

Trends in online consumer fraud: A data science perspective

Felix Soldner¹, Bennett Kleinberg^{2,1}, Shane Johnson¹

¹Department of Security and Crime Science & Dawes Centre for Future Crime, University College London, UK

²Department of Methodology and Statistics, Tilburg University, The Netherlands

Abstract

The cost of fraud is increasing both in scale and cost. Current approaches to countering it by law enforcement, the commercial and non-governmental sectors include manual and time-consuming actions. Given the volume of the problem, this makes it hard to address. Data science offers the opportunity to automate some of these approaches and hence conduct analyses at scale. For example, to commit common forms of fraud, such as fraudulent billings and non-deliveries, offenders typically attract and interact with victims using publicly available websites. This presents the opportunity to “scrape” data and analyze it automatically using techniques from the data scientist’s toolbox. Additional insights about how to tackle fraud could also come from the dark web, which harbors markets for counterfeits and fraud-enabling goods. This chapter provides an overview of current online consumer fraud-schemes, how they are being tackled by different sectors, and the role of dark web markets in such fraud. We discuss current data science approaches employed in the academic literature and some of the challenges researchers face. The chapter closes with a discussion of how online consumer fraud might be detected and prevented using these techniques.

Introduction

Following the advent of the Internet, the interaction between sellers and consumers has increasingly shifted from a face-to-face interaction towards an online environment. With that shift, new forms of consumer fraud have emerged (Rusch, 1991). Such fraud, as well as cybercrime as a whole, seems to be increasing, generating financial and psychological costs to society, including direct monetary losses, criminal revenues, and opportunity costs (Ablon et al., 2014; R. Anderson et al., 2013; van Wegberg et al., 2018). For example, a report by the FBI (2018), which examined complaints and messages from Internet crime victims in the US, found a steady increase for all types of cybercrime, with an estimated total monetary loss of \$2.71 billion in 2018. Additionally, estimates suggest that consumer products and services account for 42.6% of all financial fraud incidents, with the Internet being the dominant solicitation method (30%) in the US (DeLiema et al., 2017). Although accurate estimates of online consumer fraud are difficult to compute due to the lack and granularity of data, estimates such as those above give a sense of the scale of the problem. In terms of prevention, difficulties come from the nature of the online environment, which enables fraudsters to operate from anywhere on earth and reach individuals at scale, making traditional ways of manually detecting and preventing such fraud inefficient, and complicate the actions for law enforcement to bring fraudsters to justice (Herley, 2010).

The chapter is organized as follows. We will first examine what online consumer fraud is, revisiting the definition and common fraud schemes. This review is followed by a discussion of some of the current approaches (non-) governmental and commercial institutions take to detect and prevent online consumer fraud. We will then discuss darknet markets – internet platforms that, amongst other things sell both legal and illicit goods while providing anonymity (Christin, 2013) – and how online consumer fraud is being facilitated by them. The chapter closes with a discussion of how methods from data science can be applied to support the detection and prevention of online consumer fraud as well as possible future research directions.

Common online consumer fraud types on the surface web

For this chapter, we focus on online consumer fraud committed on the surface web (e.g. on eBay, Amazon, etc.) and define fraud in this context as transactions for which the consumer is

deceived in some way and is unaware that fraudulent activity is taking place. We limit our attention to fraud such as unauthorized billing, the non-delivery of goods ordered or services from legitimate or fraudulent websites, or fraud that requires the theft of their credentials (e.g. via phishing websites) (Anderson, 2019). Such types of fraud can often be deployed at scale (as discussed in the chapter by K. Anderson (2021)) and require little sustained attention from the fraudster. These types of offenses differ from other types of fraud such as romance or advanced fees scams, which require sustained effort from the fraudster and typically include activity that cannot be or would be more difficult to automate. We focus our attention on the former because we believe that these types of offenses – which generally involve only one or no interaction with a victim – will be easier to address using approaches from data science. Data science approaches often rely on automated detection systems, which are difficult to deploy if a personal conversation between individuals is taking place (with the exceptions of e-mail spam filters). Furthermore, to successfully implement automated detection methods, the required data needs to be accessible and usable, as in open or public domains, such as online marketplaces. Thus, we will discuss common online consumer fraud types, which are more suitable to be addressed by data science methods.

Fraudulent billings

One of the most common fraud schemes involves billings for products or services that the customer did not agree to (K. Anderson, 2019). They can occur on websites where a customer believes they are making a one-time order and payment but are in fact signing up for a subscription of some kind. Similarly, free trials for which a subscription is about to terminate can be extended without the approval of – but at a cost to – the customer. It has been estimated that unapproved billings through fraudulent websites accounted for a monetary

loss of \$48 million in the USA alone in 2018 FBI (2018). However, such estimates are based solely on complaints made to the FBI and – given that most (online) crime goes unreported (up to 85%) (Office for National Statistics, 2020) – provide only a partial picture. Different forms of such web pages exist but include those that masquerade as pages created by legitimate companies or governments, with consumers being directed to them in various ways including false advertisements on social media platforms, e-mails, text messages¹, and so on. These pages generally mimic the appearance of the official websites in a convincing manner which contributes to consumers being lured in².

Non-deliveries

Another common form of online consumer fraud is non-delivery fraud. In the simplest form of this type of offending online, the buyer and seller transact in an online marketplace, but while the seller receives payment, the product is never sent to the buyer (K. Anderson, 2019). For some online platforms (e.g. eBay), the buyer has the option of getting the money back from the payment system (e.g. PayPal) as well as writing a negative seller review and flagging the fraud quickly, which can make it difficult for offenders to sustain their activity using the same seller account over a prolonged period (Bauerly, 2009). However, this does not prevent fraud from taking place in the first place. According to the FBI, this type of offending accounts for monetary losses of more than \$180 million per annum in the USA (FBI, 2017, 2018).

In other variations of this offense, the seller convinces the customer to pay for the goods or service outside of the marketplace on which it was advertised (e.g., using a cheque, cash, or fake escrow). This type of fraud involves more effort on the part of the seller because

¹ For example, <https://www.actionfraud.police.uk/alert/fake-dpd-messages-lead-to-over-200000-in-losses-since-june>.

² For an example, <https://www.gov.uk/government/publications/phishing-and-bogus-emails-hm-revenue-and-customs-examples/phishing-emails-and-bogus-contact-hm-revenue-and-customs-examples>.

it requires them to convince the customer to transact outside of the original platform. However, where successful, the online platform cannot confirm purchases, which makes the flagging of the fraud more difficult, allowing fraudsters to maintain their account for longer.

In other instances of this type of offense, the product may be misrepresented, with the customer receiving an item of lower value, they may receive the packaging but not the actual item, or the product may be wrongly delivered on purpose. In the latter case, the seller has proof that an item was delivered and can shift the blame towards the delivery service, which makes it more difficult to resolve the problem (Abdallah et al., 2016). With the same techniques of product misrepresentations, not only small buyers can be affected, but also companies, making large-scale industrial purchases, as in the case of painted stones sold as copper for \$36m (Harper, 2021).

Commonalities in online fraud schemes

The above-mentioned fraud schemes differ in how they are executed, but they all typically involve or require the presence of open or publicly available websites. Since these websites can be observed and data scraped from them, this presents the opportunity to do something about them using techniques from data science. Moreover, since such methods can be automated, with appropriate planning, they might be addressed in a scalable way. Such techniques can also be applied to other types of fraud not mentioned above, such as fraudulent computer repairs (Miramirkhani et al., 2017) or fraud-related phishing websites, which can be prompted through pop-ups or false advertisements on any webpage (K. Anderson, 2019, 2021; Christin et al., 2010).

Current strategies to prevent and resolve online consumer fraud.

