

We Are at An Extreme Point Where We Have to Go All in On What We Really Believe Education Should Be About

Dan McQuillan, Goldsmiths, University of London, UK, D.McQuillan@gold.ac.uk, Juliane Jarke, University of Graz, Austria, juliane.jarke@uni-graz.at, Teresa Cerratto Pargman, Stockholm University, Sweden, tessy@dsv.su.se

Please cite this interview as follows:

McQuillan, D., Jarke, J. & Pargman, T.C. (2023). We Are at an Extreme Point Where We Have to Go All in on What We Really Believe Education Should Be About. In: *Postdigital Science and Education*. <https://doi.org/10.1007/s42438-023-00433-5>

Keywords: artificial intelligence, resistance, fascism, education, ChatGPT, EdTech, postdigital, futures

Introduction

Dan McQuillan is a Lecturer in Creative and Social Computing at the Department of Computing at Goldsmiths University of London. He has a PhD in Experimental Particle Physics. Prior to academia, he worked as Amnesty International's Director of E-communications. His research focuses on the resonances between forms of computational operation and their specific social consequences, especially in relation to machine learning and Artificial Intelligence (AI). His latest book is *Resisting AI: An anti-fascist approach to artificial intelligence* (McQuillan 2022).

Juliane Jarke is a Professor of Digital Societies at the University of Graz, Austria. Her research attends to the increasing importance of digital data in education, the public sector, and for ageing populations. She has co-edited *The Datafication of Education* (Jarke and Breiter 2019), *New Perspectives in Critical Data Studies: The Ambivalences of Data Power* (Hepp, Jarke, and Kramp 2022), and *Algorithmic Regimes: Methods, Interactions and Politics* (forthcoming).

Teresa Cerratto Pargman is a Professor of Human-Computer Interaction with a focus on education and technology studies. Teresa's research seeks to contribute critically to the study of the increasing digitalization and datafication of the education sector, and reflect on how emerging digital technologies come to disrupt established educational practices while constituting new ones.

About the Conversation

In this interview, Juliane Jarke and Teresa Cerratto Pargman discuss the implications of Artificial Intelligence (AI) for the postdigital future(s) of education with Dan McQuillan. They start off with an introduction to the main ideas of Dan's recent book (McQuillan 2022) and ask why we need to resist AI in education. Dan argues that the answer to the question is partly based on our understanding of education, what it means to us and how we imagine its future. He points to the harmful effects of AI and the narratives that perpetuate and boost its use in education and beyond. In the last part of the interview, Dan considers ways of resisting AI in education and sketches alternative educational futures. The interview took place end of March 2023 through an online video conference system and was subsequently transcribed and edited.

Why Do We Need to Resist AI?

Juliane Jarke and Teresa Cerratto-Pargman (JJ and TCP): We'd like to start with a question about the book that you've recently published: *Resisting AI: An Antifascist Approach to Artificial Intelligence* (McQuillan 2022). Why do we need to resist AI and what makes AI fascist?

Dan McQuillan (DM): I want to take AI seriously. I want people to take AI seriously. There is a lot of misinformation and hype about what AI can do and can't do, so the book first explains what AI actually does. In the following chapters, I unpick and expand on that. Why exactly what AI does is entangled with social relations in ways that become increasingly harmful? Just looking at the way AI works, you would say it's harmful because it's quite brittle and it's pretty

unreliable. It breaks, and it tends to break in ways that have more harmful impacts on people who are already vulnerable.

Even as a tool, AI is already a problem. But then AI doesn't really come from outer space or the future. It started within society and that's where it operates. AI has a facility for being adopted by already existing technologies such as bureaucracies and institutions. In their hands AI becomes an instrument that intensifies some of the harmful and cruel things that already happen within those institutions to vulnerable people [see e.g., Spade (2015) on 'administrative violence' and Redman and Fletcher (2021) on 'violent bureaucracy']. Bureaucracies and public institutions add to AI's capacity to scale and accelerate those things. So that's all a bit grim.

