

Jeremy H. Kidwell University of Birmingham, Lecturer in Theological Ethics

Presentation for Christian Academic Network 13:45 10 March 2018 Birbeck University

Introduction

It is often the case with any mildly philosophical discussion of technology that we can tend to focus on the exotic, futuristic, remote - and there's nothing wrong with this, but I have tended to find as an ethicist that we think differently about matters which are strictly abstract than those possibly banal things which are embedded in our ordinary daily lives. So we prefer to talk about nuclear weapons rather than landmines or handguns. Don't get me wrong here - there are very clearly Christian and theological reasons to focus on "the margins" but it is worth separating out ways of thinking which are shaped from matters of justice from those which are shaped by matters of curiosity.

I've generally found, as an ethicist, that it is better to focus our attention on those banal domestic things - that our reflection is more embedded, potentially a bit more pragmatic, and also more visceral. Dare I say it is more incarnational? In any case, I want to bring this approach to ethics to our topic for today, and in so doing, focus not on androids (though we aren't far from having these in our homes) or unmanned aerial vehicles, but on the algorithms which we engage with every single day. And as people like Sherry Turkle remind us - this way beyond "frequent" - our lives are saturated by engagement with algorithmic platforms. The average person touches the screen of their device more than 200 times a day.

So let us dive for a few moments into the ethical questions which arise with regards to our social uses of machine-learning algorithms. My ultimate goal here is to try and foreground some ways that Christian Theology might also be brought into this discussion.

Orientation - Domestic Algorithms in 2018

First, let's discuss a bit about where algorithms currently intersect our daily lives:

The most ubiquitous algorithms are probably those which parse out information from Google searches and those which drive our Facebook / Instagram feeds.

Google's "[PageRank](#)" algorithm has about 4.5 billion active users. Similarly, Facebook has written an algorithm they call "[Edge Rank](#)" which organises, sifts and filters, users news feeds. Upwards of 80% of time users spend on the internet may be filtered in some way by Facebook algorithms, whether on facebook or on one of the other platforms they've acquired, such as instagram. It's interesting to note that not all social media platforms use algorithms - twitter tried very briefly to sift and sort their users twitter feeds and there was such an outcry that the reversed the action. So what you get is an unmediated and unfiltered experience. Counter that with facebook - each user's newsfeed and instagram stream are not presented chronologically, some things "rise" to the top and others are suppressed.

These are not the only two contexts where algorithms are mediating our experience of the world, and it is important to note that as "big data" becomes the vogue, various arms of government and public services are increasingly turning to AI, that is, machine

learning algorithms, to make their work more efficient - particularly education and police, but also local authorities and other front-lines managers of public infrastructure. The other area where we implicitly feel algorithms in due to their increasingly significant place in the financial sector. The most publicly visible sign of this is the periodic occurrence of “flash crashes” like the crash on May 6, 2010 which has now been extensively studied by finance and computer science scholars. This crash provide a visceral revelation of the risks inherent in digitized financial markets and high-frequency algorithmically driven trading activities when US markets lost (and mostly recovered) more than a trillion dollars of value in the space of a half hour. Subsequent analysis has led to enhanced regulation, but there have been subsequent high-profile flash crashes in the USA, Britain and Switzerland.

Because the word “algorithm” may evoke in our imaginations the idea of a relatively simple mathematical equation, it is important to emphasise the complexity of these algorithms, they are absolutely AI. This is underlined by the fact that facebook stopped using the EdgeRank system in 2013 in favor of a machine learning algorithm that, as they suggest, takes into account more than 100,000 individual factors. I'm not going to break down the specifics of how these two algorithms work, but there are a few things that are useful to acknowledge at this stage in our reflection on AI.

Constructive Technology Criticism

First lets set some ground rules. I will borrow here from Sara Watson's terrific piece on Constructive Technology Criticism, where she sets some ground rules for journalists talking about technology, but they apply here as well.

First - lets avoid “moral panics,” those situations where people just assume that some new technology is evil because it is new and unexpected. As I am sure you all know, Christians were very uncomfortable at first about television, and radio before that, musical instruments, on on back it goes.

Second - let us be wise in seeing through “progress narratives” - same thing applies here. The bright shiny new thing can seem like the antidote to all the worlds problems. Other Christians were very excited about what television would do for evangelism, and radio, and the printing press, and so on.

The point here is to pause, take a deep breath and take a clear eyed look at what is going on in a critical, that is, careful way.

Finally, and perhaps most importantly, let's not blame the technology. As Sara observes: “it is people who design and build technology”.

Christian responses

Now, let us reflect on the matter of machine learning algorithms in everyday life.

It is tempting to take all those “ground rules” as given, reflect on the apparent banality of something like facebook and just assume that there really isn't anything for us to talk about. It's just a bit of fun, right?

This is where the Christian critic of technology ought to think of themselves as doing prophetic work. Here, I am thinking of Jeremiah, when he is asked in Ch. 7 -

Stand in the gate of the LORD'S house, and proclaim there this word, and say, Hear the word of the LORD, all you people of Judah, you that enter these gates to worship the LORD. Thus says the LORD of hosts, the God of Israel: Amend your ways and your doings, and let me dwell with you in this place. Do not trust in these deceptive words: "This is the temple of the LORD, the temple of the LORD, the temple of the LORD"

Jeremiah is written a bit like a monologue - so we do not get to hear much about how his message got across, except implicitly. But I like to try and imagine what it might have been like for him - like when he shared this message? Do you think people were receptive? I doubt it. I doubt they were even angry, they probably just had a laugh and went on about their everyday work in the temple, doing everyday things, that in their familiarity seemed pretty harmless.

