inducing' or 'fascism-producing' (Eley, 2016, cited in Malm and The Zetkin Collective, 2021) it has to appear to be beyond the capacity of traditional systems to solve. But this is only one side of the equation; the other is the decision of the dominant social class to invoke fascistic forces as a way to preserve their existing power. Historical fascisms have never actually come about through revolution but by the decision of the existing elites that they needed it as a prop for a collapsing hegemony

(Paxton, 2005). So, as far as AI is concerned, we need to be aware of both dynamics: the forms of crisis under which AI emerges and for which it is seen as a potential solution, and the aspirations of elites to use AI as a way to maintain existing political and cultural privilege.

So, the starting point for an anti-fascist approach to AI is an alertness to its operation as a technology of division, to its promotion as a solution for social crisis, and to its use to prop up power and privilege. The argument is not that the only problem with AI is the potential to enable fascist or authoritarian politics; there are many immediately harmful aspects of AI, as we shall explore in the coming chapters. But it is warning of fascism as a political possibility that shouldn't be ignored, and an assertion that any tendency to facilitate a shift in that direction should help to shape our response to AI as a whole. An anti-fascist approach is not simply one that opposes fascist tendencies but one that actively works towards structural alternatives to the conditions that give rise to the possibility of fascism in the first place.

In effect, AI acts as a kind of 'metapolitics', a term which some elements of the modern far right use for the process of shifting what's politically acceptable by shifting the culture that's upstream of it. Our concern with AI is not that it is fascist per se but that, because of its core operations, it lends itself to 'fascization', or solutions operating in the direction of fascism, and it is these that we need to be alert for as we go through the book. Likewise, having an anti-fascist approach to AI means being alert to these tendencies before they can bear fruit; it means countering any sign of such metapolitics by substituting in its place a project for a better society.

From machine learning to mutual aid

Having laid out, in Chapters 1 to 4, the reactionary politics of AI and the inability of reformist regulation to restrain it, we use Chapter 5 to scope out an alternative approach. AI's exclusions have roots going all the way down through our social structures and our ways of knowing. Fortunately, we don't have to invent a remedy for this from scratch because there are already

perspectives and practices that will help us to overcome these exclusions. In Chapter 5 we start with feminist standpoint theory, which undermines the absolutist form of scientific authority that AI tries to cloak itself with. Feminist and decolonial critiques of science can help change AI's approach to generating knowledge in ways that prioritize marginalized perspectives.

One of the fundamental positions set out in Chapter 5 is that boundaries are always constructed and what matters most is the forms of relationality that are at work in constructing those boundaries. One of the most toxic tendencies of socially applied AI is to naturalize and essentialize structural differences as part of an 'us and them' politics of inequality. Looking at AI from this different perspective allows us to understand it as an apparatus that helps produce aspects of the world through the exclusions it sets up, and suggests ways that we can interrupt this through horizontal forms of intervention. Chapter 5 articulates a collective approach to problem solving so as to open up new possibilities beyond the predictions of AI, in particular by shifting the focus from statistical inference to mutual care.

Of course, it's all very well having an alternative ethics and epistemology but what we really need are ways to turn these into tactics. Chapter 6 asks what practices can enact an alternative AI, and what forms of organization we require. The chapter proposes that the social tactic that goes with an ethics of care is mutual aid, and that the action-oriented commitment accompanying it is solidarity. It argues that mutual aid and solidarity are the basis for opposing precarity and overturning AI-driven states of exception. It looks at the stirrings of dissent within the AI industry itself among workers who already see how things are going wrong, and suggests self-organized worker's councils as a way to generalize a transformation from within. It extends this approach beyond the workplace through the model of the people's council as a form of constituent counter-power, one that assembles workers and communities into social movements capable of interrupting AI and pushing for transformative change.

Understanding AI not as some futuristic tech that has appeared in the present, but as a product of historical social processes, allows us to learn lessons from history about how best to deal with it. In the same way that Chapter 5 uses critiques from

the history of science to challenge AI's claims to authority, the proposals for worker's and people's councils in Chapter 6 draw from a long historical pedigree of political struggle against injustice and authoritarianism. One of the historical struggles against top-down technological transformation that has particular lessons for AI is Luddism. Chapter 6 looks at the similarities between Luddite times and the present day, in relation to the combination of social crisis and new forms of automation, and recovers from Luddism a sense of militancy and a commitment to the common good.

Overall, it is argued in Chapter 7, these radical perspectives can be gathered under the rubric of an anti-fascist approach to AI. This is partly about the early recognition of the threat posed by AI and having the determination to tackle it directly, but it goes beyond refusal to become a reorientation towards alternatives. Acknowledging that the roots of the problem lie in the status quo means actively pushing for a better world, one in which, by refusing computational exclusions and states of exception, we can centre the practices of mutual care. Resisting AI is significantly about restructuring the conditions that give rise to AI.

Chapter 7 draws the book to a close by setting out some sustainable directions for our technical apparatuses. It draws on historical and contemporary movements, like socially useful production and solidarity economies, to illustrate the wider idea of structural renewal and its relevance to the question of AI. Of particular importance here are the ideas of the commons and commonality, both in terms of the desirability that our apparatuses should contribute to the common good, and in terms of the specific role that 'commoning' can play in the transformation of techno-social systems. *Resisting AI* helps to illuminate a way forward for tech under the conditions of the coming global crisis.