I mention words like fascism because I think that AI has a particular direction and steers a certain way. It's coming into the world at a time when the existing systems are already experiencing a series of shocks and transformations or are in some ways degrading. The capacities of AI, again, already have tendencies towards exclusion and isolation and sort of designations of disposability that reflect their quite deep roots in the mathematics of original Victorian eugenics and the fetishization of what we call Artificial General Intelligence (AGI). All of this stuff works as a kind of chorus behind a system that is already seeking to sort of withdraw the social or healthcare benefits to the minimum requirements to survive, and to specifically allocate that withdrawal across different groups and demographics.

I feel that AI is helping to take things further towards these kinds of far-right solutions. I believe that the system is already grasping for a way to shore up its own failing and power (Harvey 2005), and bringing forth very much the eugenics that already underlie the system (e.g., Allen 2001; Benjamin 2019). We saw this during the Covid-19 pandemic: the assassination of certain parts of the population. The acceleration towards disposability is already, if you like, latent within the system. I think AI brings that out.

The first half of the book, or maybe even two thirds, talks about why AI is a problem. I try then to build back up again and say: Okay, if there's an urgency to resist AI because it essentially offers a kind of fascist solution which far too strongly resonates with what is at the same time already rising in our politics, then what do we do? In the last part of the book, I try to pull that around and ask: How can we resist AI in that kind of epic, more abolitionist sense? It's not just about saying no to AI. It's about saying: Where do we go with this? What alternatives could be possibly reconstructed and on what grounds?

JJ and TCP: Fascist movements, among others, have always been and are invested in redesigning education to serve their needs. To what extent does the use of AI in education support the causes of fascist movements?

DM: That's a really great question, which I haven't really dwelt on, but I immediately understood when you said it: How vital it is for these kinds of movements to shape the world understandings of young people from as early as they possibly can? I think the fascist threat of AI comes in from two directions. One is the more blatant and obvious, which is AI systems and their generalised powers being adopted by bad people. I'm sure we could imagine many ways in which AI could be explicitly mobilised to shape a sort of ideological content within an educational setting. The other side of the fascist face of AI is reflected in amplifying pre-existing inclinations, experiences, interactions, and even psychodynamics amongst people and their relations. I think that's where AI is very powerful because it's part of the apparatus that helps to shape subjectivities.

AI helps to produce us in a way — like every other technology or infrastructure. Our sense of self, our personhood, and our perspective on life, is shaped through interacting with these different systems. And what AI does, I think, is to give a very particular flavour to those interactions. In my book I cite Hannah Arendt (2006) and her ideas of thoughtlessness. Thoughtlessness manifests as the inability to critique instructions, the lack of reflection on consequences, and a commitment to the belief that the correct ordering is being carried out. It is the product of a certain kind of apparatus — of a certain arrangement of ways of knowing, cultural values, and institutional arrangements. This thoughtlessness is very much the kind of belief in AI: in its correct ordering of people without any sort of critical questioning.

Once you introduce AI across all forms of interaction in an educational space, whether it's knowledge finding, the kind of discourse that is supposed to happen between people in education, or the management and optimisation of education itself, then, I think, the things that are so very innate to the way AI works become a danger as they will come to dominate even more. They will shape the kinds of people that the system produces. And a system soaked in AI will be much more likely to 'produce' people who are less likely to put up a fight, if you like, to a system that asserts itself ideologically towards fascism in some way.

JJ and TCP: The way in which AI-based systems classify and categorise people becomes normalised. As children grow up, they come to believe that it's a normal thing to do in terms of how to interact with the world.

DM: Absolutely, and that's what I'm really concerned about: that political parties or actors which are close to government or in government get the hands on the levers of quite large technical systems that know a lot about pretty much everybody. Actually, my main concern is the normalisation that comes before that. If we take AI as a normalisation machine, what is it normalising?

The Educational Harms of AI

JJ and TCP: That brings us to our next question: What other harms may AI inflict in education?

DM: That's a pretty tough question, isn't it? At the moment, I'm spending half my time to organise meetings for colleagues and students to talk about ChatGPT and GPT4; what does it mean, how are students going to write their essays with chatbots, and so forth. The ways it seems to alarm educators in the immediate sense are interesting and I can understand the sudden concern because I'm a lecturer myself: I teach classes, I mark assignments. But, often, the conversations seem to completely skip what are far more profound ethical and political issues that come with the adoption of AI. The educational harms of AI are many, and we can see them hurtling towards us right now.

