serene risc Digest Spring 2016

Cutting Edge Research Summaries for Policy-Makers and Practitioners

Do national anti-botnet initiatives actually work ?

Short-term national botnet eradication initiatives are not effective. Long term planning for ICT skill increases and unlicensed software reduction is required.

page 5

How effective are big data technologies for Intrusion Detection ?

Machine learning and data mining technologies offer promise for system misuse and intrusion detection but are currently limited by a lack of training datasets.

page 6

Does multi-stakeholder Internet governance really favour equal representaion of diverse interests?

Multi-stakeholder governance requires that influence and input between government, private and civil sectors be balanced on a global scale before it can realize its promise.

page 7

Can we secure employees using their own apps with a plugin ?

Users bringing their own applications to work creates security issues. Solving these issues isn’t simple as users will not accept add-on security at the expense of utility.

page 8

Are people really happy to give up their privacy online ?

Although there are differences between privacy attitudes and behaviour, this is perhaps due to poor decision-making and not reasoned choice.

page 9

Does behavioural evidence analysis work for forensics ?

Applying behavioural evidence analysis (BEA) to digital evidence can help investigators to understand offender characteristics and authenticate digital evidence.

page 10

Is it possible to automatically assess the personality of people on Facebook ?

Language-based assessments using text from social media are capable of representing personality variance.

page 11

Do Android apps send secret messages and is that ok?

Covert communication is common in most popular free applications for Android and they cost users privacy, bandwidth and power consumption.

page 12

Can we predict outbreaks of hate online ?

Social media can act as early warning systems for deviance beyond an event. Practitioners need to focus on interventions early on to encourage a faster and more widespread de-escalation.

page 13

Can social media provide useful intelligence to first responders in a crisis ?

There are a small number of tweets in a crisis that provide useful information for responding agencies. Focusing on these can enhance intelligence gathering and decision-making processes.

page 14

Do national anti-botnet initiatives actually work ?

Botnets are major security threats for information systems and many countries have costly national initiatives to fight them. This study evaluates the impact of national anti-botnet initiatives for multiple countries. A similar pattern for all countries emerged: botnet activity shows rapid growth followed by a period of stability before a gradual decline. The data reveals no meaningful variance between the cleanup rates in countries with national anti-botnet initiatives and those without. These findings raise doubts about the effectiveness of national anti-botnet initiatives. Botnets cannot be eradicated by short-term national initiatives. A long-term view needs to be promoted for effective botnet cleanup.

How effective are big data technologies for Intrusion Detection ?

Intrusion Detection Systems identify unauthorized manipulation of information systems. They can use machine learning or data mining methods to recognize malicious activity and identify deviations from normal behaviours. Although both methods use similar statistical techniques, they are different in application. Machine learning requires a specific goal. Data mining focuses on the discovery of previously known properties in the data. Criteria such as accuracy, complexity, classification time and outcome understandability need to be taken into account when considering the application of each method. The cyber domain has peculiarities that make these methods harder to use. Data mining and machine learning methods must learn how to detect new intrusions frequently. Unfortunately, cybersecurity research on model retraining is limited. Investment in accurate and well-labeled data for model development could help create more efficient cyber threat detection tools.

Does multi-stakeholder Internet governance really favour equal representaion of diverse interests?

Governments, the private sector and civil society are all stakeholders in Internet governance. The United Nations and the World Summit of Information Systems are successful examples of multi-stakeholder governance. Multi-stakeholder governance is commonly thought of as the best model for Internet governance. An examination of the current situation from a multi-stakeholder approach finds that government and civil society are not equally represented. These findings contrast the misconceptions that government involvement impedes progress and that civil society is equally included in multi-stakeholder governance. On the contrary, there is a need for representation of broader interests in government, and civil society lacks the funding necessary to exert its interests fully. The current state of governance favours those who created the Internet. For multi-stakeholder governance to be effective, all governing elements must be represented equally. Equal collaboration between governments, civil society and the private sector is necessary. This allows for better representation of interests. Although some multiple stakeholder models are successful, they may not be universally effective. Discussion and debate are required to increase effectiveness in multi stakeholder governance.

Can we secure employees using their own apps with a plugin ?

Many companies permit the use of personal devices (BYOD) and also personal applications in the workplace, like cloud and web-based services. These applications allow employees to choose their own software to get their work done, but raise new issues for securing networks. Personal applications may be useful, but may not be secure. An always-on, webmail encryption plug-in is a promising solution, but it is not flawless. Although the always-on plug-in approach is secure, it compromises usability. Encryption solutions limit searchability and complicate organization of messages and accounts. Service users wish to be able to search and organize their messages, all while remaining secure. In order for security protocols to be successfully adopted, they must take into consideration features that are essential to service users. There is a need for comprehensive policies to better secure the use of personal device and applications in a professional setting, all while respecting usability.

Are people really happy to give up their privacy online ?

Studies show that privacy is a primary concern for citizens of the digital age, but this worry about privacy is sometimes not reflected in behaviours creating a contradiction between privacy attitudes and actions; this is often referred to as the privacy paradox. Research that attempts to explain this phenomenon has found inconsistent results, perhaps because of the variation in individual privacy concerns. Personal information is not homogeneous and individual attitudes towards privacy vary depending on the type of personal information involved. Future research on the privacy paradox could take into account the diversities of personal information types, privacy concerns and harms. As the causes and implications continue to elude understanding, decision-making should not be based on the assumption of a privacy paradox. Future research matured by the work done to date can provide the clarity required for good policy.

