

Quit Playing Games With My Heart: Understanding Online Dating Scams

JingMin Huang, Gianluca Stringhini[†], and Peng Yong

University College London[†]

jhua8590@uni.sydney.edu.au, g.stringhini@ucl.ac.uk, pengyong20@picc.com.cn

Abstract. Online dating sites are experiencing a rise in popularity, with one in five relationships in the United States starting on one of these sites. Online dating sites provide a valuable platform not only for single people trying to meet a life partner, but also for cybercriminals, who see in people looking for love easy victims for scams. Such scams span from schemes similar to traditional advertisement of illicit services or goods (i.e., *spam*) to advanced schemes, in which the victim starts a long-distance relationship with the scammer and is eventually extorted money.

In this paper we perform the first large-scale study of online dating scams. We analyze the scam accounts detected on a popular online dating site over a period of eleven months, and provide a taxonomy of the different types of scammers that are active in the online dating landscape. We show that different types of scammers target a different demographics on the site, and therefore set up accounts with different characteristics. Our results shed light on the threats associated to online dating scams, and can help researchers and practitioners in developing effective countermeasures to fight them.

1 Introduction

Online dating sites have become a popular solution for users to meet people and start relationships. The most popular dating sites have between 15 and 20 million active members, and the revenue of the whole online dating industry in 2012 was estimated to exceed one billion dollars [4]. As it happens for any popular online service, online dating sites attract cybercriminals too. This should not surprise, since online services are commonly plagued with spam [19] and malware [6]. Online dating sites, however, have a very different purpose than common online services: meeting people in real life and possibly starting a relationship. For this reason, such services attract more advanced scammers than other online services, who exploit the vulnerable emotional state of online dating users for financial gain [28]. As an example, scammers commonly set up fake accounts on an online dating site, start interacting with a user on the site, and then lure her into sending them money, for example to pay for the flight needed to meet in person [36]. Such scams are similar in spirit to the infamous “419 scams,” in which scammers ask their victims to send them a sum of money to establish trust and then promise to transfer a very large sum to them [17, 18, 23, 25]. Online dating scams, however, are more insidious than “419 scams,” because they target emotionally vulnerable people looking for love.

Compared to traditional malicious activity on online services, the one happening on online dating sites shows three main differences. The first difference is that malicious activity on online services (social networks, blogs, webmail services) is typically run in large-scale campaigns [13, 16]. As a result of this, malicious activity is automatically generated and can be detected by leveraging similarities across the same malicious campaign [26, 37]. When dealing with online dating sites however, this assumption does not hold anymore: scammers are usually real people, writing personalized messages to their victims [36]. The second difference is that, unlike traditional spam and malware attacks, online dating scams can develop over a long period of time. Scammers typically exchange many messages with their victims to win their trust, before performing the actual scam [36]. In some cases, the scam is performed once the victim and the scammer meet in person. The third difference is that, unlike other online services, online dating sites are designed to put in contact people who do not know each other. For this reason, the concept of unsolicited message, which is the core of traditional anti-spam systems, does not have any meaning when applied to online dating sites: all messages received on such sites are in fact unsolicited.

The landscape of online dating scams is widely unstudied by the research community. Previous research in this field focused on describing single scam schemes, and relied on descriptions of single incidents instead of performing large-scale measurements of the phenomenon [28, 36]. In this paper, we present the first comprehensive study of scams targeting online dating sites. We analyze the accounts used by scammers that have been identified over a period of one year on a large online dating site in China. Since the operators of the dating site do not want the name of the service to be disclosed, we will refer to it as `DATINGSITE` in this paper. We discuss the different types of scams that we identified, showing that the threats that online dating users are exposed to are usually different than the ones that are faced by the users of traditional online services (such as Online Social Networks).

Given the different nature of the threats that users face on online dating sites, current systems that detect malicious activity on online services are not enough to protect the users of such sites. This paper aims at providing the research community with insights on how online dating scammers operate, on the types of threats that users face on such platforms, and on typical traits and behaviors of the accounts that are used by scammers to perform their operations. We hope that our observations will shed some light on the problem of online dating scams, and help researchers and online dating sites operators develop better detection methods to keep their users safe.

In summary, this paper makes the following contributions:

- We discuss the threat model associated with scammers operating on online dating sites, outlining the differences between this type of malicious activity and the one that is found on other online services.
- We analyze more than 500,000 accounts used by scammers on a popular online dating site, and provide a taxonomy of the most prevalent online dating scams. In particular, we identified four types of scams. Cybercriminals performing different types of scams present a different *modus operandi* in interacting with victims, and a different level of sophistication.

- We provide detailed statistics and case studies on the detected scam accounts. We show that different types of scams target different demographics on the site, and that specific scam schemes have a higher success in receiving attention by the users of online dating sites.

2 Background and Problem Study

In this section, we first describe online dating sites in general, giving an overview of the functionalities that are typically offered by these sites to their users. Then, we describe the online dating site that we analyzed in this paper.

2.1 Online Dating Sites

There are a wealth of online dating sites on the Internet. Some of them cater to audiences with a specific ethnic or cultural background (e.g., `christianmingle.com`), while some others are targeted at all types of users (e.g., `match.com`). Some sites just aim at making people meet, while others have specific types of relationships as a target (for example marriage).

In general, the first thing users have to do after signing up on an online dating site is setting up a profile. The profile is what other users see, and having a complete and well-written one influences the first impression that possible matches have of the person [22]. Users are encouraged to add personal pictures and a description of themselves to the profile. In addition, people can add information about their favorite activities and hobbies. Users are required to add their sexual preference as well, and can specify the age range of the people they would like to meet.

All the information that the user inputs is processed by a matching algorithm. The algorithm compares the information on the user's profile with the one on the profiles of possible matches and displays to the user the profiles of people that she would probably like. The user can then review these suggestions and contact those people with whom she wants to start a conversation. Some sites allow users to browse all profiles on the site, while others restrict them to only see those profiles that were highly ranked as possible matches for her [1].

