New Me: Understanding Expert and Non-Expert Perceptions and Usage of the Tor Anonymity Network

Kevin Gallagher, *New York University;* **Sameer Patil,** *Indiana University;*
Nasir Memon, *New York University*

https://www.usenix.org/conference/soups2017/technical-sessions/presentation/gallagher

This paper is included in the Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017).

July 12–14, 2017 • Santa Clara, CA, USA

ISBN 978-1-931971-39-3

New Me: Understanding Expert and Non-Expert Perceptions and Usage of the Tor Anonymity Network

Kevin Gallagher
New York University
2 Metrotech Center
Brooklyn, NY 11201
kevin.gallagher@nyu.edu

Sameer Patil
Indiana University &
New York University
901 E. 10th Street
Bloomington, IN 47408
patil@indiana.edu

Nasir Memon
New York University
2 Metrotech Center
Brooklyn, NY 11201
memon@nyu.edu

ABSTRACT

Proper use of an anonymity system requires adequate understanding of how it functions. Yet, there is surprisingly little research that looks into user understanding and usage of anonymity software. Improper use stemming from a lack of sufficient knowledge of the system has the potential to lead to deanonymization, which may hold severe personal consequences for the user. We report on the understanding and the use of the Tor anonymity system. Via semi-structured interviews with 17 individuals (6 experts and 11 non-experts) we found that experts and non-experts view, understand, and use Tor in notably different ways. Moreover, both groups exhibit behavior as well as gaps in understanding that could potentially compromise anonymity. Based on these findings, we provide several suggestions for improving the user experience of Tor to facilitate better user understanding of its operation, threat model, and limitations.

1. INTRODUCTION

In the past, it was often sufficient to exclude one's name from an interaction to protect one's identity. The information age, however, requires more advanced means to achieve anonymity [21]. Many anonymity systems have risen to meet that demand. These anonymity systems play a vital role in the lives of society's important actors, such as journalists, activists, dissidents, law enforcement agents, and individuals for whom the disclosure of identity could lead to severe consequences. Further, these systems provide a means to assert important civil liberties, such as privacy, freedom of expression, etc. In fact, the use of such tools by the general population has experienced a large rise[1] in the aftermath of Edward Snowden's revelations of mass surveillance activities of the National Security Agency (NSA). However, incorrect use of these systems can lead to deanonymization which, in turn, can lead to a variety of consequences, ranging from slight embarrassment to imprisonment and, in extreme circumstances, death. Additionally, the strength of the anonymity system depends on the number of indistinguishable users [3]. As a result, when a user deanonymizes him- or herself, he or she weakens the strength of the anonymity network as a whole. Proper understanding and use of anonymity tools, therefore, play an important role in ensuring accurate and effective achievement of anonymity via the system. Yet, there has been surprisingly little research that looks into how users understand and conceptualize the underlying operation of these systems. We aim to address this gap.

Though many anonymity systems with different technical details and threat models [2, 34] exist, currently the most popular anonymity system is Tor [4]. Tor is a low-latency network that provides anonymity when performing tasks such as Web browsing. It works using the concept of onion routing [8], which routes traffic through multiple volunteer-run nodes, removing a layer of encryption at each node. By default, the number of nodes is set to three, which is the minimum number of nodes required to achieve anonymity. With a circuit of three or more hops, each node knows the identities of only the immediate predecessor and successor. As a result, no node knows both the source and the destination of a message. When the traffic arrives at the last node, or the 'exit node,' the plain text of the message is forwarded to the destination. In addition, Tor provides support for Onion Services,[2] which allow a server and a client to contact each other without knowing each other's IP addresses.

Tor is used globally by a wide variety of people. According to estimates by the Tor Project, Tor averages around 1,750,000–2,000,000 unique users each day from all around the world.[1] Many of these individuals are located in nations with oppressive regimes. For instance, recent estimates of Tor usage show that between June and August 2016, daily number of users of Tor in Iran ranged from 10,500–12,000.[1] Individuals based in such countries use Tor to access information from sources forbidden or censored by their nations and to pass along information about abuses of their governments to parties who can publish it without fear of retribution. Additionally, the Tor network provides the underlying platform for chat programs, such as ricochet,[3] and file shar-

[1] https://metrics.torproject.org

[2] Onion Services were previously referred to as Hidden Services, a term still used by some users.

[3] https://ricochet.im

Symposium on Usable Privacy and Security (SOUPS) 2017, July 12–14, 2017, Santa Clara, California.

ing programs, such as SecureDrop,[4] that allow journalists to communicate with sources confidentially and anonymously. As mentioned earlier, Tor is used by ordinary citizens seeking to escape ubiquitous surveillance [26] and censorship. Therefore, inaccurate use of Tor that leads to deanonymization holds potential for great individual and societal harm.

