MIMS Capstone Project

Sounds Phishy

Protecting Consumers Against Phone Phishing

Michelle Chen & Ashish Sur

Advisor: Steve Weber

May 2019

Table of Contents

Acknowledgements

We'd like to thank our advisor Steve Weber for being the person who was always there to listen to our problems and troubleshoot our roadblocks. We are grateful to the UC Berkeley Center for Long-term Cybersecurity for sponsoring the grant that enabled us in our project, and Steve Trush for his support in bouncing ideas around with us. Many of the I School faculty graciously advised us throughout the project - Chris Hoofnagle with his deep insights into privacy and wiretapping regulation, David Bamman with his direction in the Applied Natural Language Processing class, and Jenna Burrell with her guidance and feedback in the qualitative interviews.

We could not have gone as far as we had without the support from the UX Research class team (Melrose Huang, Yi Gai, Anna Waldo, and Brandon Fang) and their user research efforts and findings, as well as the Applied NLP class team (Tanvi Sunku and Eva Wu) and their situational power detection model. Finally, we'd like to thank all of our research participants, and especially thank the members of MIMS 2019 and 2020 whose brains we picked whenever we were in need.

Problem Overview

In 2018, imposter scams were the top consumer complaint to the Federal Trade Commission (FTC)[1]. Phone calls were the most common contact method reported, and the total financial loss to imposter phone scams totaled $429 million, with a median loss of $840. Imposter phone scams, or phone phishing attacks, typically result not only in financial loss but also identity theft. Posing as figures of authority, attackers used a variety of social engineering tactics to build trust with their target, by using the target's human vulnerabilities in psychological biases and limited knowledge to gain access to personal information and/or financial assets. Yet, most consumers are unaware of the severity of sophisticated phishing attacks that occur over phone. Contrary to popular belief, younger adults have been found to be more likely to fall victim to phone scams[2]. It is of great concern that the number of unwanted phone calls continues to rise every year, with a 39% increase from 2018 to 2019[3].

Currently, there is a gap in effective resources and tools for consumers, in particular, to educate and protect themselves against phone phishing. Most research, trainings, and products have been centered around raising awareness and protection efforts for businesses[4]. Consumers have to seek out such information on their own, but it is scattered across multiple sources. Most people are not incentivized to educate themselves until after they learn of a scam from their social network, or they fall victim themselves. Furthermore, even when they think they know about scams, people still succumb to phone phishing. While call spam-blocking apps such as Truecaller and Hiya are effective in warning consumers of suspicious phone numbers, these solutions do not address detecting phone scams once the call is picked up[5].

---

[1] (2019). Retrieved from https://www.ftc.gov/system/files/documents/reports/consumer-sentinel-network-data-book-2018/consumer_sentinel_network_data_book_2018_0.pdf

[2] https://firstorion.com/new-irs-scams-survey-shows-millennials-most-likely-to-fall-victim/

[3] https://truecaller.blog/2019/04/17/truecaller-insights-2019-us-spam-phone-scam-report/

[4] Krombholz, K., Hobel, H., Huber, M., & Weippl, E. (2015). Advanced social engineering attacks. *Journal of Information Security and Applications*, *22*, 113–122. https://doi.org/10.1016/j.jisa.2014.09.005

[5] https://www.truecaller.com, https://hiya.com/

Project Goals

Our capstone project consists of two main goals:

1) To understand phone phishing and consumer phone behaviors from multiple sources: existing cybersecurity literature, consumer anti-phishing resources, victims of imposter phone scams, everyday consumers, and social engineers.

2) To explore a real-time imposter phone scam detection application aimed to identify and warn users of phone phishing, with consideration towards privacy and security.

Background Research

Phone Scams, Fraud, and Social Engineering

Phone phishing scams come in a variety of flavors, from tech support scams to free trial offers to law enforcement impersonations[6]. The number of phone scams have drastically risen in the past few years, with a projected nearly 50% of mobile phone traffic to be phone scams in 2019[7]. These calls are conducted either by a live person or a pre-recorded message, known as a robocall. For our project, we are interested in sophisticated phone phishing scams that involve speaking to a live person. The term voice phishing, or vishing, is when scammers call potential victims and use social engineering over the phone to collect personal information and/or convince them to relay financial assets[8].

The Financial Fraud Research Center (FFRC) compiled a review of consumer financial fraud research[9]. There are several key insights from the report that are relevant to our project. First, profiling research has shown that there is no such thing as a "typical" scam victim. Any human person can fall victim to a scam. However, for each individual scam type (such as an IRS scam), there are certain kinds of people who are more susceptible. As to why people succumb to fraud, there are three categories of tactics that scammers use against their victims: imitation, deception, and threat. Fraud imitates legitimate services and processes in attempt to build credibility. This is closely tied with deception; as the review states, "To understand fraud is to understand deception — its persuasive strategies and practical methods" (pg. 30). Burgard & Schlembach (2013) present a framework of cyber fraud, consisting of three stages: getting hooked on, staying attuned, and cooling out[10]. At the core of a scam is the deceiver's fabrication of a different reality, and the framework details the victim's process of going from the scammer's fabricated "reality" to actual reality.

---

[6] https://www.consumer.ftc.gov/articles/0076-phone-scams
[7] https://firstorion.com/nearly-50-of-u-s-mobile-traffic-will-be-scam-calls-by-2019/
[8] Yeboah-boateng, E. O., & Amanor, P. M. (2014). Phishing, SMiShing & Vishing: An Assessment of Threats against Mobile Devices. *Journal of Emerging Trends in Computing and Information Sciences*.
[9] Deevy, M., Lucich, S., & Beals, M. (2012). Scams, Schemes, and Swindles: A Review of Consumer Financial Fraud Research. Retrieved from Financial Fraud Research Center website: http://longevity.stanford.edu/2012/11/19/scams-schemes-and-swindles-a-review-of-consumer-financial-fraud-research/
[10] Burgard, A. & Schlembach, C. (2013) Frames of Fraud: A Qualitative Analysis of the Structure and Process of Victimization on the Internet. *International Journal of Cyber Criminology*, 7(2), 112-124

Persuasion is a key element in deception and social engineering. Researchers have found that the tactics of persuasion are the similar between fraud and legitimate marketing[11]. Bullée et. al (2015) conducted a literature-based dissection of social engineering scams[12], and analyzed the scam scenarios and attack steps according to Cialdini's six principles of persuasion: reciprocity, social validation (commitment), liking, consistency, authority, and scarcity[13]. Out of the six persuasion principles, they found that authority was most commonly used in their dataset of social engineering scams.

One poignant insight from the FFRC's review is the prevalence of under-reporting financial fraud. People's behaviors do not match what they say they would do - in a national survey on white collar crime victimization, nearly 95% indicated that they would report an incident while only 21% of actual victims indicated having reported their victimization[14]. Reasons why people fail to report include embarrassment and a perception of little benefit gained from reporting. The majority who do report indicated that they do so in order to help protect others.

Current Landscape of Anti-Phone Phishing Solutions

There are numerous trusted resources to learn about phone scams, such as FTC's Scam Alerts and coverage by local and national news organizations[15]. Better Business Bureau manages BBB Scam Tracker℠, where consumers and businesses can report scams. From Reddit[16] to independent websites and blogs[17] to support forums[18], many post about their phone scam incidents online, and many more comment on their posts. However, all this information spread across multiple sources and not in an easily absorbable format for consumers to learn and analyze the latest phishing tactics. While word of mouth is another common and effective solution in spreading awareness, phone scams nonetheless continue to be a threat to consumers.

Current technical solutions and applications typically involve identifying suspicious or classified phone numbers before a call has been answered. Mobile phone providers

---

[11] Pak, K., & Shadel, D. (2007). *The Psychology of Consumer Fraud*.
[12] Bullée, J.-W. H., Montoya, L., Pieters, W., Junger, M., & Hartel, P. (2018). On the anatomy of social engineering attacks-A literature-based dissection of successful attacks: On the anatomy of social engineering attacks. Journal of Investigative Psychology and Offender Profiling, 15(1), 20–45. https://doi.org/10.1002/jip.1482
[13] Cialdini, R. B. (2001). The Science of Persuasion. Scientific American, 284(2), 76–81.
[14] National White Collar Crime Center, 2000
[15] https://www.consumer.ftc.gov/features/scam-alerts
[16] https://reddit.com/r/scams
[17] http://scamawareness.org/, https://www.hoax-slayer.net/
[18] https://forums.att.com/, https://productforums.google.com/forum/#!topic/play,

are starting to tag certain calls as “Scam Likely”[19]. Third-party apps like Truecaller and Hiya maintain a crowdsourced database of scam phone numbers. Our UX Research class team conducted a thorough competitive analysis of the existing technologies in the anti-phone phishing space[20]. One of their main takeaways was that scam alert-based systems that rely on a database of known scam numbers are the most popular amongst users currently. These apps have a simple setup procedure where users do not have to additional steps to be notified of potential scam calls after the initial onboarding process.

Notably, none of the anti-phishing apps address content of the phone conversation. The closest technical solution seems to be Google’s Call Screen, which takes a step towards detecting the motives of a caller before a human call recipient enters the call. When a user receives an incoming call, they have the option to Screen Call and allow Google Assistant to “pick up” the call and interact with the caller.

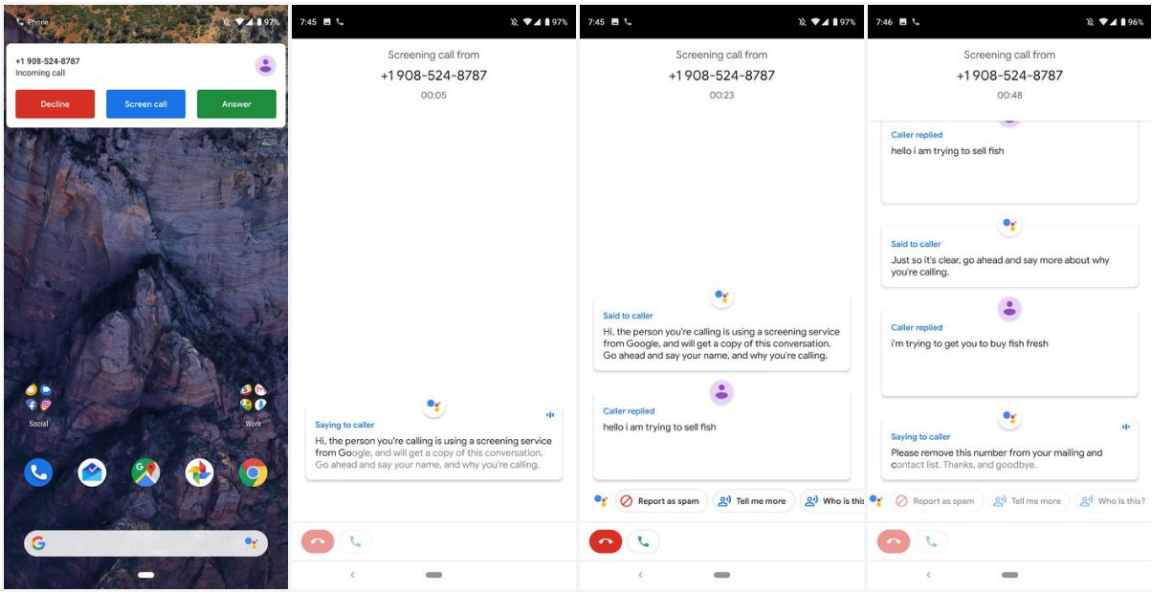

Retrieved from https://www.digitaltrends.com/mobile/google-call-screening-how-to-use/

The call begins with a disclosure that the call recipient is using Google’s Call Screen service, and the Assistant asks the caller to identify themself and their reason for calling. While the conversation is going on, the user can see a transcript of the call between the caller and the Assistant. Google’s Call Screen is the only application that we encountered during our research that addresses the problem of detecting phone scams after a call has been picked up.