A major difference between online dating sites is the subscription price: unlike online social networks, creating a profile on an online dating site is usually not free, and the user has to pay a monthly subscription to use the functionalities of the site. A subscription to a popular online dating site ranges from \$13 and \$24 per month [3]. On the other hand, a handful of online dating sites (for example `okcupid.com` [2]) offer free subscriptions, and their websites feature advertisements, similarly to what happens on traditional online social network sites.

The amount of effort required to create an online dating profile influences the way in which cybercriminals use these services. Intuitively, the high price of subscription to most of these websites makes it unsuitable for spammers to create fake accounts in bulk. Similarly, the high amount of information needed to create a believable profile on the free online dating sites limits the effectiveness of mass-created fake accounts. For this reason, miscreants use online dating sites to perform more advanced scams,

which rely on personal interactions and social engineering. We will describe the types of scams that we identified on the online dating site that we analyzed in Section 4.

2.2 Case Study: A Large Chinese Dating Site

We performed our analysis on a large online dating site in China. For confidentiality reasons, we will refer to it as DATINGSITE in this paper. DATINGSITE has more than 10 million users, which gives it a comparable user base to the most successful online dating sites worldwide.

DATINGSITE presents all the elements typical of online dating sites that we described. After registering, users have to set up a profile, including information such as their age, gender, education, marital status, etc. Users can then browse other users' profiles and contact people they like.

Unlike most online dating sites, users can create a profile on DATINGSITE for free. This fact makes it a particularly convenient platform for scammers, who can set up their accounts at no cost. To keep their users safe from scammers, DATINGSITE deployed a number of detection mechanisms that are able to flag possible scam accounts. Because the false positives of such systems are higher than what is considered acceptable in a production system, and blocking a legitimate account by mistake would be very negative for the dating site's reputation, DATINGSITE employs a team of experts that vet the flagged accounts, deciding which ones actually belong to scammers. If an account is detected as controlled by a scammer, the profile is "frozen" until the user confirms her identity and is forbidden from contacting other profiles on the site. In this paper, we analyze the accounts flagged as belonging to scammers by these human specialists over a period of one year.

2.3 Threat Model: Online Dating Scams

As we mentioned earlier, the main difference between traditional online services and online dating sites is that the latter are designed to put in contact people who have no connection whatsoever. In this context, the concept of *unsolicited message*, which is a strong indicator of maliciousness on other online services, has no meaning: all messages are "unsolicited," but users are happy to receive them instead of being annoyed by them. For this reason, we need to go beyond considering any unsolicited message as malicious and formulate a more advanced threat model.

In this paper, we consider an online dating user a scammer if he/she is using the service to take advantage (often economic) of another user. A scammer will set up one or more accounts on the online dating site, and interact with the users of the site. We call such accounts *scam accounts*. Online dating scams can be more or less sophisticated. In some cases, the scam accounts are just advertising goods or services, similarly to traditional spam (for example escort services). In this case the scam content (for example the contact information of the escort agency) is sent to the victim very early, possibly in the first message that is exchanged. In some other cases, however, scammers are more sophisticated, and establish a long-distance relationship with the victim before performing the actual scam. In many cases, the scammer tries to convince the victim to continue the conversation on a different medium, for example Skype. This is

an additional reason why online dating scams are difficult to detect: often the online dating site administrators do not see the scam happening, because the scammer and the victim have moved to a different way of communicating.

In the rest of the paper we first describe the way in which we collected a set of more than 500,000 scam accounts on DATINGSITE. We then present a taxonomy of the scam accounts that we observed, and discuss some typical characteristics of such accounts that could be used for detection.

3 Methodology

Given the difficulty of automatically detecting advanced scam accounts, online dating sites employ customer-service specialists who manually review suspicious profiles and suspend the ones that belong to scammers. These customer-service specialists are experts in detecting scammers, and therefore they can reliably assess the maliciousness of an account. However, it is unfeasible (and intrusive) for these specialists to analyze every single profile on the site and assess its maliciousness. For this reason, specialists are aided by automatic programs that narrow down the number of possible scammers as much as possible. Ideally, these detection systems should have high accuracy, so that the human specialists can quickly decide whether a profile belongs to a scammer. Their accuracy, however, is not high enough to justify a completely-automated scam detection system – this is due to the complex nature of online dating scams as we previously discussed. In addition, the cost of false positives for the company is very high: a user having his/her account suspended by mistake would leave the site and move to a competitor, and even ask for a refund in case of a paid dating site.

In the following, we briefly describe the four detection systems that help the customer service specialists at DATINGSITE in detecting scam accounts. Two of the authors of this paper worked on the development of such systems. Because these systems resemble, in large part, anti-spam systems that have been proposed by the research community over the years, we do not claim that they are novel, and we include them for the sake of completeness, and to give the reader a better idea of how the dataset of scam accounts used in the rest of the study was collected.

3.1 Behavioral-based Detection System

The goal of scammers on online dating sites is very different from the one of legitimate users: while legitimate users want to get to know new people, and possibly start a romantic relationship, scammers seek vulnerable and gullible victims, with the purpose of extorting money from them. For this reason, the behavior of accounts controlled by scammers is likely to show differences than the one of legitimate users. To capture these differences, DATINGSITE developed a detection system that models the typical behavior of scam accounts (as opposed to legitimate accounts). This system looks at two types of account characteristics. The first one are profile traits that scam accounts typically show (as we will see in Section 6, specific types of scam accounts pose as a particular demographic to appeal a particular type of victim). The second type of characteristics are related to the typical behavior of scam accounts. Such characteristics include the

number of conversations initiated simultaneously, the time waited between the creation of the profile and the first message, and the fraction of the received messages to which the account replies. This system is similar to other anti-spam systems that have been proposed by the research community [13, 24, 31].

3.2 IP Address-based Detection System

Systems that detect automated misuse on online services often look at the reuse of IP addresses by miscreants [20, 30]. This element is not as important when dealing with online dating scammers, because, as we will show in Section 6.2, they typically do not use high levels of automation for their operations. However, IP address reuse is still an useful indicator to detect those scammers with a lower level of sophistication, who create a number of different profiles to advertise their businesses. Similarly, some sophisticated scammers might create a handful of profiles, and access them from the same IP address or network. For these reasons, DATINGSITE deployed a system that flags multiple accounts accessed by the same IP address as possibly malicious.