Due to its popularity and importance as an anonymity tool, we focused on Tor to investigate people's use of anonymity systems along with their understanding of the threat model and system operation. Specifically, we addressed the following research questions:

1. Why do people use Tor?
2. How well do users understand the underlying operation of the Tor system?

We tackled the above research questions by conducting semi-structured interviews with a diverse sample of 17 Tor users. Based on an analysis of the interview responses, we make the following contributions:

- We describe user perceptions and practices regarding Tor, an anonymity tool of growing individual and societal importance.
- We uncover and describe important differences in how experts and non-experts understand and conceptualize Tor. Specifically, we show that gaps and inaccuracies in non-expert understanding of the operation and threat model of Tor could lead to a sense of more or less privacy and security than is actually the case.
- We suggest solutions that can improve the Tor user experience and boost adoption by non-experts, many of whom are in vulnerable situations and/or serve as society's important actors.

In the next section, we summarize prior research on the usability of Tor as well as that of privacy and security tools, in general. We then outline the method we used to conduct our study along with the details of participant recruitment and a description of the sample. Next, we describe our findings followed by a discussion of the insight that emerged. We proceed to apply the insight to suggest a number of potential improvements to Tor and other related aspects. We conclude after pointing out important limitations and avenues for future work.

2. RELATED WORK

There is a vast body of work on the technical aspects of Tor, such as attacks, defenses, case studies, etc. [10,11,27,33]. In contrast, our focus is on Tor users and their user experience. In this regard, we first present existing research that specifically targets the user experience of Tor, followed by a summary of the literature on user experience considerations in privacy and security tools, in general. We highlight the lack of research attention to studying Tor users and their motivations and practices.

2.1 User Experience of Tor

As Dingledine and Mathewson [3] observed, user-centered security [35] is important for anonymity systems since improving the user experience attracts more users, which strengthens the network as a whole. To this end, studies of the user experience of Tor have covered software and network operation, user interface, and external factors.

One of the first studies regarding the user experience of anonymity systems introduced latency 'shocks' into the anonymity network *'AN.ON'* over a one month period. A latency shock occurred every 105 minutes and lasted 15 minutes [15]. The results showed that the number of users who leave an anonymity network because of latency is linearly related to the amount of latency, for latency periods lasting less than 60 seconds. Fabian et al. [6] applied metrics from the literature to investigate and quantify such losses in usability caused by the latency within the Tor network. When compared with direct connections, they found that the median load time for a Web page over Tor was 5 times higher and Domain Name System (DNS) requests were 40 times slower. Based on these measurements, they postulated a request cancellation rate of 74%, leading to potential user frustration when using Tor. Given the negative impact of latency on the user experience, understanding and fixing the causes of latency within the Tor network is an important ongoing concern of Tor developers [5].

Other studies have examined the user experience of the various user interface elements of Tor. Clark et al. [1] performed a cognitive walkthrough of four configurations of the Tor software, performing four tasks in each of the configurations. They proposed user interface changes based on the difficulties encountered in completing the tasks. Norcie et al. [23] tried to identify the challenges experienced by individuals in adopting and using Tor, beginning with the step of installing the software. Their study of 25 undergraduates found that 64% of the participants faced various problems in installing and using the Tor Browser Bundle to perform the given tasks. These problems included difficulties finding and downloading the installation program, issues with decompressing the installation file, confusion in distinguishing between the Tor Browser Bundle and Firefox, latency, etc. In a follow-up study, Norcie et al. [22] evaluated the effectiveness of their proposed interface solutions aimed at fixing the problems uncovered in their initial study. They found statistically significant usability improvement in the case of most issues. Similarly, Lee et al. [17] examined the usability of the Tor Launcher that configures Tor connections. They found that the Tor Launcher interface required users to understand technical terms and did not provide appropriate and adequate feedback, thus leading to frustration and errors. They further showed that interface changes to the Tor Launcher were effective in addressing these challenges.

In a different vein, Khattak et al. [14] investigated how the Tor user experience is affected by the actions of external parties. Specifically, they looked at how Tor users are treated at the application as well as the network layer. They discovered that 1.3 million IPv4 addresses and 3.67% of the Alexa top 1,000 websites offered degraded services to Tor users or blocked them altogether.

2.2 User Experience of Privacy and Security Tools

At a more general level, researchers have devoted attention to the user experience of various commonly used privacy and security tools and mechanisms. We highlight the most salient findings in this domain pertaining to expert and non-

[4] https://securedrop.org

expert understanding and behaviors.

Leon et al. [18] studied 9 tools designed to limit or prevent online behavioral advertising and found significant usability problems in all of them, making it difficult, if not impossible, for users to make meaningful opt-out choices. Wash [31] and Wash and Rader [32] described variations in user mental models regarding viruses and hackers and explained that user decisions to follow security guidance from domain experts were influenced by the specifics of these mental models. Ion et al. [9] found that security non-experts deferred or ignored installing software updates, did not employ two-factor authentication, and did not use a password manager. They suggest that better messaging and usability are required to address the lack of adoption of common security tools. Similarly, Kang et al. [12] reported large differences in the complexity of the mental models of tech savvy participants and others. Yet, they found no link between technical knowledge and attempts to control online privacy. McGregor et al. [20] focused on journalists, a user group that often encounters situations that require anonymity, for sources as well as themselves. Journalists from the US and France indicated resorting to ad-hoc security approaches due to the lack of comprehensive and usable tools and reported difficulties in authenticating sources using existing tools.