Online consumer fraud takes place in an environment in which many different industries and sectors converge. Commercial enterprises, non-governmental and governmental organizations adopt different strategies to deal with such fraud, and these are outlined in Table 1. The commercial sector aims to maximize profit by deploying brand protection methods and removing fraudsters from “seller” webpages. This is mostly achieved by actively searching for violations manually or using semi-automated approaches (Ganguly, 2015; Pointer Brand Protection, 2019; Yellow Brand Protection, 2019). In turn, non-governmental organizations mostly focus on trying to support (vulnerable) customers by informing them about different types of fraud to heighten their awareness of them. The goal with such strategies is to increase the likelihood that customers will spot fraudulent activities or listings while they are online and act accordingly (Beals et al., 2015; Deevy & Beals, 2013; M. DeLiema et al., 2019; Peaston, 2019; Stanford Center on Longevity, 2019). Included recommendations to online shoppers often include (but are not limited to) the careful inspection of reviews, ratings, or details about the shop (e.g. physical location). The strategies to increase customer awareness requires guidance to be up to date, for consumers to be exposed to it, and for them to act upon it. While this approach is probably useful, a lot has to go right for it to work and it should be noted that awareness does not equate with effectiveness (in prevention). It would be highly valuable to conduct work examining strategies which aim to increase fraud awareness, to determine their success. Law enforcement agencies take a more reactive approach by responding to complaints and intelligence about fraud, which are investigated to find, prosecute and subsequently deter fraudsters (FBI, 2018; Intellectual Property Office, personal communication, 2019; Trading Standards UK, personal communication, February 26, 2019, personal communication, May 17, 2019) (see Table 1.)

Sector	Motivation	Strategy	Actions/Controls	How to resolve fraud
<p>Commercial industry</p> <p>(Ganguly, 2015; Pointer Brand Protection, 2019; Vistalworks, 2019; Yellow Brand Protection, 2019)</p>	<ul style="list-style-type: none"> Maximizing profit 	<ul style="list-style-type: none"> Hiring brand protection companies 	<ul style="list-style-type: none"> Searching for brand violations (manually, semi-automatically – “AI”) in product/service domain Reacting to complaints 	<ul style="list-style-type: none"> Contacting fraudsters Notifying fraud to company/platform pursuing legal actions
		<ul style="list-style-type: none"> Creating internal divisions to detect fraud (e.g. eBay, Amazon) 	<ul style="list-style-type: none"> In-person visits Test purchases Customer flagging systems Automated monitoring of item sales, transactions, and product view ratios 	<ul style="list-style-type: none"> Removing sellers from the webpage
<p>Non-governmental (e.g. Cifas)</p> <p>(Beals et al., 2015; Deevy & Beals, 2013; M. DeLiema et al., 2019; Peaston, 2019; Stanford Center on Longevity, 2019)</p>	<ul style="list-style-type: none"> Protect (vulnerable) customers 	<ul style="list-style-type: none"> Using surveys to estimate fraud prevalence and to identify vulnerable customers Creating fraud avoidance guidelines 	<p>Recommending individuals to scrutinize:</p> <ul style="list-style-type: none"> seller feedback and comments price and seller history 	<ul style="list-style-type: none"> Increase the likelihood of consumers spotting suspicious listings
<p>Law Enforcement (e.g. FBI, Trading Standards)</p> <p>(FBI, 2018; Intellectual Property Office, personal communication, 2019; National Audit Office, 2016; Raine et al., 2015; Trading Standards UK, personal communication, February 26, 2019, personal communication, May 17, 2019)</p>	<ul style="list-style-type: none"> Deter/detect criminals, reduce the negative financial impact on society 	<ul style="list-style-type: none"> Respond to fraud complaints, intelligence about fraud Prioritizing fraud types and cases depending on financial damage 	<ul style="list-style-type: none"> Manual investigation of received intelligence Cross-referencing of product and seller information 	<ul style="list-style-type: none"> Find, identify and deter fraudsters Inform affected platform

Table 1. Summary of how different sectors address online consumer fraud

Current strategies used to combat online consumer fraud by the commercial sector and law enforcement are heavily reliant on identifiable information about online users and vendors. In a non-anonymous space, such as the surface web, identifying individuals is generally feasible, as long as users register for accounts using names, e-mails, or other information that can be used to identify them. In addition, the activity of their computers can be tracked through their IP address, making geolocation possible, unless they use techniques to obscure this. Thus, authorities have some tools at hand to move against fraudsters. However, these approaches become extremely difficult on the dark web, in which users and vendors have anonymity and no trackable IP address. With this anonymity, the dark web provides a supportive space for online consumer fraud, in which individuals not only openly sell counterfeits but also sell fraud enabling services (e.g. highly-rated eBay accounts, defrauding strategies, etc.). Therefore, it is important to examine the dark web to explore the prevalence of consumer fraud-enabling products and services in this space and to find methods of combating it. To do so, we will first describe what the dark web is, how it can be accessed, and how users sell goods as well as interact with each other on it. As will be seen, exploring each of these aspects provides potential insight as to how to tackle consumer fraud on the surface web.

The Dark Web

The Internet can be segmented into three sections which are characterized by how it can be accessed and how users communicate with each other: the surface web, the deep web, and the dark web (Figure 1) (Biddle et al., 2003; Mansfield-Devine, 2009).

The surface web is the Internet we usually encounter and includes platforms such as eBay, YouTube, and news sites. Put differently, it contains the online content that is indexed by openly accessible search engines (Bergman, 2001; He et al., 2007) such as Google's.

The deep web represents the part of the internet that is not indexed by search engines and cannot easily be accessed. Content is often password-protected or restricted in other ways (e.g., requiring authentication). Online banking, webmail, or paywall content would fall within this category. Historic estimates (Bergman, 2001; He et al., 2007) suggest that the deep web is 400-550 times larger than the surface web in terms of the amount of information (data) and web pages stored. The deep web is believed to be the fastest-growing category of the internet, but it is difficult to estimate its current size precisely, as it is not openly indexed. For both the surface and deep web individual users and servers are not automatically anonymous.

The dark web represents a small portion of the deep web, on which users and hosts are anonymized. Anonymization can be facilitated through different technologies, such as The Onion Router (Tor) (The Tor Project, Inc., 2020) or the Invisible Internet Project (I2P) (The Invisible Internet Project, 2020). Such methods hide the identities of users by sending their data through a network of computers and servers that use protocols to conceal their IP addresses and serve as relays (Gehl, 2018). Hosts using this system can provide web pages called "onion sites". Navigating to an onion site requires an exact address, as they are not indexed by search engines and the administrative markers, such as ".com", ".net", or country codes, such as ".de" (Germany), ".us" (United States) are replaced by ".onion" (Ghosh, Porras, et al., 2017). While the darknet was not developed with malicious intent in mind, the safety of anonymous communications makes illegal activities less dangerous for the perpetrators as some of the usual information that is used to track fraudsters is removed (see Figure 1).

The Tor network, which is perhaps the most widely known, can be accessed through the Tor client which functions like a regular browser. The ease of entering the dark web makes it accessible to many individuals, some of whom will want to engage in fraudulent activities. Estimates of the number of “.onion” web pages range from 3700 to 32,000, with a maximum of around 13,000 exhibiting prolonged activity (Ghosh, Porras, et al., 2017; Gray, 2019; Lewis, 2016).