Is it possible to automatically assess the personality of people on Facebook ?

The text accumulating in social media is a source of rich data on human psychology. For example, millions of Facebook users regularly express emotions and attitudes with updates to their status and in messages exchanged with friends. This social media language is an especially good match for psychological study for several reasons; social media contains publicly available material that was written in natural social settings and can be retroactively accessed. This information can be leveraged to create a fast, valid and stable method for personality assessment online. By examining the Facebook content of over 66,000 volunteers, the authors built a language model using an open-vocabulary approach to language analysis. The language from social media can be used to correctly assess participants` key personality traits, providing a cost-effective alternative to questionnaires and self-reports, and allowing assessment of psychological characteristics when other options are impractical.

Does Behavioural Evidence Analysis work for forensics ?

Peer-to-peer (P2P) file sharing networks have created opportunities for the easy sharing of wide range of digital content. Unfortunately this has included Sexually Exploitative Imagery of Children (SEIC) or child pornography. Behavioural Evidence Analysis (BEA), a common practice in conventional criminal investigations although rarely applied in the digital realm, can assist the investigator, forensic practitioners and prosecutors in producing a more accurate and complete reconstruction of the crime. Drawing on digital evidence from fifteen archived cases of SEIC, the authors examined and analyzed the suspect's digital devices, attempting to reconstruct the suspect's activities in obtaining and sharing SEIC. The analysis of files on the computer revealed indicators of suspicious activity, signature behaviours and psychological characteristics of the suspect. The BEA technique can establish a bridge between the behavioural and technical aspects of digital evidence by enabling a more detailed reconstruction of evidence that can inform sentencing and prosecution.

Do android apps send secret messages and is that ok ?

Mobile applications often lead a double life, performing functions seen by the user while also sending and receiving data without the knowledge of the user. Overt communication contributes to the application functionality and is anticipated by the user, while covert communication is hidden and unexpected from the user's point of view. Nearly half of all connections from the top 500 popular Android applications available from Google Play are covert. When the covert connection statements were disabled, 63% of applications ran with no effect on the user-observable application functionality. While other applications did show some effects on functionality, disabling connections deemed covert leaves the delivered application experience either completely intact or with only insignificant interference. Covert communications present the user with unanticipated costs such as potential privacy breaches, bandwidth charges, power consumption on the device, and the unexpected presence of continued communication between the device and remote organizations. Covert communication can impair the transparency of device operation, silently consume device resources, and ultimately undermine user trust in mobile applications. These findings show that covert communication can be identified and removed.

Can we predict outbreaks of hate online ?

Specific events can influence the prevalence and severity of crimes with a prejudicial component. In particular, terrorist acts have been found to function as "trigger" events that "validate" prejudicial sentiments and tensions. However, the connection between these "trigger" events and the production of cyberhate on social media remains largely anecdotal. Drawing on 427,330 tweets from a fifteen day period, Williams and Burnap study the shape and size of cyberhate on Twitter, following the Woolwich terrorist attack in 2013. The event amplified deviant social reactions, such that an increase in cyberhate is evident within the first few hours following the terrorist event. Beyond the initial impact stage, the duration and diffusion of cyberhate is more inhibited, as the focus shifts to wider implications and issues of the event. Deviant reactions, such as cyberhate, can form part of the social response to a trigger event. Social media is a source of meaningful insights into social processes and can act as an early warning system for the amplification of deviance beyond an event itself, allowing practitioners to focus on de-escalating tensions.

Can social media provide useful intelligence to first responders in a crisis ?

In a crisis, Twitter may seem like a valuable source of information. However, previous research suggests that information shared on Twitter during a crisis varies from one crisis to another. To better understand the information shared on Twitter during a crisis, Olteanu et al. collected a sample of tweets from 2012 to 2013, covering 26 crises. The data suggests that responses depend on the spread and type of the crisis. Widespread crises result in more responses of caution and advice than localized crises. Human-induced crises have similar responses when compared to other types of crises. During a crisis, an overwhelming majority of tweets are from outsiders repeating information. Less than half of all tweets for a crisis contain important information and only 9% are eyewitness accounts. Although only a small proportion of tweets in a crisis contain valuable information, there are recognizable patterns. Automated programs may be able to work with these patterns to deliver real time information on particular crises.

Post-Mortem of a Zombie: Conficker Cleanup After Six Years

Botnets are a major security threat for information systems. A botnet - from 'robot' and 'network' - is a group of computers that communicate with one another, often for malicious tasks and without the knowledge of the owners of the machines. Interrupting and cleaning up botnet operations has been proven unsuccessful in mitigating their expansion. Various large-scale national anti-botnet initiatives have emerged during the recent years in different countries including Germany, Australia and Japan. The mitigation strategy of these costly national campaigns includes collecting data about infected machines, notifying Internet service providers (ISPs) of botnets in their networks and providing technical support as well as information to end-users.

This study evaluates the impact of national anti-botnet initiatives against Conficker, one of the most widespread botnets of all time. First detected in 2008, Conficker spread through a security vulnerability in Microsoft Windows. Asghari et al. used Conficker traffic data collected from 62 different countries over 6-year period to understand what factors affect the rate of cleanup for infected machines and whether countries with anti-botnet initiatives have better cleanup rates.