2.3 Tor Users

While several of the studies mentioned above focused on the Tor *system*, very few of them attempted to understand Tor *users*. McCoy et al. [19] analyzed the traffic from an entry guard and an exit node under their control and found that a disproportionate number of users of the Tor network hailed from Germany, Turkey, and Italy. Additionally, they uncovered that notable amounts of sensitive information was sent as plain text over insecure protocols. In contrast to such indirect indicators of Tor user practices, we present accounts of Tor use obtained directly from the users themselves. Additionally, we discuss user motivations for adopting Tor and describe user understanding of Tor operation and threat model.

A recent survey reported that 34% of a sample of American adults who were aware of government surveillance programs took steps toward protecting their online information from the government [24]. Yet, only 2 of these people reported using anonymity software such as Tor, highlighting the huge gap between the expressed need for anonymity systems and their adoption in practice. Based on interviews of 17 current Tor users, we suggest user experience improvements that could help broaden its adoption.

3. METHOD

To address our research questions, we conducted semi-structured interviews with individuals who reported using Tor. The subsections below describe how we recruited participants and provide the details of our study protocol. The protocol was approved by New York University's Institutional Review Board (IRB).

3.1 Recruitment

Recruiting Tor users for such a study is difficult because only a small proportion of the population uses Tor. Moreover, Tor users are likely privacy conscious and, as a result, may be unwilling to discuss their attitudes and behaviors, especially pertaining to their use of Tor. Therefore, we cast a wide net and utilized multiple channels to seek study participants. Such an approach was also aimed at increasing the diversity of the sample. Specifically, we advertised the study on the Tor community of Reddit,[5] the 'Et Cetera Jobs' category of Craigslist for the New York City area, and mailing lists and bulletin boards at New York University. When describing the study on Reddit's Tor community and at the university, we mentioned that the research was regarding Tor. In contrast, on Craigslist, we stated that we were studying software use, without specifying our focus on Tor. This dual strategy was adopted partially to overcome the difficulties of attracting Reddit and university participants for a general software study and partially to include participants with varying levels of familiarity and experience with Tor. Our Craigslist advertisement directed potential participants to a brief online screening questionnaire (see Appendix A). Along with age, gender, and email address, the questionnaire asked about the use of 14 technologies and online services, with 'anonymization software' as one of the options in the randomly ordered list. Those who indicated using anonymization software were contacted to ask if they had ever used Tor.

3.2 Participants

We set up interviews with the individuals who reported having used Tor and expressed willingness to participate in the study. Overall, we interviewed 17 participants (5, 2, and 10 via Reddit, university channels, and Craigslist, respectively): 10 males, 5 females and 2 who preferred not to reveal their gender. Apart from ensuring that each participant was above the age of 18, we did not collect age information in order to respect the privacy and anonymity of the participants.[6] Participant occupations covered a spectrum of technical sophistication from penetration tester to fitness trainer. As a token of appreciation for participating in the study, we offered each participant a $20 gift card for Starbucks. Many participants declined the reward, likely to preserve their anonymity.

3.3 Study Protocol

Prior to participation, we provided the participants with information on the purpose of the study along with the procedures followed for handling the collected data. Specifically, we stated that we would not collect any personally identifiable information and would treat all responses as anonymous and confidential.

After obtaining informed consent for participation (and optionally for audio recording the conversation), we interviewed the participants one-on-one using a semi-structured interview protocol (see Appendix B). When possible, interviews with the participants local to the New York City area were conducted in person at New York University. Others were interviewed via phone or conferencing software, with the exception of one participant interviewed via email[7] and

[5] https://reddit.com/r/Tor

[6] Based on the responses to the screening questionnaire and our interactions with the participants, we estimate the age range to be 21–50.

[7] The questions were sent to the participant in an initial email, with subsequent emails used to ask follow-up questions as necessary.

Tasks	Internet Service Provider	Government and Law Enforcement	Target Web site or Service	Advertising Networks
Browsing a Web site	Can see one is using Tor	Can potentially see one is using Tor	Can see some Tor user is visiting the site	Can see some Tor user is visiting the site
Reading email	Can see one is using Tor	Can potentially see one is using Tor	Can access identity and data, but not IP	Can see some Tor user is visiting the site
Receiving an advertisement	Can see one is using Tor	Can potentially see one is using Tor	Can see some Tor user is visiting the site	Can see some Tor user is visiting the site

Table 1: An empty version of the above table was presented to the participants during the interview. Participants were instructed to fill out the cells indicating which information about them they believed the corresponding entities could access when they performed the listed tasks with the Tor Browser Bundle. The above table shows the correct answers derived from the Tor Project documentation [30].

two others interviewed using a text chat program.[8] The first author conducted all interviews.

Each interview consisted of several open-ended questions. At the beginning, the participants were asked general questions about their occupation to make them feel at ease and establish rapport. After the introductory questions, the interview delved into the participants' use of Tor, beginning with how they discovered Tor and covering the details of why, where, when, and how they used Tor. We further asked the participants to describe their understanding of how Tor works.