[19] https://www.usatoday.com/story/tech/2019/01/21/scam-likely-calling-you-robocall-busters/2586105002/
[20] Third party anti-scam apps examined were: Truecaller, Hiya, TrapCall, Robokiller

User Research

Qualitative Interviews

People

Consumers who had fallen for a phone phishing attack; total of 5 interviews.

Goals

We aimed to understand the experiences of phone phishing victims and how they were manipulated from their perspective. We sought to learn the strategies that attackers used to build the targets' trust towards them and how they perceived the attackers. We were also interested their attitudes towards digital security before and after their phone scam experience, as well as any changes in the ways they perceive and act towards phone calls. Finally, we wanted to initially test out our working idea of a real-time scam detection tool. See Appendix A for a copy of the interview guide.

Findings

Scamming the Victim

Building Trust
Attackers used a variety of tactics to manipulate the target into trusting them. Displaying sympathy towards the victim and their circumstances was one tactic that scammers used to build trust, and furthermore establish distance between the victim and truth of the situation. Besides establishing rapport with the victim, the interviewees shared a number of strategies that scammers used to distract the target and mask their real intentions. The scammers recreated seemingly professional scenarios such as transferring the target to their "supervisor" when urged or advising the victim to protect them from further harm. One individual was told to purchase a new phone because the attacker said that the "real perpetrator", who is using her identity to commit the crime she has been accused with, might be tracking her calls. Upon later reflection, she believed the reason for the

new phone may have been a distraction to avoid any issues with potential anti-scam apps on her current phone.

The attackers also used information they had on the targets to strengthen their perceived credibility. One participant mentioned at first he was very suspicious of the caller, but he started to believe in the scam after the imposter provided specific and accurate personally identifiable information, such as the social security numbers of his family members. As he stated, “It's just the sheer amount of information they had on me - it just put doubt in my mind”. Attackers additionally used information they had phished from the target earlier and used the information against them. For example, one individual entered personal information into one of the “Embassy’s” websites in order to verify her identity, and then the scammers used her contact information to send her an “official” court notice email.

Taking Advantage of Human Biases

The attackers also took advantage of human psychological limitations and biases. First of all, the kinds of problems that they presented to victims were urgent issues with serious consequences if not addressed immediately, such as visa issues, late taxes, or criminal warrants. A target’s vulnerable state of mind further assists the scammer, such as engaging when sleep deprived and unable to think clearly and critically. Limited judgment comes not only from stress, but also familiar circumstances. One individual mentioned that his call proceeded just like his previous monthly service calls from his computer provider. In these calls, the real service providers required access to his computer to investigate an issue. He had expected the call since they came regularly every month. The familiarity of the situation caused him to not think of the call as something different than what he had encountered many times in the past.

Lack of Knowledge on Situation

Due to lack of knowledge on how actual processes in the particular domain typically work, victims relied on the perceived expert in the situation - the scammer. Scammers explained the situation and consequences in legal or technical terms that the target was unfamiliar and uncomfortable with as a non-expert. Therefore, the targets could not fully evaluate what was happening. Furthermore, they were under the impression that the scammer was a figure of authority in the situation, and so they defaulted to their so-called expertise.

Lack of Knowledge on Scammers Capabilities

Attackers also succeeded due to the lack of knowledge in the ways scammers are able to manipulate the technology involved. The scammers used caller ID spoofing, where the phone number of the real organization is displayed on the phone's caller ID, to deceive and build credibility towards themselves. Even when the target tried to validate the phone number online, it matched the calling number and seemingly verified the caller. Additionally, unfamiliarity with the capabilities of other technologies used during the scam had the same effect, such as not being aware of how simple it is for scammers to build an impersonating website that looks like an official website.

Detecting the Scam

Red Flags in Retrospect

From the accounts of the interviewees, there are a variety of barriers that act to conceal the reality of the scam from the victim. These barriers arise not only from the scammers' manipulations, but also from the victims' internal beliefs and mental limitations. Due to the barriers in place to prevent the victims from fully realizing what exactly was going on, they failed to notice red flags that they identified retrospectively. A common red flag mentioned was the mismatch between the Western name the imposter provided the target, and the scammer's accent. Other red flags mentioned were the attackers saying not to call the spoofed number, questionable email addresses, discounts for seniors, etc. It was easier for the victims to see the bigger picture afterwards. It went unsaid during the interviews that the victims later spent significant time mulling over what had happened and why they had fallen for the scam. There was a sense of exasperation that they should have seen the red flags.

Point of Detection / Threshold

All of the interviewees clearly recalled the moment when something the attacker said something or requested them to do something that triggered them into think critically about the situation. These moments include contradictory information provided by the scammers, out of character requests like purchasing gift cards, or strange activity like emails disappearing from their inbox. When the scammer makes a seemingly outlandish request, the victim's alert system finally starts up and begins processing the situation in a different way. All of the barriers blocking the reality of the scam then begin to break down. Unfortunately, this can occur after the victim has already transferred personal information or money to the attacker.

Previous Views on Phone Scams

All of the interviewees stated that they had never heard of anything like the scams they experienced, and they never expected to encounter anything like them. Strikingly, they knew scams existed and have gotten a suspicious call or email previously. They've read about the problem and knew it was something to be on guard for, but according to one of the interviewees, "these guys were good enough that I got scammed". Before experiencing scams themselves, there was also an attitude of disbelief that people get scammed by these situations. It may be that even though scams are known on a colloquial level, but the level of sophistication and severity is not well understood. One interviewee said a few days after her incident, someone told her about voice phishing, and that's when everything clicked for her. Before the scam, she always thought she was an alert person and always had to look after her sister, who's the kind of person "to easily believe people". Therefore, people seem overconfident with their ability to overcome a phone phishing scam before experiencing one for themselves.

Views on Security after Scam Incident

After experiencing the scam, there is a realization that these situations occur and they happen in different forms. There is a sense that a lesson has been learned, and now they are more cautious with anything related to financial and personal information. There is also an assumption that scammers will target them given that their information is likely accessible due to the increase in data breaches. The interviewees mentioned taking on additional security measures like using a password manager or changing passwords more frequently. Never answering unfamiliar calls and using an anti-scam call database app like Truecaller were other changes in people's phone security behaviors.

Potential Solutions to Imposter Scam Problem

When asked about potential solutions to addressing this problem, the targets made several different points. One interesting finding was an unstandably strong motivation to catch the scammers and prevent others from suffering the same fate. The interviewees wanted an easy way to collect information about the scammers and report it to law enforcement. There is a belief that only law enforcement can really fix the issue of phone scams by being able to handle the root cause - the scammers. Additionally, companies who are being impersonated need to provide better warnings to their consumers. In regards to how consumers can protect themselves, one interviewee mentioned the best way is by telling the story and letting people know about the scam and how serious it is.

As for the real-time scam detection phone app, the interviewees were curious about the idea. However, there were multiple concerns around privacy and technical limitations, like draining phone battery life. One person mentioned that she probably would not have used it before her scam experience, and she felt like she's learned her lesson since she's experienced it. However, a couple were enthusiastic about the real-time scam detection idea if it actually works well in accurately detecting scams.

Consumer Phone Behaviors Survey

People

Consumers with mobile phones; total of 148 respondents.

Goals

For the survey, we sought to learn about people's current mental models of phone scams, in terms of their attitudes towards phone scams, knowledge of caller ID spoofing, and protections against phone scams. Additionally, we were interested in people's current phone call behaviors and how they trust phone calls. A copy of the Consumer Phone Behaviors can be found in Appendix B.

Findings

Phone Scam Mental Models

We found that the majority of people believe phone scams (defined as "a caller pretending to be a trusted source and tries to trick you into revealing your personal information and/or sending your money to them") to be a moderate to significant problem in their life.

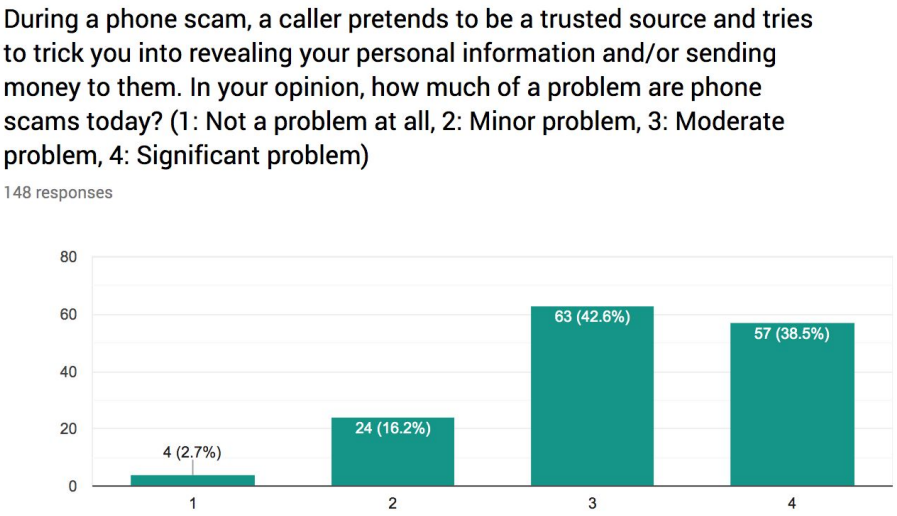

However, 72% of the respondents are using no apps or services against phone scams. Some mentioned having a built-in scam likely alert system provided by the device or their network provider. We also asked respondents to rate on a scale how true they believe it's possible for Caller ID to always display the real phone number that is calling. We found that the answers were fairly distributed. Although the highest response (30%) was that it is "very untrue" that Caller ID can always be trusted, 70% of the respondents indicated answers that range from "somewhat untrue" to "very true". Therefore, there was a concerning number of people who either absolutely do not know about call spoofing or hold uncertainty about it.

