BIG DATA SURVEILLANCE
MAY 2016 MEETING

Big Data Against Terrorism
Stéphane Leman-Langlois

The promise of Big Data is that, through advanced analysis, *added value* can be found in the maximisation of existing and/or new troves of seemingly unrelated data. From atom to molecule, to sheet metal to car, the algorithms will assemble the hidden – not *secret* but buried under masses of other data – parts of the whole to reveal the true nature of the object.

In the realm of national security big data analytics aim, for the most part, at *preemption* (stopping an attack at the preparation stage, for instance via communications analysis, following purchases of dangerous materials or behaviour revealing the "casing" of particular targets) and *prevention* (intervening further upstream before violent action is being considered, for instance trough network disruption or through the identification of persons "at risk"). Unfortunately, for the most part the analytical strategies to extract "actionable" information from large data sets have not succeeded. There are many reasons for that.

Hurdles for big data analytics in terrorism prevention
Rarity: terrorism, even if one uses the widest possible definition, is vanishingly rare in the West. In Canada it is rare enough to be considered statistically nonexistent. This of course means that there simply is no big data *on* terrorism, and that if trends continue, there never will be. This has two main consequences, depending on the type of analysis that is undertaken. If the aim is to identify *patterns* of terrorist or "pre-terrorist" (early preparation) activity, the patterns hunted for must come from an distillation of what terrorists, in general, have done when planning past attacks. This distillation is vulnerable in multiple ways, among which are sizeable gaps in our knowledge of terrorist biographies, reliance on data sources that are already filtered by common (usually wrong) knowledge about terrorism (such as media stories, police reports, etc. not to mention the dozens of cases involving non-trivial intervention by undercover agents in the plotting, see Mueller, 2011), deep variations in what is recorded caused by varying definitions of what constitutes "terrorism" and so on. Yet this is what the Total (later, *Terrorism*) Information Awareness (TIA) program planned to achieve in the early days after 11 September 2001 (Brodeur and Leman-Langlois, 2006). Other than the practical problems resulting from unreliable terrorist patterns, the TIA project was flawed in its very logic and, though aimed at prediction, was entirely structured around past acts of terror.

If the aim is to look at clustering and non-normal behaviour the need for prior patterns disappears. The target moves from the needle to the haystack. However in that case two equally insurmountable roadblocks are in the way. Clustering techniques are usually based on profiles and there are no reliable terrorist profiles. The only way to design a profile that might represent a reasonable proportion of terrorists, say, 30%, is to use extremely wide socio-demographic criteria such as age group, gender, education or socioeconomic status. In other words, to catch a significant proportion of terrorists we have to also falsely identify tens of thousands of citizens. Clustering may also look for non-normal or anomalous behaviour, but those are notoriously difficult to define. They are usually illustrated with actual cases of large amounts of fertilizer delivered to a suburban bungalow or the purchase of multiple one-way airline tickets. One approach is to start by a clustering method, first sorting out the base rates for various behaviours in order to identify statistical outliers. Neural networks are especially adept at automatically forming and testing hypotheses about regularities, through machine learning. This takes much subjective, cultural and simply erroneous assumptions, theories and classifications out of the loop. However, both terrorist and non-terrorist behaviours are extremely diverse and do not fall in clear categories, meaning that there is no way to identify the extents of the intersection between the "abnormal" and the "terrorist" sets. Because of the rarity of terrorism, and the profusion of "abnormal" behaviours, the number of false positives is again likely to skyrocket.

Finally, if the analysis of big data is aiming at finding networks, or links between persons (graph analysis), it quickly runs into a definitional problem. One common strategy of network analysis is to link communications through metadata. When sender, recipient(s), time of day, duration of call and frequency of communications are fed into the algorithm, various types of network analysis statistical techniques can be used which, combined with social psychology theory, can reveal the role of individuals within a group within a reasonable margin of error. The proper "big data" way to start the process is to access large amounts of bulk communications and attempt to identify "dark networks" that have common characteristics not shared by most telecom users (mostly time of day of the calls, duration, and localisation; statistically, they also tend towards low "link density" or fewer contacts among members). The problem of this analytical strategy is that there is no base rate, in other words we don't know how many dark networks are perfectly legitimate (as would be the case of political dissidents for instance). Every statistical study only uses confirmed illegitimate networks in its analysis (for instance, Xu and Chen, 2008, who in fact treat "dark" as synonymous with "illegitimate").

In reality most link analysis approaches begin not with big data but with very small data: once a person is labelled "terrorist" his or her contacts are simply filtered out of the captured metadata. This is at once simpler and more reliable, but rests heavily on the trustworthiness of the initial labelling. When successful this approach significantly narrows down the populations that should be given additional attention by high policing, national security intelligence organisations. But it is not big data analysis per se.

Math: Using big data analytics to prevent rare events is an overly optimistic scenario for other, more fundamental reasons. The first one is the Bayes theorem, which simply uses basic probability calculations to underline the impracticality of such predictions. Simply put, and only considering false positives, even a near-magical 99% accuracy algorithm would wrongly identify 1%, so 320 000 random Canadians, as suspects (out of 32 million adults).

