serene risc Digest Spring • 2017

Cutting Edge Research Summaries for Policy-Makers and Practitioners

Do Russian fake news stories work?

Fake news stories have the potential to wage serious information warfare online. It is known that the Russian state controls its mainstream media and online discussions. This control may have shaped audience perceptions of the conflict in Ukraine. Khaldarova and Pantti looked at how people receive and refute fake news stories. First, they found 10 alleged fake news stories online that originated from Channel One, Russia's main news source. They then went through StopFake, a fake news debunking website, to find counter-narratives to these stories. Finally, they found 6043 tweets sharing these stories and determined whether they expressed trust or distrust. Most people did not seem to trust the stories. All tweets from Ukraine expressed distrust, while the only comments showing trust came from Russia. Distrusting comments usually shared strong emotions of disgust and highlighted inconsistencies. These fake news stories used strategic narratives to present a planned position on an issue. They blurred the line between reality and fiction and often included emotive references to World War II. They also encouraged stereotypes and oversimplified complex issues. Although the public may be aware of strategic narratives, a factually inconsistent news environment can create confusion.

Is machine learning useful for finding malicious URLs?

Universal Resource Locators (URLs) identify webpages and content online. Usually they are benign and lead to safe content; however, attackers can use malicious URLs to send victims to unsafe webpages. There are techniques that filter and detect malicious URLs, but they need improvement. Machine learning techniques can already classify URLs according to certain criteria. Mamun et al. wanted to improve upon machine learning techniques in malicious URL detection. They created a program that could learn and understand the specific lexical features of malicious URLs, like URL length and the types of characters used. They also taught the program to detect techniques used to maliciously mask or hide unsafe URLs. Their work resulted in highly accurate classifiers capable of detecting many kinds of malicious URLs, including spam, malware and defacement URLs. Security specialists can use these machine learning classifiers to detect malicious URLs along with other techniques, like URL blacklists.

Does how you lock your smartphone affect how often you lock it?

Smartphones allow for convenient access to many important online services but store a lot of personal data in return. Many people lock their smartphones, but little is known about their locking habits. Mahfouz, Muslukhov and Beznosov wanted to see how people use their Android devices "in the wild." They looked at the locking habits of 41 people who use Android smartphones. They found that just over half of all participants locked their smartphones. Out of these participants, 68% used patterns, 23% used PINs and 9% used passwords to unlock their devices. Most participants who locked their smartphones also changed Android's default auto-lock setting. Almost all the participants who set their devices to lock immediately used patterns. Although PINs resulted in fewer mistakes, patterns were the quickest unlocking mechanism. Unlocking time appears to be important for participants, as many seemed willing to tolerate a few mistakes if it meant spending less time unlocking their device. Developers need to keep smartphone unlocking habits in mind when developing new models and features. They must prioritize convenience and speed to convince people to lock their devices.

Can your device's digital fingerprint help secure your online accounts?

Web browsers reveal a lot of information to websites about visiting devices. They share device characteristics like screen resolution, operating system and even location information. All this information creates a kind of "device fingerprint" that can actually identify the device. Alaca and Van Oorschot were curious about how to integrate device fingerprinting into existing authentication models, like passwords. They determined that strong device fingerprints are made of different kinds of characteristics. Reliable characteristics like screen resolution provide stable device identifying information. However, they are not very distinguishing, as all devices of a particular model may share them. Distinguishable characteristics, like specific hardware sensor calibration, can more confidently identify a device, but change over time. Device fingerprints must be composed of a good balance of reliable and distinguishable characteristics to hold up against different attacks. A score-based system would allow some characteristics to change while still confidently identifying the device. Ultimately, device fingerprinting can strengthen authentication without usability burdens.

Why do people follow computer security advice?

Although computer security is important, many people do not follow security advice. Why is this so? Decision-making is complex and can be a mix of rational choices and social motivations. Fagan and Khan wanted to understand what motivates people to follow or disregard computer security advice. The researchers sent a survey to 805 people from an online worker marketplace asking if they update software regularly, use a password manager, use two-factor authentication or frequently change passwords. They created groups based on whether participants followed each security practice or not. The researchers then looked at how each group perceived the benefits, costs and risks of their decision. Participants expressed that their own decision is the best for them. Their perceptions of risk and convenience differed: groups following the advice rated risk higher than groups not following the advice. Social motivations seemed to play a small role in security decision-making. Communicating the conveniences and social motivations of security practices may encourage people to follow security advice.

More warnings about phishing are always better, aren't they?

Social engineering attacks exploit trust to deceive people into directly giving sensitive information away. Awareness messaging that primes or warns people about social engineering can bring attention to the issue; however, the efficiency of primings and warnings is unknown. Junger, Montoya and Overink wanted to see if primings and warnings could prevent social engineering attacks. They asked 278 Dutch people at a shopping center to fill out a questionnaire asking for various types of personal information. They divided participants into three different groups. One group had primings throughout the questionnaire, one had a warning at the beginning and one directly shared their information without priming or a warning. The researchers found high levels of disclosure for all groups, with over 80% of participants disclosing personal information when asked. It is possible that participants were not in a security mindset while out shopping and were thus more likely to disclose. Primings and warnings were not effective at reducing disclosure. Conversely, warnings may even increase disclosure rates. Poorly designed awareness campaigns could have adverse effects. Intervention strategies must be evidence-based and attentive to audience characteristics.

Are comics better than text for security advice?

Computer security is complicated. Because humans are an important part of the security equation, it is essential to have effective teaching tools to help them. Zhang-Kennedy, Chiasson and Biddle wanted to explore effective ways to educate people on cybersecurity issues. They found that image-based communications and useful metaphors improved cybersecurity understanding. From these findings, they devised a witty, interactive online comic series. They tested the effectiveness of their comic by measuring the computer security understanding and risk behaviours of 52 participants before and after interacting with the comic. Participants had a poor initial understanding of computer security, but their understanding improved after interacting with the comic. Furthermore, the interactive aspects of the comics effectively communicated the benefits of following cybersecurity advice, prompting participants to change their behaviour. The participants self-reported positive behaviour changes, like changing their passwords and updating their software. Simplified, engaging content helps people learn more easily. Interactive comics seem to be an effective teaching tool when it comes to computer security.