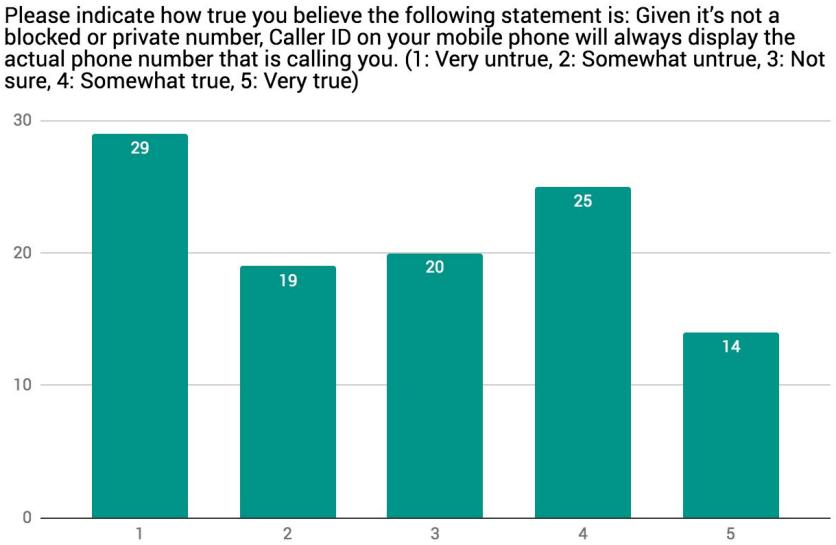

Trusting Phone Calls

In regards to reacting to a call from an unfamiliar phone number, most people indicated that they ignore the call and listen to the voicemail if there is one. Interestingly, 13% answer the call straightaway. Some responded that they answer depending on several factors, such as if the call is from a local area code, if they're not busy, or if it's an area code that they used to live in.

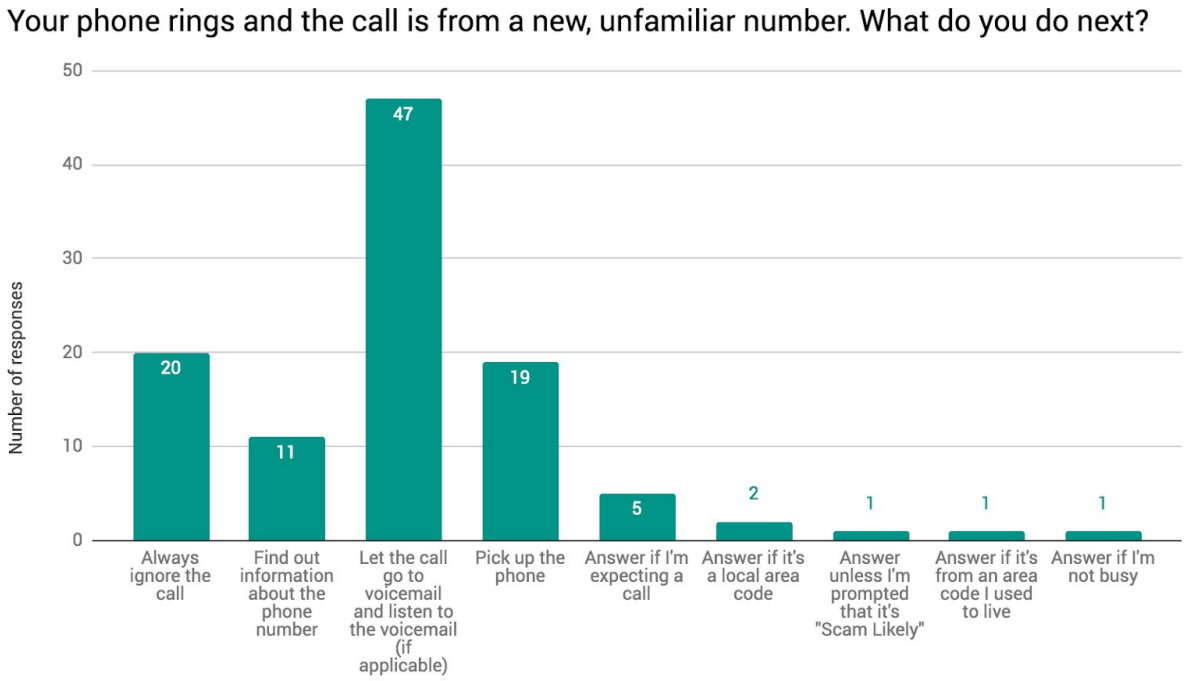

Most respondents indicated that unfamiliar phone numbers makes a phone call suspicious, followed by not knowing the person who is calling and then not recognizing the area code. Most people responded multiple characteristics. There's also a distinction between local area codes and matching area codes - matching area codes (area code that is the same as their number), are more suspicious and local area codes are less suspicious.

After picking up a phone call and being suspicious of the caller, the top action to take was hanging up and getting more information online about the caller. People also indicated fully hanging up after listening to the caller. Alarmingly, 20% of the respondents keep engaging with the caller for more information.

Design Process

Imposter Phone Scam Journey

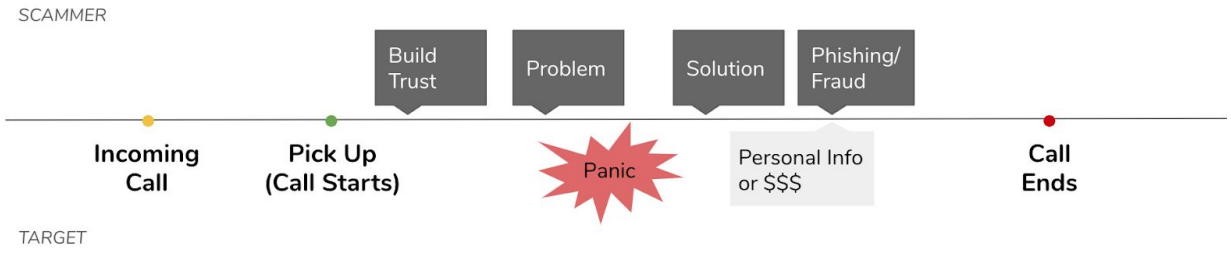

As we investigated people's experiences of imposter phone scams through the qualitative interviews and background research, a typical journey of a phone phishing scam started to emerge. It begins with an incoming call, followed by a period of time where the target makes a decision to pick up the call. The call starts when the target decides to pick up, and the conversation begins between the caller and the recipient. The scammer starts to manipulate the target and build trust as the parties establish who each other are and the purpose of the phone call. Then, the scammer presents a problem to the target, which instills panic to target. Once the target is in a state of panic, the scammer presents a solution to the target. The solution involves phishing for information and/or collecting financial assets.

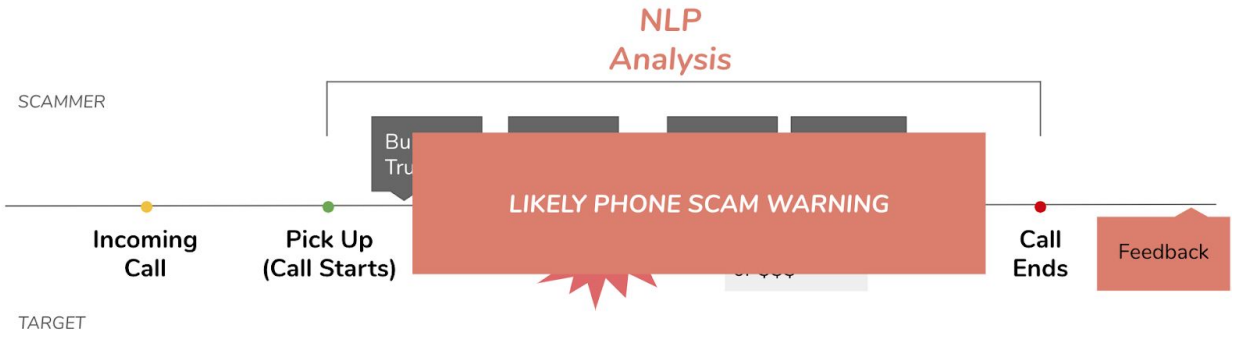

We examined opportunities for intervention within the imposter scam journey. As mentioned earlier, there are existing phone apps that intervene at the point before picking up a call. They provide an alert that an incoming call might be a likely scam. Intervention could occur after the call ends, where there could be an evaluation of the call after it occurs to see if might have been a scam call. We are interested in providing an intervention during the call, using NLP analysis of the conversation. By analyzing the conversation, we envisioned providing a warning to the target when a potential scam is detected.

Privacy Concerns

Concerns from User Research

Real-time scam detection requires access to an ongoing phone call, which is inherently an invasion of expected privacy in telephone communications. We heard concerns with privacy during the qualitative interviews with the scam victims, and the UX Research class team found significant resistance to call recording from a focus group they conducted. The focus group was generally distrusting towards a third party collecting and storing their personal data. We did not take these concerns lightly, and took a privacy-by-design approach[21] when designing the user experience.

Legal Regulation in Call Interception

The Wiretap Act of 1986, as part of the amended Electronic Communications Privacy Act, protects the transit of wire, oral, and electronic communications[22]. It regulates the interception of voice and the substance, purport, or meaning of the communication. An interception is when there is access to a communication while it is in transit, no matter if you record and store the data.

First, we needed to consider if a real-time phone scam detection app constitutes as an interception. The app does not require human access to the phone communication; during the communication, a natural language processing algorithm would be analyzing the call. Boyden (2012) argues that automated processing that not involving a human

[21] Rubinstein, I., & Good, N. (2012). Privacy by Design: A Counterfactual Analysis of Google and Facebook Privacy Incidents (SSRN Scholarly Paper No. ID 2128146). Retrieved from Social Science Research Network website: https://papers.ssrn.com/abstract=2128146
[22] 18 U.S. Code § 2511

does not violate the Wiretap Act: “The history, purpose, and judicial interpretation of the ECPA all support this reading: interception requires at least the potential for human awareness of the contents” (pg. 1)[23].

However, this is one reading of the Wiretap Act and many might argue otherwise. In order to comply with the Wiretap Act, there must be consent for the interception. Consent can come in the form of explicit consent, when the parties involved in the communication are informed that their voice may be processed, and it is implied that they provided consent when they choose to remain in the communication. Consent must be direct, meaning that it cannot be solely a notice in the Privacy Policy. Furthermore, in the U.S. eleven states require all party consent, whereas the rest require one-party consent. One exception to requiring consent for recording is CA Penal Code 631, the self-defense exception[24]. However, the exception requires that it must be known the individual is threatened within the communication. One way this can be resolved is if the call is a telemarketing call to a mobile phone, which is illegal through the Telephone Consumer Protection Act of 1991[25].