Dynamic Internet Protocol (IP) addresses change frequently, providing a single Internet connected computer with many different addresses over time. This makes a simple count of infected IP addresses an inaccurate estimate of the number of bots. To overcome this issue, an averaged number of unique infected IPs per hour has been used instead. The number has also been normalized to account for the number of Internet subscribers in each country. After these corrections, a similar pattern for all countries emerged from the data. For example, the graphs below show Conficker trends for Germany, United States, France and Russia. In each country, botnet activity shows rapid growth followed by a period of stability before a gradual decline.

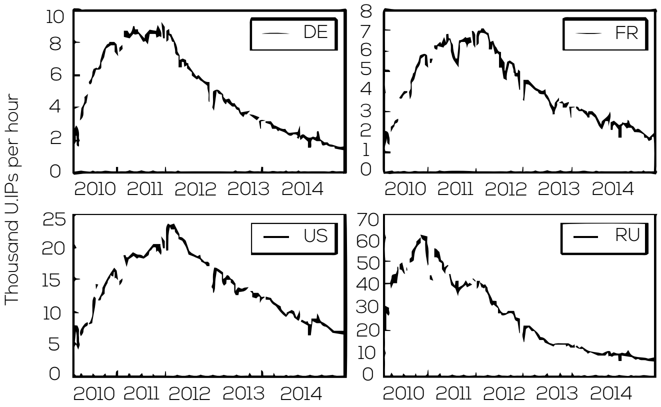

This data reveals the variance between the maximum number of bots in each country, relative to the number of Internet subscribers. Surprisingly, there was no meaningful variance between the cleanup rates in countries with national anti-botnet initiatives and those without. These findings raise doubt about the effectiveness of national anti-botnet initiatives. The study also provides a comparison of the Conficker botnet with GameoverZeus, a more recent botnet; this revealed that 15% of computers infected with GameoverZeus are also infected with Conficker. This raises the problem of further victimization of and ongoing harm to vulnerable computers. Interestingly, the number of infections per vulnerable user appeared to be related to the rate of unlicensed software use and to an index for ICT development, which incorporates factors such as a country's ICT readiness, infrastructure, and skills. These two factors explain the majority of the variation between peak infection rates in different countries.

Short-term national botnet eradication initiatives are not effective. Long term planning for ICT skill increases and unlicensed software reduction is required.

Asghari, H., Ciere, M., & Van Eeten, M. J. G. (2015). "Post-Mortem of a Zombie: Conficker Cleanup After Six Years." Presented at the 24th USENIX Security Symposium (USENIX Security 15), 1-16.

A survey on Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection

Cybersecurity is a set of technologies and processes designed to protect computers, network, programs and data from adverse cyber incidents. To do so, Intrusion Detection Systems (IDS) can identify unauthorized use, duplication, alteration and destruction of information systems. These systems employ machine learning (ML) or data mining (DM) methods to recognize known signatures of malicious activities, as well as to identify deviations from normal behaviors. This paper provides a deeper understanding of ML and DM techniques, by overviewing some popular and emerging cybersecurity methods.

Buczak and Guven (2015) review the literature on machine learning (ML) and data mining (DM) methods for cyber analytics in support of intrusion detection in both wired and wireless networks. The authors focused on highly cited papers and on recent papers presenting emerging methods.

ML focuses on classification and prediction, based on known properties previously learned from the training data. DM is the application of specific algorithms for extracting patterns from the data. ML needs a specific goal, whereas DM focuses on the discovery of previously unknown properties in the data. As both methods employ similar statistical techniques the authors label the methods as ML/DM methods. The methods discussed in the article are described in the table.

The authors conclude that the most effective ML/DM methods for the cyber domain have not yet been established. Given the richness and complexity of each method, one recommendation for each method, based on the types of attack a system is supposed to detect, cannot be accomplished. Several criteria need to be taken into account when considering different methods: accuracy, complexity, time for classifying a threat and the understandability of the outcomes for each ML/DM method. The authors also emphasize streaming capabilities are essential in the cyber domain, as cyberattacks require real-time analysis of online data.

The authors find that the cyber domain has peculiarities that make these ML/DM methods harder to use. Models need to be retrained frequently, depending on new intrusions. Yet, cybersecurity research on retraining of ML/DM models is limited, due to the scarcity of good datasets. The best dataset available so far, for testing ML/DM methods, is the 1999 corrected datasets of the Knowledge Discovery in Databases (KDD). Investing in the collection of up-to-date representative data, by for example putting sensors on networks, could foster the creation of more efficient tools that detect cyber threats in information systems.

ML/DM methods used in the cyber domain, especially for developing tools for IDSs, are still evolving. Further research is needed to develop fast incremental learning in ML/DM methods that could be used for daily updates for IDSs. Investment in accurate well-labelled datasets is the first step toward this goal.

ML/MD method	**How it can classify malicious network activity**
Artificial Neural Networks	Inspired by the brain, this method creates layers of artificial neurons capable of computing their inputs to generate a classifying output.
(Fuzzy) Association	Discovers previousl relationships among different data attributes providing association rules.
Bayesian Network	The maps, the variables and the relationship between them on a probabilistic graphical model.
Clustering	Finds patterns (similarities) in unlabelled data.
Decision Trees	Method with a tree-like structure that has leaves. The leaves represent the conjunction of features that lead to the classifications.
Ensemble Learning	Ensemble learning searches the hypothesis space to determine the right hypothesis that will make good predictions for a given problem.
Evolutionary Computation	Finds effective computational methods by the principle of survival of the fittest. Useful computational elements survive.