One other math problem was first raised by data scientist Jeff Jonas (Jonas and Harper, 2006) and has grown into a humour industry on the web (http://www.tylervigen.com/spurious-correlations): given enough data, fortuitous correlations of variables increase rapidly; they in fact become plentiful, and inevitable. Such associations are, therefore, meaningless knowledge despite their extremely high statistical "significance". Put another way, just as it has been said that 600 monkeys on 600 typewriters, given enough time, would eventually rewrite the works of Shakespeare, given enough data it seems that any hypothesis, however eccentric, can be proven. Given that much datamining is the identification of patterns and the algorithmic comparison of digital traces of behaviour, this problem is likely to become intractable with the ever increasing mountains of data. This is sure to shatter the hopes that big data is a direct path to reality: on the contrary, the more data, the more need for theory.

Civil rights: as obvious in the previous remarks, the costs of big data counterterrorism, in terms of civil rights can be projected as particularly high. Though civil rights are an external factor and have nothing intrinsically to do with data mining, they might be a source of resistance within the public, or at least within subgroups with sufficient political capital, who might counter efforts to build counterterrorism programs based on big data analytics.

Sources of data are already heavily influenced by the intense cultural focus on selected populations of "terrorists." Whether the original data comes from police, media or popular attention (through the analysis of twitter feeds for instance) it is a given that Arab, Muslim or Arab Muslim populations (or any other that they may be confused with, such as Sikhs) will disproportionately be included in any database, which in turn will further skew the models towards the same populations.

Big data on terrorism is a colony of totalitarian and authoritarian states. Untold amounts of "intelligence" is based on show trials, security services manipulation and politically expedient designations used to neutralise individuals, groups and populations. For instance, the famous "returnees from Albania" show trial in Egypt gave rise to hundreds of wrongful identifications in the West, including the issuance of "security certificates" in

Canada.

Police reports produce much of the information included in big data databases, especially the confidential ones being used by national security organisations. This creates a feedback loop where the data is mined by those creating it, the results are included in the database, and then re-analysed. This problem is common with all crime statistics. As previously stated this is especially worrisome with terrorism cases, since a majority result from undercover investigations. To this must be added the well established fact that police databases contain in some cases up to 40% erroneous records.

In the case of algorithmic profiling, we have seen that even the best designed system will create massive numbers of false suspects. This will certainly have the effect of swamping police and security forces, making them less, not more efficient. But it will also subject thousands of randomly selected persons to excessive scrutiny from police organisations, potentiallly at the international level.

As Lyon (2014) notes, datamining and especially clustering analysis amplify *social sorting*. Even if the technology was objectively neutral it would still be deployed in a non-neutral, heavily stratified real world and, under control of existing elites, would tend to reinforce and exaggerate the stratification.

Asides from problems that stem from the (mis)identification of persons one could include multiple "chilling effects" increasingly associated with mass surveillance. Research has shown that surveillance modifies the way we express ourselves, even if we are convinced we have "nothing to hide"(see Penney, 2016; Stoycheff, 2016). In other words, surveillance, even without consequent intervention, measurably affects our feeling of autonomy from the state apparatus, in this case in our ability, and in fact *desire*, to speak freely. In Stoycheff's results, fewer of those polled were willing to discuss NSA power if they knew the NSA might intercept their words. This reduced autonomy is certain to affect areas other than speech. For instance, if geolocation is included in police datamining one might attempt to avoid certain urban areas associated with illegal or potentially embarrassing behaviours.

In principle big data analytics are driven by algorithms: it is robotic surveillance. That form of machine watching has often been presented as free from the potential for abuse or other illegitimate uses of surveillance powers by humans (most notably at the time of the TIA debacle in the USA). In the case of human-designed algorithms there subsists a potential for misuse at the programming stage. But learning machines such as neural networks are meant to depart from their original programming and improve themselves. One hope is that this could neutralize any original bias built in at the programming stage, such as ethnic profiling, if it does not fit the data. However, as we have seen, most databases that have to do with terrorism are fundamentally biassed from the start. To the extent that robots "learn" from these data (they would also consult vast numbers of databases that are not explicitly linked to terrorism), they could exponentially increase the effects of the original bias.

Social and political factors behind the rise of big data in counterterrorism

Given the above, the accelerated expansion of big data analytics in the context of national security may be puzzling. The paradox is not unlike that which exists in the world of camera surveillance. Cameras are also accepted, if not defined, as the ready solution to almost any particular security problem, even when evidence of lack of results in similar environments, for similar goals, is available. In this case cameras nevertheless flourish because of 1) their non-threatening appearance — they are sometimes called "protection cameras" in order to underline their inoffensive nature; 2) their relative low cost, especially since maintenance or manning costs are never taken into account; 3) their "democratization," which gives them a lower socially constructed risk; 4) their assumed objectivity in showing reality; 5) the conviction that surveillance solves all misbehaviour problems; 6) the high-tech sales pitch and 7) risk mitigation approaches promoted by the insurance industry. Many of the same factors are at work with big data analytics.