The task of the prophet can go in two ways - but the underlying activity is always the same - their task is apocalyptic, revelatory. That word apokalypsis, is about uncovering, disclosure, and indeed much of Jeremiah is like this - God says, look, here is how things actually are, even though they seem to be fine. They are very, very, bad. In other cases, this can be about revealing how God is doing good work in the midst of circumstances that seem hopelessly bleak. This work of revealing, can have both positive and negative aspects.

My point here, is that as Christians we are called to take up this apocalyptic task - which might otherwise be called critique, or as Oliver O'Donovan puts it exercising "judgement".

And if we look more closely, what we find in our query about algorithms is that there are some very nasty things going on here.

One thing that quickly surfaces in critical study of these domesticated forms of AI is that our algorithms are very, very racist.

Some examples:

- <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- <https://m.huffpost.com/us/entry/970451>
- https://motherboard.vice.com/en_us/article/mg7g3y/how-to-make-a-not-racist-bot
- https://www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms#t-145078

It may be tempting to respond to these discoveries as technicians: "ah, that just means they need to write better code, there is some fine tuning needed there". But this obscures an underlying dynamic - and also begins to explain why this issue of harmful and deeply consequential discrimination has persisted in this space for a very long time. As I noted above, people make code.

As Sarah Rich observes:

all technological artifacts contain certain "prescriptions" within them...
designers can inscribe intentions into the things they build
(<https://www.theatlantic.com/technology/archive/2013/07/the-machine-zone-this-is-where-you-go-when-you-just-cant-stop-looking-at-pictures-on->

[facebook/278185/](https://www.facebook.com/278185/))

I am not suggesting that people who write algorithms are more racist than other people. The reality of the situation is that people who write algorithms are mostly white. And when you combine something that acts in subtle implicit ways like white privilege with technology that is deliberately engineered to work quietly, behind the scenes, what you have are some very subtly disguised forms of racism. I know this may be hard to accept, as it is a phenomena that is massively hidden. But this is that hard prophetic work. God loves all his children, and grieves when we marginalise and harm those who are already more vulnerable.

But there is some ethical advice for the engineers as well here. Algorithms are becoming increasingly popular for two reasons:

(1) They are tremendously powerful - enabling a person to parse through data, respond to customers, provide a service, etc. more quickly than it would otherwise be possible. Algorithms enable economies of scale which are otherwise not possible and they do this by amplifying the power of ordinary humans. As Lynn White argues, power enhancement and amplification is a fundamental feature of technology, and this is the same for waterwheels, mills, and machine guns. The ethical and prophetic response to any technology is thus to ask, "is this a situation which is morally stable enough to be worthy of such amplification?"

To put it another way, do our toilet seats really need to be connected to the internet of things?

I think that Christians have a special obligation to think carefully about the way that they inhabit power relationships. Definitely be empowered, as Paul suggests in Eph 6:13, "take up the whole armour of God" but bear in mind that this is so that - as Jesus points out twice in the gospel of Matthew - we may take up a cross.

Matt. 10:38 and whoever does not take up the cross and follow me is not worthy of me.
Matt. 16:24 Then Jesus told his disciples, "If any want to become my followers, let them deny themselves and take up their cross and follow me.

So algorithms, like all technologies, are about power and thus it is worth exerting caution before we deploy them into aspects of our domestic daily lives or in contexts where crucial decisions about welfare of others are at stake. This is not because the "tech" is immature, but because we are and the tech is so powerful.

But there is another aspect at work here, which I have already alluded to above, which is that algorithms are a bit like snails in a fishtank, in that they are very good at hiding. You only realise that you have a completely transformed ecology when you hit the lights in the morning and there they are, hundreds of them! This is exacerbated by the fact that technological development is increasingly being done out of sight, in deliberately obscured locations, often concealed by legal non-disclosure agreements.

A group of AI specialists called "AI now" released a report last year with a range of very clear-eyed ethical guidance for public agencies seeking to make use of algorithms. They suggest that:

Core public agencies, such as those responsible for criminal justice, healthcare, welfare, and education (e.g "high stakes" domains) should no

longer use “black box” AI and algorithmic systems.

I will close by suggesting that Christians should be championing such proposals, not simply because we care about our privacy (though this is a reasonable concern), or because we are anxious about the competency of our public officials, but because we are actively concerned about injustice and inequity and there is very good reason to assume that these things which are hidden from view will be made known in one of two ways:

By the revealing hand of God bringing to light very bad things which are happening on our watch, for which we will be subject to judgement.

Or in surfacing the concealed ways that God is working (through us I hope) to bring about God's kingdom, one which is known not by convenience and consumerism, but by justice and mercy.

Some further reading:

<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

<https://www.theatlantic.com/technology/archive/2013/07/the-machine-zone-this-is-where-you-go-when-you-just-cant-stop-looking-at-pictures-on-facebook/278185/>

<https://medium.com/@AINowInstitute/algorithmic-impact-assessments-toward-accountable-automation-in-public-agencies-bd9856e6fdde>

<https://marketingland.com/edgerank-is-dead-facebooks-news-feed-algorithm-now-has-close-to-100k-weight-factors-55908>

Other reading:

<https://www.howtogeek.com/290919/how-facebooks-news-feed-sorting-algorithm-works/>