Would backdoor encryption be safe?

In 1997, a group of computer scientists and security experts launched a formal debate on whether communications service providers should guarantee lawful exceptional access to all forms of data to law enforcement. This debate is still a hot topic, so many of the same researchers reconvened in 2015 to re-examine the implications of exceptional access. Since the authors did not have access to a complete statement of how governments would conduct exceptional access, they came up with different hypothetical scenarios. They argued that allowing exceptional access would pose greater security risks, hinder innovation, and have commercial and political consequences. Furthermore, it would increase operational complexity, as a large number of institutions would have to ensure authentication, organisation and information transfer. The authors also argued that beyond hindering security, more complexity also increases costs for everyone. Although exceptional access could be beneficial to law enforcement, it may also have numerous serious consequences that must be taken into consideration.

Are people paying ransomware with Bitcoin?

Cryptoware is a malicious software that encrypts valuable files on a victim's system and demands Bitcoin payments as ransom. Bitcoin is reputed for being supposedly anonymous and regulary used in cybercrime. However, it does not completely guarantee anonymity. Liao, et al. wanted to know more about CryptoLocker malware targets. They identified 795 ransom payments from data on 968 Bitcoin addresses. Because CryptoLocker targeted professionals, the researchers assumed that recipients processed payments during business hours, allowing them to assign countries to the payers by time zone. Ransom payers seemed to be mostly in the United States, Great Britain and Australia. The researchers also found relationships between CryptoLocker Bitcoin addresses and money laundering activities, such as the use of Bitcoin fog mixer. CryptoLocker was responsible for over 300 million dollars in losses and cost victims more than $300 000 in a single quarter. Not enough is known about ransomware such as Cryptolocker and more work must be done in order to develop effective countermeasures.

How can criminals trust each other online?

Criminals use many well-known risk reduction strategies in physical illegal markets, but we know less about those used in the online criminal world. Holt et al. wanted to know the different ways buyers and sellers minimize risk in online illegal markets. They looked at 1889 conversations in Russian and in English from 13 public web forums on the trade of stolen personal and financial data. The researchers identified different strategies that online offenders used in these forums. Forum members used informal strategies to minimize risk, like private messaging and using the feedback system. Administration used formal strategies, like reviewing comments and upholding forum rules. Forum members seemed to be mainly concerned with internal risks, such as dishonest sellers and buyers. Administrators are supposed to protect and ensure the quality of the forum, but can fail at fulfilling all their formal duties. Slander attacks by police that disrupt the forum harmony seem efficient, but may not work in forums with tighter regulation. Undercover forums operated by law enforcement may be more effective at collecting forum data and creating distrust between members of these online illicit markets.

Fake News: The narrative battle over the Ukrainian conflict

Fake news stories can be more than mere fibs; they can be an important tool in waging information warfare in the modern Internet-connected world. Russia devotes significant resources to control both its mainstream media and discussions online. Russia seems to be engaging in a new level of information warfare by managing both the mainstream media and Internet discussions during the Ukrainian crisis. State backed fake news stories may have shaped national and international perceptions of the conflict in Ukraine. Fake news stories can resemble propagandistic entertainment and might use shocking material, accusations, dramatic music and misleading images to support pre-fabricated ideas. These stories can form part of "strategic narratives", or planned messaging that presents a position to influence audience perception.

Khaldarova and Pantti were interested in exactly how people receive and refute fake news stories. They conducted a study in two parts. The first part involved finding alleged fake news stories. The researchers collected 339 stories in Russian and 260 in English from 'StopFake', a website dedicated to debunking fake news. They identified about 30 fake news reports in both languages that originated from Channel One; a popular, state-run Russian television network. They then selected the 10 most popular of these Channel One stories on social media. The researchers looked for strategic narratives in these stories and interpreted the StopFake debunking of these reports as counter-narratives. For the second part of their study, the researchers analysed 6043 tweets related to the selected 10 stories containing links to the Channel One reports or to the StopFake counter-reports. The researchers then determined whether people seemed to trust or distrust the news story they were tweeting about.

People in general seemed skeptical of these news stories and aware of their strategic narratives. The researchers found that 50.7% of the tweets distrusted Channel One news. Most of the tweets were from Ukraine, Russia and the United States. All the comments showing trust were from Russia, while those from Ukraine exclusively expressed distrust. The most frequently expressed emotions in distrusting comments were sarcasm and disgust with Channel One propaganda in general or at the content of a news story specifically. Those seeking to debunk fake stories included various messages highlighting the lack of evidence, inconsistencies and counter-narratives that contradicted or disproved the fake news.

Diffuse media environments contain many opinions, rapidly changing stories and response-provoking material. Strategic narratives require ongoing engagement and interaction to survive diffusion by the numerous opinions of a broad audience exposed to multiple news sources. In order to keep shaping the perception of emerging events, they aim to provoke emotional responses and blur the line between reality and fiction.

In this case, Russian news narratives often included emotive references to World War II. They focused on the connections between the atrocities witnessed by and achievements of the Russian nation in the conflict against Nazi Germany. They often used terminology referencing the war or directly linked persons with those known from the war, such as Nazi collaborators. These narratives also encouraged stereotypes and put complex issues into binary terms, such as 'the West vs Russia.'

It is possible that Russian fake news stories represent a state strategic narrative instead of factual reporting on events. Despite public awareness of strategic narratives, these environments blur the line between what is real and what is not. Even irrational and unfounded strategic narratives can have an effect. A news environment that requires constant independent fact checking may fuel confusion. The public is not a passive consumer of incredible news. It will both contribute to and contest reports to either perpetuate or dissolve narratives by contributing sources, testimony and video to an increasingly confusing mix of media.

Fake news can form part of strategic narratives but can also fuel confusion. Even when irrational and unfounded, they can blur the line between reality and fiction.