For an elicitation of the participants' understanding of the underlying operation of Tor, we asked them to engage in a drawing task as suggested by Kearney et al. [13]. Specifically, we asked the participants to draw a free-form sketch of their views and understanding of Tor, including its various front- and back-end (i.e., visible and invisible) components, processes, and actors. We stated that the sketches may include information about data flows and access controls. As they drew, the participants were encouraged to vocalize their thoughts in order to allow the collection and comprehension of the corresponding detail. Those who were interviewed via phone, conferencing, chat, or email were asked to send a picture of the drawing to the interviewer. When needed, we sought clarification and asked follow-up questions during the task. All drawings were retained for analysis.

We next asked the participants to fill out a table to capture their awareness of the threats countered by Tor (see Table 1). The table included a set of tasks along with various entities involved in those tasks. The participants were asked to indicate which pieces of information each of these entities could access when they used the Tor Browser Bundle to carry out each of the listed tasks. We encouraged the participants to think aloud when filling out the table. These answers, coupled with the responses to the other questions, allowed us to determine the participants' understandings of the potential deanonymization risks.

At the end, we asked the participants about the societal role of privacy tools, specifically in relation to contemporary national security debates and discussions in the US and Europe. We concluded the interviews with a brief multiple-choice questionnaire that used 5 questions on cybersecurity and anonymity taken from the 'Technical Knowledge of Privacy Tools Scale' from Kang et al. [12]. We chose this scale due to its topical relevance as well as short length. Participants who provided no more than one incorrect answer were marked as 'experts' with the remaining labeled 'non-experts.' These cutoffs were determined based on prior pilot testing with privacy and cybersecurity domain experts. Overall, 6 of our participants were classified as experts and the other 11 were treated as non-experts.

Most interviews lasted approximately 45 minutes. For the interviews that were audio recorded, the audio files were labeled with an anonymous identifier and destroyed after transcription. We analyzed the text of the interview responses along with the corresponding interviewer notes and the sketches collected during the drawing task. We followed an inductive process, allowing insight to emerge from the collected data. In order to avoid biasing the inductive analysis, we deferred a systematic review of the literature related to mental models of security and privacy tools until after the analysis was completed. The analysis included iterative open coding, axial coding, and selective coding [7] using the Atlas.ti software.

The first author began the three stages of coding – open, axial, and selective – right after the first interview. The coding proceeded continuously as the interviews were being conducted. During open coding, the text was coded sentence by sentence. Codes were created from the data with no initial hypotheses. For example, the sentence *"curiosity; I heard a lot of different things about it and was wondering how it works"* was labeled with the code 'being curious.' Axial coding examined the collection of codes generated by open coding and grouped related codes into categories. For instance, the codes 'feeling less watched,' 'feeling at ease,' 'evading surveillance,' and a few others were categorized under 'benefits derived from Tor use.' We further examined how frequently codes were mentioned together. Finally, in selective coding, the interactions between the categories and the codes were analyzed qualitatively and, to a smaller extent, quantitatively. The following sections describe the high level insight regarding user perceptions and understandings of Tor that emerged from the analysis.

4. FINDINGS

Unsurprisingly, we found notable and large differences between the experts and the non-experts in terms of understanding of the operation of Tor as well as the threat it counters. The experts exhibited deep knowledge of Tor's underlying operation while the views of the non-experts were simple and abstract. Notably, not all experts were free of

[8]These participants did not wish to reveal their voice and demanded a text communication channel with end-to-end encryption.

gaps in knowledge that could potentially affect anonymity during Tor use. Interestingly, the experts focused on the *technical details* of Tor operation, while the non-experts were much more likely to situate Tor within a broader *sociotechnical* landscape of purposes, actors, and values. We unpack these results by discussing the details of the participants' understanding of Tor operation and threat model, respectively.

4.1 Mental Models of Tor Operation

As mentioned above, we uncovered differences in the mental models of the experts and the non-experts pertaining to how Tor operates as a system. However, within each of the two participant groups, the models exhibited common threads.

4.1.1 Experts View Tor as a Complex Network

The experts understood Tor as a complex decentralized network used to move packets of information from one node to another. When describing how Tor works, the experts focused on network related aspects, such as connections, paths between Tor nodes, routing, etc., along with technical details, such as encryption layers. For example, one expert discussed the evolution in his understanding of Tor operation using the technical jargon of computer networks:

"When I started off I understood [Tor] pretty crudely as just kind of a way to get past state firewalls and to hide your identity from Web sites you are visiting. As I continued to use it, it's really good for NAT [Network Address Translation] traversal for example. Like, if you want to host a Web site from your home address and you're behind NAT, a Tor hidden service is a great way to give you that kind of access." (P8, Expert, Male)