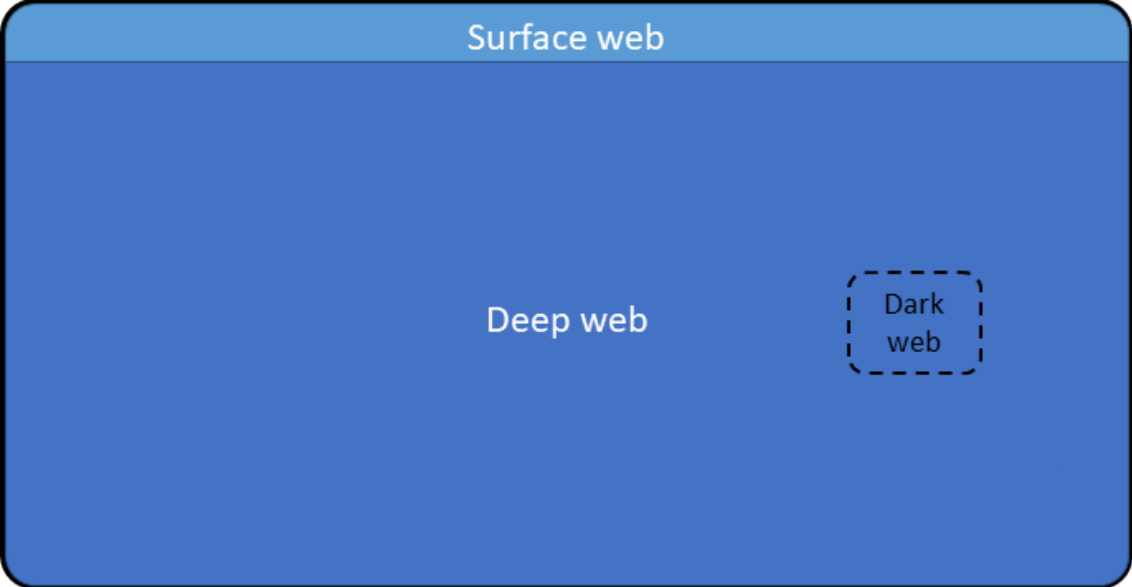


Figure 1. The World Wide Web subdivided into the surface, deep, and dark web. The domain proportions are approximations, as the deep web would dwarf the other domains to be unrecognizable.

Markets on the dark web

The dark web harbors selling platforms, which are called “black markets”, “cryptomarkets”, or “dark(net) markets”, which offer various products and services (Christin, 2013; Gehl, 2018). These platforms are not online shops but — like markets such as eBay — provide spaces for users to transact (Soska & Christin, 2015). Similarly, like eBay, platform providers receive a small margin of each monetary transaction. Darknet market transactions are anonymized through the use of cryptocurrencies, which can be obtained from online exchanges (e.g., Coinbase). Most of these currencies are based on a peer-to-peer system that does not rely on

a bank or other centralized third party (Soska & Christin, 2015); (for more information on cryptocurrencies, see the Chapter by Kamps et al. (2021) of this book). The use of cryptocurrencies as well as anonymized communications further facilitates illegal activity and makes tracking fraudsters very difficult. The Silk Road, which was the first commonly known darknet market, started operating in February 2011 and used Bitcoin (BTC) as a medium of transaction. Buyers did not pay the seller directly but used an escrow system — a form of holding area — embedded within the platform. Escrow systems allow platform operators to compute commissions as well as to supervise transactions between buyers and sellers to ensure that products are shipped (or services are provided) only after payments have been received (Christin, 2013). Thus, transactions can be completed in a highly anonymized space, in which accountability is otherwise almost absent.

Estimates suggest that for larger markets (e.g., Silk Road 1 & 2, Agora, AlphaBay), about 50-80% of all category listings (Demant et al., 2018; Europol, 2017; Soska & Christin, 2015) are for various forms of drugs (e.g. Cannabis, Amphetamines). However, a large range of other products are also typically offered including weapons, counterfeit goods, guides for malicious activities, stolen credit or debit card information, other personal data (e.g. log-in credentials), or services such as hacking-attacks, or secure hosting (Adamsson, 2017; Du et al., 2018; van Wegberg et al., 2018). By way of an example, Figure 2 shows a screenshot from one darknet market – AlphaBay, which operated for around three years (2015 – 2017).

Examining darknet markets more closely is important as doing so offers a lens through which to understand online consumer fraud from a different perspective than the surface web. However, since all communication is anonymized, it is very difficult to retrieve market or user information (see Figure 2).

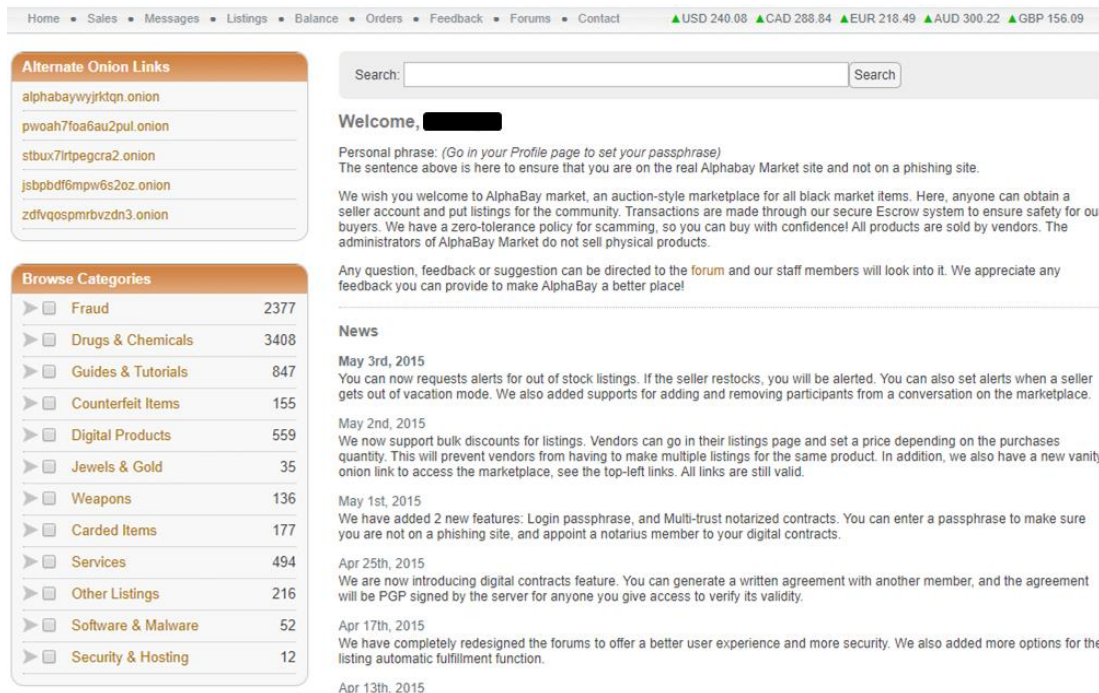


Figure 2. Screenshot of the homepage from AlphaBay (Branwen et al., 2015).

Obtaining darknet market data

Collecting darknet market data can be complicated as such data are rarely shared by the platforms. One way of collecting darknet market data is by visiting the websites as a user and retrieve it manually. However, the most commonly used method is to employ a web scraper that revisits markets at set intervals (e.g., once a day) over weeks, months, or longer (Ghosh, Das, et al., 2017). Such web scrapers retrieve information from a web page and download it to a local machine. Employing web scrapers is relatively straightforward on the open web, it is time-consuming on the dark web due to the rerouting of page requests – necessary to create anonymity – through the TOR network, which takes significantly longer than on the surface web. Moreover, since darknet market sites often have different layouts, almost all scrapers need to be coded individually for each website (Du et al., 2018; Hayes et al., 2018). Also, researchers are often faced with unexpected problems during the scraping process, which

requires a constant review of the collected data. This can entail changes in accessibility credentials, changes in the layout of the page, or downtime of the host (Ball et al., 2019; Ghosh, Porras, et al., 2017). Given that scraping is the predominant method of obtaining darknet market data, discussions of these challenges are well-rehearsed elsewhere and the interested reader is referred to studies by (Ball et al., 2019; Du et al., 2018; Ghosh, Porras, et al., 2017; Van Buskirk et al., 2014, 2015, 2016).

