

Do larger collections of child abuse material lead to longer sentences?

In most decisions, no. Judges are more concerned with whether the offender is physically abusing or on a path towards abusing a child.

5
page**Are people good at picking memorable passwords?**

Yes. People are able to tell which passwords will be hard to remember. Dictionary words, patterns and phrases are easier to remember.

6
page**What types of attacks could compromise a 911 Emergency call infrastructure?**

Attacks could target the disclosure of private information, the misallocation of resources or the unavailability of emergency services.

7
page**What are the attributes of the most effective telephone scams?**

Telephone scams impersonating an internal entity in the targeted organization are the most successful.

8
page**Can Bitcoin exchange addresses be identified in a transaction network?**

Exchange addresses can be identified through algorithmic methods based on the distribution of connections in the transaction network.

9
page**Do psychological differences make some people more vulnerable to being deceived online?**

Disposition, including motivation and personality, appears to be the key in how messages and experiences are processed.

10
page**Does the law protect the use of platform data to serve the public interest?**

No. Although platforms like Airbnb have generally allowed such access so far, they have many legal tools at their disposal to limit this use.

11
page**Are the sensitive data used to train deep learning algorithms safe from attacks?**

No, especially when attackers know which model and parameters are used, or when they are participants in the machine learning process.

12
page**Who is using the Tor network, and how?**

A landmark 2008 study showed computers in Germany, China and the US use Tor to browse the internet, but BitTorrent data accounts for most of the traffic across the network.

13
page**What is the reality of how system administrators manage software updates?**

Administrators get information from a variety of sources and prioritize functionality over security.

14
page

Do larger collections of child abuse material lead to longer sentences?

There has been a considerable increase in recent years of the number of child abuse images listed in Canadian police reports. Fortin et al. evaluated the effect of the reported number of images on the sentencing of child sexual exploitation material offenders, to determine whether larger collections are associated with more severe sentencing decisions. They analyzed 101 Quebec cases between 2002 and 2012, along with 97 other Quebec cases extracted from the CanLII database. Results suggest that judges are more concerned with whether the offender is, or is on a path towards physically abusing a child, rather than their consumption of illicit material. Image analysis, which is expensive and time-consuming, should be reconsidered if the objective is to impact sentencing.

Are people good at picking memorable passwords?

Is there any truth behind the assumption that passwords that are hard to guess are hard to remember? Alomari et al. compared passwords perceptions with electroencephalogram (EEG) data, which records brain activity. They presented their 77 study participants with different passwords, asking them to rank their memorability. Password recall was successfully linked to brain activity data. Passwords consisting of dictionary words, patterns, or phrases were perceived as more memorable, whereas passwords consisting of names or random letters and numbers were perceived as less memorable. People thus appear to have the ability to sense a password's usability upon its presentation and make decisions based on that information.

What types of attacks could compromise a 911 Emergency call infrastructure?

As the technological capabilities of 9-1-1 services continue to expand, systems that were once safely managed solely by the telephone company are increasingly connected to less well-regulated networks. Goebel et al. describe the technical structure of the 9-1-1 call-taking system in the United States in terms of its current susceptibility to breaches of confidentiality, integrity and availability. The current "Next Generation 9-1-1" upgrade effort in the US that is set to allow call centers to receive text messages, images and live video streams, could add further cybersecurity vulnerabilities to those already present in the telephone-based systems.

What are the attributes of the most effective telephone scams?

In the midst of growing concern about telephone scams, it remains unclear why people fall for them and how to combat the problem. Tu et al. ran a telephone phishing experiment on a university campus, testing a variety of scam attributes on a population of 3,000 Arizona State University staff and faculty to assess their individual effectiveness. Impersonating an internal entity was the attribute most linked to attack success. On the contrary, manipulating the type of motivation, voice production, voice accent and caller name did not result in a higher success rate. Given the success of impersonation techniques, scam prevention efforts should prioritize impersonation countermeasures.

Can Bitcoin exchange addresses be identified in a transaction network?