Khaldarova, I. & Pantti, M. (2016). "Fake News: The narrative battle over the Ukrainian conflict." Journalism Practice, 10(7), 891-901, doi: 10.1080/17512786.2016.1163237

Android users in the wild: Their authentication and usage behavior

Smartphones are an important part of everyday life for many people. They allow quick and easy access to online services such as banking, social media and even dating. Consequently, these devices can contain massive amounts of personal information. Many people use lock screens on their smartphones to protect their personal information and accounts; but exactly how they use them is unclear. The security and usability of unlocking mechanisms in controlled conditions are well understood. However, how different smartphone users choose unlocking mechanisms and how they work in the real world is still unknown.

Mahfouz, Muslukhov and Beznosov wanted to learn how people use Android smartphone devices 'in the wild.' They studied how these people lock, unlock and use their devices during their day-to-day activities. The researchers recruited 41 participants at the University of British Columbia and via social media. Participants installed an invisible monitoring tool from the Google Play store on their Android devices. The researchers monitored each device for 20 days, from December 2014 until March 2015. They collected information on the use of different unlocking mechanisms; such as number passwords (PINs), passwords and patterns. They also collected information on the 'auto-lock timeout' feature that automatically locks an unused device after a time period. Increasing the auto-lock gives attackers more time to get to an unlocked smartphone. Decreasing it or setting it to activate immediately by touching the power button reduces this window of opportunity.

Among the 41 participants, 22 locked their smartphones. Out of these 22, two locked their devices with passwords, five with PINs and 15 with patterns. Participants who locked their smartphones interacted with their devices more often and for longer than those who did not. Most of these participants also changed auto-lock time: only six did not. Seven participants increased it to 30 minutes, while nine set it to activate immediately. Seven out of the nine participants also used patterns to unlock their smartphones. Patterns take very little time and effort, so users can easily afford the inconvenience of immediate auto-lock timeout. On average, the auto-lock feature was responsible for 11% of all device locking, and most devices remained unlocked for over a minute when not in use.

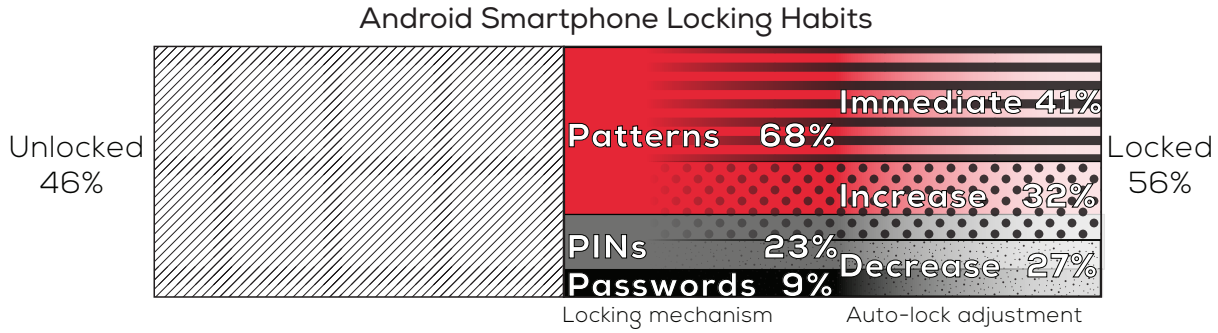

Patterns were the fastest unlocking mechanism. Participants entered patterns at a similar speed to PINs but significantly more quickly than passwords. Patterns unlocked their devices in 1.7 seconds on average, compared to 2.5 seconds for PINs and 4.1 seconds for passwords. However, PINs had the highest success rate. PINs were 5.5 times less likely to have mistakes than patterns and 6.3 times less than passwords. They also had the least amount of consecutive mistakes, but repeated mistakes were uncommon in general. Unlocking time seems to be important to smartphone users. Those who used patterns or PINs spent an average of 100 seconds each day unlocking their device compared to password users who spent 200 seconds doing so.

Participants appeared willing to tolerate a few mistakes here and there, as long as it meant they could still quickly unlock their devices. Due the low rates of repeated mistakes, consecutive errors could be an indicator of a guessing attack. Smartphone users seem to be picky about their unlocking methods and do not want to waste time. They appear to be willing to balance the effort and time needed to unlock their devices with the immediacy of automatic lock. Developers must consider the most appealing features of popular unlocking mechanisms when creating new models. They should consider models that are quick and convenient, even if they result in a few more failed attempts.

The speed and convenience of locking mechanisms appear to influence people's locking habits.

Mahfouz, A., Muslukhov, I. & Beznosov, K. (2016). "Android users in the wild: Their authentication and usage behavior." Pervasive and Mobile Computing, 32, 50-61, http://dx.doi.org/10.1016/j.pmcj.2016.06.017

Detecting Malicious URLs Using Lexical Analysis

A Universal Resource Locator (URL), or web address, identifies pages and content on the web. Attackers can use malicious URLs as a part of online attacks to harm visitors. A common technique used to filter out malicious URLs involves blacklisting known harmful webpages. However, attackers can subtly change web addresses, rendering blacklists useless. Companies often use expensive and complicated technologies to determine malicious addresses but do not provide the resulting lists freely. Heuristic-based techniques are an alternative that can identify newly created malicious websites in real-time. However, they are not infallible and could be anticipated by attackers. Because detection techniques are time and resource intensive, they are generally limited to the classification of URLs or to a specific attack. For these reasons, there is a need for approaches that can better detect and categorize malicious URLs.

Machine learning techniques offer a potential solution. Machine learning techniques can already classify malicious websites by their URL, content and network activity. Mamun et al. investigated the use of these techniques for identifying bad URLs. The researchers collected about 110 000 URLs known for spam, phishing, malware distribution and defacement, as well as benign URLs. They then created a program that could identify 79 features of the words used in the URLs and taught it to detect the implications of these lexical features. The lexical features included elements such as the length and the characters used to compose the URL. The program found five sets of lexical features that help to identify bad URLs. The researchers also looked at six techniques used maliciously to mask or 'obfuscate' harmful URLs. Finally, they created and tested malicious URL classifiers using the selected features and obfuscation techniques.

The selected features and classifiers appeared to be highly accurate. The classifiers detected 98% of spam and malware URLs and 99% of defacement URLs. These methods provide a measure of confidence on whether a URL is malicious or not.