Darknet Markets economies

The Silk Road 1 operated from 2011-2013 and was the first darknet market studied by researchers, who examined what and how much was being sold (Christin, 2013). This darknet market had between 30,000-150,000 active customers, 220 distinct product categories, and vendors who made an overall monthly revenue of \$1.2 million in 2012 (Christin, 2013). Drugs were predominately offered, and half of all products sold were shipped worldwide, mostly originating from the US (43%), followed by the UK (10%) and the Netherlands (6%). In their study, Soska and Christin (2015) estimated transactional and sales volumes through the analysis of product feedback for 32 different marketplaces for the period 2014-2015. They concluded that sales volumes typically varied between \$300,000-\$500,000 per day. With an estimation of 9,386 unique vendors across markets, 70% made less than \$1,000 during their active time on the platform, while 1% accounted for more than half of all sales (Soska & Christin, 2015).

Baravalle and Lee (2018) examined AlphaBay between 2015 and 2017 and estimated sales to reach \$79.8 million over the two years, with \$69.2 million of this attributed to drugs and chemicals alone. They also estimated that \$1.7 million of sales were generated from fraud

(e.g., the sale of fake IDs, or accounts), \$1.6 million from counterfeit items, \$1.4 million from hacking attacks or server hosting, \$748,544 from Software and Malware, \$210,000 from Guides and Tutorials about committing fraud, \$198,000 from digital products, \$35,990 from Security and Hosting, and \$125,472 from other listings. This suggested that counterfeit goods and goods and services related to fraud were well represented on the market.

A study by van Wegberg et al. (2018) investigating cybercrime commodities on eight darknet markets between 2011-2017, and estimated revenue of at least 15 million USD. The study showed that business-to-consumer rather than business-to-business transactions were taken place and that cybercrime commodity trade exhibits some overall growth.

Multinational governmental institutions, such as Europol are increasingly interested in darknet markets, particularly those associated with the drug economy, but also those that facilitate intellectual property (IP) crimes (EMCDDA-Europol, 2017; Europol, 2017). Based on data from five darknet markets (AlphaBay, Dream Market, Hansa, TradeRoute, and Valhalla), Europol estimate that fraud and counterfeits accounted for around 17% of all listings. Examining AlphaBay specifically, EMCDDA-Europol (2017) estimated that 10,000 products were counterfeit goods. Fake banknotes and IDs were the most frequent, but clothes, electronics, jewelry, software, e-books, subscriptions, and watches were also common. Fraud seemed to be well represented too, accounting for around 22% of all listings on Alphabay (Adamsson, 2017). Europol (2017) reports that vendors often sold small amounts, mostly specialized in one type of product but were present in several different markets at the same time. EMCDDA-Europol (2017) concluded that IP crimes were increasing in darknet markets, but that products are not always clearly categorized, which makes it difficult to accurately estimate volumes of offers and sales.

Vendor and user behavior

The sale of products or services on the Internet requires a level of trust between the customer and vendor. This is true on both the surface and dark web, with trust often established through feedback systems such as reviews or ratings. For most dark web markets, reviews can only be submitted after a verified purchase (verified by the escrow system) and some markets make reviews mandatory (Calis, 2018). Since everyone is anonymous and no public or physical stores exist, the review system is a particularly important feature of dark web markets, along with escrow. Although there are no findings of fraud on darknet markets, the presence of such review, escrow, and registration procedures seem to suggest that fraud occurs frequently. Such security procedures are not new but are still often missing on regular surface web platforms and it would be interesting to investigate their effects on the prevalence of fraud on the surface web.

Methods of vendor identification

Identifying individuals on the dark web is inherently difficult due to their anonymity (Gehl, 2018; The Invisible Internet Project, 2020; The Tor Project, Inc., 2020). Nonetheless, some studies have analyzed text data or product photos to attempt to identify individuals or vendors with different user names that operate across platforms (Ho & Ng, 2016; Wang, 2018; Wang et al., 2018). Photos are important in markets since they serve as a form of proof that the vendor owns the product advertised. At the same time, pictures are only one aspect of building trust and they should be identifiably different from one vendor to the next, to highlight ownership. Therefore, Wang et al. (2018) suggested that photo styles could serve as an identification of vendors across markets. To examine whether this was plausible, they used transfer learning, for which a machine learning classifier was pre-trained on a large data set

(ImageNet) and further fine-tuned by retraining it on a smaller data set (vendor-specific photos). The task for the classifier was to identify whether two or more vendors were the same based solely on the images used. The model performed well for a single market for which a subset of vendor and photo associations were known. However, the performance was more difficult to assess for matches across markets as their true association was unknown. Nevertheless, the approach offers a possible identification methodology for an otherwise anonymized environment, which, although not suggested by the authors, could also prove useful on the surface web. For example, once a fraudulent vendor is identified, other vendor accounts across platforms (e.g. Ebay, Amazon) operated by the same fraudster(s) might also be identified in that manner.

In a different study, a series of automated methods, combined with manual investigations were used to identify vendors on darknet markets (Hayes et al., 2018). The researchers collected data from vendors and their associated listings from an undisclosed collection of darknet markets. Using Maltego, which conducts automated surface web cross-referencing (Paterva, 2019), they searched for obtained (darknet) e-mail addresses, user names, and other personally identifiable information, and were able to identify some darknet market vendors on the surface web. As noted by the authors, this approach has some limitations associated with identifiers, which are similar to high-frequency words. For example, a username, which resembles a popular brand or product name (e.g., Coca-Cola), would lead to meaningless cross-referencing. Nevertheless, their results suggest that the automated identification of some individuals is possible, but can be easily disrupted when users employ countermeasures, such as adapting their identifiable information (e.g. username, e-mail address) accordingly.

Importantly, the methods outlined show possible ways of identifying vendors across darknet markets, which might also be applicable, in some circumstances, for the identification of fraudsters across markets on the surface web as well as fraudulent products between the dark- and surface-web.

The role of data science

Data science has been previously defined as *“a set of fundamental principles that support and guide the principled extraction of information and knowledge from data”* (Provost & Fawcett, 2013, p. 2). This definition captures a high-level perspective, but data science also involves *“statistics, or the systematic study of the organization, properties, and analysis of data and its role in inference, including our confidence in the inference”* (Dhar, 2013, p. 1). Thus, data science draws from many disciplines including computer science, mathematics, statistics, and importantly, the knowledge domain of the data in question (Dhar, 2013; Provost & Fawcett, 2013). In the realm of online consumer fraud, data science can be utilized to speed up or scale up manual processes, such as collecting, cross-referencing, and analyzing information from advertisements, product listings, sellers, or webpages to identify if they are likely to be fraudulent. Such methods will be reviewed in this section by looking at current and possible future applications and their associated advantages and disadvantages.

Current data science approaches in the academic literature

Current data science methods applied to online consumer fraud on the surface web mainly focus on the automated detection of specific types of fraud (Abdallah et al., 2016). For example, Pandit et al. (2007) collected eBay data and modeled the topological connections (*i.e.* the network) between sellers and buyers using transaction data. They hypothesized that

fraudsters, who commit the crime, should be heavily connected to accomplices, whose role is to boost the fraudsters' feedback ratings. Accomplices are needed as they can continue to operate once the fraudster is banned from the site, retaining any positive ratings they accumulate. As such, fraudsters would be expected to have fewer connections to honest users than they would have to accomplices. Thus, the authors describe, that fraudsters have many connections to accomplices, but the accomplices nor the fraudsters are interconnected, forming near bipartite cores, instead of cliques, which exhibit strong interconnectedness. The automated detection of such bipartite cores was evaluated on a synthetic network, containing over 66,000 nodes and 795,000 edges. The network was filled with artificial fraudster-accomplice structures of random sizes, which were detected with around 90% accuracy. An additional network was created from scraped eBay data, which contained ten known fraudsters who were identified through manual inspections and investigative media reports. All ten fraudsters were automatically detected by the authors' devised model. As the researchers discuss, the evaluation based on the analysis of the eBay dataset has limitations since no fraud detections (other than those already detected through a manual investigation) could be verified. As such, additional verification of the models' performance through a well-labeled data set would be needed to determine its applicability (Pandit et al., 2007).