Anonymity is a major feature of Bitcoin, which allows criminals to conduct illegal trading activities and evade regulation. Identifying the addresses linked to Bitcoin exchanges, where users can trade Bitcoin for fiat money, is crucial to identify transactions of interest. Liang et al. developed a reliable method to identify the addresses of Bitcoin exchanges, using only transaction and user pattern data. They tested multiple identification algorithms on a network constructed from Bitcoin transaction data and successfully differentiated exchange addresses from general nodes. This study provides new tools to detect illegal behaviors in the Bitcoin network based on transaction and user patterns instead of user identities.

Do psychological differences make some people more vulnerable to being deceived online?

What are the psychological differences which make people more vulnerable to be deceived online? Norris et al. conducted a systematic review of the literature relating to how victims respond to fraudulent communications. They extracted 1036 articles, from which they selected and categorized 34. Messages displaying high sources of credibility, requiring a quick response and responded to on a smartphone were more effective in enticing potential victims. Dispositional factors, such as motivation and personality, appear to be the key mediating factor determining how message and experiential factors are processed. This study unfortunately highlights that empirical examinations of online fraud based on established psychological theory are rare. Reviewed papers lacked coherence and consistency in selecting appropriate psychological principles explaining an increased likelihood of victimization.

Does the law protect the use of platform data to serve the public interest?

Online platforms such as Airbnb publicly display a wide range of information on their users, which has become of significant interest to researchers, journalists and civil society organizations. Scassa analyzed statute and case law in order to assess the current legal situation around the scraping of publicly accessible platform data, using Airbnb as a case study. This study highlights how the ecosystem of platform data users could potentially be denied access through a series of legal tools. Economic and power imbalances in the data access litigation process create a risk that user perspectives and the public interest are not being represented as the law evolves.

Are the sensitive data used to train deep learning algorithms safe from attacks?

Deep learning, a machine learning method inspired by the information processing of a biological brain, is being successfully applied to many types of potentially sensitive user data. Do deep learning models leak individuals' data samples from their training sets? Nasr et al. present a comprehensive framework for the privacy analysis of deep neural networks. They experimented on three different datasets, based on a novel open-box attack approach where attackers take advantage of their prior knowledge of the algorithmic model and parameters in use. Results show that deep learning models previously assessed as not very vulnerable to closed-box inference attacks could be substantially more vulnerable to open-box attacks.

Who is using the Tor network, and how?

Tor directs its users' internet traffic through a three-hop path among a worldwide network of relays in order to encrypt and conceal it from surveillance. McCoy et al. set up their own Tor router and analyzed passing data to observe who is using the service and how. Results show that while internet browsing traffic comprises the majority of connections, BitTorrent activity uses a disproportionately high amount of bandwidth. Insecure protocols which transmit login credentials in plain text, from email servers for instance, were commonly observed and created an important risk for the initiating client's anonymity. Two thirds of running routers originate from two countries and 2% of routers transport 50% of the traffic, which could compromise the network's anonymizing properties.

What is the reality of how system administrators manage software updates?

How do administrators manage updates, and what factors impact how effectively they perform those updates? Li et al. studied US administrators responsible for managing updates in their organizations, administering a large-scale survey of updating practices and conducting interviews with administrators. Findings suggest that administrators rely on various update information sources, testing and deployment strategies. Internal policies and management also play an important role. Participants demonstrated that they prioritize functionality over security.

The Effect of Child Sexual Exploitation Images Collection Size on Offender Sentencing

There were more than 150,000 reports of child sexual exploitation on the Internet between 2008 and 2015 in Canada, with close to a fourfold increase during this period. Alongside the growth in Internet speed and computer storage, there has been a considerable increase in the number of child abuse images listed in police reports. Policing organizations analyze these images to identify victims and to provide information to courts in the hope of securing longer sentences. Image analysis is not only time-consuming and expensive for police organizations but it also takes time away from other investigations. It is still unclear what impact such exhaustive analysis has on the sentencing of offenders.

Fortin et al. evaluated the effect of the reported number of images on the sentencing of child sexual exploitation material offenders. Their primary goal was to determine whether larger collections of images are associated with more severe sentencing decisions. They were also curious as to which characteristics of child abuse material collections are considered when sentencing offenders accused of possessing such material.

The authors analyzed 101 Quebec cases between 2002 and 2012. They used statistical analysis to determine relationships between the number of images, sociodemographic information, details about the criminal events and prior convictions. The authors further studied in detail the text of 97 Quebec cases extracted from the Canadian Legal Information Institute (CanLII) database.