The methods developed showed very high detection results. Security professionals can use machine learning to detect malicious URLs including defacement URLs. They can use the identified classifiers to augment blacklist techniques and increase protection against malicious URLs.

Machine learning can provide an additional tool for identifying bad web addresses.

Mamun, M.S.I., Rathore, M. A., Lashkari, A., H., Stakhanova, N. & Ghorbani, A. (2016). "Detecting Malicious URLs Using Lexical Analysis" in: Chen J., Piuri V., Su C., Yung M. (eds) Network and System Security. NSS 2016. Lecture Notes in Computer Science, vol 9955. Springer.

Device Fingerprinting for Augmenting Web Authentication: Classification and Analysis of Methods

Web browsers explicitly or implicitly reveal device information to websites. Browsers explicitly share information such as the screen resolution, operating system or local time of a device. They can also implicitly provide other types of software or hardware information. All this information combined can uniquely identify a computer or smartphone; creating a kind of "device fingerprint." Device fingerprinting is used for many purposes. For example, advertisers use device fingerprinting to track people online. The ability to uniquely identify devices means that device fingerprinting has the potential to augment existing online authentication methods; such as passwords.

Alaca and Van Oorschot looked at integrating device fingerprinting techniques into authentication methods without disrupting the user experience. They identified and classified 29 available device fingerprinting mechanisms. Most were known mechanisms and some were new. The researchers assessed how suitable and practical each mechanism would be for improving user authentication.

A good device fingerprint must contain a well-balanced combination of both reliable and distinguishable identifying characteristics. A reliable characteristic is stable and does not change, like model number or screen size. Because all devices of a particular model have the same fixed hardware, model-relatable specifications alone are not enough to distinguish one device from another. Furthermore, if a smartphone model is only available in a certain country or zone, there could be a similar relationship between the model, location and time zone. Attackers can determine a portion of device fingerprints from the relationships between commonly shared characteristics. There are other details that are more distinguishable. Characteristics such as hardware sensor calibration or specific browser settings are more random and harder to guess. However, many of these characteristics change. For example, these changes can occur after operating system upgrades. It is possible to improve fingerprint longevity by combining multiple fingerprint characteristics to create a scored level of confidence or trustworthiness. Since device fingerprints can change over time, a score based system would allow elements to change while still identifying the device.

Authentication with a Different Device Fingerprint

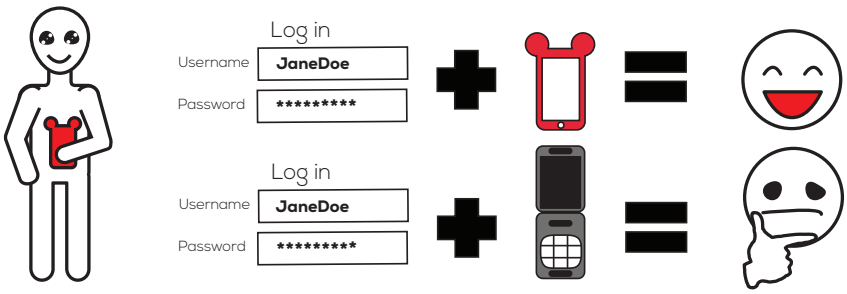

The researchers also identified different threats based on the assumption that an attacker would eventually learn the properties of a device fingerprint. They then determined how device fingerprinting would hold up against different attacks. When used to augment passwords, device fingerprinting can completely stop conventional password-guessing attacks. It can also reduce the success of password and fingerprint guessing attacks. However, it can be more difficult for device fingerprinting to protect against targeted attacks in which the attackers have device-specific information. Phishing attacks that capture both the password and certain aspects of the fingerprint are also problematic. These attacks are less likely to succeed if the device fingerprint consists of multiple, variable characteristics. Finally, if an attacker hijacks an existing session, they may easily retrieve the device fingerprint. Fingerprints should therefore be difficult to forge and used to authenticate sessions not only once at the beginning, but also multiple times throughout.

Fingerprinting has the potential to strengthen user authentication. Using multiple carefully chosen characteristics for device identification makes it more possible to identify the device in a way that is difficult to guess. Device fingerprinting mechanisms do not require additional attention from device users and seem to pose no additional usability burden. Consequently, they offer promise as an additional factor for authentication.

Fingerprinting as an additional factor has the potential to strengthen authentication.

Alaca, F. & van Oorschot, P.C. (2016). "Device Fingerprinting for Augmenting Web Authentication: Classification and Analysis of Methods." Proceedings of the 32nd Annual Conference on Computer Security Applications Conference, Los Angeles, CA, USA.

Why Do They Do What They Do?: A Study of What Motivates Users to (Not) Follow Computer Security Advice

Computer and data privacy are important, but not everyone follows best security practices. Researchers have long been curious as to what motivates individuals to follow computer security advice. Decision-making is a complicated process. People might make decisions by considering the benefits or consequences of an action. In a rational process, humans make informed choices based on the perceived costs or risks of their actions. Social motivations, such as the desire to please others, may also influence their decision. It is important to understand what motivates the decision-making process in order to help people make better security decisions.

Fagan and Khan wanted to clarify computer security decision-making. They used a short screening survey to gather 805 survey participants on Amazon Mechanical Turk; an online worker marketplace. They wanted to know if participants followed certain security recommendations and why. They looked at four common recommendations: updating software, using a password manager, using two-factor authentication and changing passwords. They then created 8 groups of 30-40 participants who completed follow-up surveys. They formed "yes" and "no" groups for each security recommendation based on whether participants said they followed the advice or not. The researchers then considered the rational and social aspects of each decision. They looked at the perceived benefit, cost and risk of following a recommendation or not.

Participants felt as if their decision brought them the most benefit. Whether they followed advice or not, users justified their decisions by saying the benefits outweighed the risks or inconveniences. Those who followed security recommendations rated the risk more highly. For them, the security benefits outweighed the cost of the inconvenience. Conversely, the convenience benefits outweighed the security risks for the people who did not follow security recommendations. They rated the convenience cost higher than those who did follow the advice. Participants felt that their current decision provided them with a greater benefit than they would get from changing. Social motivations seemed to have very little influence on computer security decisions.