Hernandez-Castro and Roberts (2015) used an automated detection method that can process data without human interventions, to try to identify illegal sales of elephant ivory on eBay. Specifically, they used the CN2 induction algorithm (Clark & Niblett, 1989), which is a supervised machine learning method. Supervised methods require a training phase, in which they learn to infer from features (e.g., metadata or writing style of the advertisements) what the corresponding label (e.g. fraudulent versus legitimate, or same versus different account holder of a vendor profile) of a data point is likely to be. The goal is typically to find a

combination of features that enable the best discrimination of the outcome labels. The labels are often obtained through manual annotations, which can be time-consuming, particularly where a large volume of labels is required. However, the advantage of such approaches is that they can uncover previously unknown patterns within the data and utilize these to classify unlabeled data. Such methods find a wide range of applications that can be utilized on the dark- as well as the surface-web, depending on the discriminatory problem. In the study by Hernandez-Castro & Roberts (2015) two former law enforcement officers annotated 1,159 product listings as selling ivory or not. Utilizing the CN2 induction algorithm (Clark & Niblett, 1989), which made decisions based on metadata parameters of the listings (e.g., item price, number of reviews), the framework was able to categorize almost all listings correctly. The advantage of the CN2 algorithm is, that it induces decision rules, which can be inspected to gain an understanding of the sale strategies of the potential fraudsters. Other supervised machine learning methods, such as random forest or logistic regressions are also interpretable, but many others, such as neural networks can be black boxes, since the decision rules are generated automatically, making it difficult to understand the inner workings of the classification system and hence the possible fraud strategies. While the system developed by Hernandez-Castro and Roberts (2015) could find suspicious online listings quickly, without manual searchers, it is difficult to assess the reliability of their current classifier, as the true labels of the listings (selling ivory or not) were not known and no inter-rater reliability score between the annotators was reported (Hernandez-Castro & Roberts, 2015). However, a similar approach of using domain experts to label instances of online consumer fraud (e.g. counterfeits on eBay) could be helpful. For example, classifiers trained using these data could then act as a pre-selection (or triaging) tool by labeling suspicious listings. Although such an approach does not eliminate the problem of falsely labeled fraudsters, it could reduce the

number of listings experts subsequently have to investigate, reducing the overall workload. The model could then be updated after each investigation to increase its performance and reduce false positives.

Other researchers also utilized supervised classification methods, which are trained and tested on different subsets of labeled data (Almendra & Enachescu, 2012; Chang & Chang, 2012; Sahingoz et al., 2019; Xu et al., 2016). For example, Chang and Chang (2012) used decision trees, to detect fraud on the Yahoo Taiwan online auction site. Decision trees continuously split the data into subgroups (e.g. [non-] fraudulent seller) based on a binary decision about a predictive feature (e.g. number of negative reviews). As an example, a seller might be labeled as fraudulent if more than half of the reviews are negative. The remaining unlabeled sellers undergo more splits based on other feature values until all data are categorized (for an introduction into machine learning methods see (Rosenbusch et al., 2021)). In the study by Chang & Chang (2012) the decision trees were trained on the history of sellers' activities and the sellers were labeled as legitimate or fraudulent, depending upon whether they had been blacklisted from the auction site or not. While the performance of the model was promising, the labeling process makes the evaluation problematic. That is, it remains unclear why users were blacklisted and whether and how many blacklisted sellers were genuine. Supervised machine learning methods have also been employed to detect phishing websites by examining their URL (Sahingoz et al., 2019) or HTML-structure (Xu et al., 2016). The approach is the same as that described above – using an annotated dataset the algorithm learns to associate the URLs or the webpage's HTML structure with the correct label to infer a decision rule. This is then used to classify new incoming data (webpages) that do not have labels. A similar approach can also be applied to classify webpages on the dark web, to

automatically identify what type of webpage it is (e.g. a marketplace versus a discussion forum), and what is being discussed or sold (Ghosh, Porras, et al., 2017).

An important aspect of machine learning is that many (but not all) supervised methods rely on the selection and crafting of informative features, which are the parameters the algorithms learn from. Thus, the goal is often to create features that are readable by the algorithm and convey a high amount of information. Machine learning methods need numerical features to work, which means that in the case of non-numerical data, such as text, the creation of features becomes more complicated as it requires the data to be converted to numerical representations. This can be achieved using Natural Language Processing (NLP) methods, which essentially work at the intersection of computer and human language, to bridge the understanding of the two domains (Jurafsky & Martin, 2019; Nadkarni et al., 2011). By utilizing NLP methods, text (e.g., characters, words, sentences) can be converted to numerical data, which can be processed and analyzed by a computer. This can be as simple as creating frequencies of words that occur in the text, but also includes numerical representations of grammatical structures or semantic meanings (Goldberg & Levy, 2014). For a more detailed overview of NLP methods and how they function, readers are referred to Goldberg (2016), Jurafsky and Martin (2019), and Nadkarni et al. (2011).

Supervised machine learning methods are powerful tools, which can uncover previously unknown patterns in large amounts of data and find usages not only on the surface web but also on the dark web. However, supervised models need well-created data sets, with suitable features and labels, which are not always easy to obtain. In particular, the labeling process is often problematic as the “ground-truth” of the labels is often unknown (see above).

Ground-truth

In the context of this chapter, ground-truth refers to the true labels of a data set. The knowledge about ground-truth or certainty of the label could be about an item on a selling platform, a confession, or anything else. In an experimental setting, in which researchers have full control, true labels can be obtained. For example, participants might be asked to test a pair of purchased headphones, followed by instructions to write an honest and fake review (with opposing sentiments) about them. However, outside of an experimental setting, and in most other cases, the ground-truth is not as unambiguous and only a level of certainty can be attributed to a particular data point. For example, how would we know if a shoe advertised on an online platform was counterfeit or legitimate? A common strategy employed by consumers is to choose a reputable seller by inspecting their ratings and reviews (Dellarocas, 2006; Houser & Wooders, 2006; Melnik & Alm, 2003; Resnick et al., 2000). To provide a better label, an expert could assess the listing, or better still, order a pair of shoes and compare them to a sample purchased directly from the manufacturer. Each method will provide an assessment of the listing with different levels of certainty about the labels (counterfeit vs. genuine) associated with an item or event.

In many cases, obtaining labels with a high degree of certainty will either be impossible or costly in terms of time, money, or both (Raine et al., 2015). In some instances officials from online marketplaces or law enforcement (e.g. Trading Standards) order products from sellers or visit the sellers' physical locations to conduct manual inspections (Ganguly, 2015; Raine et al., 2015) but this is costly. Furthermore, transparency during the labeling process is important as is the clear conveyance of the label's limitations, to ensure that a judgment about the associated certainty of the labels is possible, without needing further expert knowledge. This is especially important to law enforcement, for whom resources are limited and misallocations can be costly. The strategy to minimize mislabeling is to reduce false positives

labels (e.g., mislabeling a genuine product to be fraudulent) and false-negative labels (e.g., mislabeling a fraudulent product to be genuine). Training a model to work well on a poorly labeled dataset is of little value and consequently, the determination of the data labels on which a supervised model is trained is paramount. Therefore, the value of a model should not be judged on performance metrics alone but also on how the data were acquired and labeled for each context.

Possible future data science applications

As described above, current data science approaches often focus on supervised machine learning methods which would support a proactive approach to the detection of possible frauds on the surface web before people are affected. However, given how the labels used for these approaches are obtained, it is unclear if such detections should always warrant a lawful intervention (Abdallah et al., 2016; Almendra & Enachescu, 2012; Chang & Chang, 2012; Hernandez-Castro & Roberts, 2015; Pandit et al., 2007; Xu et al., 2016). That said, these approaches could be used to presort and prioritize cases (e.g. suspicious advertisements, sellers, webpages) which authorities could then focus on manually, making better use of limited resources.