Typically, the experts viewed the Tor network as composed of three elements: a sender, a receiver, and a path of decentralized nodes connecting the sender and the receiver. Moreover, they frequently referred to themselves as the sender who uses the network of Tor nodes to send messages to various receivers. For example, consider expert P10's sketches of Tor operation; he drew two diagrams, one depicting a connection from himself to a 'clearnet' site (see Figure 1) and another showing his connection to a Tor Onion Service (see Figure 2). In the first drawing, P10 indicated how relay information is loaded (including the possibility of a Tor bridge with obfuscation). The bottom half of the drawing shows that the traffic between the client (User) and the exit relay (Exit Node) is encrypted (green) and the traffic between the exit relay (Exit Node) and the Web site (Clearnet Site) is potentially unencrypted (red). In the second drawing, P10 showed the role of Tor Onion Service Directories, Rendezvous Points, and Introduction Points in connecting to a Tor Onion Service (Hidden Service). These drawings and descriptions present a mental model of the Tor network that demonstrates an understanding of the Tor system architecture akin to that of a Tor developer or researcher.

Other experts described Tor operation in varying levels of detail, with P10's being the most descriptive and complete. Despite differences in the level of completeness of the descriptions, all elicitations of the experts referred to the decentralized network nature of the Tor system architecture along with the role played by onion routing and encryption in the operation of Tor. For instance, the experts discussed the workings of Tor in terms of technical mechanisms, such as traffic obfuscation techniques, anti-tracking measures, latency reduction solutions, etc.

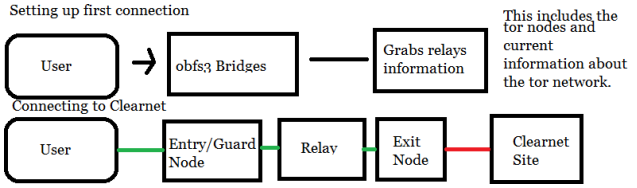

Figure 1: An expert's sketch of Tor's connection to a 'clearnet' Web site. (P10, Expert, Male)

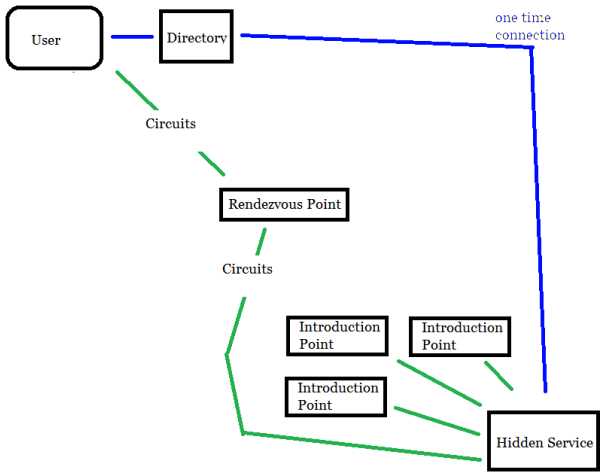

Figure 2: An expert's sketch of Tor's connection to a Tor Onion Service (Hidden Service). (P10, Expert, Male)

4.1.2 Non-experts Treat Tor as a Service

Seven of our non-experts began using Tor out of curiosity. This curiosity took different forms, with four curious about the 'Deep Web' and controversial hidden services and others about the ability to surf anonymously or bypass censorship. Similar to the experts, the non-experts viewed themselves as information senders within the Tor system. However, unlike the experts, the non-experts often treated several key components of Tor's network based architecture as an abstract and opaque 'black box' with certain inputs and outputs. Specifically, we noted that the non-experts tended to treat Tor as a 'service.' They described calling upon the Tor service to perform specific functions, such as "*bouncing signals*" (P3, Non-expert, Male) or "*providing security*" (P11, Non-expert, Male). For instance, non-expert P17 drew his model of Tor as a service that provides a "*new me,*" obscuring his identity from those he is connecting to (see Figure 3). Additionally, Figure 3 reveals that the non-experts often mistakenly understood the Tor 'service' as *centralized*, with an administrator watching over and controlling the operation of individual Tor nodes. Only one non-expert correctly mentioned the decentralized nature of Tor nodes.

Different non-experts believed that the Tor service performed different functions; some said it provided security, others mentioned it made them anonymous, and still others stated it granted them access to previously inaccessible sites and resources. These functions were seen as enabling Tor to help the user achieve specific goals and tasks. These included tasks such as visiting sites that the participant wished to

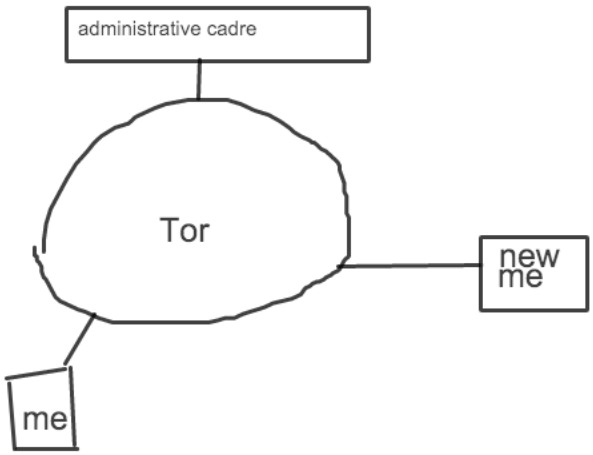

Figure 3: One non-expert's sketch describing Tor as a service with an administrative section watching over its inner workings. (P17, Non-expert, Male)