Results suggest that judges are more concerned with whether the offender is, or is on a path towards physically abusing a child, rather than their consumption of illicit material. In cases where a sexual crime occurred prior to or at the time of the offense, no effect from the number of images was observed. However, the number became relevant in the absence of sexual crimes, as the judges needed additional information to evaluate the level of culpability of the offender. In these cases, the number of images and the time spent in the creation and classification of the collection were used to separate the behaviors of an invested offender from those of the curious.

The results suggest that the number of images is not determinant in most decisions. Accordingly, the amount of time police officers devote to classifying unidentified images should be reconsidered if the reason for such effort is to impact sentencing. There may be valid reasons for detailed analysis of the information on seized hard drives (e.g. intelligence-gathering, identifying new victims, spotting new trends) but affecting sentence severity is not one of them.

Image analysis of child abuse material by the police, which is expensive and time-consuming, does generally not impact sentencing and should be reconsidered if it is the intended goal.

Fortin, F., Paquette, S., & Leclerc, C. (2019). The effect of child sexual exploitation images collection size on offender sentencing. *International Review of Law, Computers & Technology*, 33(3), 330-348.

Inside out : A study of Users Perceptions of Password Memorability and Recall

Internet users maintain an average of 25 accounts that require passwords. Good passwords are hard to remember, so to cope with the increasing number of passwords they are expected to remember people tend to reuse them. Is there any truth behind the assumption that passwords that are hard to guess are hard to remember? This study addresses this assumption by physiologically measuring the extent to which the perception of password memorability is in conflict with "password strength".

Alomari et al. compared passwords perceptions with electroencephalogram (ECG) data, which records brain activity. They presented their 77 study participants with different passwords, asking them to rank their memorability. The brain activity data, recorded through a wireless headband, was combined with the reported ease of recall. This was then used to classify passwords and predict what would be perceived as more or less memorable. The authors also investigated password recall using brain activity data, asking participants to memorize two passwords for a period of 8 to 10 days. Participants were also asked to describe which characteristics make passwords more memorable.

The study consisted of three parts:

RECALL PREDICTION: Participants were asked to recall two algorithmically generated passwords based on seed words they provided. Participants were tested three times: on the first, second, and eighth days of the study.

PERCEPTION OF MEMORABILITY: Participants were asked to rank the memorability of 15 blocks of five real passwords extracted from password leaks. All passwords were 12 characters long and were evaluated and ranked by a strength estimator program.

A SURVEY: Participants were asked what makes a password more or less memorable.

The electroencephalogram test results were linked to how hard people thought passwords would be to remember. Password recall was successfully linked to brain activity data with an average accuracy of 81 percent on the first-day session. Predicting recall over longer periods of time appeared to be challenging, with poor performance on days 2 and 8. Password characteristics had a notable effect on how participants perceived their memorability. Passwords consisting of dictionary words, patterns, or phrases were perceived as more memorable, whereas passwords consisting of names or random letters and numbers were perceived as less memorable. Interestingly, password length was not a concern for memorability. The responses about what made a password memorable were generally consistent with the rankings of password memorability during the experiment.

People seem to know if a password will be easy to remember and consequently can make decisions based on that information. The characteristics of passwords and their strength had a substantial effect on the perception of memorability, even when participants had no information as to how strong the passwords actually were. This provides firm support that passwords that seem easier to remember, such as those containing words, are actually easier to remember than those containing random strings of characters; which can impact the utility of chosen passwords.

Passwords that seem easy to forget, are probably actually more forgettable.

Alomari, R., Martin, M. V., MacDonald, S., Maraj, A., Liscano, R., & Bellman, C. (2019). Inside out-A study of users' perceptions of password memorability and recall. *Journal of Information Security and Applications*, 47, 223-234.