Other areas in which data science methods could alleviate the current workload of law enforcement are in gathering, pooling, and analyzing information about reported items, services, sellers, or webpages (National Audit Office, 2016; Raine et al., 2015; Trading Standards UK, personal communication, February 26, 2019).

Data science can also be used in other disciplines, such as psychology, to investigate fraud on a more individual level. For example, machine learning models could be utilized to uncover underlying patterns of how and why certain people fall victim to fraud. In turn, the

same approaches could be used to uncover patterns associated to individuals acting as a fraudster. However, discussing such use cases in more detail is out of the scope of this chapter, which aimed to highlight more direct applications, usable by practitioners or law enforcement.

Automation of manual tasks

An initial step in alleviating the workload would be to automate the standard processes of gathering and pooling information relevant to a reported online fraud. (e.g. reported from a consumer or a company). Information gathering and pooling could be addressed by using rule-based systems that could automatically collect information (e.g. telephone numbers, e-mail addresses, prices, product- or seller names) from webpages. Additional automated web searches could supplement and cross-reference the existing information, including whether they exist and in what capacity they are registered (e.g., does the physical address of the seller exist, and is the business legally registered with the relevant authority).

Uncovering unknown patterns in existing data

Existing data could be analyzed using unsupervised machine learning methods that do not rely on labeled training data. Here, the idea is to cluster the data based on inherent traits to find common features or properties that may not be directly apparent (Arthur & Vassilvitskii, 2006; Ester et al., 1996; Liu et al., 2012). For example, a consumer-reported incident of fraud or counterfeit items on shopping platforms might be clustered into several groups. These groups could be based on similarities in their descriptions, prices, item locations, and so on. Such similarities could point towards meaningful associations, which may open up new leads in an investigation. For example, where items are identified as sharing features such as similar descriptions, seller locations, and images, but that are advertised by sellers with different

usernames, might point to a common seller that wishes to conceal their identity and fraudulent activity.

Understanding online markets at scale

It is important to understand the online markets in which consumer fraud occurs as it places the reported fraud cases into context and may inform new fraud detection strategies. Data science can help in accumulating, structuring, and analyzing data for online markets, such as eBay, Amazon, Alibaba, or any darknet market. By devising scraping methods for these markets, large amounts of data for products and sellers can be collected. Through NLP methods it is also possible to operationalize and measure text characteristics and styles from titles, descriptions, reviews, and so on. Understanding markets with such data will (we hope) enable us to update our beliefs about common product and seller characteristics, and determine what constitutes “normal”, and understand how fraudsters might behave.

Utilizing darknet markets to identify consumer fraud on the surface web

Identifying illegal activity on the surface web is a challenge since most products and services are advertised as being genuine or legal. In contrast, on the dark web, most of the services are, by definition, illegal, and generally advertised with no effort made to conceal this. This raises an interesting possibility — could information found on the dark web serve as a source of intelligence to determine what illegal activity is or might happen on the surface web. Although categories on dark markets surrounding consumer fraud only represent a small portion (17%) of what is sold or traded (EMCDDA-Europol, 2017; Europol, 2017), they are often labeled as such. Thus, such information could serve as a source of ground truth data for illegal products and services on the surface web, where most fraudulent sales occur.

Moreover, to increase profits, it appears that vendors sell on multiple darknet markets *and* the surface web (Europol, 2017), suggesting that cross-domain referencing could be of value.

As far as we are aware, to date, research has not examined this possibility of using such data to inform fraudulent product identification on the surface web. The focus of such cross-referencing would be the identification of counterfeit goods, which is the main category of items on the dark web that also appears on the surface web. However, fraud-related categories, such as personal credentials, guides, and tutorials may also be of interest. For example, if hacked seller accounts (e.g., eBay), or specific fraud strategies are frequently advertised, their occurrence could represent an indicator of security breaches, which might be worth investigating. To examine the feasibility of such a cross-referencing approach, historical darknet data could be compared to law enforcement datasets of online consumer fraud for the same periods. In the case that such connections can be made, further work could focus on monitoring current darknet markets and cross-referencing them with current listings on the surface web. Since counterfeits on darknet markets are mostly clearly labeled and contain many product details (e.g. name, price, pictures, and description), a similar automated cross-referencing approach could be taken as has previously been described in studies (see above) that have sought to identify individual vendors (Hayes et al., 2018). Identifying fraud (enabling) services, such as hosting phishing sites or well established eBay accounts from the dark web on the surface web, might not be as easy to identify, as they often do not have similar listed details as fraudulent products before the sale (Thomas et al., 2015). However, depending on the level of details provided by vendors offering such services, these listings could also serve as an informative tool for researchers seeking to identify illegal activity on the surface web.

Conclusion

The costs associated with online consumer fraud are significant and the detection and their prevention using data science approaches have seen some progress. However, they are not yet well integrated or used outside of the academic literature. Using and implementing data science methods is complicated by a range of factors, such as the interplay of different sectors and jurisdictions that are affected, seller anonymity, and data labeling. A pressing issue is the volume of fraud complaints, most of which cannot be dealt with appropriately due to the lack of personnel and the highly time-consuming manual processing of such cases. Data science approaches would offer some help in automating and speeding up this (currently) manual work. Furthermore, darknet markets are increasingly utilized to sell fraudulent products online, such as counterfeits, or fraud enabling services. However, it is still unclear to what extent and how fraudulent transactions on these markets are related to those on the surface web. Nevertheless, such dark web listings could serve as an informative tool for detecting and identifying fraudulent behavior on the surface web.

References

- Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90–113. <https://doi.org/10.1016/j.jnca.2016.04.007>
- Ablon, L., Libicki, M., & Abler, A. (2014). *Markets for Cybercrime Tools and Stolen Data: Hackers' Bazaar*. RAND Corporation. <https://doi.org/10.7249/RR610>
- Adamsson, H. (2017). *Classification of illegal advertisement*. Uppsala University.
- Almendra, V., & Enachescu, D. (2012). A Fraudster in a Haystack: Crafting a Classifier for Non-delivery Fraud Prediction at Online Auction Sites. *2012 14th International Symposium on Symbolic*

and Numeric Algorithms for Scientific Computing, 233–239.

<https://doi.org/10.1109/SYNASC.2012.21>

Anderson, K. (2019). *Mass-Market Consumer Fraud in the United States: A 2017 Update* (p. 153).

Anderson, K. (2021). Mass-Market Consumer Frauds: What the Data Show. In *A Fresh look at Fraud: Theoretical and Applied Approaches*.

Anderson, R., Barton, C., Böhme, R., Clayton, R., van Eeten, M. J. G., Levi, M., Moore, T., & Savage, S. (2013). Measuring the Cost of Cybercrime. In R. Böhme (Ed.), *The Economics of Information Security and Privacy* (pp. 265–300). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-39498-0_12

Arthur, D., & Vassilvitskii, S. (2006). *k-means++: The Advantages of Careful Seeding*. 11.

Ball, M., Broadhurst, R., Niven, A., & Trivedi, H. (2019). *Data Capture and Analysis of Darknet Markets*. 15.

Baravalle, A., & Lee, S. W. (2018). Dark Web Markets: Turning the Lights on AlphaBay. In H. Hacid, W. Cellary, H. Wang, H.-Y. Paik, & R. Zhou (Eds.), *Web Information Systems Engineering – WISE 2018* (Vol. 11234, pp. 502–514). Springer International Publishing. http://link.springer.com/10.1007/978-3-030-02925-8_35

Bauerly, R. J. (2009). ONLINE AUCTION FRAUD AND EBAY. *Marketing Management Journal*, 19(1), 134–144.

