Reflecting on Engagement: Malware and a Social Scientist.

Andrew Dwyer – CDT Yearbook 2019

It was a freezing January morning in Abingdon, 10km south of Oxford, when I entered Sophos' headquarters. A photograph was hastily snapped for my ID badge by the reception staff and I was escorted to its malware analysis laboratory, SophosLabs. There I met several analysts, shown to my computer and desk, and an analyst whizzed through set-up. A comment from my first hour has stuck with me: "you know how to do everything on the command line, right?" Instead of saying no, I spent the remainder of the day baffled. I felt completely out-of-depth.

My thesis explores the spaces of malware; how it interacts on computing and social environments, and its impact on geopolitics and international relations. For seven months between January and July 2017 I used an underutilised methodological technique, (auto)ethnography, that blends self-reflection and learning with observations, taking-part, and conducting interviews. In this (reflective) piece, I outline an interdisciplinary approach to the study of malware through imbrication of a social scientist in a malware analysis laboratory. I first explore how its germination in the CDT broadened the horizons of a geographer, offer insights into some my emergent thinking, the engagement with an *industrial* partner, my research method on 'becoming-analyst' in SophosLabs, and conclude with some thoughts on this experience.

Germination

I joined the CDT in 2014 with no extensive experience, in-depth knowledge, or training in cybersecurity. I arrived after studying (human) geography at Durham University. Unlike many geography programmes however, I did an unconventional combination of in-depth theory, geopolitics, security and the study of emergencies. This set me up with a broad appreciation of the interface between data, computer science, and security that I sometimes still refer to for inspiration. The initial year of CDT training (ranging from more comfortable

options in ethics and international relations, to the less comfortable study of high integrity systems and security architectures) pushed me to explore different directions and consider something 'technical.' Without this, I doubt I would have engaged in research on malware at all. I cannot overstate the importance of this initial component of the DPhil. This provided a bridge, at least for me, to traverse what I felt was insurmountable. I am sure I would have found this cross-disciplinary engagement at few other places. This gave me the confidence, essential knowledge, and tools to at least consider developing a project that co-constituted two very different worlds.

The possibility to open-up entering and working with different disciplines is most beneficial not only through scheduled teaching, but through chatting and conversing with your cohort. Though I had not intended to return to geography it has provided a foundation in which to build a distinctive project. This blends elements of computer science, international relations, media archaeology, geography, and anthropology together whilst invoking the critical insights that criss-cross science and technology studies. In an initial meeting, my supervisor and I arrived at malicious software as something interesting to pursue. Primarily to critique current directions of travel, on the distinction between *good* and *bad* software, how this binary becomes operationalised in the media, who determines what is malicious or not, the impact that this has on how we talk about cyber-attacks, and how it constructs politics and international relations. It may appear odd to dwell on these initial aspects, but it explains how I subsequently engaged and why drawing together people from a variety of different social science positions (not social science as a category on its own!) can be beneficial to the study of cybersecurity.

Engagement with a Malware Analysis Laboratory

Malware Analysis Laboratories (MALs) are some of the most understudied spaces in cybersecurity today; perhaps because they are seen to 'work' and are not directly related to what are often the more immediate concern of security operation centres (SOCs) for governments and corporations. However, the endpoint security industry is a multi-billion-dollar market that provide security protection to a large number of the world's population and industries. Yet, it is not only how many the endpoint security industry cover, but that MALs have become the private enterprises that contour and define the vast amount of what is considered malicious. I in no way claim that this is the *only* site where software is designated as malicious, but that it is a crucial site for most - as the detections (whether static signatures or advanced behavioural techniques) are used in a variety of other security spaces. Therefore how, who, and what is considered malicious is condensed within MALs. They define, within societal discourse, what is malicious in one context *or ecology*, and in another not. This is where I think geography and cybersecurity can have a productive conversation - where environments, societal context, as well as political discourse shape security response.

As an example of the benefit of working in the MAL, I was able to determine that malware, does not exist *a priori*. This may seem an odd statement to make where there are clear malicious actors who wish to cause damage and *harm*. In claiming that malware does not exist before certain moments in time and space, I do not attempt to remove the important ethical and criminal actions that certain individuals, groups, and states take. Malware, without human expectation, is ultimately software - there is emotion to it, no badness. There may be certain software that I would define as parasitical - that it attempts to extract, monitor, reroute, and sometimes take over other pieces of software (which is different to how computer science would define it). However, that is not malicious. It is only when it coalesces with certain environments that it becomes as such. Therefore, certain materialities need to exist; whether that be a certain operating system patch, an internet connection, a command and control server, a human intervention, among others, to enable software to *emerge as malicious*. I perhaps take a slightly different approach to others - but one that I think follows what Lockheed Martin do in their *Cyber Kill Chain*®[i]- where the stages represent an interpretation that malware does not exist as an object moving around with maliciousness embedded. It can be 'killed' through preventing connection, 'killing' maliciousness. Hence, the role of the MAL becomes ever-trickier. It must try and assess both the multiple environments and the software that it operates in. Many of the apparent *mal*ware that I encountered in the MAL failed to execute or failed due to implementation errors. Clearly even considering something as *failing* shows how dominant human expectation is! Therefore, the work of the MAL is the attempt to define the contours of what the abnormal (and by extension, malicious) attributes of software are. How these combine in certain temporalities and environments is what allows malware to be malware. This involves selection and decision about what software should be allowed and what should not.