Although it can be argued that we are not legally required to disclose consent for the real-time phone scam detection app, we decided that obtaining consent for real-time analysis of users’ phone calls would be in our best interest in avoiding legal violations and providing transparency to our users. Therefore, when the phone scam detection feature is enabled, we generate an audible consent disclosure to all parties on the line: “The person you are calling is using a screening service and will receive a transcript of this conversation”.

Data Processing and Storage

Data processing and storage is another issue we considered as part of our privacy-by-design approach. We wanted to ensure that we only collect sufficient data required for scam detection and do not retain unnecessary data. In order to perform natural language processing for scam detection, we determined that we do not need access to the entire audio recording. Rather, we just need the content of the conversation. Therefore, we decided to convert the audio conversation into text and solely analyze the text when applying the scam detection model during a phone call.

---

[23] Boyden, Bruce E., Can a Computer Intercept Your Email? (March 26, 2012). Cardozo Law Review, Vol. 34, No. 2, 2012; Marquette Law School Legal Studies Paper No. 12-05. Available at SSRN: https://ssrn.com/abstract=2028980 or http://dx.doi.org/10.2139/ssrn.2028980

[24] CA Penal Code 631

[25] https://www.fcc.gov/document/rules-and-regulations-implementing-tcpa-act-1991

In order to build a phone scam detection model, we need data to train the model and improve it over time. As detailed in a later section on building the NLP model, there are no easily accessible call data to build this model. Asking users to donate their call transcripts in the name of science was our only way of creating a dataset. The user will have the option to donate any transcripts, which would be transferred securely over to a secure server. Thus, only donated calls will be accessible and processed externally to improve the phone scam detection model. If the user decides not to donate a transcript, the transcript will be deleted and not retained (not even locally).

Even though the user chooses whether or not to donate the call transcript, there might be personal information included in the content of the conversation. Hence, we came up with the idea of using basic redaction techniques so we are not exposed to the user's personal information as much as possible. Additionally, although we can ensure privacy of phone call data at the server where we store the transcripts using encryption, it is still a possibility for hacker to hack into the servers.

Design Workflows

Onboarding

Due to the concerns around privacy uncovered in our user research, we made extensive considerations around user control and defaults. We wanted to ensure that users are properly protected from phone scams, but not at the expense of relinquishing control over their own privacy preferences. Within scam detection, there are three key features: the act of applying scam detection to a call, the sensitivity of the scam detection model, and the method of communicating a detected scam. We could design these features to be controlled by the user to a certain degree, and our question was - to what degree?

An initial idea was allowing users to influence the scam detection feature according to their “proficiency level” during onboarding.

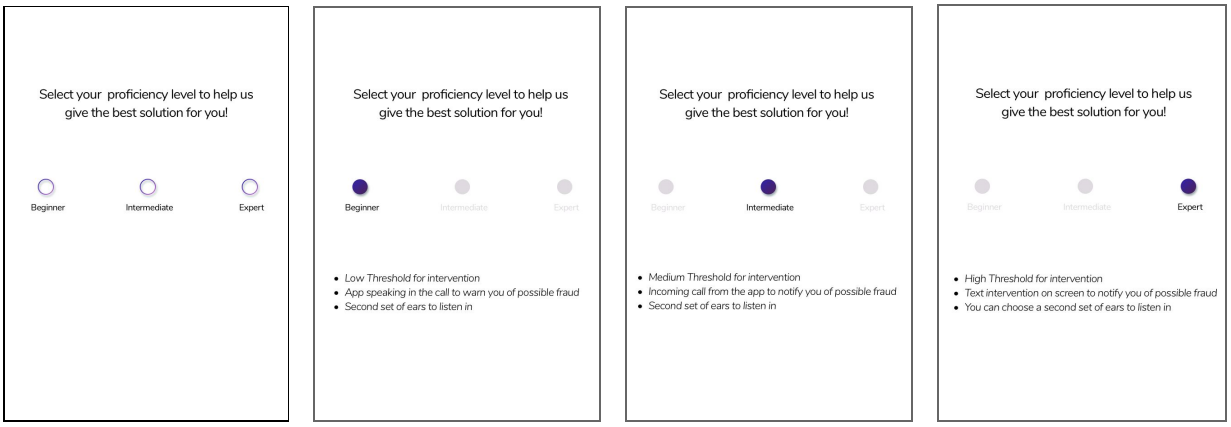

However, we determined this design would have a high potential to leave users unprotected due to the uncertainty of assessing one’s proficiency level with phone scams. As our user research and background research showed, people are very confident in their ability to detect scams until they fall victim to one. Therefore, there was considerable tension between providing user control over determining their proficiency level versus forcing all users into one standard level. In the end, we decided to focus on the design goal of keeping things simple rather than adding complexity with various settings. Therefore, we moved the most important control - turning on or off the scam analysis feature - into the phone call flow. As for the detection sensitivity and notification

method, we decided to keep it standard across all users for now, and leave it as a possibility to iterate on this design if usability testing shows a need for different people to have different scam detection thresholds and/or alert interventions.

Onboarding was still an important place to convey critical information to the user. The goal of the app is to detect phone scams without any excessive features, like any typical utility app. In order to keep the app simple, there are not many opportunities to communicate to the user. One of the important findings from the UX Research class team is the need to understand what the app does and why does it do it, specifically around transcribing phone calls. Therefore, we determined that onboarding would be the most fitting place to communicate this information to the user.

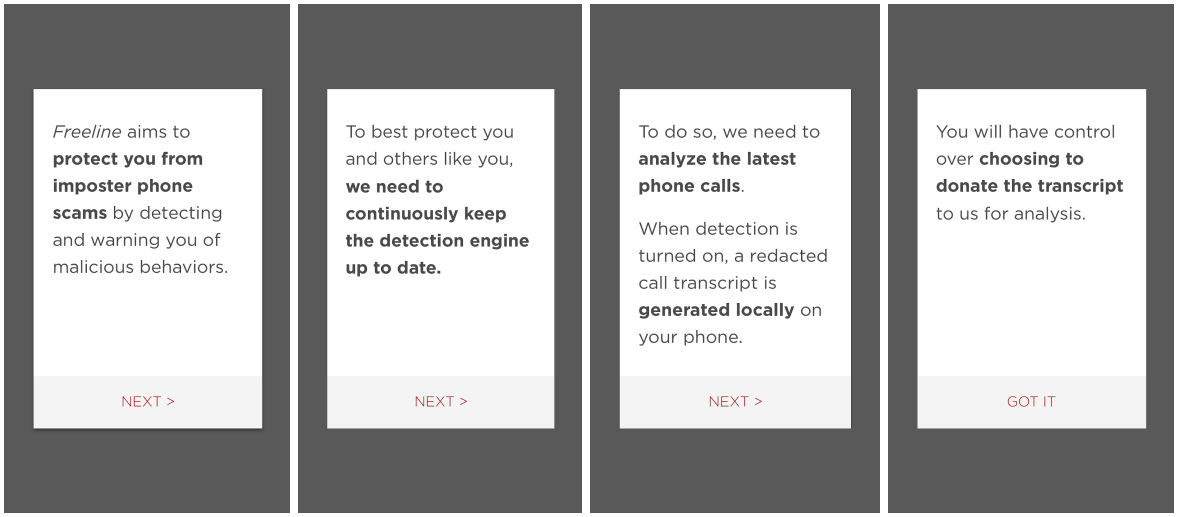

After the user first installs and opens the app, onboarding takes the user though an explanation of the app's main goal: to protect the user against imposter phone scams. It goes onto explain that in order to best protect its users, it must continuously maintain up to date with the latest phone scams by analyzing new phone calls. We emphasize that a redacted call transcript is generated locally on the phone, not on any external service. The user then has the choice to donate the transcript to assist in improving the scam detection model.

Scam Detection User Flow

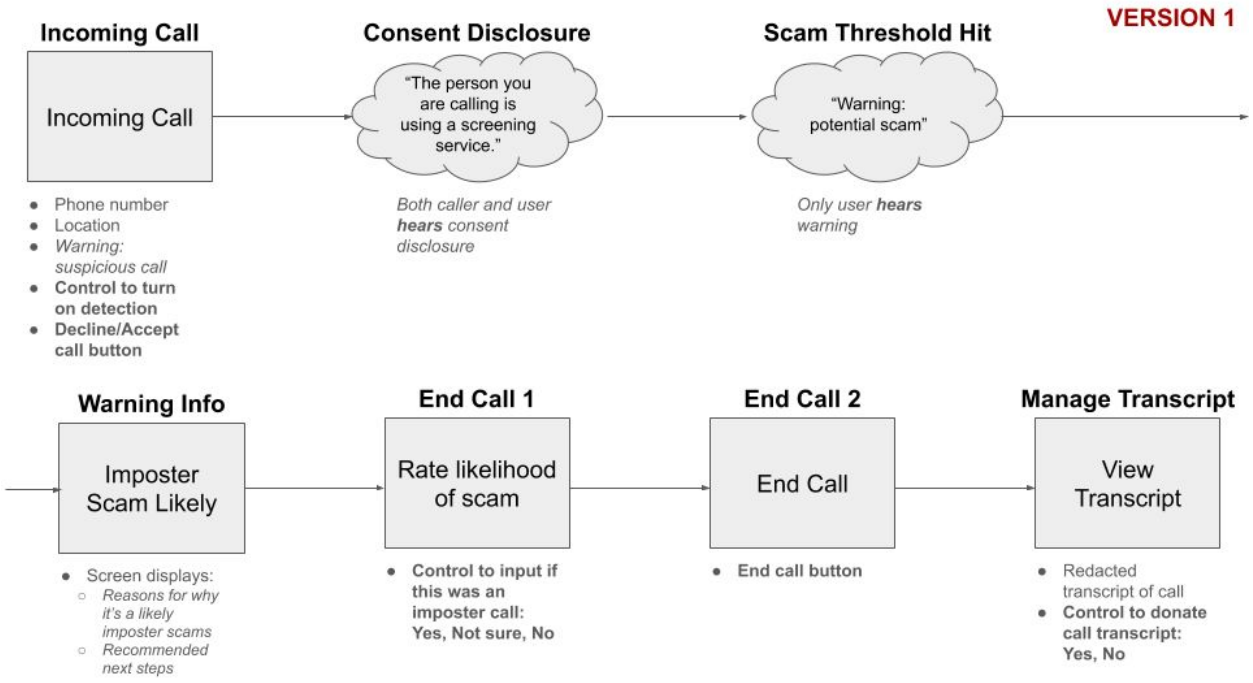

In our first version of the phone scam detection user flow, the incoming phone call screen displays typical call details such as the phone number and location of the call, and the button controls to accept or decline the call. There is an additional suspicious call warning, and a control to turn on the detection analysis. If the detection analysis is turned on, there is a consent disclosure heard at the beginning of the call once it has been picked up by the user. As the phone conversation proceeds, the call is analyzed in the background. However, once the scam detection threshold is hit, there is a audible warning to the user only that there is a potential scam. On the phone screen, additional details about the scam alert is displayed, such as reasons for why it is a likely a phone scam and recommended next steps. Before user can end the call, they must rate the likelihood of the call being an imposter scam call to help us label the call. Afterwards, they can end the call and view the redacted transcript of the call. At this point, there is a control to choose whether to donate the transcript.