Machine Learning and Data Mining technologies offer promise for system misuse and intrusion detection but are currently limited by a lack of training datasets.

Buczak, A., & Guven, E. (2015). A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection. Communications Surveys & Tutorials, IEEE, 1-26.

Power Plays in Global Internet Governance
Regulation by Risk of Harm

It is commonly thought that as the Internet has multiple stakeholders a multiple stakeholder model is the best method of governance. Is this correct or is it just a comfortable assumption that fits the view of those who currently hold power? The truth is important because the Internet represents current power relations and diverse, competing interests.

Historically, the private sector has been the default coordinator of Internet governance. In the multi-stakeholder model of Internet governance, governments, the private sector and civil society are the primary actors. The multi-stakeholder governance model is favored by the United States of America and has received considerable support. The United Nations and the World Summit on Information Systems are examples of successful multi-stakeholder initiatives. However, very few experts have critically analyzed the practical implementation of the multi-stakeholder approach. Carr presents multi-stakeholder governance in the context of liberalism: a paradigm and philosophy grounded in equality and liberty.

An examination of the current situation using a multi-stakeholder approach finds an unequal inclusion of government and civil society. For multi-stakeholder governance to be effective these two groups should be included equally. If these groups are represented equally, multi-stakeholder governance has the ability to accommodate diverse interests, to seek expertise on an issue, and to develop a more inclusive view.

There are two misconceptions about the current state of Internet governance. First, the idea that government should be removed from Internet governance so that it does not impede progress is faulty. There is a need for representation of broader interests of the different states in the light of their interpretations of an optimal Internet. Secondly, there is also an assumption that civil society is equally included in multi-stakeholder governance. Currently, civil society is unable to compete for equal inclusion because the sphere lacks the funding and influence to exert its interests fully.

These misconceptions point to some challenges in multi-stakeholder governance. Without diverse discussion and debate multi-stakeholder approaches are inefficient. To effectively deliver on the promise of true multi-stakeholder governance requires the combination of all governance elements so they can represent their interests, from the maintenance of state security, through the preservation of industry to the upholding of ethical practices.

MYTH Government should be removed from Internet governance so that it does not impede progress
There is a need for representation of greater interests of the different states in the light of their interpretation of an optimal Internet.

MYTH Civil society is equally included in multi-stakeholder governance
Currently, civil society is unable to compete for equal inclusion because the sphere lacks the funding and influence to exert its interests fully.

Although there are examples of successful multi-stakeholder initiatives, it is a pure assumption that the model is universally effective. The practical application of a balanced multi-stakeholder approach for Internet governance requires considered, careful and conscientious implementation. Multi-stakeholder governance favors the agendas of those whom created the Internet, as evidenced by the overwhelming representation of the private sector. Just because multi-stakeholder models work for those who created the Internet, it does not mean that the approach works for everyone or is effective. Governments should seek private sector and civil society participation in governance. Additionally, the private sector should also seek government and civil society participation when searching for a well-rounded approach to solving a problem.

Multi-stakeholder governance requires that influence and input between government, private and civil sectors be balanced on a global scale before it can realize its promise.

Carr, M. (2015). "Power Plays in Global Internet Governance." Millennium-Journal of International Studies, 43(2), 640-659.

Seamless And Always-on Security in a Bring-Your-Own-Application World

The evolving world of information technology raises new issues for securing networks. Many enterprises offer a Bring-Your-Own-Device (BYOD) environment, which supports users who choose to take their own technology, like laptops and smartphones, to work and use them on a secured network. More recently cloud- and web-based services have allowed users to move applications seamlessly between computing environments. In a Bring-Your-Own-Application (BYOA) context companies encourage the use of consumer products such as personal applications in the workplace. This is often an extension of BYOD, one which also permits personnel to choose their software for getting their work done. Existing institutional and corporate policies for BYOD may not be enough to manage, and may even conflict with, this emerging BYOA practice.

Hecht et al. examined an always-on approach to usable security that could be applied in a BYOA environment. The goal was to test a non-intrusive and easy to use portable encryption method which showed the security of communications using an open and closed lock metaphor. Seventy-two participants downloaded and used the alternative encryption plug-in for webmail clients, such as Google's Gmail, on the Chrome browser. The researchers collected public user data from these participants and followed up with interviews of the 29 users who could be reached to provide feedback.

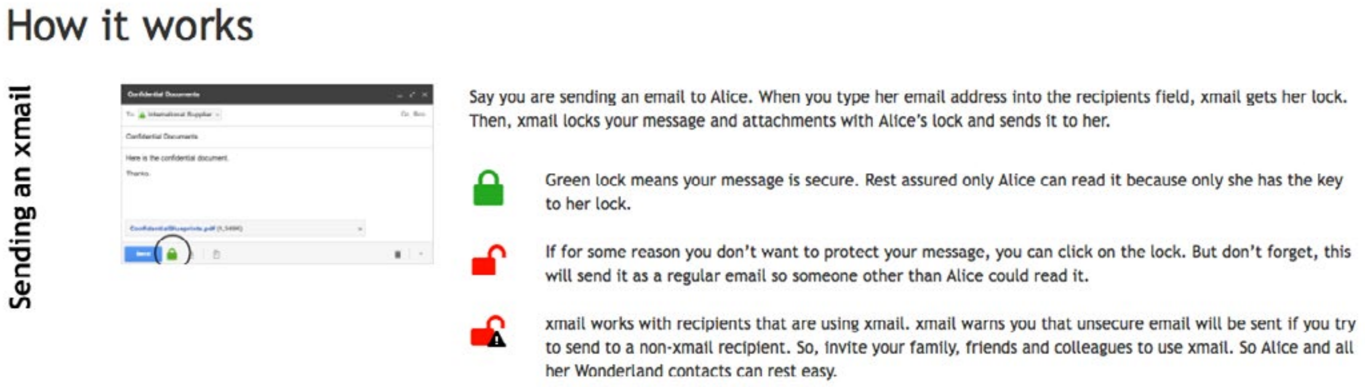
How it works

Sending an xmail

Confidential Documents

To: International Supplier

Confidential Documents

Here is the confidential document.