Hacking 9-1-1: Infrastructure Vulnerabilities and Attack Vectors

9-1-1 call centers are a critical component of communications infrastructure. They accept emergency calls, dispatch field responders and provide callers with emergency medical instructions before their arrival. Although the original implementation of basic 9-1-1 services relied mostly on telephone technology, enhanced 9-1-1, deployed in the 1980s, introduced significant new elements aimed at identifying the caller's telephone number and location. As the technological capabilities of 9-1-1 services continue to expand, systems that were once safely managed solely by the telephone company are increasingly connected to less well-regulated networks. The interconnected technologies of 9-1-1 infrastructure are undoubtedly valuable for those needing emergency services, but they could also introduce new risks and potential vulnerabilities.

Goebel et al. describe the technical structure of the 9-1-1 call-taking system in the United States in terms of its current potential for compromising emergency services. The identified vulnerabilities were categorised as risking confidentiality, integrity and availability. This refers to the unauthorized disclosure of private information, the alteration and misdirection of system functionality and the provoked unavailability of emergency services respectively.

Breaches of Confidentiality: Surveillance of the 9-1-1 system could provide attackers with valuable information. Metrics such as call volumes and response times could be used to maximize damage during an attack.

Breaches of Integrity: Attackers can misdirect responder resources using text-based accessibility services provided for the deaf. These services reduce the amount of potentially identifying information provided and are prohibited from keeping records. Attackers can also generate false location or caller information using spoofing applications. These misdirections of resources can delay emergency response and increase collateral damage.

Breaches of Availability: Telephone denial of service attacks fill up the 9-1-1 call center phone lines, preventing the public from reporting emergencies. These attacks have already happened. They were executed by a malicious script released via social media and by automatic dialer devices. Ransomware attacks, where malicious software renders a computer unusable until a ransom is paid to the attacker, could be part of the next generation of attacks targeting availability.

The 9-1-1 infrastructure identifying the location of the caller has also evolved to incorporate wireless cellular providers. For example, the system incorporates cell tower information to better identify caller location. Voice-over-IP providers such as Skype present a special problem for 9-1-1 location services. VoIP services providers need to maintain a database of their subscribers with the appropriate location information, usually based on billing information.

Although all 9-1-1-focused attacks to date have been linked to criminal hackers, they could soon be employed by state-sponsored actors and terrorists. The current "Next Generation 9-1-1" upgrade effort in the US that is set to allow call centers to receive text messages, images and live video streams, could add further cybersecurity vulnerabilities to those already present in the telephone-based systems. Vulnerabilities could be mitigated by frequent software patching, network segmentation and dedicated cybersecurity response plans and teams.

Technological enhancements to the 9-1-1 services infrastructure are progressively introducing new vulnerabilities which need to be quickly mitigated.

Goebel, M., Dameff, C., & Tully, J. (2019). Hacking 9-1-1: infrastructure vulnerabilities and attack vectors. *Journal of medical Internet research*, 21(7), e14383.

Users Really Do Answer Telephone Scams

Telephone fraud is a significant and growing problem. Decreasing costs and automation have made the telephone an attractive medium for disseminating unsolicited information. Telephone phishing can be more convincing than other forms of fraud because it presents an enticing harmony of both visual and audible cues. In the midst of concerns with telephone scams, however, it remains unclear why people fall for them and how to combat the problem.

Tu et al. ran a telephone phishing experiment on a university campus. Researchers reviewed more than 150 real-world telephone scam samples to identify their main attributes. They then tested variations of those attributes on a population of 3,000 Arizona State University staff and faculty to assess their individual effectiveness. Contacts were randomly chosen from the university telephone directory, and then evenly separated into 10 experiment groups. Participants were called on their work phone with a pre-recorded message, according to one of ten sets of specific experimental characteristics. They were then prompted to enter the last four digits of their Social Security number. This led to a debriefing announcement and a survey asking if they had been convinced by the scam. The experiment was carefully designed in collaboration with the university's Institutional Review Board because of its involuntary participation.

Impersonating an internal entity was the attribute most linked to attack success. The experiment where participants were shown a caller ID displaying a faked internal department name was the most effective, with 10 percent (31/300) of recipients entering the last digits of their Social Security number. Manipulating the area code also had a small but noticeable impact on attack success. On the contrary, manipulating the type of motivation, voice production, voice accent and caller name did not result in a higher success rate. Across all 10 experimental groups, 5 percent of participants (148/3000) entered at least a digit when requested to enter their Social Security information. To account for the possibility of recipients entering a fake Social Security number, the authors removed the participants who subsequently stated that they were unconvinced by the scam during the survey process.