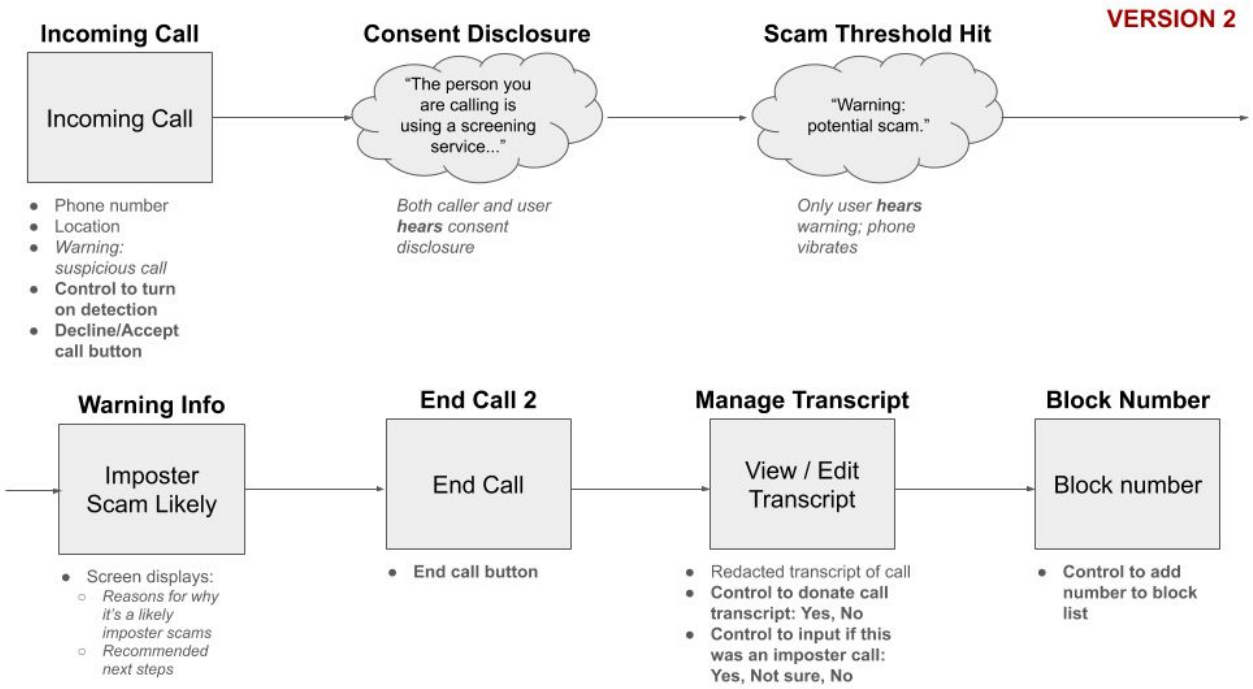

After receiving feedback from the UX Research class team on our first version of the user flow, we decided to make a number of changes to the scam detection user flow. In case the user does not hear the audible warning alert, we add a vibration as part of the notification to further get the user's attention. We also modified the feedback and transcript flows so that the user can edit the transcript to remove any additional information that they do not wish to share besides personally identifiable information from the redaction model. We moved the feedback mechanism so that the user is not forced into providing feedback that may not be used if they choose not to donate the transcript about the call. Finally, we added a mechanism to easily block the number due to user research feedback in which users want to immediately address any known unwanted callers.

Freeline App Screens

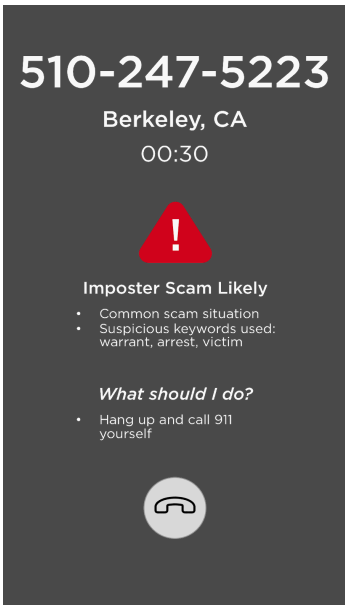

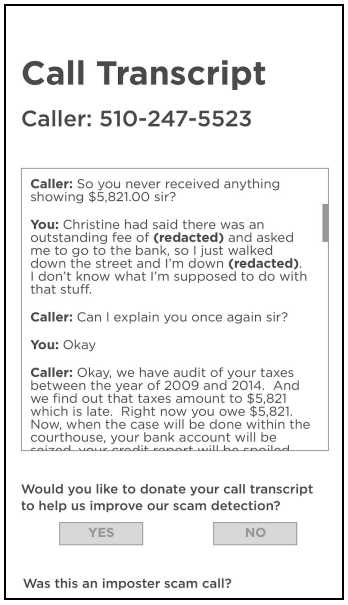

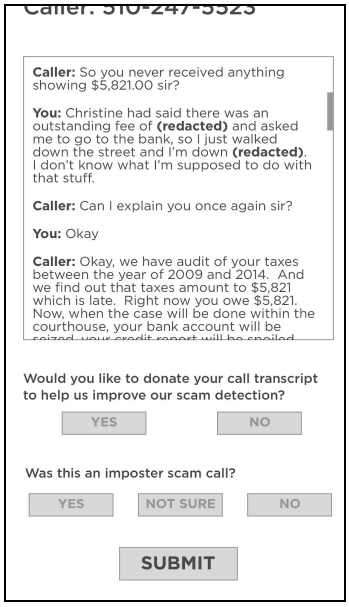

Natural Language Processing Model

Background Research

While conducting background research on the subject, we found some notable research papers on applying machine learning to the phone phishing problem. Bordjiba et al. used complaint data about phone calls (scams and not scams) to determine phone numbers that were spoofed and groups that were involved in attacks[26]. Edwards et al. worked on detecting persuasion in fraudulent emails[27]. Marzouli et al. have done a data-driven analysis of phone fraud landscape using a honeypot technique to identify that most spam calls originate from a few perpetrators[28].

Authority and Social Power

Amongst various social engineering tactics, Bullee et al contend that "authority" in social attack is a major driver to subjugate people into revealing personal or financial information, or performing some call to action[29]. They found that 80% of scam callers use authority to persuade their targets.

Lam et al.'s "Power Networks: A Novel Neural Architecture to Predict Power Relations" indicates that neural network models are able to predict power relations between people in an email conversation[30]. Situational power is defined as the power or authority one individual holds over another in the context of a task or a specific situation. Power differentials are apparent between people in interactions, but it is difficult to measure or detect using machines. Natural language analysis of written communication has been promising to detect other types of power like organizational power.

---

[26]Bordjiba, H. E., Karbab, E. B., & Debbabi, M. (2018). Data-driven approach for automatic telephony threat analysis and campaign detection. Digital Investigation, 24, S131–S141. https://doi.org/10.1016/j.diin.2018.01.016

[27] Edwards, M., Peersman, C., & Rashid, A. (2017). Scamming the Scammers: Towards Automatic Detection of Persuasion in Advance Fee Frauds. Proceedings of the 26th International Conference on World Wide Web Companion - WWW '17 Companion, 1291–1299. https://doi.org/10.1145/3041021.3053889

[28] Marzuoli, A., Kingravi, H. A., Dewey, D., & Pienta, R. (2016). Uncovering the Landscape of Fraud and Spam in the Telephony Channel. 2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA), 853–858. https://doi.org/10.1109/ICMLA.2016.0153

[29] Bullée et al. (2018)

[30] Lam, M., Xu, C., Kong, A., & Prabhakaran, V. (2018). Power Networks: A Novel Neural Architecture to Predict Power Relations. ArXiv:1807.06557 [Cs]. Retrieved from http://arxiv.org/abs/1807.06557

Literature review showed that in Gilbert (2012) showed that analysis of words and phrases that are strong predictors of workplace hierarchy[31]. Prabhakaran et al. (2012) did initial work on perception of situational power in and concluded that it does not correlate with hierarchical power within an organization. This can be understood from an example that there can be a situation where an employee knows more about a project than his/her manager and hence, if in a conversation about working of the project, the employee will have much more power than the manager, even though he/she is lower in rank.

To sum up the above papers into a theory, we hypothesize that being able to detect “authority” or “superiority” in language used in phone phishing calls could be a leading indicator to detect a scam when considered alongside other information about the call (frequency of the call, time of call, etc).

As a limitation of only using authority for scam detection, there are “service call” scammers who are more docile while pretending to help you fix your computer. These scammers are less authoritative over the phone. So, they won’t be picked up as as scam calls using the authority model.

Authority also works in different ways within various social relationships. Let us consider some cases. For example, it is not just that an authoritative figure will always be a scammer. A caller can be your mother, your boss or even the authorities (police, etc). Now, mothers and bosses can be frequently callers over a long period of time. So, given the history of such callers, they should be classified as normal calls even if they are ‘authoritative’.

So, in the case of a new caller who is not on the contact list, who speaks authoritatively and may end up calling very frequently over a short period of days, those calls would be considered as scam. Calling at odd times of the day/night can also be considered a feature.

Data Collection Strategy and Issues

The basic premise of the phone scam detection app is to classify an incoming call into two classes, namely scam and non-scam. Therefore, the ideal dataset would

[31] Gilbert, E. (2012). Phrases that signal workplace hierarchy. Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work - CSCW ’12, 1037. https://doi.org/10.1145/2145204.2145359

comprise of phone transcripts that contain scam and normal calls. As we did not originally find such datasource, our initial data gathering idea was to scrape web-forums to understand behavior of people being scammed, and featurize those descriptions.