Thanks.

ConfidentialBlueprints.pdf

Say you are sending an email to Alice. When you type her email address into the recipients field, xmail gets her lock. Then, xmail locks your message and attachments with Alice's lock and sends it to her.

🔒 Green lock means your message is secure. Rest assured only Alice can read it because only she has the key to her lock.

🔓 If for some reason you don't want to protect your message, you can click on the lock. But don't forget, this will send it as a regular email so someone other than Alice could read it.

🔓⚠ xmail works with recipients that are using xmail. xmail warns you that unsecure email will be sent if you try to send to a non-xmail recipient. So, invite your family, friends and colleagues to use xmail. So Alice and all her Wonderland contacts can rest easy.

The always-on, plug-in approach is a promising solution but some issues remain. Users were successful in exchanging emails securely and were able to understand the lock and key metaphor used in the experiment. However, users expected greater flexibility in the implementation of the security features, due to a variety of email practices. In particular BYOA security solutions must permit users to:

- Keep a complex system of email accounts.
- Maintain email mobility across devices.
- Search protected messages.
- Decide the security status of emails.

This study demonstrated that users will not compromise usability for security in their webmail usage. For example, users want the ability to search all messages, even protected ones. This, however, conflicts with the encryption provided by the experimental plug-in. To be successfully adopted, security protocols must match these types of expectations for usability. This study demonstrated the need for comprehensive institutional and corporate BYOA policies that can meet evolving day-to-day practices such as use of popular webmail clients to manage all emails.

Users bringing their own applications to work creates security issues. Solving these issues isn't simple as users will not accept add-on security at the expense of utility.

Hecht, P., Fels, S., & Anacleto, J. (2015). "Seamless And Always-on Security in a Bring-Your-Own-Application World." Presented at the Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, 2019-2024.

Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox

Studies show that privacy is a primary concern for citizens of the digital age. However, this worry about privacy is sometimes not reflected in related behaviours creating a contradiction between privacy attitudes and actions; this is often referred to as the privacy paradox. Many individuals are willing to reveal personal information in return for some benefit even though they recognize the importance of privacy. Research that attempts to explain this phenomenon has found inconsistent results.

Kokolakis examines the research on the privacy paradox to illustrate the diversity of results and research methods used to understand this complex phenomenon. There is both evidence that supports and evidence that challenges the idea of a privacy paradox. Researchers have turned to various disciplines in search of theories that can contribute to the conceptualization of the phenomenon and the investigation of probable explanations. Most studies focus on individual aspects of the privacy paradox phenomenon. This has led to a need for integrative studies based on comprehensive theoretical models that take into account the diversity of personal information types and privacy concerns.

Any discussion on privacy is highly contextual and interpretive. Personal information is not homogeneous and individual attitudes towards privacy vary depending on the type of personal information involved. Some types of privacy concerns may have a stronger influence on attitudes and behaviours than others. Furthermore, much of the research on the privacy paradox suggests that individuals do not understand privacy in the digital world. While giving your name and address to a stranger on the street seems unsafe, individuals consistently disclose that same information online.

Kokolakis provides a number of recommendations for future research on the privacy paradox. As surveys can be unreliable data for actual user online behaviour, the author suggests avoiding self-reports of behaviour. Further research should also take into account the diversities of personal information types, privacy concerns and harms. Each individual will have varying perspectives on the importance of their information and how information disclosure could be harmful.

A privacy paradox is a new concept that is not yet understood. Any research should take into account that privacy is a highly contextual phenomenon. The samples chosen for study should be as representative of the public as possible because the privacy paradox not only concerns users of all ages but it also may affect them differently. Studies of the relationship between privacy behaviours and privacy awareness campaigns or improving privacy enhancing technologies could add depth to the privacy paradox phenomenon research.

There is evidence of a difference between privacy attitudes and behaviour. However, the extent, dynamics, cause, nuances and implications are far from being understood. Decision-making should not be based on the assumption of a privacy paradox and our current primitive understanding of this complex phenomenon. Only future research matured by the work done to date can provide the clarity required for good policy.

Privacy Calculus Theory	Individuals make privacy decisions as rational agents – perform a cost-benefit analysis of expected loss of privacy and the potential gain for disclosure
Emotional Rewards	Users experience emotional rewards of belonging to a community, rewards that are concrete and immediate, which override the abstract, calculated risks of data misuse
Optimism Bias	Individuals have a tendency to believe that they are less at risk of experiencing a negative event compared to others
Affect Heuristic	Users tend to underestimate the risks of information disclosure when confronted with a user interface that elicits positive affect
Bounded Rationality	People do not have access to all necessary information and lack the cognitive ability to calculate privacy risks and disclosure benefits
Indeterminacy	The outcome of a decision making process is determined at the time the decision is made but not prior to it, altering user preferences indeterminately

Although there are differences between privacy attitudes and behaviour, this is perhaps due to poor decision-making and not reasoned choice.