For instance, the devaluation of the idea of education, particularly critical education. The cutting-edge large language models (LLMs), such as ChatGPT - I would think of it like fostering anti-critical pedagogy. It is an attack on the idea of an education that educates people to question things. AI also puts educators in the hands of Big Tech. You know, very few organisations on the planet can afford to train a large language model. If we thought the centralisation of Google classroom is bad, that's nothing compared with the power asymmetries induced by embracing large language models.

The narrative that I have experienced in workshops that I've joined from other institutions is to try and pick up a sense of the public discourse that goes like this: 'ChatGPT has happened and it's the way forward and we should embrace it and we should learn to be creative about our assessments'. I'm not against being creative about assessments, but the idea of embracing ChatGPT seems incredibly toxic. I mean, do people not read how these things are made? They depend on eye watering – this is their own term – costs of carbon emissions. They are built on massive exploitative 'ghost labour'; crowdsourced and outsourced labour that follows the patterns of colonial relations (e.g., Bender et al. 2021).

Large language models are ethically radioactive and the idea of embracing them at the centre of education seems really unwise even before you get to an analysis of what it actually does or doesn't do to people's capacity to think for themselves or to embrace the sort of friction in writing something or to answer questions or to do research or anything else. It's a pretty intense time to be thinking about what large language models are going to do to education. What is the long plan? I suppose, for EdTech and the consultancies, education, as it is currently known, is just not really needed; it is basically 'not fit for purpose'. Education to them is really dispensable because that's not what the young people of today really need to know. What they need—weirdly—are the skills to be better at helping with AI. It's very circular. You know, what we need today are young people who can do AI so we can build a bigger economy and strategy around the AI so we can have more AI.

JJ and TCP: The discourse around LLMs and ChatGPT is very popular these days but then AI is also very prominent in terms of writing out risks of education. The ways in which learning analytics promise to predict students at risk and then ask educators to act upon, aligns very much with the ways in which individuals are categorised according to what's desirable and effective education.

DM: Right! Before large language models came along, one of the other smaller dissonances at the college where I work, was exactly about an automated attendance system that is sold on this basis. You know that it's going to have a kind of traffic light model of students at risk, e.g., visualising the risk of a student failing a course with red (high risk), yellow (average risk), and green (low risk).

AI systems are powerful because they are performative; they produce the things that they claim to measure. A troubled student is a great example of the performative power of AI. The measures that are taken to supposedly deal with a student at risk can become a kind of self-reinforcing feedback loop with potentially lifelong consequences. If we take the lid off the box and look at the alleged reasoning that's going on inside these systems, they are really quite awful, reductive, and error prone misreadings of the world. So they are very, very harmful!

The Fragility of Artificial Intelligence

JJ and TCP: In your book you also describe the fragility of AI, the way it manifests in practice, and the consequent problems with applying AI as a solution to social problems. More precisely, by fragility, you write in the book:

The statistical nature of machine learning means it assumes that the distribution of data on which the algorithm has been trained covers the spread of occurrences in the wider world. Any shift in the underlying distribution (new behaviours, unexpected events) can throw a spanner in the works. There have been cases where a self-driving car has collided with a tow truck at the side of the road because its training data didn't include a statistically significant representation of a tow truck. ... Any AI in the real world is going to be faced with unexpected examples, whether it's navigating the chaos of traffic or deciding on unique immigration applications. AI, it seems, is both powerful and fragile. (McQuillan 2022: 28)

How do you think this fragility of AI should be reflected in the current discourse about using large language models in education?

DM: I think that the large language model discussion is how we come to experience their fragility, which is kind of festered in large language models for which the industry term weirdly is 'hallucination' (see also further McQuillan 2023). And that's what the industry itself uses, the idea that they make stuff up because they're not doing language, they're doing an optimised imitation of language. So, they can do a peer review of a paper without having ever read it because they can fake it really well, according to the statistical model of peer use that they've already learned. So they're fragile.