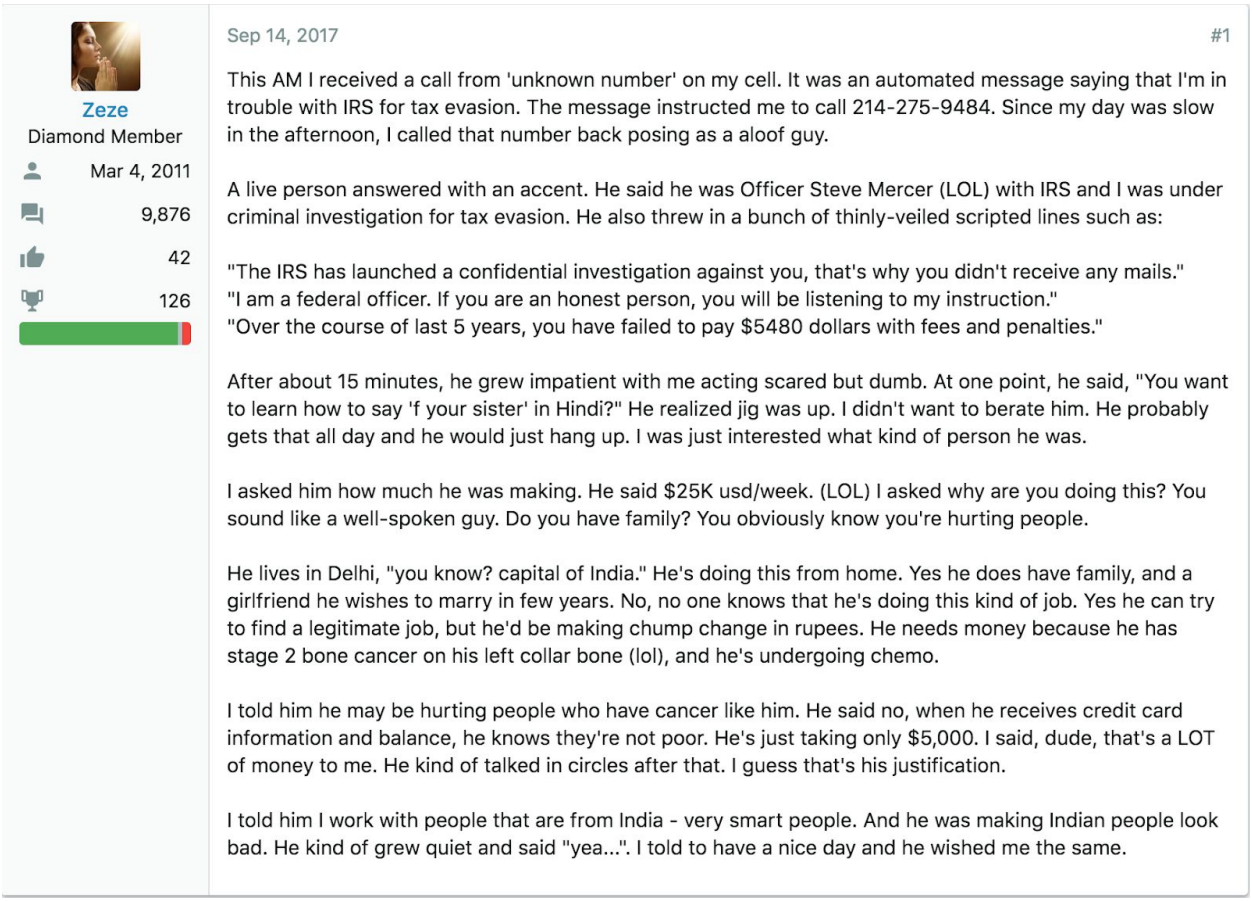

Zeze
Diamond Member
Mar 4, 2011
9,876
42
126

Sep 14, 2017 #1

This AM I received a call from 'unknown number' on my cell. It was an automated message saying that I'm in trouble with IRS for tax evasion. The message instructed me to call 214-275-9484. Since my day was slow in the afternoon, I called that number back posing as a aloof guy.

A live person answered with an accent. He said he was Officer Steve Mercer (LOL) with IRS and I was under criminal investigation for tax evasion. He also threw in a bunch of thinly-veiled scripted lines such as:

"The IRS has launched a confidential investigation against you, that's why you didn't receive any mails."
"I am a federal officer. If you are an honest person, you will be listening to my instruction."
"Over the course of last 5 years, you have failed to pay $5480 dollars with fees and penalties."

After about 15 minutes, he grew impatient with me acting scared but dumb. At one point, he said, "You want to learn how to say 'f your sister' in Hindi?" He realized jig was up. I didn't want to berate him. He probably gets that all day and he would just hang up. I was just interested what kind of person he was.

I asked him how much he was making. He said $25K usd/week. (LOL) I asked why are you doing this? You sound like a well-spoken guy. Do you have family? You obviously know you're hurting people.

He lives in Delhi, "you know? capital of India." He's doing this from home. Yes he does have family, and a girlfriend he wishes to marry in few years. No, no one knows that he's doing this kind of job. Yes he can try to find a legitimate job, but he'd be making chump change in rupees. He needs money because he has stage 2 bone cancer on his left collar bone (lol), and he's undergoing chemo.

I told him he may be hurting people who have cancer like him. He said no, when he receives credit card information and balance, he knows they're not poor. He's just taking only $5,000. I said, dude, that's a LOT of money to me. He kind of talked in circles after that. I guess that's his justification.

I told him I work with people that are from India - very smart people. And he was making Indian people look bad. He kind of grew quiet and said "yea...". I told to have a nice day and he wished me the same.

Retrieved from:
https://forums.anandtech.com/threads/so-i-called-back-an-irs-scam-call-talked-to-this-guy-in-india.2518768

It was observed that forum posts contain a lot of narration with snippets of the actual transcript. Although a lot was learned about the modus operandi of the scammers and reaction from victims by reading through similar forums, it became apparent that each forum post is very different from each other and creating a common methodology for scraping data was determined to be difficult.

With this issue on hand, we moved to looking for complete transcripts. Limited success was achieved in finding written scam transcripts as they were very few in number (~20). However, this low number wasn't ideal for processing using NLP methods. Our next approach was to ask a national mobile network provider to partner for call

transcripts. The response was it would take months to scrub personal information from the data and draft a non-disclosure agreement, which was not suitable for the project.

On continuation of our search, the National Cellular dataset (v2.3) was discovered. It looked promising as it contained voice samples of calls from 2336 speakers in English from locations throughout the United States. On further inquiry into the data and the data collection strategy, it became clear that these were free-flowing responses to predefined prompts. Given the prompts, all the phone conversations were normal calls, and there were no scam calls in the dataset.

At the same time, it was also discovered that there are a large number of YouTube videos of conversations during scam calls, which we were able to transcribe in a short time using an online transcription service. We considered merging these two datasets, transcripts from videos and the National Cellular dataset. But considering that for a classification model to predict classes with confidence, it should not have interference from the source of the data. Both the classes, scam and non-scam, need to be present in the same dataset. The idea of merging the two datasets was then abandoned.

Given the difficulty in attaining a dataset of call transcripts the next thought was to use something similar to conversations- emails - to create a data model, and then transfer and test that model on speech data. Therefore, to work around the data issue, our ultimate strategy was to train data that is similar to speech (text, in this case) and then test the model with speech data later on. We plan to use the same Enron email dataset[32] as the Lam et al. (2018) paper to model authority using email conversations labelled as “subordinate” and “superiors”[33]. Our goal here is to detect situational power of an individual using email threads, in which the individual is an active participant.

As future work, this model would be used alongside call features such as frequency, time of the day, whether the caller was in the contact list, call history, etc as these are also features that users go by when deciding to pick up a call.

NLP Model and Evaluation

The NLP model works very well with the Enron email dataset. We had extracted lexical and non-lexical features and created models that use either or both of them. Our approach was to use simpler methods first to generate a baseline model, and then use

[32] http: /www.cs.stanford.edu/~vinod/data/PowerAnnotations_V1.0.zip
[33] More information about the dataset available here: https://www.cs.cmu.edu/~./enron/

more complex methods such as a convolutional neural network (CNN) and Bi-directional Long Short Term Memory model (BiLSTM).

The body of the email was used as the lexical features for the data model. For the non-lexical features, we extracted annotation present in the dataset that indicated whether the email was written for requesting information, to inform, make a commitment, etc. Next, we also extracted how many questions were asked by a sender and how many of them were answered by the receiver in a particular email thread. These were done by counting the number of Forward and Backward links present in the email. They are called F-links/SFlinks and Blinks respectively.

A template of the emails in the dataset with features:

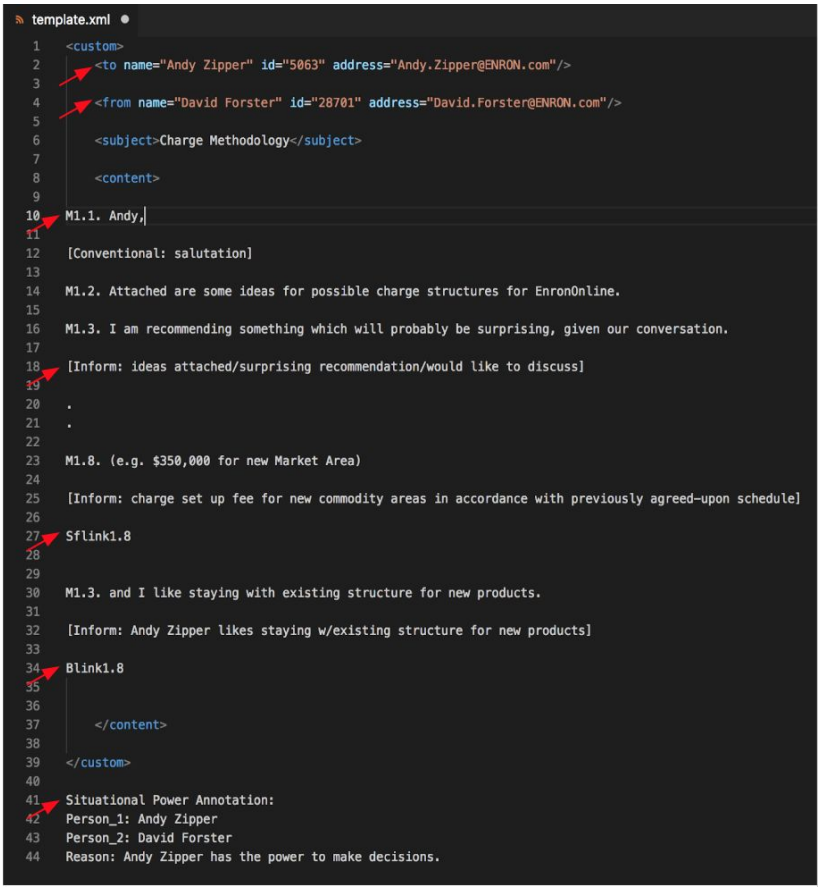

With logistic regression using just the body of the emails as input, our accuracy came up to 63% with an F-score of 0.63 in finding who has situational power in the email thread.

When the same input was used with a CNN, a better accuracy of 92% was achieved. The F-score in this case was 0.91

Next, we used a CNN with a bidirectional LSTM. It gave us a better result of 98.43%.