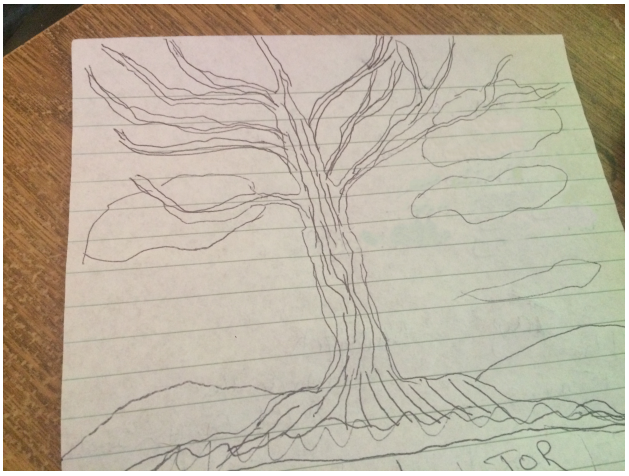

Figure 4: One non-expert's sketch depicting Tor as the Tree of Knowledge. (P9, Non-expert, Female)

conceal from the spouse, accessing geographically restricted content, circumventing content restrictions of filters and firewalls, etc.

While all non-experts described Tor as an abstract service, some descriptions exhibited more technical sophistication than others. For example, one participant mentioned that Tor may assign a new IP address, showing some understanding of the role of an IP address as an identifier.

"And then like IP address . . . I don't know . . . does Tor jumble up your IP? Maybe, perhaps it does, perhaps it doesn't. Perhaps it gives you a new IP." (P2, Non-expert, Female)

Three non-experts mentioned cryptography, even though they did not understand the role it played in the operation of Tor. Two non-experts mentioned 'signal bouncing' without explaining how it was accomplished.

"Like, the signal gets split up among other things, that would be cool if that happens, not too sure how that works, but I don't have an extensive knowledge of that." (P2, Non-expert, Female)

"As far as I'm aware the way it works is it bounces your signal around a lot . . . To various countries and such." (P3, Non-expert, Male)

It should be noted that there was a large degree of uncertainty among the non-experts about their understanding of the operation of Tor. While some non-experts were confident in their answers, five seemed unsure that their understanding was accurate or complete. For instance, when P2 was asked to clarify her idea that Tor performs *"signal dispersion,"* she replied that it works with *"cryptography,"* admitting that she did not know what that meant, indicating confusion between terminology and operation. Other non-experts simply stated that they did not understand how Tor worked, but knew that it did.

"It's one of those things where I know it works, it exists." (P9, Non-expert, Female)

Five non-experts described their understanding of Tor operation through metaphors. For example, one non-expert stated that it worked just like a faucet: *"if one turns the handle, the water appears"* (P9, Non-expert, Female). P11 clarified his sketch of Tor by equating it with Fort Knox, through which his traffic passed in order to become secure. This demonstrates his conception of Tor as a central service meant to secure, rather than anonymize, his traffic.

"Let's say I am like a circle. I am a circle on the left side. Inside of the circle I have for example, let's say I have my laptop and this for example would be down, hanging down. And on this it says it's my computer or it's my laptop. So that's on the left side and above that for example you can put any human picture and I give it a face. In the middle for example you will have a wall like Fort Knox and that would be in the middle obviously with no face because . . . it's not human and on the right side it is also a circle and that would be another computer with another human face." (P11, Non-expert, Male)

Many metaphors utilized by the non-experts described the ideologies and the values that the participants believed Tor stands for. For instance, P14 referred to Tor as the Statue of Liberty, bringing liberty to those who use it.

"Yes, the Statue of Liberty on Ellis Island. So just to describe to you we can probably explain it as giving us liberty to watch what I need, you know, and so at the same time [providing] freedom." (P14, Non-expert, Male)

Another example is a non-expert drawing Tor as the Tree of Knowledge (Figure 4), granting access to many different branches of knowledge.

"Tor ties in with the Tree of Knowledge for the simple reason that it's one of the best forms of confronting knowledge. Because there's no filters really on Tor." (P9, Non-expert, Female)

Apart from underscoring the non-expert treatment of Tor as a service, these metaphors also reveal that the non-experts often viewed Tor as a tool for social good. This aspect was mentioned in multiple non-expert interviews, with the participants discussing Tor as a tool used by activists, journalists, and ordinary citizens for communicating freely without surveillance, bypassing state censorship, and achieving em-

powerment in civic engagement.