It may be tempting to assume that people who choose not to follow security advice simply need to be better educated on the benefits of doing so. However, it is important to consider how people perceive the value of these benefits. Communicating convenience may be a greater motivator than risk for those not following advice.

Differences in Security Mindsets

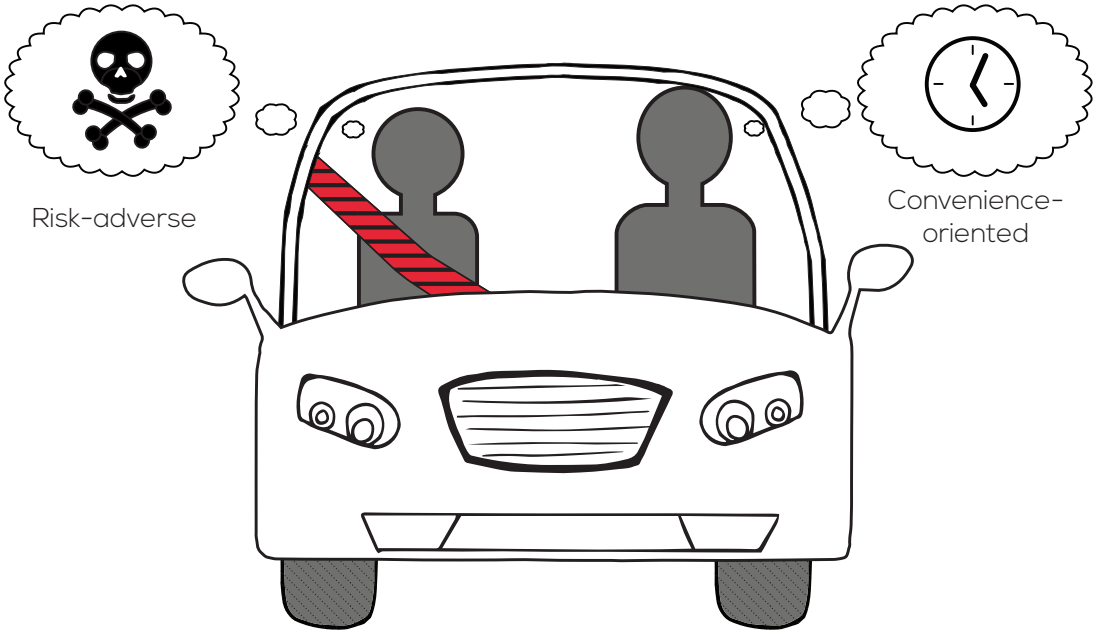

Not everyone is motivated by recommendations focused on risk. Convenience and social reasons might have more effect.

Fagan, M. & Khan, M. M. H. (2016). "Why Do They Do What They Do?: A Study of What Motivates Users to (Not) Follow Computer Security Advice." Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS 2016), Denver, CO, USA.

Priming and warnings are not effective to prevent social engineering attacks

Social engineering attacks exploit social interaction science to efficiently deceive people so that they directly give information away. Concepts of trust are firmly embedded in culture and tricksters are nothing new. This raises questions about how to protect the public against social engineering attacks. Awareness messaging can bring attention to social engineering and online dangers. However, it is unclear whether these awareness techniques are actually effective in protecting people.

Junger, Montoya and Overink wanted to see if awareness messaging such as priming and warnings effectively prevent social engineering attacks. The researchers distributed three slightly different questionnaires to a total of 290 Dutch people at a shopping centre. All of the questionnaires asked participants to disclose their email address, a part of their bank account number and the details of their last online purchases. The one questionnaire exposed (primed) participants to the topic and included questions on phishing and cybercrime issues. The second questionnaire started with explicit warnings not to give out personal information. The remaining questionnaire only asked for the email, banking and purchase information. Around 90 people answered each version, forming three groups. The researchers then calculated the final risk score of each group based on the disclosed personal information.

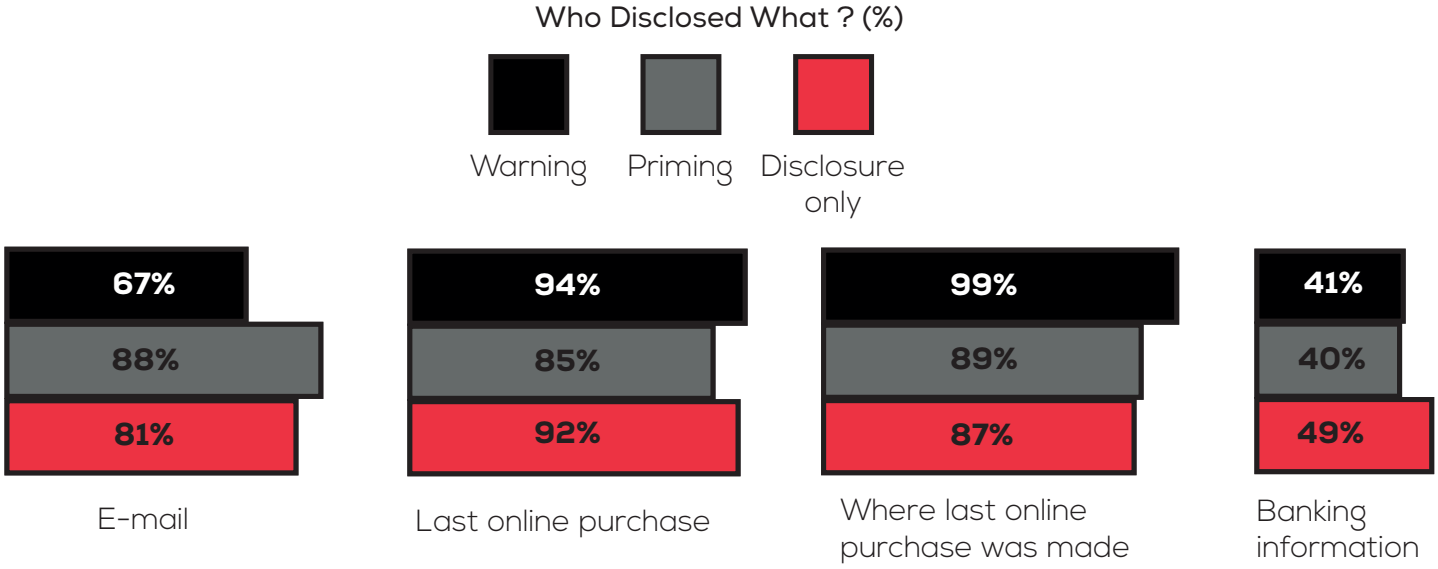

The researchers found relatively high levels of disclosure across all groups. Around 80% of all participants disclosed personal information when asked. These results are not shocking, as people generally trust each other. The person delivering the survey was also a young, friendly-looking man; which may have increased disclosure rates.