Kokolakis, S. (2015). "Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon." Computers & Security, 1-13.

Behavioural Evidence Analysis Applied to Digital Forensics: An Empirical Analysis of Child Pornography Cases using P2P Networks

Peer-to-peer (P2P) file sharing networks have created opportunities for the easy sharing of wide range of digital content. Unfortunately this has included Sexually Exploitative Imagery of Children (SEIC) or child pornography. As law enforcement agencies develop innovative investigative techniques to deal with the computer-facilitation of crimes, some more traditional techniques can also be leveraged in the digital realm. Among these, Behavioural Evidence Analysis (BEA) is a common practice in conventional criminal investigations. Studies have only recently included BEA to examine digital evidence. Along with the technical examination of digital evidence, it is important to learn as much as possible about the individuals behind an offence, the victim(s) and the dynamics of a crime. BEA can assist the investigator, forensic practitioners and prosecutors in producing a more accurate and complete reconstruction of the crime.

Mutawa et al. set out to examine the utility of BEA in investigating criminal cases that involve the possession and dissemination of SEIC through P2P networks, and to increase the understanding of the benefits of BEA in the interpretation of digital evidence. BEA has been very useful in conventional criminal investigations such as homicides, sex offences and rape, but has not been explicitly used to investigate SEIC cases.

The study was conducted with archival digital evidence from fifteen cases of SEIC offences made available by the Dubai Police. The authors used a deductive approach to analyze the individual cases separately and apply BEA strategies to the examination of each case. The authors examined and analyzed the suspect's digital devices, attempting to reconstruct the suspect's activities in obtaining and sharing SEIC. The analysis of files on the computer (e.g. Internet history files, recently accessed files, time stamps, deleted files) revealed indicators of suspicious activity, signature behaviours and psychological characteristics of the suspect.

Whereas conventional criminal offences often have common characteristics and behaviours, this study finds that computer-facilitated SEIC offenders did not share a common demographic profile. While this study suggests that it is not possible to construct a single profile of SEIC offenders, the results do suggest a type of offenders who are mainly viewers, downloaders and sharers of SEIC through P2P file sharing networks.

Using BEA to investigate computer-facilitated crimes can:

1. Assist the investigator in assessing the reliability of digital evidence and the strength of conclusions
2. Produce a more detailed reconstruction of evidence that can inform sentencing and prosecution in court
3. Assist in mapping and understanding offending behaviour and the dynamics of offences

For example, the location of SEIC files can indicate offender intentions. If a suspect has hidden or categorized SEIC files, this can indicate active participation in the offence.

Utilizing BEA in cases of SEIC offences can better equip forensic practitioners and prosecutors to take advantage of a behavioural interpretation of evidence. This technique can establish a bridge between the behavioural and technical aspects of digital evidence by enabling a more detailed reconstruction of evidence that can inform sentencing and prosecution.

Applying Behavioural Evidence Analysis (BEA) to digital evidence can help investigators to understand offender characteristics and authenticate digital evidence.

Mutawa, N. A., Bryce, J., Franqueira, V. N. L., & Marrington, A. (2015). "Behavioural Evidence Analysis Applied to Digital Forensics: An Empirical Analysis of Child Pornography Cases using P2P Networks." Presented at the 10th International Conference on Availability, Reliability and Security, 293- 302.

Automatic Personality Assessment Through Social Media Language

The text accumulating in social media is a massive source of rich data. Consider, for example, the many millions of Facebook users, regularly expressing emotions and attitudes with updates to their status and exchanging messages with friends. This data represents a potential to increase psychological research substantially. Social media language is an especially good match for psychological study for several reasons; social media contains an unprecedented amount of publicly available written language, material that was written in natural social settings and can be retroactively accessed. In addition, users disclose information about themselves at unusually high rates. If researchers can develop models from the data on social media, this information could be leveraged to create a fast, valid and stable method for personality assessment online.

Park et al. approached this study with the goal of building a model for accurately predicting personality using the written language data from social media. This primary research study used participants drawn from users of myPersonality, a third-party application on the Facebook social network.

Using the large sample size of over 66,000 participants, the authors built a language model using an open-vocabulary approach to language analysis. Open-vocabulary methods extract numerous and rich features from language samples, including single words or topics. The authors built a predictive model of personality within a sample of Facebook users, each of whom volunteered samples of their language and completed personality tests. The results of the language-based assessment (LBA) were compared with self-reporting questionnaires, informant reports and external criteria, such as information from online profiles. This method improved the systems accuracy over other language-based predictive models.

The language from social media can be used to correctly assess participants` key personality traits. LBAs are capable of capturing true personality variance and suggests that the language in social media can be harnessed to create a valid and reliable measure of personality. As well, compared with self-report questions, LBAs are fast, cheap and have low levels of participant burden. Once the model is created, the application of the model to a new user’s language data only takes seconds. Using computational techniques for LBA can reveal new layers of psychological richness in language and avoid the inherent biases in self-reports.

Language-based assessments offer a cost-effective alternative to questionnaires and self-reports, allowing assessment of psychological characteristics when other options are impractical. LBAs may also enable new approaches to studying geographic and temporal trends, comparing regional psychological differences and trends, and within-person variation over time and across locations. As LBAs can be generated retroactively, this approach can give researchers, investigators and situational awareness analysts the ability to provide a clearer portrait of the mentality behind online actions.