AI is an interesting machine. It's doing something remarkable. However, whatever it's doing is so far from what the hype would have us believe that its doing. The danger is to overlook the fact that large language models are simply completely incapable of having any understanding of the world whatsoever. Unleashing a machine like that with sort of awesome scale and power, absolute lack of understanding, and an inevitability of making mistakes, is going to have the same kind of harmful effects than unleashing lots of self-driving cars that will crash into things.

And, I think there's another layer to it, which is kind of more sinister in a way. Even when it's clear and obvious that large language models get things wrong, they will harm people. These technologies will be pushed through, they will continue to be imposed, because they are seen as the only real answer to the 'problems' of the education system. These problems are nested within the problems of society, and they appear to be so huge and apparently insurmountable in the frame of reference that we are given, that there seems just no chance to refuse AI, even if we know that it's

got problems. That's what people say in these meetings about ChatGPT: 'It's here now. We have to learn to live with it. We have to adapt to it.' I think this is wrong. Of course, we can say no and we should say no to these systems in education, and by my reading, pretty much everywhere else as well. Otherwise, the errors will just be seen as necessary, collateral damage and we will have to spend a lot of our time making sure AI kind of works.

If we do not resist AI, we will have to make sure that whatever it is that AI is doing is kind of working because we've had to change our educational system so much to fit in around this dysfunctional, harmful machinery. This is particularly bitter for me: in order to make LLMs appear to work in the real world, everybody around them has to invest a huge amount of time into and into plugging the gaps.

AI, Feminisms and an Ethics of Care

JJ and TCP: In your book (McQuillan 2022) you discuss feminist concepts such as ethics of care. In our commentary to this Special Issue, we have made an argument about why we need a redefinition of what we understand innovation to be; an understanding of innovation that centres around care and community rather than dominance and competition (Macgilchrist et al. 2023). How could resistance to AI look like in education?

DM: The relations of care in education are actually the most underpinning aspect. No institution, no school or college, works without care in its broader sense. People supporting each other, looking out for each other as much as they can, within the constraints of a framework that can only exist because people are prepared to put in the unpaid or invisible time to keep it working, so that other people can sort of march in and deliver the visibly value adding bits of education or whatever it is.

I think, in the first instances, resisting AI in education is trying to look at what's already there, what's already existing. We have learned our lessons very quickly, during the pandemic. At least in the UK, the narrative was all about 'essential workers'. And the essential workers were all the people that nobody valued at all before the pandemic. The cleaners, the low paid nurses, the delivery people, or the people collecting the rubbish, who provide normally—in feminist terms—'invisible labour' that props up society. Their labour was kind of momentarily made visible and the hierarchies were narratively inverted and their work was seen to be 'essential'. Of course, we went back to normal and all of that is completely forgotten.

Perhaps a starting point for resistance then is the turning towards those relations that already exist within a space and that are already relations of care. But it kind of means in a way a tactical turn: you're explicitly turning towards those people who are already the most marginalized in some sense. One of the tools I tried to apply in the book as a way of inverting the sort of harmful effects of AI is standpoint theory (Harding 1998) or sort of 'situated knowledges' (Haraway 1988). The idea is that if AI is the thing that seems to be amplifying some of the most toxic or damaging things in society, then what is its opposite? What are the things that are left out of the abstractions? It's exactly the marginal viewpoints, the undervalued perspectives.

Again, what I'm concerned to try to do is to transform an ethics of care into a sort of politics or at least into some kind of counter power. I think this is related to my passion for the transformative potential of people getting together and saying: well, actually, we share a problem. The approach of a general assembly or a council of some kind and people simply getting together and saying: you know, we're the ones with a stake in this, and, no matter our differences and dissonances, we have to figure this out together and get some kind of a response out of our plurality. That's very much a counter, an inversion to the way AI deals with rendering everything commensurable, rendering everything tradable, in a universal uniform and an algorithmic way.

Alternative Futures

JJ and TCP: This is connected with what you mention in your book about the possible alternatives that we could construct when resisting AI. We wonder if you could say more about these alternatives. What are these alternatives you are speaking about?

DM: Where I got to in ‘Resisting AI’ was not the specifics as I was trying to map out—for people and also for myself—a sense of direction. I started out with some aspirations towards thinking about what an ‘AI for the people’ would look like, how could we use AI well, and I came out with the realisation that this would be impossible.