3.3 Photograph-based Detection System

Previous research noted that spammers that are active on Online Social Networks typically reuse profile pictures for their fake accounts [31]. For this reason, it makes sense to look for profiles that share the same profile picture to find possible fake accounts on those platforms. We observed the same trend happening on online dating sites. Although scams are usually run by humans and do not have the scale of spam campaigns on other online services, scammers often use the same picture (typically of an attractive young woman or a handsome middle-aged man) for a multitude of accounts. DATINGSITE deployed a system that is able to detect duplicated profile pictures and flag those accounts as possibly malicious. This system is based on a perceptual hashing algorithm [39], to be able to detect images that look the same, but have been re-encoded.

3.4 Text-based Detection System

Content analysis is a popular way of fighting spam [11, 29]. Such systems typically identify words that are indicative of malicious and benign content and leverage them for detection. Content analysis systems work well in automatically detecting spam. In our specific case, they can be useful in detecting scammers who employ a low level of sophistication (for example, who send a copy of the same message to many other profiles). For this reason, DATINGSITE developed a system that looks for keywords in messages that are typical of scam accounts, and flag them as malicious.

Because of their differences, each detection system is more suited to detect different types of scam accounts. For example, the photograph and text based systems are more suited to detect accounts that are part of large scale campaigns and share the same picture or the same profile text. On the other hand, the behavioral-based detection system works best in detecting advanced scammers, who set up their own accounts and exchange a long series of manually-crafted messages with the victim before performing the actual scam.

4 Description of the Scam Account Dataset

Previous work showed anecdotes and case studies of scams on online dating sites [28, 36]. In this paper, we provide the first large-scale measurement of online dating scams, performed over a period of 11 months between 2012 and 2013. Our analysis is based on a set of scam accounts detected by the systems described previously on DATINGSITE. In particular, we analyze a dataset composed of the scam accounts detected on DATINGSITE during the period of observation. In total, the dataset is composed of 510,503 scam accounts.

Note that, since the detections performed by the detection systems were manually vetted by human analysts, we are confident that they were correct, taking aside potential mistakes made by the analysts. This “clean” set of malicious accounts allowed us to come up with a taxonomy of online dating scammers. In the next section we discuss this taxonomy in detail.

5 A Taxonomy of Online Dating Scammers

With the help of the human analysts employed by DATINGSITE, we further analyzed the scam accounts in our dataset. We identified four types of scammers that infest the site: “*Escort Service Advertisements*,” “*Dates for Profit*,” “*Swindlers*,” and “*Matchmaking Services*.” Although some of the scammer schemes that we identified are specific to the Chinese culture and society, we think that this taxonomy holds also for dating sites located in other countries. In the following, we describe the types of scams that we identified. In Section 6 we analyze the scammers that we detected in the wild more in detail, breaking them down by scam type, and we investigate the different demographics that they use to attract victims, as well the different strategies that they use.

Escort Service Advertisements. As we said, the bulk creation of accounts with the purpose of spreading spam is not a viable solution for cybercriminals operating on online dating sites. However, this does not prevent miscreants from creating a number of accounts and sending unsolicited advertisements to the users of the dating site. Although this type of activity is not as predominant as it is on other media (such as online social networks [31] and email [33]) we still observed it on DATINGSITE. It is interesting to note that, since the main purpose of an online dating site is to connect two people who do not know each other, the concept of *unsolicited messages* has little meaning on such platforms. For this reason, it is more difficult to fight spammers on online dating sites than it is on other services.

On DATINGSITE, the vast majority of unsolicited advertisements promote escort agencies. In total, we detected 374,051 accounts of this type during our study. In this scheme, cybercriminals operate as follows: first, they create a profile belonging to a young woman (often including an attractive picture). Then, they start contacting other profiles, including the contact information of the escort agency in their messages. Note that considering accounts that contact a large number of profiles as malicious is usually not enough. In fact, legitimate accounts that contact many profiles, trying to establish a contact, are not uncommon on online dating sites; previous work showed that 78% of the messages sent on online dating sites never receive a reply from the other person [14]. We confirm this low turnaround of messages in Section 6.2.

Dates for Profit. As we said, the purpose of using online dating sites is to meet people in person and possibly start a relationship. This opens additional possibilities for scammers, who can leverage the fact that their victims will want to meet them in person and perform the actual scam once this happens. This element has no parallel on other online services. Even in advanced “419 scams,” which rely on social engineering, the only “physical” interaction between the scammer and the victim is the transfer of a sum of money.

On DATINGSITE, we observed an interesting trend in scams that exploit in-person meetings. Owners of establishments such as cafes and restaurants would hire girls to create profiles on the online dating site, contact multiple victims, and ask them to meet in person at that particular establishment. Such establishments are usually very expensive, and it is customary in China for males to pay for food and drink on dates. Therefore, the owners of the establishments can make a considerable amount of money out of these scams. Obviously, after this first encounter, the victim is never contacted again. To put this type of scam in perspective, a meal at most of the rogue establishments can cost between \$100 and \$2,000. This amount of money is similar to what is gained by traditional “419 scams” [12]. However, the success rate of this type of scam is much higher, because the scammer leverages the desire of the victim to meet an attractive woman. In addition, it is likely that the victim will never realize that he has been scammed, since the date really happened, and the victim possibly had a good time. Therefore, there is a low chance for the owner of the establishment to get caught. In total, we detected 57,218 accounts of this type during our study.

This particular type of scam appears to be a popular scheme for scammers on DATINGSITE, as we will show in Section 6. Although the cultural setting of China might make such scams more likely to be successful, we do not see a reason why similar scams should not be happening on dating sites based in other countries too.