ID spoofing is a very effective feature in telephone scams. Given the success of impersonation techniques, scam prevention efforts should prioritize impersonation countermeasures. Technical solutions such as caller ID authentication are recommended to provide users with early impersonation warnings. Feedback from survey participants indicates that vigilance was an important reason for not falling for a scam. Education and awareness campaigns are thus also recommended as scam countermeasures.

ID spoofing is a very effective feature in telephone scams. Prevention efforts should accordingly prioritize impersonation countermeasures.


Tu, H., Doupé, A., Zhao, Z., & Ahn, G. J. (2019). Users really do answer telephone scams. In 28th USENIX Security Symposium (USENIX Security 19) (pp. 1327-1340).

Bitcoin Exchange Addresses Identification and Its Application in Online Drug Trading Regulation

Bitcoin is a decentralized digital currency using blockchain technology to build a public ledger for transaction security and coin control. Anonymity is a major feature of Bitcoin, which allows criminals to conduct illegal trading activities and easily evade regulation. Anonymity and decentralization, two inherent properties of Bitcoin, make it difficult for regulators to monitor and investigate illicit transactions, for instance drug sales through online darknet markets. Identifying the addresses linked to Bitcoin exchanges, where users can trade Bitcoin for fiat money, is crucial for regulation. Exchanges provide the only channel that links people with virtual Bitcoin addresses, and their addresses can be used to identify transactions of interest.

Liang et al. developed a reliable method to identify the addresses of Bitcoin exchanges, using only transaction and user pattern data. They downloaded Bitcoin transaction histories from July 3 to 9, 2018, collecting 3,100,000 unique addresses and 1,350,000 transactions. Using the data, they then constructed its corresponding transaction network. The structure of this network was analysed with multiple algorithmic methods to allow a comparison of the outcomes and determine their reliability. Results indicate that the addresses of exchanges in the transaction network are identifiable and notably different from general addresses in the distribution of their connections. For instance, the most common number of connections was three for general addresses and one for exchange addresses. General addresses also mostly ranged from one to five connections, while more than half of all exchange addresses displayed more than five connections. The proposed identification algorithms tested appeared to be effective.

While previous research efforts have been attempting Bitcoin address de-anonymization, this study provides new tools to detect illegal behaviors in the Bitcoin network based on transaction and user patterns instead of user identities. This research provides a new basis for regulating online darknet markets, hopefully mitigating the public health implications of illicit drug trade.



Top 5 Bitcoin and Exchange nodes compared by frequency of Network Connections


Bitcoin exchange addresses can be identified through algorithmic methods based on the distribution of connections in the transaction network.

Liang, J., Li, L., Zeng, D., Luan, S., & Gan, L. (2019). Bitcoin Exchange Addresses Identification and Its Application in Online Drug Trading Regulation. In PACIS (p. 49).

The Psychology of Internet Fraud Victimization: a Systematic Review

Internet-based fraud is a fast-growing crime. Although fraud studies often refer to psychological explanations such as impulsiveness and loneliness, their actual use of established behavioral theories and methods is often limited.

Norris et al. conducted a systematic review of the literature relating to how victims respond to fraudulent communications. They restricted their search to empirical examinations of established psychological theories, from peer-reviewed journals, conference presentations and book chapters in English. The authors then categorized the 34 articles they obtained in three categories of decision-making cues.



MESSAGE FACTORS: How fraudulent messages are framed to hook their target and maximize their enticement potential.

EXPERIENTIAL FACTORS: How knowledge of internet scams and computer experience can help make individuals more resilient to fraud.

DISPOSITIONAL FACTORS: How personality and cognitive ability are related to fraud susceptibility.

The review underlined the importance of the content of scam messages and of the device on which they are received. Messages presenting as from a credible source, requiring a quick response and responded to on a smartphone were more effective. Computer experience and expertise also appeared to lead to resilience. People with higher levels of security knowledge and better email processing abilities were less susceptible to phishing attempts. However, it was also clear that whilst message content and internet experience have some predictive ability, they cannot alone explain how a person becomes a victim of web-based fraud. Dispositional factors, such as motivation and personality, appear to be key in determining how messages and experiences are processed. Unfortunately, research has not settled on a common set of psychological principles explaining an increased likelihood of victimization.