4.2 Threat Model Addressed by Tor

During the interviews, we attempted to discover the participants' understandings of the threat model of Tor. We discovered misunderstandings of the following threats to anonymity on Tor:

1. **Client side scripting:** Client side scripting may place users at risk. For example, Flash code running outside the browser's control can be used to deanonymize users. Similarly, various vulnerabilities in JavaScript running within the browser can be exploited for deanonymization.
2. **Browser fingerprinting:** When Tor users use a browser other than the Tor Browser Bundle over the Tor network, the browser sends information to visited sites, such as installed add-ons, version, etc. Since the number of people with matching sets of information is likely to be low, the browser fingerprint lowers anonymity, with the worst case being unique identification.
3. **Side channel leaks:** Information provided by users to third parties external to Tor, such as login credentials, credit card numbers, or even language choice, can be used to denanonymize the user to varying degrees. In addition, if users do not ensure the use of encrypted connections, their information can be accessed by the exit node on their Tor circuit.
4. **Node operation:** Tor nodes are independently owned by volunteers. As a result, data flowing within the network is not controlled or seen by any single party, including the owners and the operators of the Tor Project itself. As a result, it is possible for malicious actors to run Tor nodes with the aim of attacking users who utilize the node (which is typically an exit node).

Similar to the operational details of the Tor system, the experts and the non-experts differed in the understanding of the threat model that Tor addresses. We discuss each in turn.

4.2.1 Experts Mostly Grasp the Threat Model

The experts showed a reasonably accurate understanding of the threat that Tor attempts to counter. Importantly, they understood that Tor is not a *complete* solution for all potential anonymity related issues and additional steps may be needed to achieve the desired level of anonymity. For instance, when filling out Table 1, expert responses revealed that they understood the complexities of the different browsing tasks and situations. These complexities are tied to the threat model of Tor. For example, all experts understood that logging into a Web site could deanonymize them. When asked whether the email service could access any information when reading email using the Tor Browser Bundle, the response of one expert demonstrated his understanding of the limits of Tor's protection:

"Yes they do, because you have an account with them. Assuming you've provided personal information, they kind of know who you are and, you know, what you've sent, but they still don't know where you are. You've still obscured your IP address." (P1, Expert, Male)

Further, all experts mentioned that the traffic exiting a Tor exit node may not be encrypted, again demonstrating the limitations of the protection Tor provides.

"Between the laptop and the entry I will write a little note that says ISP can see that I'm using Tor. And then between the entry and the middle I'm going to say 'encrypted traffic.' And then between the middle and the exit I'm going to say 'encrypted traffic' and between the exit and the Web site I'm going to say 'ISP can see requests, but not the originator."' (P1, Expert, Male)

In addition, many experts understood that the threat model of Tor allows a certain number of compromised Tor nodes, and that some of the Tor nodes might be a threat.

"Nodes may be owned/0wned[9] by governments." (P5, Expert, Unspecified gender)

Though the experts understood the threat model, all but two of them neglected to mention the Tor Browser Bundle as a part of the Tor system, mentioning only its network elements, such as the nodes, and security elements, such as encryption. Two experts configured their own Web browsers or used other non-standard ways to connect to the Tor network to receive Web content. This makes them vulnerable to fingerprinting attacks mentioned above, thus leading to potential deanonymization. Moreover, using a Web browser other than the Tor Browser Bundle is complicated and could lead to mistakes such as DNS leaks caused by a misconfigured browser resolving DNS requests independent of Tor. One expert stated that he used *wget* (an alternative tool for Web content retrieval) over Tor, which has a similar effect if the user does not anonymize the USER-AGENT string.[10]

4.2.2 Non-experts Conflate Threat Models

Unlike the experts, the responses of the non-experts revealed a lack of consensus regarding the threats that Tor addresses. While some non-experts possessed a complete understandings of the Tor threat model, five believed that Tor provided more security than it actually does. For instance, one non-expert believed that Tor was a tool for protecting sensitive data, such as credit card numbers, in transit on the Internet.

"It's going to something and entering my credit card or some kind of financial or some Web site where I don't want them to have my information because they're going to follow me." (P13, Non-expert, Female)

Another non-expert believed that Tor kept one anonymous from one's email provider, even when logged into the service. Other non-experts, however, held the view that Tor did not offer complete protection, with four claiming that Tor is effective for privacy protection from entities such as advertising networks, but not from governments and ISPs. Two others believed that the Tor Project has access to all traffic on the Tor network and could provide it to governments and law enforcement agencies. One non-expert argued that the Tor Project does not provide such access only because doing so would be counter to their goals.

"I know that I'm not doing anything dangerous but they don't know that, so I can see why the government would want to

[9] 0wned here refers to the computing slang term indicating a device being taken over and controlled by an external party, with or without the knowledge of the device owner.

[10] A USER-AGENT string is a line of text containing information about the browser or the program.

have access to that kind of thing. Or maybe they can receive alerts from Tor saying 'hey this person is suspicious by your standards' . . . but that's bad business, so. . . " (P4, Non-expert, Female)

Two non-experts claimed using Tor to circumvent geographical restrictions imposed by Web sites, such as Hulu, Netflix, etc. Yet, many of these sites run Adobe Flash or JavaScript, which can not only deanonymize users but also leave them vulnerable to injection attacks from malicious Tor exit nodes.

In general, the non-experts operated with incomplete, and sometimes inaccurate, understanding of the Tor threat model, often conflating it with other threat models that Tor is not designed to address. These gaps and inaccuracies could lead to a sense of more or less anonymity and privacy than is actually the case.