Big data is becoming a popular meme in the national security area, first and foremost because most modern organisations, whether public, private or hybrid, are already collecting large

amounts of data on every facet of their day-to-day operations, with exponentially more to come in the near future, and that most are trying to keep afloat by improving their data centres or moving them to the cloud. Almost every data centre management software vendor, as well as all cloud providers now sell their services with “big data analytics” of some sort.

At the epistemological level, big data is presented as a transparent channel to objective reality: the more data, the more reality. Under this view, most other forms of knowledge about humans and their various environments are flawed, mainly because they rest on an imperfect scientific process involving theorisation, partial empirical observation and an epistemologically weak, or disputed, logical link connecting the two (deduction, induction, falsification or abduction, all demonstrated as useless in philosophy of sciences 101 classes). This scientific process is vulnerable to subjective and cultural biases, as its history of infamous mistakes demonstrates, and should be replaced by the direct contact with reality and its workings that the era of big data analytics now offers us. What defines “big data” is not its size, speed, variety, etc. but the expectation, on the part of its users, that it crosses the threshold where the data becomes the thing, the map becomes the land.

As noted above, the idea that there could be a form of machine-led “privacy,” where algorithmic surveillance is freed from human misuse, was a cornerstone of the defence of the TIA program in early 2002. It was also an important response offered by the NSA to the Snowden revelations. Since we can assume that the future of all forms of technical surveillance will in fact be big data analytics, because of their expansion and because of the costs and unreliability of human monitoring, machine-led privacy is bound to become the leading response to claims of loss of privacy.

Perhaps the most powerful factor of acceptance of big data surveillance is its claim to provide security. For the time being this is more a promise than an actual result, as security successes are extremely few. Notably, when asked to demonstrate the usefulness of its various surveillance programs, the USA NSA was only capable of providing some 50 examples, worldwide, for over 15 years of existence of the programs. Considering the astounding resources engaged, this is a rather meagre return on investment. What is more, under scrutiny the 50 turned out to be more like a single instance of NSA datamining playing a secondary role (Cahall, Sterman, Schneider and Bergen, 2014). Nevertheless, the discourse of security continues to be very strong and debates about mass surveillance that followed the Snowden papers all centred on security, and often with no questioning of whether programs labelled “security” actually produce increases in security.

The simple invocation of “security” is a nearly indisputable safe conduct for all forms of surveillance. In the case of catch all big data surveillance, it also breaks down in a series of more specific benefits. In terms of prevention, big data promises for example to identify those who are “radicalizing” by digesting massive numbers of tweets or web searches. It also promises to follow the connections between individuals and to track persons of interest from one computer trace to the other. In a society where even an extremely limited terrorist attack triggers loud demands for more policing and more tracking of risky persons, these technologies are easy to sell. That they don’t produce results matters far less; that they *might* work and that they “don’t hurt” is often explicitly invoked in their legitimation. Finally, when all else fails, big data can be used to identify active terrorists and target them with drone strikes.

These few explanations do not entirely solve the paradox of big data for counterterrorism. As it quickly conquers the world, big data analytics will certainly continue to fail at predicting acts of terror but in all likelihood these failures will be ignored. This makes it a fascinating object of study and a powerful window on the social representations of mass surveillance, security and terrorism.

References

Brodeur, Jean-Paul and Stéphane Leman-Langlois, 2006. “Surveillance-Fiction : High And Low Policing Revisited,” K. Haggerty and R. Ericson, *The New Politics of Surveillance and Visibility*. Toronto: University of Toronto Press, 171-198.

Cahall, Bailey, David Sterman, Emily Schneider and Peter Bergen, 2014. “Do Nsa's Bulk Surveillance Programs Stop Terrorists?” *New America*, https://www.newamerica.org/international-security/do-nsas-bulk-surveillance-programs-stop-terrorists/.

Chen, Tingting and Sheng Zhong, 2009. Privacy-Preserving Backpropagation Neural Network Learning. *IEEE Transaction on Neural Networks*, 20(10), http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=5223520.

Lyon, David, 2014. Surveillance, Snowden, and Big Data: Capacities, consequences, critique. *Big Data and Society*, 1(13), http://bds.sagepub.com/content/1/2/2053951714541861.

Mueller, John (ed.), 2011. *Terrorism since 9/11: The American Cases*. Mershon Center, Ohio State University.

Penney, Jonathon, 2016. Chilling Effects: Online Surveillance and Wikipedia Use. *Berkeley Technology Law Journal*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2769645.

Stoycheff, Elizabeth, 2016. Under Surveillance: Examining Facebook’s Spiral of Silence Effects in the Wake of NSA Internet Monitoring. *Journalism and Mass Communication Quarterly*, 1(16) http://m.jmq.sagepub.com/content/early/2016/02/25/1077699016630255.full.pdf?ijkey=1jxrYu4cQPtA6&keytype=ref&siteid=spjmq.