Given that we are able to predict which person in a conversation has power with a pretty high degree of confidence, the next step would be to test this on annotated conversational data from scam and normal calls for each of the developed models. We have received critical feedback that the accuracy seems to be really high and may not generalise well to real life data, which can well be the case as the model maybe learning the data in the corpus too well but may not generalise well on external data. This is yet to be tested and is on our list on our list of future work.

Engineering

Because of the data access issues, we added capturing data from users as a priority to our app design and concentrated on testing the technical feasibility to build our real-time phone scam detection application. Github information can be found in Appendix C.

We have started developing an application using React Native based on our app design. The main features that we wanted to test building were the following:

- Capturing calls
- Transcribing the call from speech to text
- Redacting sensitive data from the transcript
- Donating the transcript using a secure, encrypted API
- Storing donated transcripts
- Providing consent disclosure during the call
- Notifying user of scam detection warning

The following describes the data flow at a high level. As the call proceeds, a chunk of the call audio is sent to the Speech-to-Text (STT) engine. Once the STT engine spits out text for each audio chunk, these chunks are then pieced together on-device into a message. We decided against streaming the audio to the STT engine because that is known to drain the phone battery.

The scam detection NLP model is based on ‘message-level’ detection, which means it processes one speaker’s segment of speech at a time. By segment, we mean when one participant starts talking and ends talking. For example, see the below fictitious scam call. Each segment of speech should be ingested into the scam detection model one at a time:

> *Person 1: “I am calling from the Internal Revenue Service. You are under investigation for under reporting of your taxes for the year 2012. We have launched and investigation and this call is to gather more information and talking about next steps for you.”*
>
> *Person 2: “Wait, what? I filled my taxes using XYZ tax consultants from Boston. What is wrong with my taxes?”*
>
> *Person 1: “Sir, you have under reported them by $3000.”*
>
> *Person 2: “That is not possible!”*
>
> *Person 1: Sir, is 123-456-7890 your SSN id?*
>
> *Person 2: That is correct.*

The redaction will happen after the call ends. The redaction model does not need to be fed the text segment-by-segment, although that can be done. Redaction at a call level would save processing power as we won’t have to piece together the redaction at the end.

Architecture Diagram

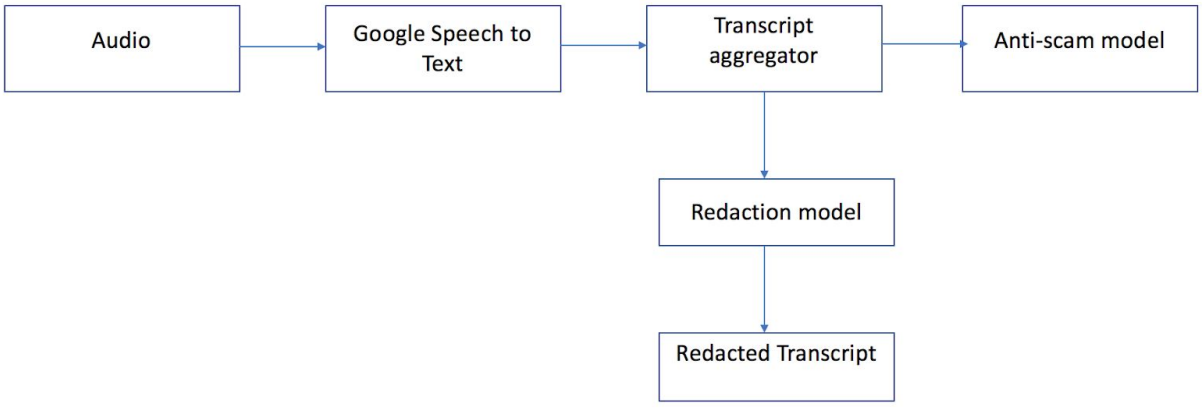

Android Call Capturing

We built a simple Android service using React Native for call capturing. Call recording is a complex task, so the design uses a very simple microphone as a source. The goal is to use a broadcast receiver and services together with React Native.

In simple terms, an Android service is a code kept in background and receiver is a code that will execute itself when some of registered events is fired. So in our case the receiver will call our service when something happened like a new incoming call is started or a phone state is changed (for example, a phone is picked up). This call will be recorded and passed on the Google Cloud Speech-to-Text API which transcribes the audio with caller annotation.

The current code is limited by permissions in Android 9 where the background applications do not have access to important sensors, such as the microphone. Therefore, this code is developed for any device below Android SDK 28 (below Android 9). For Android 9, the application would have to be overhauled into a foreground application that would essentially be a phone calling application because Android 9 requires the application to be in the foreground.

Call Transcription - Google Speech to Text

Google Cloud Speech-to-Text enables developers to convert audio to text by applying powerful neural network models developed by the firm. The API is low-latency and the return time is low enough for ingesting the text into the NLP model in near real-time. Also, the user would expect the transcript to be generated as soon as they end the call.

We decided to use this API because 1) we did not have capacity or expertise to build a speech to text engine, 2) text can be compressed for analysis as compared to audio, and 3) Google is a reliable product in terms of data privacy and security.

Redaction - Named Entity Recognition

To generate trust amongst users, a basic redaction strategy was developed. The users could see that sensitive information such as names, locations and some personal information is redacted on the transcript and not transmitted to us for any processing.

After the transcription is generated locally, the entire transcript will be run through a NLP model that detects named entities (names of people in the conversation, organizations and locations) and scrubs or deletes them from the transcript. Alongside that, we will use regular expressions to remove addresses, social security, tax and bank account numbers. The transcript would appear with these terms redacted when it is displayed to the user at the end of the call. See Appendix D for a link to more details with the NER model.

The following is an example of NER for names and geolocations. B- stands for beginning and I- statings for tag of the second word onwards of the detected entity.

```
text="Jenna works at the coffee shop at Microsoft and she works out of San Francisco"
text2 = "James works at the Apple store in near that office"
text3 = "My name is Jaime Lannister. I am from London in the UK"
predict(text, model, rev_tags)
```

```
Jenna   B-person
works   O
at      O
the     O
coffee  O
shop    O
at      O
Microsoft       B-facility
and     O
she     O
works   O
out     O
of      O
San     B-geo-loc
Francisco       I-geo-loc
```

```
predict(text2, model, rev_tags)
```

```
James   B-person
works   O
at      O
the     O
Apple   B-product
store   O
in      O
near    O
that    O
office  O
```

```
predict(text3, model, rev_tags)
```

```
My      O
name    O
is      O
Jaime   B-person
Lannister.      O
I       O
am      O
from    O
London  B-geo-loc
in      O
the     O
UK      B-geo-loc
```

Conclusion

Throughout this project, we faced many challenges given the complexity of our problem and project goals. We discovered valuable insights to our original mission in protecting consumers against phone phishing, as well as learned many lessons along the way. We were able to study the phone phishing scam experience from a variety of perspectives, design a real-time phone scam detection app with consideration towards user, legal, and engineering requirements, build an authority detection NLP model as a proxy to a phone scam detection model, and test the implementation feasibility of our app design.

We have a number of key takeaways from our process in exploring a solution to the phone phishing scam problem for consumers:

Design Tradeoffs

We came up against various tradeoffs in the values toward privacy and security between the values of everyday consumers and cybersecurity professionals. As experts in phone phishing, we understood the severity of the imposter phone scam problem and the difficulties for consumers to overcome the social engineering tactics of sophisticated scammers. However, when we presented our idea of the real-time phone scam detection app and its requirement to tap into phone calls in order to work, we were often met with skepticism. From the user research, some reactions were full on repulsion to the fact that their phone conversations would be analyzed, even thought it would be by a machine and not a human. Therefore, it was difficult to continue with the project knowing that we would need to compromise people's preference for privacy in order to provide security to them.

We made the decision to proceed with the design because we as security researchers understand the problem more than the typical consumer, and we attempted to design our app with privacy and security in mind. However, we believe researchers and security professionals must consider the question: at what point does the desire for privacy by the user outweigh the need for protection? A follow up question is who should make this decision - the user or the security expert?

Connecting with Experts

We had considerable challenges in getting time with experts. Although we had success in hearing back from experts with great interest about the project, almost every meeting never occurred. We reached out to campus law enforcement, social engineering professionals, and other researchers working on phone scam detection. We additionally hoped to connect with a team at a major mobile provider company, but we were also unsuccessful due to confidentiality concerns and protocols. It seems that everyone we contacted were enthusiastic or at least intrigued with the project, but understandably our project is not a high priority of theirs. We believe the challenge to collaborate with security professionals on the problem of consumer phone phishing must be noted, as it might be hindering the progress on addressing this problem.

A Really Hard Problem to Solve

The major challenges we faced in solving the phone phishing problem through NLP were:

1. Collection of phone transcript data
   Phone calls are a very private affair and people are unwilling to part with their call data. Hence, these kinds of datasets are close to impossible to find. The closest to free flowing speech we have come across was the National Cellular dataset, but even that is not a natural conversation that two normal people would have over a cell phone. The companies that have this data are either too wary of confidentiality breaches or are government affiliated, example DARPA[34].

2. Limitations to security
   Our goal was to design a process that can detect scams with the resources we had. The Google API of voice-to-text transcription works great, but technically someone can intercept the audio file segments being sent to the API because the audio cannot be encrypted when it is being sent to the API and get processed.

Almost all the problems we faced ensue from the fact that we need user data into a single remote location for training on models. However, what if that is not needed? What

[34] Shen, W. (n.d.). Active Social Engineering Defense (ASED). 15.

if there is a way of training and testing models on device and passing on the updates to a global model that can in-turn deploy updates to devices? Federated learning[35] is a possible answer. Federated learning at scale[36] is an ongoing research area. Google, as of mid-May 2019, has demonstrated it for its GBoard and a new Captioning product.

Beyond the issues with solving phone phishing with NLP, the problem still remains extremely challenging as we are going against human psychology - there are so many types of scams; people fall victim to scams in different ways; and scammers will keep updating their methods. Hence any machine needs to be flexible enough to tackle new attack vectors. It is very difficult to build such a model, and finding holes and breaking the algorithm is easier. In one conversation with a professional social engineer, he mentioned that attackers have always been able to work around any system.

Our solution, in its current state, can become a good data collection application. It is still limited by the above privacy and security concerns from becoming a usable product at scale. Our future work includes to continuing to implement the different features and connecting them together. We hope to move towards releasing the app in a controlled environment to generate the initial data corpus, which then can be used as the seed model for federated learning.

Overall this project was an exercise in addressing the issue of phone phishing for consumers. We gathered insightful information about imposter phone scams from a variety of sources, and we believe that our work has uncovered better understandings toward the challenges in solving this problem. We hope that the course of this project has raised more awareness around the dangers and severity of phone phishing.