Beals, M., DeLiema, M., & Deevy, M. (2015). *Framework for a taxonomy of fraud* (p. 40). Financial Fraud Research Center; Stanford Center on Longevity; FINRA Investor Education Foundation. <http://162.144.124.243/~longevl0/wp-content/uploads/2016/03/Full-Taxonomy-report.pdf>

Bergman, M. K. (2001). White Paper: The Deep Web: Surfacing Hidden Value. *Journal of Electronic Publishing*, 7(1). <http://dx.doi.org/10.3998/3336451.0007.104>

Biddle, P., England, P., Peinado, M., & Willman, B. (2003). The Darknet and the Future of Content Protection. In J. Feigenbaum (Ed.), *Digital Rights Management* (Vol. 2696, pp. 155–176). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-44993-5_10

- Branwen, G., Christin, N., Décary-Héту, D., Andersen, R. M., StExo, El Presidente, Anonymous, Lau, D., Sohlz, Kratunov, D., Cacic, V., Whom, McKenna, M., & Goode, S. (2015). *Dark Net Market archives, 2011-2015* (2015-07-12). <https://www.gwern.net/DNM-archives>
- Calis, T. (2018). *Multi-homing sellers and loyal buyers on darknet markets*. Erasmus University.
- Chang, W.-H., & Chang, J.-S. (2012). An effective early fraud detection method for online auctions. *Electronic Commerce Research and Applications*, 11(4), 346–360.
<https://doi.org/10.1016/j.elerap.2012.02.005>
- Christin, N. (2013). Traveling the silk road: A measurement analysis of a large anonymous online marketplace. *Proceedings of the 22nd International Conference on World Wide Web - WWW '13*, 213–224. <https://doi.org/10.1145/2488388.2488408>
- Christin, N., Yanagihara, S. S., & Kamataki, K. (2010). Dissecting one click frauds. *Proceedings of the 17th ACM Conference on Computer and Communications Security - CCS '10*, 15.
<https://doi.org/10.1145/1866307.1866310>
- Clark, P., & Niblett, T. (1989). The CN2 induction algorithm. *Machine Learning*, 3(4), 261–283.
<https://doi.org/10.1007/BF00116835>
- Deevy, M., & Beals, M. (2013). *The scope of the problem* (p. 46). Financial Fraud Research Center; Stanford Center on Longevity; FINRA Investor Education Foundation.
- DeLiema, M., Fletcher, E., Kieffer, C. N., Mottola, G. R., & Pessanha, R. (2019). *Exposed to scams. What separates victims from non-victims?* (p. 24). Stanford Center on Longevity, Federal Trade Commission, FINRA Foundation, International Association of Better Business Bureaus, BBB Institute for Marketplace Trust.
- DeLiema, M. I., Mottola, G. R., & Deevy, M. (2017). *Findings from a Pilot Study to Measure Financial Fraud in the United States*. Stanford Center on Longevity; FINRA Investor Education Foundation. <https://www.ssrn.com/abstract=2914560>
- Dellarocas, C. (2006). Reputation Mechanisms. In *Handbook on Economics and Information Systems* (p. 38). Elsevier.

- Demant, J., Munksgaard, R., & Houborg, E. (2018). Personal use, social supply or redistribution? Cryptomarket demand on Silk Road 2 and Agora. *Trends in Organized Crime*, 21(1), 42–61.
<https://doi.org/10.1007/s12117-016-9281-4>
- Dhar, V. (2013). Data science and prediction. *Communications of the ACM*, 56(12), 64–73.
<https://doi.org/10.1145/2500499>
- Du, P.-Y., Zhang, N., Ebrahimi, M., Samtani, S., Lazarine, B., Arnold, N., Dunn, R., Suntwal, S., Angeles, G., Schweitzer, R., & Chen, H. (2018). Identifying, Collecting, and Presenting Hacker Community Data: Forums, IRC, Carding Shops, and DNMs. *2018 IEEE International Conference on Intelligence and Security Informatics (ISI)*, 70–75.
<https://doi.org/10.1109/ISI.2018.8587327>
- EMCDDA-Europol. (2017). *Drugs and the darknet: Perspectives for enforcement, research and policy*. Publications Office of the European Union.
- Ester, M., Kriegel, H.-P., & Xu, X. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining, Portland, OR, AAAI Press*, 226–231.
- Europol. (2017). *INTELLECTUAL PROPERTY CRIME ON THE DARKNET*. Enforcement Cooperation.
<https://www.europol.europa.eu/publications-documents/intellectual-property-crime-darknet>
- FBI. (2017). *2017 Internet Crime Report* (p. 29). Federal Bureau of Investigation.
- FBI. (2018). *Internet Crime Report*. Federal Bureau of Investigation.
https://pdf.ic3.gov/2018_IC3Report.pdf
- Ganguly, P. (2015, January 8). How e-retailers such as Flipkart, Amazon are keeping the fake products at bay. *The Economic Times; New Delhi*, 2.
- Gehl, R. W. (2018). Archives for the Dark Web: A Field Guide for Study. In Lewis Levenberg, T. Neilson, & D. Rheams (Eds.), *Research Methods for the Digital Humanities* (pp. 31–51). Springer International Publishing. https://doi.org/10.1007/978-3-319-96713-4_3

- Ghosh, S., Das, A., Porras, P., Yegneswaran, V., & Gehani, A. (2017). Automated Categorization of Onion Sites for Analyzing the Darkweb Ecosystem. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '17*, 1793–1802. <https://doi.org/10.1145/3097983.3098193>
- Ghosh, S., Porras, P., Yegneswaran, V., Nitz, K., & Das, A. (2017). ATOL: A Framework for Automated Analysis and Categorization of the Darkweb Ecosystem. In *AAAI Workshops* (p. 9).
- Goldberg, Y. (2016). A Primer on Neural Network Models for Natural Language Processing. *J. Artif. Intell. Res.(JAIR)*, 57, 345–420.
- Goldberg, Y., & Levy, O. (2014). word2vec Explained: Deriving Mikolov et al.'s negative-sampling word-embedding method. *ArXiv:1402.3722 [Cs, Stat]*. <http://arxiv.org/abs/1402.3722>
- Gray, H. (2019). *Dark Web Map*. <https://www.hyperiongray.com/dark-web-map/#zoom=0.8521016982969332&x=0.5064520330563047&y=0.572866049039204>
- Hanoch, Y., & Wood, S. (2021). Introduction. In *A Fresh look at Fraud: Theoretical and Applied Approaches*. Routledge.
- Harper, J. (2021, March 10). Trader gets painted stones instead of \$36m of copper. *BBC News*. <https://www.bbc.com/news/business-56330378>
- Hayes, D., Cappa, F., & Cardon, J. (2018). A Framework for More Effective Dark Web Marketplace Investigations. *Information*, 9(8), 186. <https://doi.org/10.3390/info9080186>
- He, B., Patel, M., Zhang, Z., & Chang, K. C.-C. (2007). Accessing the deep web. *Communications of the ACM*, 50(5), 94–101. <https://doi.org/10.1145/1230819.1241670>
- Herley, C. (2010). The Plight of the Targeted Attacker in a World of Scale. *The Ninth Workshop on the Economics of Information Security (WEIS) 2010, Harvard University, USA.*, 12.
- Hernandez-Castro, J., & Roberts, D. L. (2015). Automatic detection of potentially illegal online sales of elephant ivory via data mining. *PeerJ Computer Science*, 1, e10. <https://doi.org/10.7717/peerj-cs.10>