This study highlighted the rarity of empirical examinations of online fraud based on established psychological theory. The majority of the 1036 papers initially extracted discussed other fraud types (corporate, academic), did not focus on individual factors or did not include at least one established and testable theory. The review of the 34 included papers then indicated no clear agreement on a set of applicable psychological principles.

The evidence currently supporting online fraud literature is limited and too often anecdotal, which can lead to misleading myths. Additional research with a theoretically and practically informed agenda is necessary in this important and growing field, so that we can provide individuals most at risk of fraud with targeted preventative measures. Mood as a factor in how people think about potentially fraudulent messages could present an overlooked area that could prove fruitful.

Dispositional factors, such as motivation and personality, appear to be the key mediating factor determining how message and experiential factors are processed.

Norris, G., Brookes, A., & Dowell, D. (2019). The psychology of internet fraud victimisation: A systematic review. *Journal of Police and Criminal Psychology*, 34(3), 231-245.

Ownership and control over publicly accessible platform data

Airbnb, publicly displays a wide range of information such as unit descriptions, prices and locations, hosts and guest reviews. This information is often of interest for public serving groups, like governments and researchers interested in housing. In a data-driven economy, questions arise as to who has the right to control such data, and the ways such control can be exercised.

Scassa analyzed statute and case law in order to assess the current situation around the scraping of publicly accessible platform data. The research involved a detailed analysis of Airbnb's documentation; the platform; studies and reports relying on scraped Airbnb data; as well as regulations from Canada and the USA.

A thriving data ecosystem has arisen around Airbnb's publicly accessible data that provide valuable insights into a range of issues such as tourism, tax avoidance and rental discrimination. Civil society organizations, journalists, researchers and various businesses all make use of Airbnb data to serve the broader public interest.

Access to this data isn't guaranteed as platforms have several legal tools they could use to potentially cut off this supply of information to actors in their data ecosystem.

Ownership

Intellectual property rights: Airbnb does not claim copyright on member content but may have a copyright on the overall compilation of the site's content as its host and compiler.

Rights to use: Although there are currently no records of lawsuits against Airbnb data scrapers, the availability of fair use defenses remains uncertain, especially for small organizations with little resources.

Chattel rights: In some data-scraping cases, plaintiffs have argued that scrapers are engaging in a trespass to chattels, interfering with their web servers.

Contractual or Technological Restrictions

In at least two cases, courts have held that contractual terms of service that prohibit scraping may provide a basis for breaches of contract liability.

Anti-circumvention provisions, now found in most copyright statutes, could provide an additional recourse.

Privacy law

To the extent that platform data includes personal information, data protection laws may impose further restrictions on the collection, use and disclosure of these data.

Legal Tools available to platform operators to restrict open use

Economic and power imbalances in the data access litigation process create a risk that consumer perspectives and the public interest are not being represented in the law as it evolves. An ecosystem approach is particularly useful to address the reality that non-commercial users are unlikely to pursue issues in court because of high litigation costs. It is important to keep in mind the diverse system that relies on data coming out of these companies and not just hear the claims of major corporate players.

Platform data has become vital source of information for many organizations serving the public interest, yet this function of data access is not recognised currently in law.

Scassa, T. (2019). Ownership and control over publicly accessible platform data. Online Information Review, 43(6), 986-1002. doi:10.1108/OIR-02-2018-0053

Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning





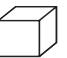



Deep learning, a machine learning method inspired by the information processing of a biological brain, is being successfully applied to many types of potentially sensitive user data. Massive datasets of user speech and images, medical records, financial data and location data points are provided to those systems to train their algorithms, with concerning privacy implications. How secure are those data records, especially in a federated learning setting where training data is distributed among multiple parties?

Nasr et al. present a comprehensive framework for the privacy analysis of deep neural networks, using a novel Open-box membership inference approach. Previous research has been focused on closed-box attacks, a scenario where an attacker's observations are limited to the model's output and intermediate calculations remain hidden. Open-box inference attacks, where attackers take advantage of their prior knowledge of the algorithmic model and parameters in use, are more accurate than their closed-box equivalent. They also correspond better to many real-world federated learning scenarios where participants themselves could be potential attackers.