Language-based assessments using text from social media are capable of representing personality variance.

Park, G. H., Schwartz, A., Kosinski, M., & Stillwell, D. (2015). “Automatic Personality Assessment Through Social Media Language.” Journal of Personality and Social Psychology, 108 (6), 934-952.

Covert Communications in Mobile Applications

Mobile applications often lead a double life. They perform functions that their users see. However, it is quite common for Android applications to perform some functions and send and receive data without the knowledge of the user. Overt communication contributes to the application functionality and is anticipated by the user, while covert communication is hidden and unexpected from the user's point of view. Generally, this covert communication does not deliver any tangible value to the user and the user cannot opt-out from sharing data without uninstalling the application.

Rubin et al. set out to identify application functionality that is hidden from the user. This primary research is focused on understanding the extent of covert communication in applications downloaded from popular application stores. The authors developed a static analysis technique that can automatically identify covert communication connections. Once identified, those connections were disabled and the researchers performed a usability assessment to identify any change in application functionality. This technique correctly identified all but two covert connections, performing with an average precision of 93.2% suggesting that the static analysis method proposed in this paper could be useful for detecting covert connections. The technique is highly scalable and provides actionable output that can be used for disabling covert communication in a majority of cases. Furthermore, disabling these connections had little significant effect on application functionality.

The results from this study reveal that nearly half of all connections from the top 500 popular Android applications available from Google Play are covert. When the covert connection statements were disabled, 63% of applications ran with no effect on the user-observable application functionality. While other applications did show some effects on functionality, disabling connections deemed covert leaves the delivered application experience either completely intact or with only insignificant interference.

This study provides evidence of the prevalence of covert communications in the 500 top-popular free applications on Google Play for Android. Covert communications present the user with unanticipated costs such as potential privacy breaches, bandwidth charges, power consumption on the device, and the unexpected presence of continued communication between the device and remote organizations. Covert communication can impair the transparency of device operation, silently consume device resources, and ultimately undermine user trust in mobile applications. These findings show that covert communication can be identified and removed.

Covert communication is common in most popular free applications for Android and they cost users privacy, bandwidth and power consumption.

Rubin, J., Gordon, M., Nguyen, N., & Rinard, M. (2015). "Covert Communication in Mobile Applications." Presented at the 30th IEEE/ACM International Conference on Automated Software Engineering.

Cyberhate on Social Media in the Aftermath of Woolwich: A Case Study in Computational Criminology and Big Data

Specific events can influence the prevalence and severity of crimes with a prejudicial component. In particular, terrorist acts have been found to function as "trigger" events that "validate" prejudicial sentiments and tensions. These social responses contribute to the overall impact of an event. Contemporary online spaces, such as social media, offer new public spheres for expression and amplification. Cyberhate has recently been identified as a social problem on these platforms, one that must be addressed. However, the connection between these "trigger" events and the production of cyberhate remains largely anecdotal.

Williams and Burnap study the shape and size of cyberhate on Twitter, following the Woolwich terrorist attack in 2013. The sample consists of 427,330 tweets about the event, from a fifteen day period following the Woolwich attack. Due to the large sample size, the authors developed an interdisciplinary method called computational criminology. This method allowed them to analyse the data using advanced computing techniques. A supervised machine classifier that learned the features of hateful tweets toward minorities was employed.

Findings demonstrate that the event amplified deviant social reactions. An increase in cyberhate is evident within the first few hours following the terrorist event. Beyond the initial impact stage, when immediate responses to the event occur, the duration and diffusion of cyberhate is more inhibited. Cyberhate is limited in the reaction stage, where focus shifts to wider implications and issues of the event. The figure below shows the "half-life" of cyberhate tweets and retweets, evidenced by the rapid de-escalation immediately following the attack.

The study demonstrates that deviant reactions, in this case cyberhate, can form part of the social response to a trigger event. The study illustrates how social media presents a rich new form of data, which assisted by computational methods, can extract meaningful insights into social processes. Social media can act as early warning systems for the amplification of deviance beyond an event itself. Practitioners need to focus on interventions within the impact stage to encourage faster and more widespread de-escalation.