I think ours is a time when ideas of education have over decades and decades been gradually sort of paired down and salami-sliced in order to fit in with the overwhelming and powerful neoliberal ideologies that are about optimisation and productivity. We have so far always found a rear-guard action in a sense that we can carry on doing this, but we can still preserve something of our idea of education within this space that we're trying to keep. There is an ongoing resistance in education towards marketisation and neoliberal values and I think maybe AI marks the end game of that kind of resistance. In the face of AI, the sort of gradual retreat that we have practiced in the past is no longer a viable resistance tactic. Maybe we're at an extreme point where we have to sort of—it seems like the wrong phrase to use, but—put our money where our mouth is and go all in on what we really believe education should be about. And that isn't clearly possible in an immediate way in any existing school or university, the places where education happens.

I took part in the Anti-University festival in the UK a few months ago, which is exactly about trying to take learning outside of the constraints of disciplines, and extremely expensive and debt-ridden learning spaces, and make it accessible to everybody. And I think, maybe, through the self-organisation of people at every grassroots level, there's a possibility of building a basis for something really different. I'm not talking about a sort of long-term aspiration; I'm talking about trying to do educational things that fit in with what we really believe education can be about. Maybe it's possible to do some of that inside of school or university. Maybe it's only possible to do that on the picket line. Or maybe it's only possible to do that literally outside of the university. Maybe we have to find each other outside these spaces and start to build what we think is a real form of education, even if they're small and parallel for the time being, and face up to the idea that the system that we are currently working in simply needs to be replaced.

I think the idea of replacing the education system would find a lot of resonance in the general political moment. The neoliberal answer has completely run out of promise. It is dead on its feet, it's a zombie-world system. So if there's going to be any alternative to that, then it's a ground-up one. Schools in New York are running classes for kids to gain a critical perspective on AI. That seems like a great and immediate thing: to gain critical AI literacy. But I think, honestly, that chances of immediately reversing things in the educational establishments are pretty slim. My political understanding leads me to the idea of small concrete alternatives, different education experiences, that can possibly be scaled to federating rather than through re-imagining the top-down institutions.

References

- Allen, G. E. (2001). Is a New Eugenics Afoot? *Science*, 294(5540), 59-61. <https://doi.org/10.1126/science.1066325>.
- Arendt, H. (2006). *Eichmann in Jerusalem: A Report on the Banality of Evil*. London: Penguin Classics.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? . In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 610-623). New York: Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>.
- Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge, UK and Medford, MA: Polity.
- Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3), 575-599.
- Harding, S. (1998). *Is Science Multicultural?: Postcolonialisms, Feminisms, and Epistemologies*. Bloomington, IN: Indiana University Press.
- Harvey, D. (2005). *A Brief History of Neoliberalism*. Open University Press.

- Hepp, A., Jarke, J., & Kramp, L. (2022). *New Perspectives in Critical Data Studies: The Ambivalences of Data Power*. Cham: Palgrave MacMillan. <https://doi.org/10.1007/978-3-030-96180-0>.
- Jarke, J., & Breiter, A. (Eds.). (2019). *The Datafication of Education*. London: Routledge. <https://doi.org/10.4324/9780429341359>.
- Macgilchrist, F., Allert, H., Cerratto Pargman, T., & Jarke, J. (2023). Designing Postdigital Futures: Which Designs? Whose Futures? *Postdigital Science and Education*. <https://doi.org/10.1007/s42438-022-00389-y>.
- McQuillan, D. (2023). Predicted benefits, proven harms. How AI's algorithmic violence emerged from our own social matrix. *The Sociological Review*. <https://thesociologicalreview.org/magazine/june-2023/artificial-intelligence/predicted-benefits-proven-harms/>
- McQuillan, D. (2022). *Resisting AI: An Anti-fascist Approach to Artificial Intelligence*. Bristol: Bristol University Press.
- Redman, J., & Fletcher, D.R. (2021). Violent Bureaucracy: A Critical Analysis of the British Public Employment Service. *Critical Social Policy*, 42(2), 306-326. <https://doi.org/10.1177/02610183211001766>.
- Spade, D. (2015). *Normal Life: Administrative Violence, Critical Trans Politics and the Limits of Law*. Durham, NC: Duke University Press.