Swindlers. Online dating sites are seen as a valuable resource for mid-aged people who so far have not found a partner for life. This type of person is particularly vulnerable to scams, and scammers take advantage of it. We observed multiple instances of scammers contacting mid-aged men and women, and establishing a long distance relationship with them. Over time, the trust that the victim has in the scammer would grow. After some time, the scammer will ask the victim to send him a sum of money to help in some task. The scammer might be in financial trouble or need money to buy a plane ticket to come visit the victim. At this point in the relationship, the victim is very likely to trust the scammer and send him the required money. After sending the requested amount of money, the victim will never hear back from the scammer. We call these type of scammers “*Swindlers*.” Swindlers are a big problem on online dating sites. Similar scam schemes have already been studied by previous work [28,36]. In total, we detected 43,318 accounts of this type during our study.

A particular type of swindlers, which are very peculiar to the culture of certain areas of China, are what we call “*Flower-basket Swindlers*.” In this particular type of scam, the scammer creates a fake profile, usually of an attractive mid-aged man, contacts a mid-aged woman, and starts exchanging messages with her. Over time, the victim will think that she is developing a romantic relationship with the scammer, and will start trusting him. After a while, the scammer will start implying that he wants to marry the

victim, and that his parents need some proof of the victim’s “good will.” At this point the scam happens. It is customary in those areas of China to send baskets of flowers to newly-opened shops, as gifts and to wish good luck. The scammer will then pretend to be opening a new shop, and will ask the victim to send some baskets of flowers to the shop. He will also give her the contact information of a florist in his area, who is actually one of the scammer’s accomplices. As for other swindler scams, once the victim buys the flowers, she will never hear from the scammer again. These shipments of flowers can be very expensive and cost up to \$20,000 to the victim.

Another type of swindlers that we observed on DATINGSITE are the so-called “*lottery-ticket swindlers*.” This type of scams is more similar to traditional “419 scams.” The miscreant establishes a romantic relationship with the victim, by pretending to be a successful businessperson in a foreign country, and lures the victim into sending him a sum of money, as part of a rewarding financial operation.

Swindlers are present on all online dating sites and are not specific to DATINGSITE or the Chinese culture in particular. On the other hand, flower-basket swindlers seem to be leveraging a specific Chinese tradition, and this type of scam is probably peculiar to Chinese online dating sites.

Previous research showed that scams perpetrated by email need to narrow down the number of potential victims that will reply to their messages, because each message requires some effort on the scammer part to reply to it [21]. For this reason, the scammer sets up a story that is hardly believable to a general audience, making sure that whoever responds is already prone to fall for the scam. In the case of online dating swindlers, this does not happen. The scammer starts a long-lasting conversation with the potential victims, and the purpose of the scam is revealed very late in the relationship. A partial selection of users that are more likely to fall for the scam is performed by setting up profiles that appeal a specific audience, such as mid-aged divorced women. We will discuss more details on these strategies in Section 6.1.

Matchmaking services. Before online dating sites were popular, people used to sign up for matchmaking services, which would find them a partner. We observe that agents from these matchmaking services are active on online dating sites too. They typically create fake profiles of attractive people, lure users into handing out some contact information, and then contact them directly, advertising the matchmaking agency. Although no financial scam is involved in the operation of matchmaking services, users are annoyed by them, and online dating sites see them as a competitor that is illicitly using their site to gain customers. We detected 35,916 accounts of this type during our study.

In the next section we provide more details on the scam accounts that we detected in the wild.

6 Analysis of the Scam Account Dataset

In this section we first analyze the different demographics and profile characteristics of the various types of scam accounts that we detected, comparing them to the legitimate accounts that were active on DATINGSITE during the analysis period. We then study the *modus operandi* of the different types of scammers, providing new insights on how these cybercriminal crews operate.

6.1 Demographics of Different Scam Account Types




Fig. 1. Gender distribution of the different types of scam accounts, compared to legitimate accounts.

As we mentioned, the different types of online dating scammers have very different goals, and cater to a very different pool of victims. For this reason, the accounts that they use to perpetrate their scams show very different characteristics, and impersonate people with different demographics. Not only, but the population of malicious accounts for each scam type shows reasonably homogeneous characteristics, which are very different from the general population on the dating site. In this section we analyze the demographics of the different types of scam accounts that we detected.

We first analyzed the gender distribution of the different types of scam accounts. As Figure 1 shows, legitimate accounts on DATINGSITE are mostly male, with a 65% - 35% ratio compared to female accounts. Different types of scams, on the other hand, aim at attracting victims from a specific gender, and therefore their gender ratio is dramatically skewed. For instance, “Dates for Profit” scammers try to lure men looking for a date with an attractive lady to go to a particular establishment; therefore, the accounts performing such scams are almost entirely female. A similar situation holds for “Escort Services” scam accounts.

“Swindler” accounts, on the other hand, cater to an older audience composed mostly of divorced or widowed ladies. For this reason, this type of scam accounts is prevalently male. A similar reasoning goes for “Matchmaking Services” scam accounts. These accounts mostly try to lure older women into handing out personal information such as email addresses and phone numbers, and their gender distribution is in line with the one of the legitimate user population of DATINGSITE.

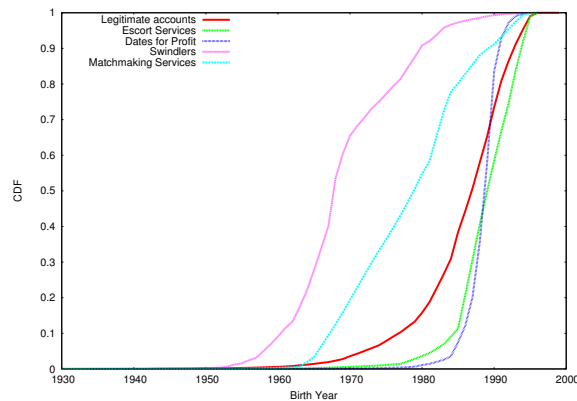


Fig. 2. Cumulative Distribution Function (CDF) of the birth years of the different types of scam accounts, compared to legitimate accounts.