Nasr et al. experimented all attack types on three different datasets consisting of color images, online shopping records and hospital discharge records. The authors analyzed and exploited in their attacks the privacy vulnerabilities of the stochastic gradient descent algorithm, which is the de facto standard for training artificial neural networks.

Results show that deep learning models previously assessed as not very vulnerable to closed-box inference attacks could be substantially more vulnerable to open-box attacks. The DenseNet model for instance, tested with a 54.5% closed-box inference accuracy (50% being the baseline for a random guess), returned a white-box attack accuracy of 74.3%. In a federated learning setting, where the training data is distributed among multiple parties, adversarial participants were demonstrated able to successfully run active membership attacks against other participants, pushing the algorithm to leak other parties' data.

Findings highlight the vulnerability of potentially sensitive training data used in deep learning neural networks. They suggest additional precautionary measures, especially in a federated network context.

 Closed Box Only output visible	 Passive The output visible	 Stand-alone Data in one place	 Supervised attack trained on known target data
 Open Box Everything visible	 Active everything visible	 Federated Data shared between parties	 Unsupervised attack trained not knowing target

Sensitive data used to train deep learning algorithms we considered safe are in fact vulnerable to inference attacks. Additional precautionary measures are needed.

Nasr, M., Shokri, R., & Houmansadr, A. (2019). Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In 2019 IEEE Symposium on Security and Privacy (SP) (pp. 739-753). IEEE.

Shining Light in Dark Places: Understanding the Tor Network


Tor is a popular privacy-enhancing system which conceals and protects its users' internet usage from network surveillance. It directs Internet traffic through a carefully constructed three-hop path among a worldwide network of volunteer relays, using a layered encryption strategy.

McCoy et al. set up their own Tor router and analyzed passing data to observe who is using the service and how. They deployed their router from December 15-19, 2007 and from January 15-30, 2008, capturing sending address information and identifying which protocols were used. The authors implemented a logging detection technique, identifying malicious Tor routers targeting insecure protocols to capture usernames and passwords.

While internet browsing traffic comprised an overwhelming majority of the connections observed, the BitTorrent protocol (a peer-to-peer protocol used to download large files) used a disproportionately high amount of bandwidth in the Tor network. Insecure protocols which transmit login credentials in plain text, from email servers for instance, were also commonly observed. In addition to potentially compromising those accounts, plain text credentials allow malicious Tor routers to trace back all traffic on the same Tor circuit back to the now identified client. The authors found one router logging plain text email traffic, and observed it attempting to login using a name and password pair they purposefully provided.

The vast majority of clients originated in Germany, with China and the United States providing the next largest number of clients. Germany and the United States together contributed nearly two thirds of all running routers. From the perspective of the researchers' relay, the top 2% of all routers in the Tor network transported about 50% of the traffic, while the bottom 75% together transported just 2%.

Results suggest that tunneling insecure protocols like email over Tor presents an important risk to the initiating client's anonymity. Location diversity in the distribution of Tor routers, although desirable to enhance privacy, is difficult to guarantee because of their high concentration in a few countries. Since the vast majority of Tor traffic is handled by a very small set of routers, an adversary controlling a set of the highest performing routers would be able to conduct traffic analysis and defeat the network's anonymizing properties. Incentive programs to encourage volunteers to run routers in under-represented countries should be investigated accordingly.



Tor's lack of diversity in its router location and bandwidth distribution could compromise client anonymity.

McCoy, D., Bauer, K., Grunwald, D., Kohno, T., & Sicker, D. (2008). Shining light in dark places: Understanding the Tor network. In International symposium on privacy enhancing technologies symposium (pp. 63-76). Springer, Berlin, Heidelberg.

Keepers of the Machines: Examining How System Administrators Manage Software Updates

Information systems administrators serve as 'keepers of the machines' entrusted with keeping computers updated and running smoothly. Failing to patch known software vulnerabilities can lead to devastating consequences, as critical infrastructure becomes potentially subject to crippling attacks. Conversely, deploying updates in the context of a large organization can lead to serious problems. How do administrators manage updates, and what factors impact how effectively they perform those updates?