The results suggest that neither the priming questions nor the warnings were effective in reducing disclosure. Surprisingly, those who received a warning were actually more likely to share information about their last online shopping locations. It is possible that participants did not view the priming or warnings as personally relevant or pay attention to them. They may not have given priority to security when someone interrupted them while shopping to fill out a questionnaire. This may also have resulted in greater disclosure.

There is a pressing need for greater education on social engineering attacks. More work on the effectiveness of cybersecurity awareness campaigns is needed. It is important to determine educative priorities as teaching everything at once may not increase awareness. Poorly designed awareness or warning campaigns could actually produce adverse effects. Certain specific warnings may increase disclosure of some types of personal information and require extra attention. Effective interventions should be evidence-based and sensitive to the particular characteristics of the audience.

Poorly designed awareness or warning campaigns could actually produce adverse effects.

Junger, M., Montoya, L. & Overink, F.-J. (2016). "Priming and warnings are not effective to prevent social engineering attacks." Computers in Human Behavior, 66 (2017), 75-87, http://dx.doi.org/10.1016/j.chb.2016.09.012

The Role of Instructional Design in Persuasion: A Comics Approach for Improving Cybersecurity

The administration of systems plays an important part in security. This means that people using, and by necessity administering, home computers must make various security decisions. However, computer security is complex and home users are generally untrained and uninterested. Although it is important to make security inherent to computers, people may always be a part of the security equation. Even if we make security easier to use, it remains essential to educate home users to make better security decisions. Advancing this cause requires that we know how people understand security, including different kinds of attacks and protection methods. We also need to know what kinds of educational material can efficiently alter people's perceptions and behaviours to improve their security.

Zhang-Kennedy, Chiasson and Biddle looked at ways to effectively educate people on home computer security and persuade them to improve their security habits. First, they designed security infographics to test how well certain types of images and metaphors contributed to learning. They discovered helpful metaphors that improved cybersecurity understanding. They found image-based communications to be a more effective teaching tool than text alone. Based on these findings, they designed a witty, interactive comic series. The researchers used different instructional design principles to come up with interactive comics about password-guessing attacks, malware protection and mobile online privacy. They then used eye-tracking technology to optimize the comics and make adjustments to the flow and length as necessary.

To test the effectiveness of their comic, the researchers recruited 52 study participants at Carleton University to interact with their comic. They interviewed the participants to find out how much they knew about the three security areas that the comics address. They then used questionnaires and interviews to test the participants' security understanding before and after interacting with the comic. One week later, they looked to see if the participants had changed their behaviour.

Initially, the participants had a poor understanding of computer security. Many had difficulties distinguishing between different kinds of attacks. Others thought they were unlikely targets of cyberattacks. This may affect their ability and motivation to practice safe behaviour online. The comics followed four principles that link instructional design and persuasive technology principles.

1. Simple designs that support easy reading and navigation make persuasive learning tools more effective.
2. Interactive moments of reflection allow the readers to pause and reflect on their learning progress.
3. Characters, story, humour and conversational language increase immersion.
4. Metaphors that build on existing knowledge help readers create strong mental models of what they are learning.

After interacting with the comics, the participants knew more about cybersecurity. The simplicity of the comics reduced the cognitive effort required to understand the complex security content. Furthermore, the interactive aspects of the comics provided insight into the benefits of following their security recommendations. These aspects effectively persuaded participants to follow the advice given in the comics. Overall, the participants self-reported more positive behaviour changes after interacting with the comics. They reported changing their passwords, updating their security software and sharing advice from the comics with family and friends.

Simplifying content through images and metaphors in an entertaining comics-based approach can help people overcome the challenges of learning about cybersecurity. Furthermore, this type of approach may even be helpful in other fields. Projects that aim to teach, inform and modify behaviour could benefit from an evidence-based immersive comics approach.

Evidence-based, interactive comics are an effective teaching tool for cybersecurity.

Zhang-Kennedy, L., Chiasson, S. & Biddle, R. (2016). "The Role of Instructional Design in Persuasion: A Comics Approach for Improving Cybersecurity." International Journal of Human–Computer Interaction, 32:3, 215-257, doi: 10.1080/10447318.2016.1136177

Keys under doormats: mandating insecurity by requiring government access to all data and communications

There is global debate about requiring communications service providers to guarantee law enforcement legal access to all data, including encrypted data. This debate highlights the difficulties involved in balancing privacy and security, as the privacy of citizen's communications on private networks is questioned by governments when those communications are vital to security investigations. This is not a new debate. In the 1990s, the United States government proposed a law requiring that all encrypted data storage and communication systems be designed for "exceptional access" by law enforcement agencies for security purposes. In response to this proposal for exceptional access, a group of computer scientists and security experts conducted a study in 1997 on the potential ramifications. The authors found that it would have been too difficult and expensive to implement. Many of the same authors reconvened in 2015 to examine the modern incarnation of the debate in the current operational and technical environment.