The purpose of the different types of scammers is also reflected in the different ages of the profiles that they use. As Figure 2 shows, 50% of the legitimate users on DATINGSITE are 26 years old or older. “Dates for Profit” and “Escort Services” accounts list themselves as slightly younger than average, with an average age of 25 and 24 respectively. In particular, 20% of the “Escort Services” profile list their age as between 20 and 18 (which is the minimum age allowed on DATINGSITE). In comparison, only 1% of the normal users is younger than 20 on the site. Because they target older victims, “Swindler” profiles are a lot older on average: 50% of these accounts pretend to be at least 46 years old, with only 3% of them younger than the average age on the site. “Matchmaking Services” accounts are older than average as well, although not as much as “Swindler” accounts: 50% of them are older than 35, while 18% of them is younger than the average age on DATINGSITE.

The final element for which scam accounts differ from regular users is marital status. As Figure 3 shows, more than 85% of the regular users on DATINGSITE are single. Scam accounts mostly list themselves as single as well, probably because they cater to the majority of the users on the site, and in general it is a more neutral connotation. Figures 5, 6, and 7 show this trend for “Dates for Profit,” “Matchmaking Services,” and “Escort Agencies” scam accounts respectively. A notable exception are “Swindler” scam accounts. As it is shown in Figure 4, these accounts prevalently list themselves as widowed or divorced. The reason for that is that these accounts need to build a believable story, as we mentioned in Section 4, and being widowed is a good starting point to establish trust and a long-lasting relationship with an emotionally fragile person, and eventually steal her money.

To recap, in this section we confirmed that stereotypes do exist for each type of scammer: while “Dates for Profit” and “Escort Agency Services” mostly present themselves as young, single females, “Swindlers” present themselves as mid-aged, widowed men.




Fig. 3. Fraction of the marital status for legitimate accounts on DATINGSITE.




Fig. 4. Fraction of the marital status for "Swindler" scammers.




Fig. 5. Fraction of the marital status for "Dates for Profit" scammers.




Fig. 6. Fraction of the marital status for "Matchmaking Service" scammers.




Fig. 7. Fraction of the marital status for "Escort Service Advertisement" scammers.

6.2 Strategies Used by Different Scam Account Types

Another interesting aspect is looking at the different strategies used by different types of scammers to reach victims. In the previous section we analyzed how different types of scammers set up their accounts to look appealing to a different audience, such as older women or younger men. After setting up a profile, a scammer needs to attract victims. This is probably the most challenging part, because replying to a number of users who are not really convinced of the authenticity of the scam profile is time consuming for the scammer. As previous research noted, it is very advantageous for a scammer to make sure that the accounts that either reach out to him or that will reply to his first message are likely to fall for the scam [21]. In this section we show that some types of scammers are more successful in doing this.

In general a scammer has two strategies while trying to attract victims: he can contact users on the dating site himself, or he can make his account so appealing that a number of potential victims will contact him themselves. To understand the typical strategies put in practice by scammers, we first looked at the number of accounts contacted by scam accounts. As Figure 8 shows, scam accounts typically contact many more accounts than legitimate users on DATINGSITE do. Legitimate accounts contact a small number of profiles: from our observations, 50% of them contacted at most four profiles, while only 20% accounts contacted more than 27 profiles. Conversely, the majority of scam accounts contacted 100 or more profiles during their activity on DATINGSITE. This shows that scam accounts are on average a lot more aggressive in contacting other people on the dating site than regular users. Note that these numbers are a lower bound, because the scam accounts in our dataset were shut down by the support people at DATINGSITE, and therefore would have contacted more victims if they were left free to act.

Instead of contacting their victims, scammers can wait for users to contact them. This approach has the advantage that whoever contacts a scam account is more likely to

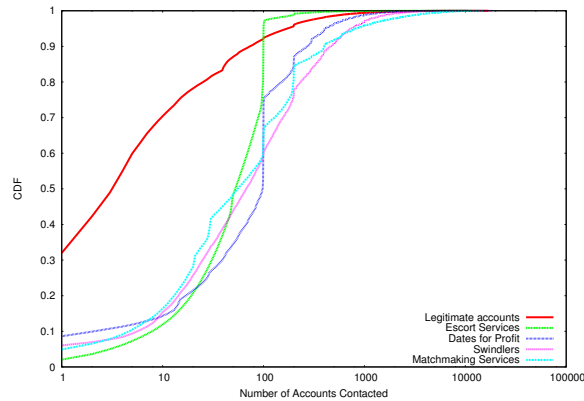


Fig. 8. Cumulative Distribution Function (CDF) of the accounts contacted by the different types of scammers, compared to legitimate users.

be genuinely interested in him/her. On the scammer side, all he needs to do is setting up a profile that is appealing to potential victims. Figure 9 shows the CDF of the number of messages initiating a conversation received by different types of accounts on DATINGSITE. As it can be seen, 50% of legitimate accounts were contacted at least once by another user on the dating site. We observe a two-fold distribution in the success of scam accounts. “Escort Services” and “Matchmaking Services” are on average less successful than regular accounts in receiving attention from users on the site, although these numbers are not dramatically lower than what we observe for legitimate users. The reason for this lower turnaround might be that these profiles look phony, and they do not appear believable to many users. Also, these accounts are usually easier to recognize by the human specialists than more advanced types of scam accounts, and are therefore “frozen” quickly, so that possible victims cannot contact them anymore. This is in line with previous research, which showed that social network users are fairly good in realizing whether a profile is fake, although their detection rate is not perfect [34]. “Swindlers” and “Dates for Profit” scam accounts, on the other hand, attract more victims on average: 70% of these accounts were contacted by at least a user on the site, while 20% of the “Swindler” accounts were contacted by at least 90 potential victims. All together, this shows that advanced scammers are remarkably successful in setting up profiles that appeal their potential victim audience, and, unlike the less sophisticated types of scams, their profiles are considered legitimate by many users.

We then analyzed how often scam accounts reply to the messages they receive from potential victims. Figure 10 shows the Cumulative Distribution Function of the number of replies that different types of scam accounts sent as responses to messages they received. As it can be seen, legitimate users rarely reply to the messages they received: only 14% of the users ever replied to a message on the site¹, and more in general only

¹ Note that this number takes into accounts also those accounts that never received a message at all on the site.