It is not only the power to determine what the vast majority of software *are* that makes the MAL so important. Nor is it only whether software is malicious or not; but how it becomes communicated or *curated* through blogs, with journalists, through conversations with governments, and providing crucial technical (and sometimes geopolitical) analysis. As one of the better studies of malware and international relations outlines – these reports provide some of the best analyses that feed into these debates[ii]. This combination as both the broker of information and its technical expertise gives it immense (if not particularly directed) political power. Though there is broad agreement in society over what is deemed unacceptable behaviour - such as the stealing of credentials, holding a computer to ransom, or attacking critical national infrastructure, there are many *grey* areas. This includes advertising, monitoring, and grey hat tools that are used by pen testers. It is in these spaces that the fallacy of malware as some sort of *malicious* object that exists *a priori* falls apart. Not only in the spatial distribution that is required through certain ecological alignments that allow something to happen, but in human expectation. There is a double tactic at work here - between first, whether a piece of software

operates *as the author intended* (through certain alignments), and second whether these should be deemed as malicious at all.

You may ask why I have even considered what may seem to be either a) pointless, or b) overly fluffy. However, it is precisely what I developed above that necessitated a tight engagement with industry. Social science cannot develop the in-depth technological knowledge from afar, or solely through interviews. This is why Sophos was such an important engagement for me to not only learn about malware analysis and detection, but to also be able to research it at the same time. Though I may have been referred to as an *intern*, it is difficult sometimes for social scientists to offer something '*productive*' for an industrial partner. Though it is something beneficial to assist with the design of a product, generate better human-computer interactions, or develop a policy or strategy; this, importantly, is not the sole aim of social science. Critical work, that is able to take a step back and reflect, is essential for a disciplinary acquis to develop. Some may position this as a service to a body of knowledge. I am not naïve enough to think this will convince many industrial partners; after all, commercial concerns are often too strong for this. So, I offered myself as an intern of sorts to train-up (the value for me being clear), and then to be a productive member of the Sophos team, generating a different form of value through the detections I wrote. Indeed, the training and subsequent time 'working' for Sophos allowed for a social scientist to engage properly with the technical material of the MAL whilst hopefully being able to give back to Sophos in a way they desired. This is not always, and should not, be possible – it depends on each project. Industrial engagement can be useful for both sides; but it is tricky to negotiate for a social scientist who wishes to maintain a critical distance to talk of what they have experienced. I have great freedom in what I can say – I would not have engaged otherwise. I am not sure whether there is a resolution (or whether one is beneficial due to the potentially richer projects that emerge from this friction).

In light of this tension, my research engaged with ethnography, that draws on initial insights from anthropology. My interpretation of a broad brand of ethnographic research is (auto)ethnography; that broadly focuses on self-reflection and immersion of oneself into a 'field.' This allowed me to engage seriously with malware analysis whilst being reflective of broader issues; on how anti-virus works, on the impact of malware, and the politics that emanates from these spaces. My in-depth research over these seven months in the Sophos MAL has provided insight in some theoretical discussions highlighted above that blend with my direct experience; such as when something becomes malicious and under what conditions, how malware and its interpretation becomes disseminated outside of the MAL, and on such questions on malware as a tool and on the *cyberweapon*. Yes, the daily grind of commuting an hour each way, 5-days a week, along with writing my research diary to keeping up with the requirements of a DPhil were draining. There were moments when I wanted to quit. Though I can now reflect on my whole time with a glow of appreciation, and in fact, the frequent enjoyment I experienced as I grew more confident and comfortable – I do not wish this to blind the (literally) dark early days. Getting to grips with the intricacies of technologies is not easy for a social scientist – and I have great sympathy for my colleagues who often receive snide comments that they don't 'get' the technology. It is not easy – and sometimes not worth it. As with any serious engagement the 'other way' to social science, it takes a sinking in to material which is hard – interdisciplinary work does not come without its sacrifices.