- Ho, T. N., & Ng, W. K. (2016). Application of Stylometry to DarkWeb Forum User Identification. In K.-Y. Lam, C.-H. Chi, & S. Qing (Eds.), *Information and Communications Security* (Vol. 9977, pp. 173–183). Springer International Publishing. https://doi.org/10.1007/978-3-319-50011-9_14
- Houser, D., & Wooders, J. (2006). Reputation in Auctions: Theory, and Evidence from eBay. *Journal of Economics & Management Strategy*, 15(2), 353–369. <https://doi.org/10.1111/j.1530-9134.2006.00103.x>
- Intellectual Property Office. (2019). *Annotation task* [Letter to Felix Soldner].
- Jurafsky, D., & Martin, J. H. (2019). *Speech and Language Processing An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition—Third Edition draft*. <https://web.stanford.edu/~jurafsky/slp3/>
- Lewis, S., Jamie. (2016, July 3). *OnionScan Report June 2016—Snapshots of the Dark Web*. Mascherari Press. <https://mascherari.press/onionscan-report-june-2016/>
- Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2012). Isolation-Based Anomaly Detection. *ACM Transactions on Knowledge Discovery from Data*, 6(1), 1–39. <https://doi.org/10.1145/2133360.2133363>
- Mansfield-Devine, S. (2009). Darknets. *Computer Fraud & Security*, 2009(12), 4–6. [https://doi.org/10.1016/S1361-3723\(09\)70150-2](https://doi.org/10.1016/S1361-3723(09)70150-2)
- Melnik, M. I., & Alm, J. (2003). Does a Seller's eCommerce Reputation Matter? Evidence from eBay Auctions. *The Journal of Industrial Economics*, 50(3), 337–349. <https://doi.org/10.1111/1467-6451.00180>
- Miramirkhani, N., Starov, O., & Nikiforakis, N. (2017). Dial One for Scam: A Large-Scale Analysis of Technical Support Scams. *Proceedings 2017 Network and Distributed System Security Symposium*. Network and Distributed System Security Symposium, San Diego, CA. <https://doi.org/10.14722/ndss.2017.23163>
- Nadkarni, P. M., Ohno-Machado, L., & Chapman, W. W. (2011). Natural language processing: An introduction. *Journal of the American Medical Informatics Association*, 18(5), 544–551. <https://doi.org/10.1136/amiajnl-2011-000464>

- National Audit Office. (2016). *Protecting consumers from scams, unfair trading and unsafe goods* (HC 851). National Audit Office; Department for Business, Energy & Industrial Strategy.
- Office for National Statistics. (2020). *Nature of fraud and computer misuse in England and Wales: Year ending March 2019* (p. 29). Office for National Statistics.
- Pandit, S., Chau, D. H., Wang, S., & Faloutsos, C. (2007). Netprobe: A fast and scalable system for fraud detection in online auction networks. *Proceedings of the 16th International Conference on World Wide Web - WWW '07*, 201. <https://doi.org/10.1145/1242572.1242600>
- Paterva. (2019). *Paterva Home*. PATERVA A New Train of Thought.
<https://www.paterva.com/index.php>
- Peaston, S. (2019). *The Fraudscape* (p. 19). Cifas.
<https://www.cifas.org.uk/secure/contentPORT/uploads/documents/Cifas%20Fraudscape%202019%20Full%20Digital%20Report%20.pdf>
- Pointer Brand Protection. (2019). *Online Brand Protection With An Impact*. Pointer Brand Protection.
<https://pointerbrandprotection.com/>
- Provost, F., & Fawcett, T. (2013). Data Science and its Relationship to Big Data and Data-Driven Decision Making. *Big Data*, 1(1), 51–59. <https://doi.org/10.1089/big.2013.1508>
- Raine, J., Mangan, C., & Watt, P. (2015). *THE IMPACT OF LOCAL AUTHORITY TRADING STANDARDS IN CHALLENGING TIMES* (p. 148). The Department for Business, Innovation and Skills and The Trading Standards Institute.
- Resnick, P., Kuwabara, K., Zeckhauser, R., & Friedman, E. (2000). Reputation systems. *Communications of the ACM*, 43(12), 45–48. <https://doi.org/10.1145/355112.355122>
- Rosenbusch, H., Soldner, F., Evans, A. M., & Zeelenberg, M. (2021). Supervised machine learning methods in psychology: A practical introduction with annotated R code. *Social and Personality Psychology Compass*. <https://doi.org/10.1111/spc3.12579>
- Rusch, J. J. (1991). *The “Social Engineering” of Internet Fraud*. 12.
http://www.isoc.org/isoc/conferences/inet/99/proceedings/3g/3g_2.htm

- Sahingoz, O. K., Buber, E., Demir, O., & Diri, B. (2019). Machine learning based phishing detection from URLs. *Expert Systems with Applications*, 117, 345–357.
<https://doi.org/10.1016/j.eswa.2018.09.029>
- Soska, K., & Christin, N. (2015). Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem. *Proceedings of the 24th USENIX Security Symposium*, 17.
- Stanford Center on Longevity. (2019). *Financial Fraud – Stanford Center on Longevity*.
<http://longevity.stanford.edu/2017/03/29/safeguarding-clients-from-financial-fraud-and-exploitation/>
- The Invisible Internet Project. (2020). *I2P Anonymous Network*. <https://geti2p.net/en/>
- The Tor Project, Inc. (2020). *The Tor Project | Privacy & Freedom Online*. <https://torproject.org>
- Thomas, K., Huang, D. Y., Wang, D., Bursztein, E., Grier, C., Holt, T. J., Kruegel, C., McCoy, D., Savage, S., & Vigna, G. (2015). Framing Dependencies Introduced by Underground Commoditization. *Proceedings (Online) of the Workshop on Economics of Information Security (WEIS)*, 24.
- Trading Standards UK. (2019, February 26). *Interview about possible data science solutions* [Personal communication].
- Trading Standards UK. (2019, May 17). *Interview about Trading Standards areas of operations, and current issues concerning online crime* [Personal communication].
- Van Buskirk, J., Naicker, S., Bruno, R. B., Breen, C., & Roxburgh, A. (2016). *Drugs and the Internet*.
https://www.drugsandalcohol.ie/20369/1/NDARC_Drugs&TheInternet_Bulletin1.pdf
- Van Buskirk, J., Roxburgh, A., Farrell, M., & Burns, L. (2014). The closure of the Silk Road: What has this meant for online drug trading?: Editorial. *Addiction*, 109(4), 517–518.
<https://doi.org/10.1111/add.12422>
- Van Buskirk, J., Roxburgh, A., Naicker, S., & Burns, L. (2015). A response to Dolliver’s “Evaluating drug trafficking on the Tor network.” *International Journal of Drug Policy*, 26(11), 1126–1127.
<https://doi.org/10.1016/j.drugpo.2015.07.001>

- van Wegberg, R., Tajalizadehkhoob, S., Soska, K., Akyazi, U., Ganan, C., Klievink, B., Christin, N., & van Eeten, M. (2018). Plug and Prey? Measuring the Commoditization of Cybercrime via Online Anonymous Markets. *Proceedings of the 27th USENIX Security Symposium (USENIX Security'18)*, 19.
- Vistalworks. (2019). *Vistalworks—Keeps online shoppers safe from harm*. Vistalworks.
<https://vistalworks.com>
- Wang, X. (2018). *Photo-based Vendor Re-identification on Darknet Marketplaces using Deep Neural Networks* [Master Thesis]. Faculty of the Virginia Polytechnic Institute and State University.
- Wang, X., Peng, P., Wang, C., & Wang, G. (2018). You Are Your Photographs: Detecting Multiple Identities of Vendors in the Darknet Marketplaces. *Proceedings of the 2018 on Asia Conference on Computer and Communications Security - ASIACCS '18*, 431–442.
<https://doi.org/10.1145/3196494.3196529>
- Xu, J.-C., Shin, K., & Liu, Y.-L. (2016). Detecting Fake Sites based on HTML Structure Analysis. *Proceedings of the 6th International Conference on Communication and Network Security - ICCNS '16*, 86–90. <https://doi.org/10.1145/3017971.3017980>
- Yellow Brand Protection. (2019). *Anti-Counterfeiting | Yellow Brand Protection*.
<https://www.yellowbrandprotection.com/services/anti-counterfeiting>