Fig. 9. Cumulative Distribution Function (CDF) of the number of messages received by different types of scammer accounts, compared to legitimate users.

7% of the messages that are sent on DATINGSITE ever receive a reply. This result is in line with what was shown by previous research [14]. Looking at scam accounts, we see a clear two-fold distribution: “Escort Services” and “Matchmaking Services” reply even less than regular users. The reason for this is that these scams are typically composed of a single spam-like message, and do not involve message exchanges with their victims. In addition to that, these accounts show the tendency of being quickly abandoned by the scammers, or being shut down by the DATINGSITE operators. “Swindlers” and “Dates for profit,” on the other hand, are a lot more likely to start conversations with their victims: around 50% of these accounts sent replies to potential victims after being contacted. The fact that they do not reply to about half of the messages they receive, on the other hand, shows that scammers are carefully choosing which users to engage with, favoring the ones that are more likely to fall for the scam. This shows that these scams are more sophisticated and require a longer set up (and a longer exchange of messages). For this reason, these types of scammers take the time to reply to the messages they receive, crafting them in a way that it will make them appealing and believable to potential victims.

Finally, we studied the way in which different types of scammers connected to their accounts. Previous work showed that malicious accounts on social networks are typically controlled by botnets [19,31], which are networks of compromised computers acting under the same cybercriminal. Botnets provide a convenient way for cybercriminals to control their malicious accounts, and run large-scale campaigns on online services. To investigate the *modus operandi* of the connections performed by different types of scammers on DATINGSITE, we studied the number of scam accounts of the same type that were accessed by single IP addresses. Although observing multiple accounts connecting from the same IP address is not suspicious per se, and it is usually an artifact of Network Address Translation (NAT) networks, having a high number of accounts accessed by the same IP address is highly suspicious. Figure 11 shows the CDF of the

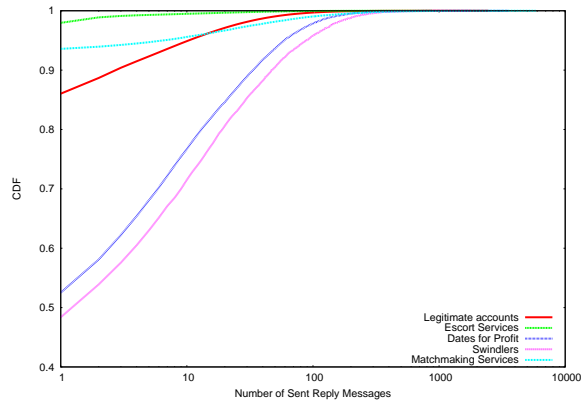


Fig. 10. Cumulative Distribution Function (CDF) of the replies sent by the different types of scammers in reply to messages they received, compared to legitimate users.

number of different scam accounts accessed by a single IP address on DATINGSITE. As it can be seen, legitimate accounts are typically accessed by only one IP address during a period of one week (85% of them). Although it is more common to have IP addresses that access many accounts, the ratio of scam accounts accessed by multiple IP addresses is low too: it is rare to have IP addresses that accessed more than ten scam accounts, and the majority of them accessed a single IP address. This does not give us conclusive evidence that scammers are using botnets to perform their malicious activity on DATINGSITE. The reason for this is that online dating scams require quite a bit of interaction by the scammer to succeed, and are more difficult to automate than traditional spam or phishing operations. Reports by the human experts at DATINGSITE, as well as by law enforcement, show that an exception to this trend are “Escort Services” scam accounts, whose accounts are usually controlled by bots. Still, these bots are typically used in isolation or in small numbers, and therefore we have no indication that these types of scammers are taking advantage of botnets.

7 Discussion

Online dating scams are a relatively new problem and present very unique aspects, compared to traditional Internet threats. In this paper we described in detail the insights that we obtained by analyzing the scam accounts detected on a large online dating site over a period of one year. Many points are open and will be the focus of future research. In this section, we first discuss the efforts needed to secure online dating sites from scammers, both from the policy and the technical side. Finally, we introduce some ideas that we are planning to pursue as future work.




Fig. 11. Cumulative Distribution Function (CDF) of the number of other accounts of the same scam type that were accessed by a single IP address over a period of one week, compared to legitimate users.

7.1 Scammers are Perseverant

According to the human experts at DATINGSITE, each individual scam on the site is usually set up and performed by a small number of people, composed of two to five individuals. For example, a typical “Date for Profit” scam will involve the owner of a restaurant, a girl who agrees to go on a date with the victim, and optionally a third person who interacts with the online dating site. Similarly, a “Flower-basket swindler” scam is performed by a person interacting with the victim on the dating site and by an accomplice who pretends to own a flower shop. Optionally, other accomplices will pretend to be family members of the scammer, and contact the victim to test her “good will.” This is quite different than what we observe in traditional cybercriminal schemes, in which there are many actors involved, and a complex economy exists behind these operations.

Based on the reports by the DATINGSITE human experts, there are a number of scammer groups that keep using DATINGSITE to perform their scams, even after their fake accounts are detected and blocked. This happens because the scams that these groups perform on the site are particularly remunerative for them. The perseverance showed by online dating scammers teaches researchers two lessons: first, better detection systems are needed, to make sure that the scam accounts are detected and shut down quicker, making life more difficult for the scammers. We hope that some of the insights provided in this paper, such as the typical demographics and behaviors showed by scam accounts will help practitioners in developing more effective detection techniques. Second, similar to other cybercriminal operations, it is difficult to get rid of scammers by just securing the online service. In many cases, an action from law enforcement is needed to prosecute those scammers who are most dangerous. Law enforcement operations are typically difficult to coordinate when we deal with cybercriminal schemes that are spread across multiple countries, but should be more feasible in the case of DAT-

INGSITE, in which the site, the victims, and the scammers all reside in the same country. Similar law enforcement measures might get more complicated for online dating sites that have a more international user base, such as `match.com` or `okcupid.com`.