An obstacle for the new study was the lack of a complete statement of how governments would conduct exceptional access. Accordingly, the authors used different potential scenarios to illustrate how exceptional access would work. The authors argued that allowing exceptional access would pose greater security risks, hinder innovation, and have commercial and political consequences. Exceptional access would jeopardize current security best practices because it would inevitably lead to more opportunities for hackers to intercept communications. Creating exceptional access would also increase operational complexity. A large number of institutions would have to securely and safely negotiate attacks on the authentication, organisation and information transfer of lawful information access. Increasing the number of stakeholders required would increase the number of targets for criminals. If a third party were trusted with the required information that a law enforcement agency would need to access private communications, then a malicious insider could take advantage of that trust. The authors argued that complexity generally hinders security and increases costs for everyone involved. For example, the United States government would have to increase staffing to accommodate their requirements. Software companies would also be laden with extra costs to engineer their software along the mandated guidelines. Allowing the government exceptional access also limits how communication services providers can construct their software. It also puts communication providers at a competitive disadvantage with firms in other countries.

Managing Global Scale Exceptional Access

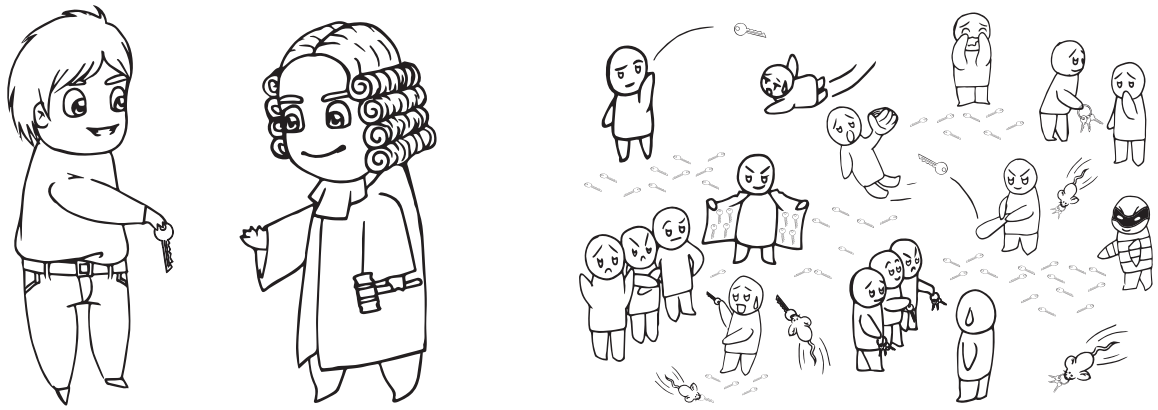

Expectation Reality

These arguments go a step beyond just debating if law enforcement agencies should or should not have access to private communications. They deal with the practical realities of exceptional access. Creating exceptional access, although beneficial for law enforcement agencies, could result in additional costs and security risks. The issues highlighted by the researchers are exacerbated by the growing amount of services and connectivity and related cybersecurity issues. Consequently, if governments and law enforcement agencies are going to mandate exceptional access, they should be aware of the possible consequences.

The organizational practicalities of implementing exceptional access to encrypted communications makes them impossible to secure, even if the technology works.

Abelson, H., Anderson, R., Bellovin, S. M., Benaloh, J., Blaze, M., Diffie, W., Gilmore, J., Green, M., Landau, S., Neumann, P. G., Rivest, R. L., Schiller, J. I., Schneier, B., Specter, M. A., & Weitzner, D. J.(2015)."Keys under doormats: mandating insecurity by requiring government access to all data and communications." Journal of Cybersecurity, 1(1), 69-70.

Behind Closed Doors: Measurement and Analysis of CryptoLocker Ransoms in Bitcoin

CryptoLocker is a family of ransomware using Bitcoin as a payment method. This type of malicious software encrypts all valuable files on a victim's system. Once all the files are encrypted, the program lets the victim know about the infection by sending a message. In order to decrypt the files, the victim is asked for money. The data will remain encrypted until the ransom is paid. Usually, these ransoms are paid using Bitcoin, which is a decentralized cryptographic currency. Known for its pseudo-anonymity, this digital currency gained popularity over the past few years and is used in different types of cybercrime, such as financial fraud. Money transfers in Bitcoin are virtually impossible to reverse and difficult to trace. All confirmed transactions made in Bitcoin are visible to the public in something called a blockchain. The transactions are issued with a Bitcoin address instead of with a name. Initially, this intended to preserve the user's anonymity, because nothing was supposed to relate a person to their Bitcoin address. But, with the existence of the blockchain, this technique has been proven ineffective in ensuring anonymity.

Liao, et al. wanted to understand who were the targets of the CryptoLocker malware. They used information collected from forums online to gain insight into this new digital threat and a better knowledge of ransomware criminal enterprise strategies.

During their study, they discovered a total of 968 Bitcoin addresses used by the CryptoLocker operators. Initially, they found two victims had posted addresses on Reddit, a social news networking service. These addresses were considered as seeds. The remaining addresses were found by analyzing Bitcoin transactions linked to these seeds. The researchers then collected data about the payments and identified 795 ransom payments. CryptoLocker targeted professionals and so the researchers assumed that the payments were processed during business hours. This allowed them to assign a country of origin to the payers. For example, the period from 9:00 a.m. to 5:00 p.m. corresponds to 17:00 UTC to 1:00 UTC in the Pacific Time Zone, which is UTC-08:00. Therefore, a payment made between 9:00 UTC and 17:00 UTC (UTC±00:00) is more likely to have come from Great Britain.

The results indicate that the top three countries from which ransoms were paid are the United States, Great Britain and Australia. In addition to this, the researchers were able to extract eight communities from the cluster of 968 Bitcoin addresses using an algorithm to detect communities in a network. It provided a better understanding of the CryptoLocker financial infrastructure. In some of the communities identified, they discovered relationships between the Bitcoin addresses used with CryptoLocker and those involved in money laundering activities. They also found the use of Bitcoin fog mixer; a service that mixes and randomizes outgoing transactions for a percentage fee. This kind of service can assist cybercriminals to launder the stolen Bitcoin. In the center of one of the communities, they were able to identify an address belonging to BTC-e, one of the largest currency exchange services for Bitcoin. The losses related to CryptoLocker within four months were estimated at $310,472.38. Furthermore, they identified possible connections between different forms of crime involving Bitcoin; such as phony black market scams and ransomware.