Emergence

In this section I outline my experience in the MAL, where malware is not simply a technical product, an object that moves through space undisturbed, as a direct extension of the human hacker. It emerges in context. This is because software is performative – that is, it unfurls (think of its assembly through a processor), at particular space-times – meaning its interactions cannot be conditioned or predetermined completely before they occur. Now, this does not mean that there are no structures, or clear mathematical logics, that underpin computing. However, these are only parameters of engagement – where there is a variety of activity possible within those parameters. In this, I follow the work of Luciana Parisi[iii] who argues that environments are productive spaces of encounter, where there is a more-than-human agency that produces

what we would term as maliciousness. Humans who are not the only ones who are important. The 'more-than' element respects that there are more actors in cybersecurity than humans; they are part of the story but not the total extent, as we all know.

In my experience it was clear there is little sense of a wholly *rational* reverse-engineer, or malware analyst:

> "It's based on experience, intuition, gut feel, risk. So, blend all those things and you make a decision about something… things like machine learning can really help because we base… We base our decisions on the data we have available. As I said, 10 years ago we had no data."
>
> -Joe, Senior Malware Analyst

Though there are many extremely useful technologies that support the work of MALs, much of it still comes down to gut instinct. Something I recognise *doing it*. The value of being in the MAL, through shadowing, chatting, being frustrated, and becoming-analyst, was the development of what social scientists may call an *affective* engagement. That is, an emotional relationship with the objects and subjects that we come into contact with. Only (auto)ethnographic engagement allows this deep experience to come to light in a sustained manner – and can explain why automated technologies maintain high false positive rates (above an acceptable threshold for users). In MALs, automation blends with conventional analysis and human gut instinct – which wraps the societal expectation of maliciousness (instead of the parasitical that I believe automated technologies are more working upon) into the mix of detection. This produces detections that are always suboptimal (though acceptable to be used due to the mathematical logics behind it).

That gut instinct and insight took time to develop. It took at least three months to even feel vaguely confident enough to broach the topic of writing 'real' detections after playing with training samples. After April 2017, and significant training, I moved into a phase where I started to settle in to producing detections for different malware *families* that used a variety of tools and techniques that I broadly say come under a *pathological* logic. That is, there are attempts to render malware knowable through a complicated 'gaze'[iv], where there is a medical-like sensibility in the MAL; through dissection through disassembly, through monitoring behaviours in sandboxes, to big data processing, and to categorise and develop new detections through machine learning. I in no way claim this is a *bad* thing – but merely note this is how we approach malware, which has implications for how we see it. That is, we assume it has an origin source in the author, where intentionality flows in certain ways from them, to the fact that we speak of families and rarely of specific instances according to environmental variables. Most malware has some sort of reference to its environment; even if it could be regarded as inconsequential – such as a file path. Yet it is not something that is immediately apparent in our discussions outside of MALs. Sophos permitted a research project, through (auto)ethnography, that allowed for a broad study of malware without particular aim or goal. I went in to the MAL with some vague ideas of what happened, some of them radically overturned, and the findings of my research are situated there, but can be combined with other research to argue for broader movements in malware and cybersecurity.

Reflections

Gaining access to sites of security are a crucial research issue for social scientists. As part of the CDT in Cyber Security, I was aided in gaining access that would be rare outside – and is a debt I owe to the centre. This makes interdisciplinary centres for social scientists a crucial component of cybersecurity. However, we are not here to provide assistance to a certain product, service or computer science research area (though at times, this can be exceptionally productive). The 'value' (in scare quotes for a reason) of social science cannot be frequently tied to a research output in the sense that we would expect from our computer science and engineering colleagues. Indeed, there is no singular social science to speak of, as much as we do not speak of science as a totality. Though I do not elaborate on exactly what I found in the MAL here apart from some insights and observations, nor my subsequent findings that will be forthcoming, I hope I have provided an overview of the importance of exploratory social science work – that helps to (re)construct familiar notions in cybersecurity; that I hope can be productive not only in geography but

elsewhere. Indeed, it is also to thank Sophos for their time and commitment to me, who were very supportive, and assisted in any way they could.

[i] Lockheed Martin, 2018. Cyber Kill Chain®. [online] Available at: <https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html> [Accessed 14 Jun. 2018].

[ii] Balzacq, T., & Cavelty, M. (2016). A theory of actor-network for cyber-security. European Journal of International Security, 1(2), 176-198. doi:10.1017/eis.2016.8, p. 192.

[iii] Parisi, L., 2017. Computational logic and ecological rationality. General Ecology: The New Ecological Paradigm, pp.75–99.

[iv] This draws extensively from Michel Foucault's use of the term in his 1963 book, *The Birth of the Clinic*.