7.2 Future Work

Aside from law enforcement efforts, online dating scams can be made less successful by improving the effectiveness of detection systems deployed on the online dating site. As we showed in this paper, detecting the more advanced types of scammers, such as “Swindlers,” is challenging, because the activity of these accounts is not automatically generated, and their scams unfold over a long period of time and several exchanges of messages with the victim. To detect the majority of these scam accounts, DATINGSITE administrators decided to have a better recall over precision for their anti-scam systems, and to employ a team of human experts to vet such detections.

In this paper, however, we showed that scam accounts show differences from legitimate accounts in the way they set up their profiles, in the way they select the people they chat with, and in the way they interact with the site. Such differences could be the base for developing effective detection techniques that could work with minimal human interaction. In addition, future work could include the use of stylometry to detect message language that might be typical of scams [5].

We are aware that the insights provided in this paper are only the first step into understanding and fighting scammers on online dating sites. In particular, although many of the scams that we identified are likely to happen on other online dating sites too, our analysis is limited to DATINGSITE, and therefore might be biased toward the Chinese culture. In the future, we are interested in studying online dating sites based in other countries. In particular, online dating sites with an international audience might attract scammers that are located in different countries than the victim; tracking these schemes might be challenging, but could give interesting insights on how to fight the phenomenon. In addition, we are interested in studying the underground economy behind these scams, such as tracking how much money is actually stolen as a consequence of these schemes. We would also like to study online dating sites that require users to pay a monthly subscription. We suspect that the population of scammers on such sites will be heavily skewed towards more advanced schemes.

8 Related Work

The popularity of online dating sites has attracted a wealth of research. Fiore et al. performed an analysis of the demographic characteristics of online dating users, as well as of the characteristics of the messages that those users exchange [14]. Chen et al. presented a work that analyzed the social network of the users on an online dating site [10]. Hitsch et al. looked at what characteristics influence the chance of an online dating message to receive a reply [22].

Our work is the first comprehensive study on online dating scams. A few previous papers provided case studies of single scam schemes happening on online dating sites [27, 28, 36], while Wang et al. observed that miscreants crawl online dating sites,

and use the crawled pictures to set up fake online dating accounts, which they then use for their scams [35]. Unlike these preliminary works, our paper provides a comprehensive study of the scams that happen on a large online dating site, and gives real-world numbers on how prevalent such scams are.

Online dating scams that try to lure a victim into sending money to the miscreant have many points in common with the so-called “419 scams.” This type of scams has been widely studied by the research community [12, 17, 18, 21, 23, 25]. Despite the similarities, in this paper we showed that online dating scams are very different from these schemes, because they typically involve a long exchange of messages between the scammer and the victim before the actual scam is performed.

Many systems have been presented to detect spam and malicious activity on social networks. Some systems analyze the characteristics of social network accounts, looking for sign of mass-created fake accounts [7, 16, 24, 31, 40, 42]. Other systems look at the social network structure of accounts, or at how the messages posted by them spread, looking for anomalies indicative of a malicious profile [8, 9, 15, 32, 38, 41]. These systems work well in detecting automated activity finalized at spreading malicious content. However, they are often ineffective in detecting advanced scams, such as the ones presented in this paper. The reason is that such scams require a lot of manual effort by the attackers, and often times span through long periods of time (and multiple message exchanges between the victim and the scammer).

9 Conclusions

In this paper, we analyzed the problem of online dating scams. We analyzed a set of ground-truth scam accounts on a large online dating site, and provided a taxonomy of the different types of scams that we observed. We showed that different types of scammers targeted different audiences, and have a very different success rate in attracting victims. In particular, we showed that advanced scammers are more successful than regular users in getting attention by other online dating users, which can potentially turn into scam victims. This paper is the first large-scale measurement of online dating scams, and it sheds light on these operations and on the challenges that researchers face in fighting them. Future work will focus on developing better detection techniques to block even the stealthiest scammers, and studying new types of online dating scams.

Acknowledgments

We thank the anonymous reviewers for their comments. We would also like to thank Ali Zand, Adam Doupé, Alexandros Kapravelos, Ben Y. Zhao, and Christo Wilson for reviewing an early draft of this paper. Your feedback was highly appreciated.

References

1. eHarmony. <http://www.eharmony.com>, 2013.
2. OkCupid. <http://www.okcupid.com>, 2013.
3. Online dating sites pricing. http://www.nextadvisor.com/online_dating/compare.php, 2013.

