Policing the internet for extremist material

Josh Cowls discusses the Oxford Internet Institute’s report on the complexities of balancing security and privacy online.

The internet serves not only as a breeding ground for extremism, but also offers myriad data streams which potentially hold great value to law enforcement. The report by the Oxford Internet Institute’s (OII) Ian Brown and Josh Cowls for the VOX-Pol project Check the Web: Assessing the Ethics and Politics of Policing the Internet for Extremist Material explores the complexities of policing the web for extremist material, and the implications for security, privacy and human rights. Josh Cowls discusses the report with blog editor Bertie Vidgen. The views given here do not necessarily reflect the content of the report, or those of the lead author, Ian Brown.

Interviewer: Josh, could you let us know the purpose of the report, outline some of the key findings, and tell us how you went about researching the topic?

Josh: Sure. In the report we take a step back from the ground-level question of “what are the police doing?” and instead ask, “what are the ethical and political boundaries, rationale and justifications for policing the web for these kinds of activity?” We used an international human rights framework as an ethical and legal basis to understand what is being done. We also tried to further the debate by clarifying a few things: what has already been done by law enforcement, and, really crucially, what the perspectives are of all those involved, including lawmakers, law enforcers, technology companies, academia and many others.

We derived the insights in the report from a series of workshops, one of which was held as part of the EU-funded VOX-Pol network. The workshops involved participants who were quite high up in law enforcement, the intelligence agencies, the technology industry, civil society and academia. We followed these up with interviews with other individuals in similar positions and conducted background policy research.

Interviewer: You highlight that many extremist groups (such as Islamic State) are making really significant use of online platforms to organize, radicalize people, and communicate their messages.

Josh: Absolutely. A large part of our initial interest when writing the report lay in finding out more about the role of the internet in facilitating the organisation, coordination, recruitment and inspiration of violent extremism. The impact of this has been felt in Paris and Beirut in 2015, and many other places worldwide. This report pre-dates these specific developments but was written in the context of these sorts of events.

Given the internet is so embedded in our social lives, I think it would have been surprising if political extremist activity hadn’t gone online as well. Of course, the internet is a very powerful tool and in the wrong hands it can be a very destructive force. But other research, separate from this report, has found that the internet is not usually people’s first point of contact with extremism: more often than not that actually happens offline through people you know in the wider world. Nonetheless it can definitely serve as an incubator of extremism and can serve to inspire further attacks.

Interviewer: In the report you identify different groups in society that are affected by, and affecting, issues of extremism, privacy, and governance – including civil society, academics, large corporations and governments.

Josh: Yes, in the later stages of the report we do divide society into these groups and offer some perspectives on what they do, and what they think about counter-extremism. For example, in terms of counter-speech there are different roles for government, civil society and industry. There is this idea that Islamic State are really good at social media and that is how they are powering a lot of their support; but one of the people that we spoke to said that it is not the case that they are really good, it’ is just that governments are really bad!

We shouldn't ask government to participate in the social network: bureaucracies often struggle to be really flexible and nimble players on social media. In contrast, civil society groups tend to be more engaged with communities and know how to "speak the language" of those who might be vulnerable to radicalisation. As such they can enter that dialogue in a much more informed and effective way.

The other tension, or paradigm, that we offer in this report is the distinction between whether people are "at risk" or "a risk". What we try to point out is that people can go from one to the other. They start by being "at risk" of radicalisation, but if they do get radicalised and become a violent threat to society, which only happens in the minority of cases, then they become "a risk". Engaging with people who are "at risk" highlights the importance of having respect and dialogue with communities that are often the first to be lambasted when things go wrong, but which seldom get all the help they need, or the credit when they get it right. We argue that civil society is particularly suited to being part of this process.

Ed: It seems like the things that people do or say online can only really be understood in terms of the context. But often we don't have enough information, and it can be very hard to just look at something and say "This is definitely extremist material that is going to incite someone to commit terrorist or violent acts".

Interviewer: Yes, I think you're right. In the report we try to take what is a very complicated concept – extremist material – and divide it into more manageable chunks of meaning. We talk about three hierarchical levels. The degree of legal consensus over whether content should be banned decreases as it gets less extreme. The first level we identified was straight up provocation and hate speech. Hate speech legislation has been part of the law for a long time. You can't incite racial hatred, you can't incite people to crimes, and you can't promote terrorism. Most countries in Europe have laws against these things.

The second level is the glorification and justification of terrorism. This is usually more post-hoc as by definition if you are glorifying something it has already happened. You may well be inspiring future actions, but that relationship between the act of violence and the speech act is different than with provocation. Nevertheless, some countries, such as Spain and France, have pushed hard on criminalising this. The third level is non-violent extremist material. This is the most contentious level, as there is very little consensus about what types of material should be called "extremist" even though they are non-violent. One of the interviewees that we spoke to said that often it is hard to distinguish between someone who is just being friendly and someone who is really trying to persuade or groom someone to go to Syria. It is really hard to put this into a legal framework with the level of clarity that the law demands.

There is a proportionality question here. When should something be considered specifically illegal? And, then, if an illegal act has been committed what should the appropriate response be? This is bound to be very different in different situations.

Interviewer: Do you think that there are any immediate or practical steps that governments can take to improve the current situation? And do you think that there any ethical concerns which are not being paid sufficient attention?

Josh: In the report we raised a few concerns about existing government responses. There are lots of things beside privacy that could be seen as fundamental human rights and that are being encroached upon. Freedom of association and assembly is a really interesting one. We might not have the same reverence for a Facebook event plan or discussion group as we would a protest in a town hall, but of course they are fundamentally pretty similar.

The wider danger here is the issue of mission creep. Once you have systems in place that can do potentially very powerful analytical investigatory things then there is a risk that we could just keep extending them. If something can help us fight terrorism then should we use it to fight drug trafficking and violent crime more generally? It feels to me like there is a technical-military-industrial complex mentality in government where if you build the systems then you just want to use them. In the same way that CCTV cameras record you irrespective of whether or not you commit a violent crime or shoplift, we need to ask whether the same panoptical systems of surveillance should be extended to the internet. Now, to a large extent they are already there. But what should we train the torchlight on next?

This takes us back to the importance of having necessary, proportionate and independently authorised processes. When you drill down into how the right to privacy should be balanced with security then it gets really complicated. But the basic process-driven things that we identified in the report are far simpler: if we accept that governments have the right to take certain actions in the name of security, then, no matter how important or life-saving those actions are, there are still protocols that governments must follow. We really wanted to infuse these issues into the debate through the report.

Josh Cowls *researches the impact of technology on politics, communication and the media at the Massachusetts Institute of Technology.*

The interview was first published on the Oxford Internet Institute's Policy and Internet Blog here*. This article gives the views of the authors, and not the position of the Policy and Internet Blog, nor of the Oxford Internet Institute.*

Published on:March 25, 2016