Li et al. sent a survey to 102 US system administrators in September and October 2017. Using themes identified in a series of pilot interviews, the authors then recruited and interviewed 17 administrators, most of whom had participated in the survey beforehand. About half of the study participants worked at organizations with over 500 employees. Participants typically managed large computing infrastructures; two-thirds of them maintaining more than 100 computers.

Findings suggest that administrator update workflows consist of five stages, each with their own challenges and limitations:

- 1) Learning About Updates: Participants usually rely on 5 or more sources, as update information is highly dispersed. They need to be constantly browsing the news.
- 2) Deciding to Update: Administrators prioritize security updates, but they can be bundled with feature changes which are potentially disruptive.
- 3) Preparing for Update Installation: Staggered deployment and dedicated testing setups are the two main update strategies. Staggered deployment is concerning because production machines, which are the most exposed to potential attackers, are updated last.
- 4) Deploying Updates: Administrators often depend on custom scripts and third-party managers. The need to maintain compatibility with vendors lagging behind prevented some participants from deploying automatic updates.
- 5) Handling Update Issues After Deployment: Participants dealt with update problems by simply uninstalling it, thus reverting to an insecure state. They prioritized functionality over security.

Internal policies and management could also play an important role in update decisions. Freedom to apply updates could result in ad-hoc decisions by administrators, potentially resulting in poor practices. Conversely, approval requirements could delay or prevent the application of updates. Some administrators even skipped less severe updates to avoid the hassle.

Findings suggest the following potential solutions:

- Standardized and consolidated update information in a centralized repository
- Software vendors bundling security patches separately from feature patches
- Dynamic software updating allowing for live updates without restarts or downtime
- A cultural shift at organizations to recognize the importance of expedient updates

Software updates are too often unreliable and system administrators prioritize functionality over security in their updating practices.

Li, F., Rogers, L., Mathur, A., Malkin, N., & Chetty, M. (2019). Keepers of the machines: Examining how system administrators manage software updates for multiple machines. In Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019).

serene-risc.ca

Smart Cybersecurity Network (SERENE-RISC) is a knowledge mobilization network that aims to be a cybersecurity exchange that is recognized as an open, inclusive, and unbiased forum to unlock the value of knowledge on cybersecurity risks, threats and solution for all Canadians. We are a network of networks that provides publications, tools, events, professional development and education. We are always looking for partners to collaborate to make cybersecurity better, so please contact us to find out more.

konnnect.serene-risc.ca

We are building the Canadian source for summaries of research, opinion pieces, video presentations, fraud bulletins, public awareness materials and more. You can filter the collection by type of content, click on keywords or search. The website provides lots of original content produced by us and in collaboration with partners. To save your time, we are sending an email regularly a summary of the new content to save you time and make it easier to find content. You can subscribe to this list online at: <http://konnnect.serene-risc.ca/subscribe-abonnement/>

We will be looking for contributions to this page from our community, so if you have an idea for a piece that you would like to share please let us know.

You can help by:

- Subscribing to the regular update online at: <http://konnnect.serene-risc.ca/subscribe-abonnement>
- Following @SERENE_RISC on twitter and retweeting,
- Joining the LinkedIn Group and submitting or commenting on posts, and
- Posting links to konnnect.serene-risc.ca content on other platforms you are involved with.

cybersec101.ca

Want to provide basic cybersecurity classes but don't know how or have the time to get started?

We have developed evidence-based materials that can help you quickly put on training programs. Find resource sheets, presentation videos, class tools and slides to build basic cybersecurity education and awareness sessions at. There are ten modules in French and English with materials to help from basic concepts to practical step-by-steps for better security online.

secrev.org

We host an international conference on Cybersecurity. It is an event that goes around the world to provide a zero dollar cost event so that researchers can collaborate. The event is in its fourth year in 2020. To find out more please contact us at info@serene-risc.ca



@SERENE_RISC



/serenerisc



/serene-risc

The SERENE-RISC Cybersecurity Knowledge Digest

Editor-in-Chief: Michael Joyce

Scientific Editor: Benoît Dupont

Editor: Louis Melançon

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



Université 
de Montréal