4. Online Dating Statistics. <http://www.statisticbrain.com/online-dating-statistics/>, 2013.
5. AFROZ, S., BRENNAN, M., AND GREENSTADT, R. Detecting hoaxes, frauds, and deception in writing style online. In *IEEE Symposium on Security and Privacy* (2012).
6. BALTAZAR, J., COSTOYA, J., AND FLORES, R. KOOFACE: The Largest Web 2.0 Botnet Explained. In *Trend Micro Threat Research* (2009).
7. BENEVENUTO, F., MAGNO, G., RODRIGUES, T., AND ALMEIDA, V. Detecting Spammers on Twitter. In *Conference on Email and Anti-Spam (CEAS)* (2010).
8. CAI, Z., AND JERMAINE, C. The Latent Community Model for Detecting Sybils in Social Networks. In *Symposium on Network and Distributed System Security (NDSS)* (2012).
9. CAO, Y., YEGNESWARAN, V., POSSAS, P., AND CHEN, Y. PathCutter: Severing the Self-Propagation Path of XSS Javascript Worms in Social Web Networks. In *Symposium on Network and Distributed System Security (NDSS)* (2012).
10. CHEN, L., AND NAYAK, R. Social network analysis of an online dating network. In *Proceedings of the 5th International Conference on Communities and Technologies* (2011).
11. DRUCKER, H., WU, D., AND VAPNIK, V. N. Support vector machines for spam categorization. In *IEEE transactions on neural networks* (1999).
12. DYRUD, M. A. "I brought you a good news": An analysis of nigerian 419 letters. In *Association for Business Communication Annual Convention* (2005).
13. EGELE, M., STRINGHINI, G., KRUEGEL, C., AND VIGNA, G. Compa: Detecting compromised accounts on social networks. In *Symposium on Network and Distributed System Security (NDSS)* (2013).
14. FIORE, A., AND TRESOLINI, R. *Romantic regressions: An analysis of behavior in online dating systems*. PhD thesis, Massachusetts Institute of Technology, 2004.
15. GAO, H., CHEN, Y., LEE, K., PALSETIA, D., AND CHOUDHARY, A. Towards Online Spam Filtering in Social Networks. In *Symposium on Network and Distributed System Security (NDSS)* (2012).
16. GAO, H., HU, J., WILSON, C., LI, Z., CHEN, Y., AND ZHAO, B. Detecting and Characterizing Social Spam Campaigns. In *Internet Measurement Conference (IMC)* (2010).
17. GAO, Y., AND ZHAO, G. Knowledge-based information extraction: a case study of recognizing emails of nigerian frauds.
18. GLICKMAN, H. The nigerian "419" advance fee scams: Prank or peril? *Canadian Journal of African Studies* (2005).
19. GRIER, C., THOMAS, K., PAXSON, V., AND ZHANG, M. @spam: the underground on 140 characters or less. In *ACM Conference on Computer and Communications Security (CCS)* (2010).
20. HAO, S., SYED, N. A., FEAMSTER, N., GRAY, A. G., AND KRASSER, S. Detecting Spammers with SNARE: Spatio-temporal Network-level Automatic Reputation Engine. In *USENIX Security Symposium* (2009).
21. HERLEY, C. Why do Nigerian scammers say they are from Nigeria? In *Workshop on the Economics of Information Security (WEIS)* (2012).
22. HITSCH, G. J., HORTACSU, A., AND ARIELY, D. What makes you click: An empirical analysis of online dating. In *Society for Economic Dynamics Meeting Papers* (2005).
23. ISACENKOVA, J., THONNARD, O., COSTIN, A., BALZAROTTI, D., AND FRANCILLON, A. Inside the SCAM Jungle: A Closer Look at 419 Scam Email Operations. In *International Workshop on Cyber Crime (IWCC)* (2013).
24. LEE, K., AND CAVERLEE, JAMES AND WEBB, S. Uncovering social spammers: social honeypots+ machine learning. In *International ACM SIGIR conference on Research and Development in Information Retrieval* (2010).

25. PARK, Y., JONES, J., MCCOY, D., SHI, E., AND JAKOBSSON, M. Scambaiter: Understanding targeted nigerian scams on craigslist. *Symposium on Network and Distributed System Security (NDSS)* (2014).
26. PITSILLIDIS, A., LEVCHENKO, K., KREIBICH, C., KANICH, C., VOELKER, G. M., PAXSON, V., WEAVER, N., AND SAVAGE, S. botnet Judo: Fighting Spam with Itself. In *Symposium on Network and Distributed System Security (NDSS)* (2010).
27. PIZZATO, L. A., AKEHURST, J., SILVESTRINI, C., YACEF, K., KOPRINSKA, I., AND KAY, J. The effect of suspicious profiles on people recommenders. In *User Modeling, Adaptation, and Personalization* (2012).
28. REGE, A. What's love got to do with it? exploring online dating scams and identity fraud. *International Journal of Cyber Criminology* (2009).
29. SAHAMI, M., DUMAIS, S., HECKERMANN, D., AND HORVITZ, E. A Bayesian approach to filtering junk e-mail. *Learning for Text Categorization* (1998).
30. STRINGHINI, G., HOLZ, T., STONE-GROSS, B., KRUEGEL, C., AND VIGNA, G. BotMagnifier: Locating Spambots on the Internet. In *USENIX Security Symposium* (2011).
31. STRINGHINI, G., KRUEGEL, C., AND VIGNA, G. Detecting Spammers on Social Networks. In *Annual Computer Security Applications Conference (ACSAC)* (2010).
32. STRINGHINI, G., WANG, G., EGELE, M., KRUEGEL, C., VIGNA, G., AND ZHENG, H. AND ZHAO, B. Y. Follow the green: growth and dynamics in twitter follower markets. In *ACM SIGCOMM Conference on Internet Measurement* (2013).
33. SYMANTEC CORP. Symantec intelligence report. http://www.symanteccloud.com/mlireport/SYMCINT_2013_01_January.pdf, 2013.
34. WANG, G., MOHANLAL, M., WILSON, C., WANG, X., METZGER, M., ZHENG, H., AND ZHAO, B. Y. Social turing tests: Crowdsourcing sybil detection. *Symposium on Network and Distributed System Security (NDSS)* (2013).
35. WANG, G., WILSON, C., ZHAO, X., ZHU, Y., MOHANLAL, M., ZHENG, H., AND ZHAO, B. Y. Serf and turf: crowdurfing for fun and profit. In *World Wide Web Conference (WWW)* (2012).
36. WHITTY, M. T., AND BUCHANAN, T. The online romance scam: A serious cybercrime. *CyberPsychology, Behavior, and Social Networking* (2012).
37. XIE, Y., YU, F., ACHAN, K., PANIGRAHY, R., HULTEN, G., AND OSIPKOV, I. Spamming Botnets: Signatures and Characteristics. *SIGCOMM Comput. Commun. Rev.* (2008).
38. XU, W., ZHANG, F., AND ZHU, S. Toward worm detection in online social networks. In *Annual Computer Security Applications Conference (ACSAC)* (2010).
39. YANG, B., GU, F., AND NIU, X. Block mean value based image perceptual hashing. In *International Conference on Intelligent Information Hiding and Multimedia Signal Processing* (2006).
40. YANG, C., HARKREADER, R., AND GU, G. Die Free or Live Hard? Empirical Evaluation and New Design for Fighting Evolving Twitter Spammers. In *Symposium on Recent Advances in Intrusion Detection (RAID)* (2011).
41. YU, H., KAMINSKY, M., GIBBONS, P. B., AND FLAXMAN, A. Sybilguard: defending against sybil attacks via social networks. *ACM SIGCOMM Computer Communication Review* (2006).
42. ZHANG, C. M., AND PAXSON, V. Detecting and Analyzing Automated Activity on Twitter. In *Passive and Active Measurement Conference* (2011).