[35] https://federated.withgoogle.com/; Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated Learning: Strategies for Improving Communication Efficiency. ArXiv:1610.05492 [Cs]. Retrieved from http://arxiv.org/abs/1610.05492

[36] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., … Roselander, J. (2019). Towards Federated Learning at Scale: System Design. ArXiv:1902.01046 [Cs, Stat]. Retrieved from http://arxiv.org/abs/1902.01046

Appendix

Appendix A - Qualitative Interview Guide

Interview Guide | Phone Phishing Targets

About the Incident(s):

- Can you tell me about the phone scam incident that you went through?
    - How did they engage you?
    - Did they have any information about you?
    - What actions did you take and why?
- What were ways that the attackers manipulated you?
- What made you trust the attackers?
    - Did the attackers say anything that made you trust them?
- How did you end up realizing it was a scam?
- What happened after the incident?
    - How did you feel afterwards?
    - What actions did you take to recover from the event?

Protection Against Phone Scams:

- What was your view of phone scams before the incident?
- Has anything changed after the incident?
    - In terms of your views on information security?
    - In terms of your behaviors with information security?
    - Are you more cautious of phone calls now?
    - After the incident, are you concerned that it could happen again? Have you taken any steps to protect yourself and/or others?
- What do you think could have helped you during the scam?
    - Are there things that you wish you could have known before the scam?
    - Are there things that you wish you could have done during the scam?

Real-time Detection Tool:

- What do you think about a real-time detection tool during scam calls? Basically it would be an app that records and tracks incoming calls from unfamiliar phone numbers, and raises a red flag when it detects a potential scam.
    - What are some reasons why you would use it? Some reasons why you wouldn't?

Appendix B - Consumer Phone Behaviors Survey

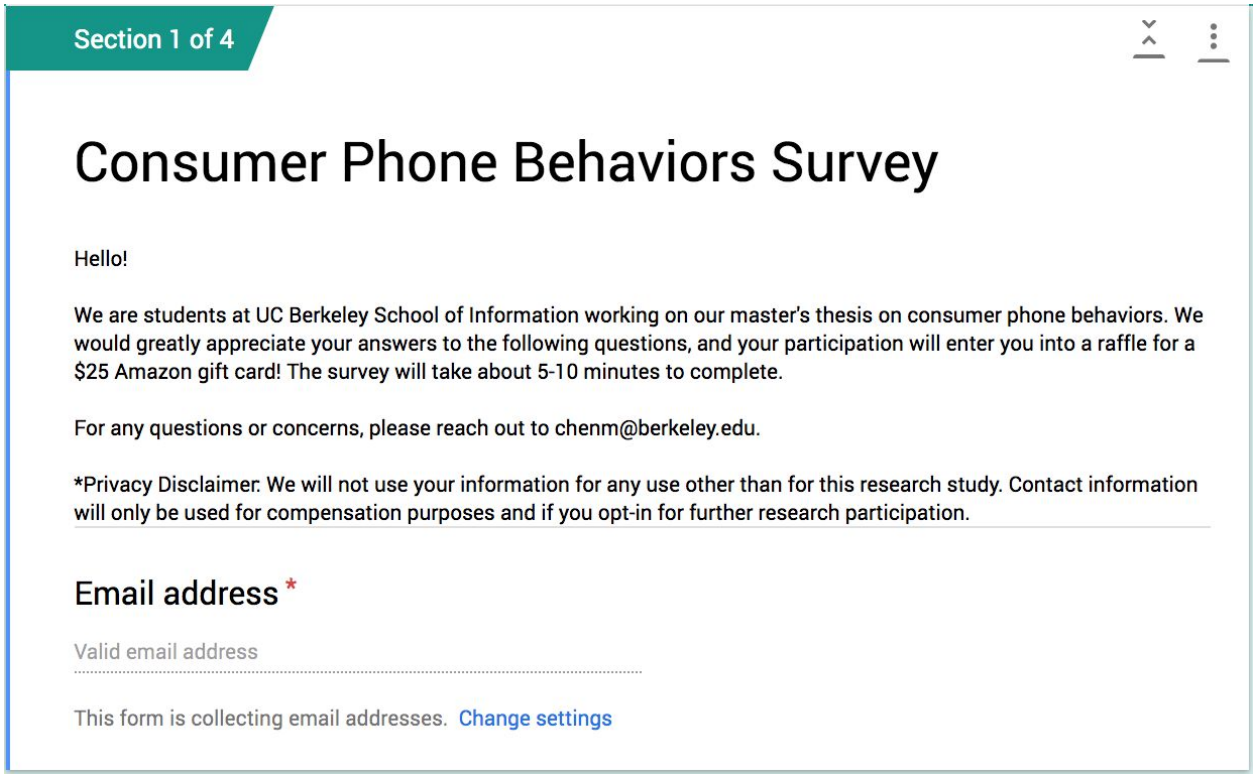

Section 1 of 4

Consumer Phone Behaviors Survey

Hello!

We are students at UC Berkeley School of Information working on our master's thesis on consumer phone behaviors. We would greatly appreciate your answers to the following questions, and your participation will enter you into a raffle for a $25 Amazon gift card! The survey will take about 5-10 minutes to complete.

For any questions or concerns, please reach out to chenm@berkeley.edu.

*Privacy Disclaimer: We will not use your information for any use other than for this research study. Contact information will only be used for compensation purposes and if you opt-in for further research participation.

Email address *

Valid email address

This form is collecting email addresses. Change settings

Section 2 of 4

Page 1 of 3

Description (optional)

During a phone scam, a caller pretends to be a trusted source and tries to trick you into revealing your personal information and/or sending money to them. In your opinion, how much of a problem are phone scams today? (1: Not a problem at all, 2: Minor problem, 3: Moderate problem, 4: Significant problem) *

	1	2	3	4	
Not a problem at all	○	○	○	○	Significant problem

Please indicate how true you believe the following statement is: Given it's not a blocked or private number, Caller ID on your mobile phone will always display the actual phone number that is calling you. (1: Very untrue, 2: Somewhat untrue, 3: Not sure, 4: Somewhat true, 5: Very true) *

	1	2	3	4	5	
Very untrue	○	○	○	○	○	Very true

Section 3 of 4

Page 2 of 3

Description (optional)

Your phone rings and the call is from a new, unfamiliar number. What do you do next? *

- ○ Pick up the phone
- ○ Let the call go to voicemail and listen to the voicemail (if applicable)
- ○ Always ignore the call
- ○ Find out information about the phone number
- ○ Other...

What makes a phone call suspicious to you? Please select all that apply. *

- ☐ I don't recognize the phone number.
- ☐ I don't recognize the area code.
- ☐ I don't know the person calling.
- ☐ I'm unfamiliar with what the caller is telling me.

- [ ] The first 6 digits of the phone number is the same as mine.
- [ ] It is not a local area code where I currently live.
- [ ] Other...

You pick up a phone call, and you are suspicious of the caller and what they are telling you. What do you do to address your suspicions about the caller? Please select all that apply. *

- [ ] I keep talking to the caller to get more information from them.
- [ ] While still on the call, I look up the caller's phone number online.
- [ ] While still on the call, I look up the caller's organization or company online.
- [ ] I hang up and get more information about the caller online.
- [ ] Other...

Which of the following apps or services do you currently use, if any? Please select all that apply. *

- [ ] Truecaller
- [ ] Hiya
- [ ] Robokiller
- [ ] Nomorobo
- [ ] Google's Call Screening (only available on Google Pixel)
- [ ] None
- [ ] Other...

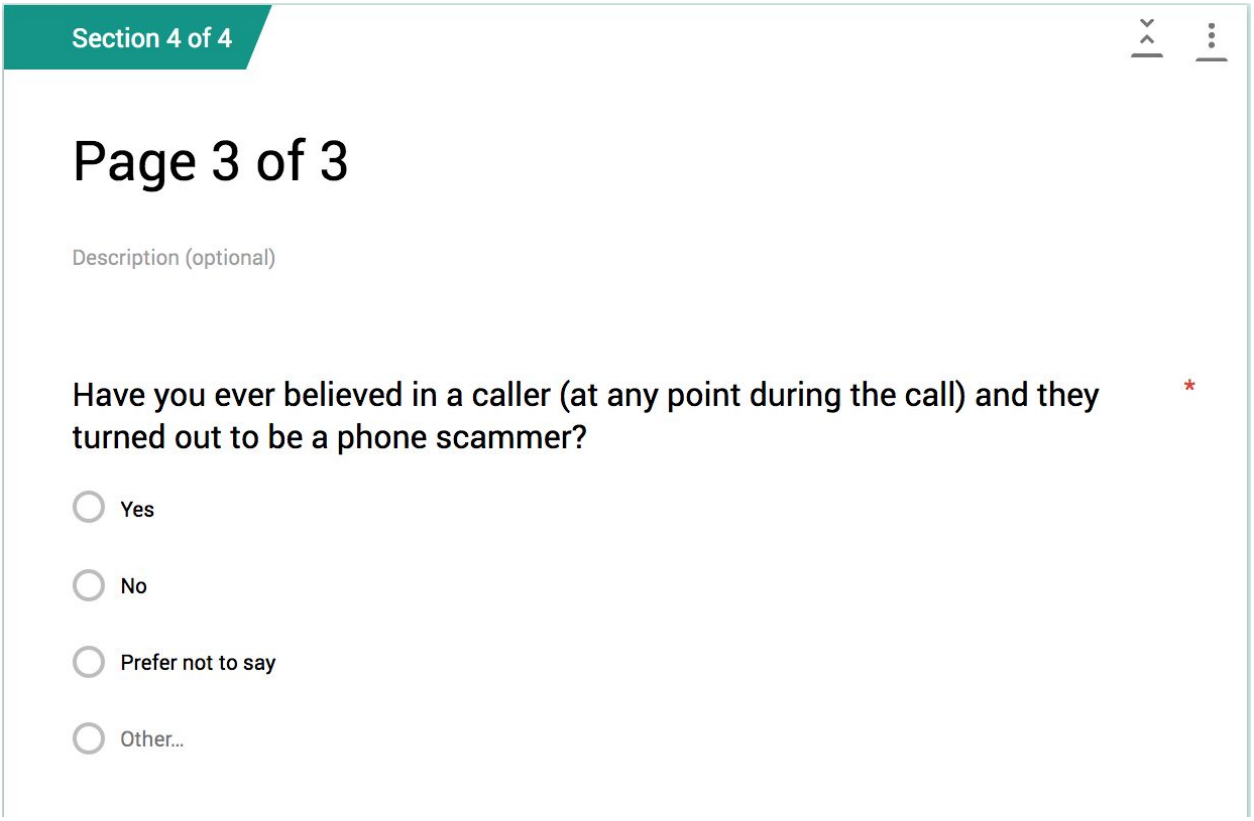

Appendix C - Github Link

https://github.com/surashish/CapstoneEngineering

Appendix D - NER Model

https://nbviewer.jupyter.org/gist/surashish/607cb3af511a9c761294ce6735ff98e4