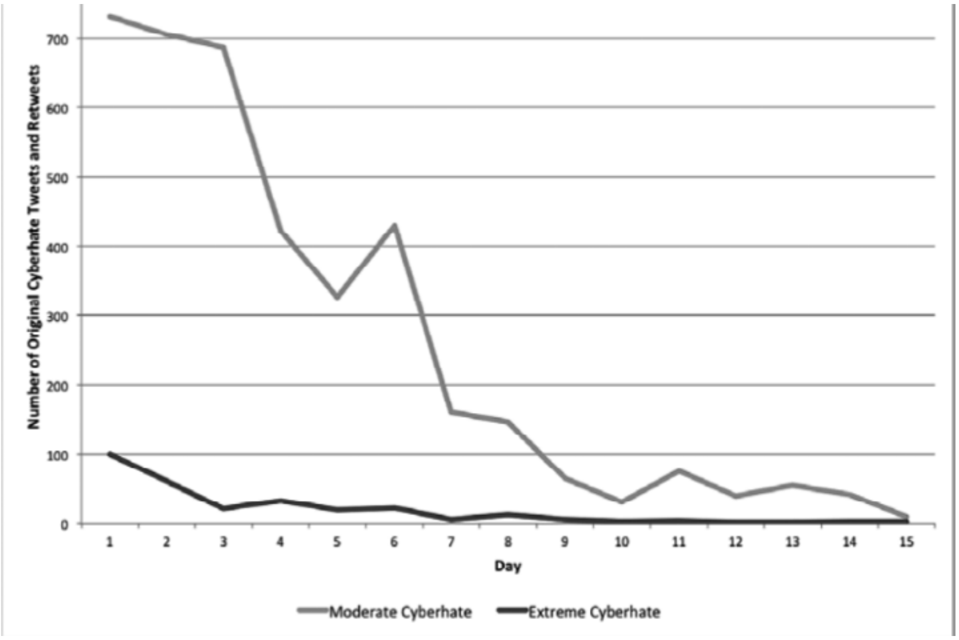

Social media can act as early warning systems for deviance beyond an event. Practitioners need to focus on interventions early on to encourage a faster and more widespread de-escalation.

Williams, M. L., & Burnap, P. (2015). "Cyberhate on social media in the aftermath of Woolwich: A case study in computational criminology and big data." British Journal of Criminology, 1-28.

What to Expect When the Unexpected Happens: Social Media Communications Across Crises

When a disaster occurs, time is limited and safety is in question, so people need to act quickly with as much knowledge of the situation as possible. Knowing where to find information is a crucial step. However, affected populations, response agencies and other stakeholders may not know what to expect from different information sources. Previous research found that information shared on Twitter varies substantially from one crisis to another, but understanding the similarities that could ease the information overload wrought by social media during crisis situations.

To build this understanding, Olteanu et al. investigate several crises in a systemic manner and with a consistent methodology. Olteanu et al. used a retrospective sample of tweets from 2012 and 2013, related to 26 events that spawned significant activity on Twitter. The data was compiled by a crowdsourcing team and then categorized by crisis dimensions; such as the type of hazard and geographic spread ,as well as by content dimensions, including the informativeness, type and source of information. The dataset (available at http://crisislex.org/) compiles on average 132 million tweets per month, amounting to roughly 38 GB of compressed data.

There are commonalities in the types of information people tend to be concerned with, given the particular dimensions of the crisis. Messages of caution and advice, and those containing information about infrastructure and utilities, were the most repeated. There are also subtle differences in Twitter responses that depend on the spread of the event over space and time. For example, when a crisis event is geographically diffused, the proportion of caution and advice tweets is higher, and when a crisis is localized the proportion of caution and advice tweets is lower. Despite these variations, the Twitter responses to human-induced crises are more similar to other human-induced crises than other types of hazards.

Eyewitness accounts as the source of information comprise, on average, only 9% of tweets for a crisis. Nearly 80% of all tweets relevant to a crisis are from outsiders reporting on the event with indirect and repeated sources of information. Furthermore, only an average of 47% of all crisis-related tweets contain the four main types of information (affected individuals, donations and volunteering, caution and advice, and infrastructure and utilities). Therefore, to use Twitter as a source of information during a crisis, a practitioner would have skip through the 80% of crisis-related tweets that are built on indirect and repeated information, then identify the 47% that concern the selected information types. This represents a relatively small number of tweets that could be of value for emergency management and response.

These findings will be of interest to members of the public, emergency managers and formal response agencies who are increasingly trying to understand how to effectively use social media during a crisis. Stakeholders who know what content and reliability to expect will not have to sift through masses of social media posts; instead they have a reasonable expectation of what they will find, and can then make more informed decisions regarding their situational assessment process. The results of this study show that there are a relatively small number of tweets that can help turn the information from Twitter into real-time information about the crisis.

There are a small number of tweets in a crisis that provide useful information for responding agencies. Focusing on these can enhance intelligence gathering and decision-making processes.

Olteanu, A., Vieweg, S., Castillo, C. (2015). "What to Expect When the Unexpected Happens: Social Media Communications Across Crises." Presented at Collaborating Around Crisis CSCW 2015, 994-1009.

SERENE-RISC Six Key Activities

Smart Cybersecurity Network (SERENE-RISC) is a knowledge mobilization network that organises six key activities intended to reach its various audiences: workshops and seminars, a knowledge brokers' forum, quarterly knowledge digests, Konnect - online knowledge-sharing platform, a public website and a professional development program.

Workshops & Seminars

April 2016
Vancouver
October 2016
Ottawa

Knowledge Brokers

Expanded Access Program

Knowledge Digest

More than 50 Summaries.
Sponsorship Opportunities Available

Konnect

More than 500 hand-selected resources on Cybersecurity including exclusive content

Website

Cyber security tips section , news on the network and Digest Archive.

Professional Development

Ask us about the Graduate Development Sessions

The SERENE-RISC Quarterly Cybersecurity Knowledge Digest

2016 Winter
Editor-in-Chief Michael Joyce; **Scientific Editor** Benoît Dupont
Editors: Emily Maddock, Masarah Paquet, Justin Anstett
Asghari, H. et al. summarised by Hanieh Moshki sponsored by SERENE-RISC
Hecht, P. et al., summarised by Hervé Saint-Louis sponsored by SERENE-RISC
Carr, M. summarasied by Alexandra Green sponsored by SERENE-RISC

Links to external sources for papers are provided for convenience only and do not represent an endorsement or guarantee of the external service provider.

Government of Canada Networks of Centres of Excellence
Gouvernement du Canada Réseaux de centres d'excellence
serene risc
Université de Montréal