Ransomware is a significant problem, with Cryptolocker alone costing victims more than $300,000 in a single quarter. The current understanding of ransomware and the role of Bitcoin in online crime is insufficient for the development of effective countermeasures.

The Bitcoin payed in ransomware is significant.

Liao, K., Zhao, Z., Doupé, A., Ahn, G. (2016). "Behind Closed Doors: Measurement and Analysis of CryptoLocker Ransoms in Bitcoin." IEEE Computer Society. Published in Proceedings of the 2016 APWG Symposium on Electronic Crime Research, eCrime 2016. doi: 10.1109/ECRIME.2016.7487938

Examining the risk reduction strategies of actors in online criminal markets

The risk reduction strategies used by criminals in illegal markets in the physical world are well documented. Their practices are shaped by their perception of risk. Online and physical environments are considerably different. Therefore, individuals participating in online illegal markets employ different risk reduction strategies. They use resources available to them to protect themselves from formal and informal risks associated with law enforcement and with their criminal peers.

Holt et al. conducted a study focusing on the different mechanisms used by buyers and sellers participating in online illegal data markets to minimize formal and informal risks. They analysed content found in a sample of 1889 conversations from ten Russian and three English public web forums on the trade of stolen financial and personal data. They only selected conversations related to the sale and trade of personal and banking information. They found three of the forums through search engines and identified the rest through these forums. Eight of the sites were entirely openly accessible, while the other five required registration to access the market sections.

The researchers identified many strategies to minimize risks for both sellers and buyers in these forums. Forum members are mainly concerned with internal risks, such as disreputable or competing sellers, disgruntled buyers or dishonest members. To ensure successful transactions, forum members use informal strategies negotiated between them for their mutual protection. Since cheated members cannot complain to the authorities because of their implication in illicit trade, forum administrators implement formal strategies to protect the forum. These strategies help preserve its credibility, proper operation and regulate transactions between members.

Informal strategies used between buyers and sellers	
Electronic payment and fake accounts	Immediate payment and rapid dissemination of goods. Avoid face to face interaction and provide privacy and anonymity.
Private messaging system	Privacy for negotiating terms of trade.
Public posting of positive or negative feedback	Help determine the reliability and reputation of a seller. Discredit and harm the reputation of dishonest sellers.
Formal strategies used by the forum administrators and moderators	
Administrator intervention	Dictate rules of participation in the forum. Review and deletion of posts.
Escrow service system	Hold payment on behalf of the buyer until the seller releases the merchandise.
Review and product testing	Ensure the quality of products and services offered.
Threat management system	Ban or block of users attempting to disrupt the market. Demand proof of any wrongdoing.

In reality, not all forums offer escrow or product testing services, and their application is often inadequate. Furthermore, administrators are not always active, forcing buyers and sellers to use feedback posting and reviews to find reliable transaction partners. Even though external risks such as legal authorities are not perceived as the main threat to illegal forums, members still use anonymity strategies to ensure their privacy.

The use of slander attacks against forums, where law enforcement attempts to disturb the market by making false complaints against other members, is well documented as an efficient strategy. But, it might not be as effective in regulated forums where administrators can ban disruptive members. Using undercover identities to create forums for stolen data could be more valuable for gathering intelligence, identifying key members, and tracking the behaviour of participants, along with building cases against entire networks. The creation of stolen data forums by law enforcement as an investigative strategy would provide sufficient evidence to build solid cases, and would furthermore create distrust amongst participants and consequently affect the supply and demand of these illicit markets.

Trust in the forum is used as a proxy for trust in other criminals. Reducing the trustworthiness of criminal forums could be effective.

Holt, T. J., Smirnova, O., Chua, Y, T., & Copes, H. (2015). "Examining the risk reduction strategies of actors in online criminal markets." Global Crime, 16 (2), 81-103, doi: 10.1080/17440572.2015.1013211

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 relationship brokers' forum, quarterly knowledge digests, Konnect - online knowledge-sharing platform, public websites and professional development program.

Workshops & Events

SERENE-RISC @ GoSec

August 2017
Montreal

Bi- Annual Workshop

October 2017
Ottawa

Knowledge Digest

More than 90 pieces of research on cybersecurity summarised.

Opportunities available to network partners to help empower future professionals with practical knowledge translation skills through mini-scholarships.

Konnect Platform

Find the presentation videos and slides from workshops and tutorials and much more.

Over 800 selected articles and research pieces on cybersecurity for and by Canadian experts.

Website

Find out the latest news, tips and job postings at:
www.serene-risc.ca
Connect with us on Twitter, Facebook and Linkedin:

@SERENE_RISC /serenerisc /serene-risc

Cybersecurity Public Awareness Tools & Library Outreach

Cybersecurity is an issue for everyone so we see a role for everyone in addressing the issue. We are aware of the tremendous role libraries play as educators and trusted sources of knowledge. We found that libraries are doing great work to support technology. They are doing such a good job that they are often the primary source of Internet access, technical support and practical information for not only the public generally, but those that might be at greater risk online. With the support of Public Safety Canada, we have developed interactive training tools using leading edge Canadian research provided by our national network of academic, government and industry partners. The training program is available free of charge to enable libraries and community centres across Canada to deliver instruction to the public on the best practices for staying safe and secure online.

The program consists of a number of flexible modules that can be either presented together as a course or combined to augment existing information technology training sessions where appropriate. The resources are freely available and include trainer guides, classroom handouts, planning tools and a multimedia website with videos, mini-quizzes etc. We are working with libraries across the country to deliver train-the-trainer sessions to enable a sustainable system of disseminating the best advice to the people who need it in a form they can access. Find out more at:

www.cybersec101.ca

The SERENE-RISC Cybersecurity Knowledge Digest

2017 Spring

Editor-in-Chief: Michael Joyce
Scientific Editor: Sonia Chiasson
Editor: Shannon McPhail

Abelson et al. *summarised by* Hayley McNorton *sponsored by* SERENE-RISC
Liao et al. *summarised by* Nedra Hamouda *sponsored by* SERENE-RISC
Holt et al. *summarised by* Chloé Majdalany *sponsored by* SERENE-RISC