4.3 Discovery and Use of Tor

We examined how the participants discovered Tor, why they used it, and how long they had been using it. There are no real distinctions between the experts and the non-experts regarding the discovery of Tor. Both groups primarily discovered Tor through news articles, and many participants reported discovering it around the time of the initial publication of the Snowden documents. Some exceptions exist, with five participants finding Tor through searches on popular search engines or hearing about it from friends. One participant discovered Tor at a conference, and two participants (both experts) did not remember how they discovered Tor. In terms of use, however, we found significant differences between the experts and the non-experts.

4.3.1 Experts Used Tor for Many Reasons

All experts reported that they used Tor more frequently and for more purposes than the non-experts. A few experts used the Tor Browser Bundle as their primary browser, using it for most tasks and reserving non-anonymous browsers, such as Google Chrome and Mozilla Firefox, only for tasks which are ill-suited for the latency Tor creates (e.g., video streaming, etc.).

"I use [Tor] primarily as my everyday browser for most of my tasks. But I use regular Firefox if I want to do something, if the Web site is blocking Tor or if I want to do something on localhost that doesn't need outside Internet access." (P1, Expert, Male)

In addition to anonymous browsing and censorship circumvention, the experts mentioned alternative uses of Tor apart from Web browsing, such as downloading via alternative means such as *wget*, circumventing NAT using Onion Services, etc.

"So if I'm at school I can use Tor to ssh into a computer on my home network and it's not a problem. I don't have to deal with all of the IP address stuff." (P8, Expert, Male)

Curiosity differed between the experts and the non-experts. The experts tended to be curious about the network and its components, rather than the information held in Onion Services.

"Pure curiosity drove me toward it. It was just a different way of distributing information systems, so it was like, hmm, if we could do it a bit differently that would be a bit better." (P10, Expert, Male)

Additionally, the experts who started using Tor out of curiosity tended to remain Tor users and become more involved in the Tor community, while the non-experts who started using Tor due to curiosity stopped using it relatively quickly.

4.3.2 Non-experts Have Specific Motivations

Although the non-experts mentioned a variety of reasons for using Tor, all but two used it only within the context of a single specific purpose. Non-expert motivations for using Tor included: satisfying curiosity regarding the content accessible via Tor, bypassing censorship, circumventing geographical restrictions imposed by Digital Rights Management (DRM), countering surveillance by governments as well as other parties such as advertisers, communicating with activists, protecting the discovery of one's visits to pornography and gambling sites, researching sensitive legal matters, etc.

The non-experts who used Tor out of curiosity tended to be more curious about the information available via Tor rather than about the operation of the anonymity system itself. Specifically, the non-experts were drawn to information available on Onion Services, also called the 'Deep Web.'

"To be honest, the Internet black market. Uh, yeah, just to access it and see what's up. Um, the 'Deep Web.' Yes, that's it, the 'Deep Web.'" (P2, Non-expert, Female)

Four non-experts believed that Tor was designed primarily in the context of their own specific use case. For example, one non-expert used Tor only when abroad in a country that censored Web sites.

"I was using [Tor] because I was living abroad and I wasn't allowed to access certain sites. . . I was looking for ways to access these sites or rather looking for ways to get around the countrywide ban." (P4, Non-expert, Female)

She stated that Tor was not very needed in the US because the US government did not block many Web sites.

"I feel like it's less relevant in the US for the average user because the US doesn't block too much. They don't block Facebook and they don't block Google or that sort of thing. Whereas within a lot of foreign countries there's a lot of content that the US would consider benign that the governments wouldn't want you to access." (P4, Non-expert, Female)

Another non-expert used it only when performing credit card transactions, believing Tor to be a tool meant to safeguard data in transit.

Similarly, the non-experts tended to focus on only one adversary while using Tor. For example, one non-expert claimed that she used Tor because she did not want the government to see that she had looked up drugs, contract killer postings, and other such information.

"I just kind of used it those few times to look on the Internet and be like 'look how much acid costs on the Internet' and then like. . . find all the Web sites that are like oh I'm a hitman and I'm going to kill the president for a few million dollars. . . I like the president, but . . . I was kind of like just lurking and seeing what's up. That was the main pur-

pose and I didn't really want to get like a knock on my door, which they do in China... So that's like something that I'd like to avoid, which I'm sure doesn't happen as frequently in the States but... I don't like want to get arrested for some unrelated incident and then have my record like.... my computer searched, and then its like 'You were looking at hitmen, what's up with that?"' (P2, Non-expert, Female)

Lastly, half of the non-experts reported using Tor infrequently or having quit using it altogether, citing a lack of need or fading curiosity for the tool as their reasons.

4.4 National Security and Tor

In contrast to other aspects, we found no major differences among the experts and the non-experts regarding the relationship between Tor and national security concerns. When asked about the morality of Tor and its role in national security, most participants stated that Tor was a trade-off between privacy and national security and acknowledged that it likely made law enforcement more difficult. Yet, all but one participant believed that Tor was a good tool and the balance between individual privacy and national security should be closer to privacy.

"On balance I think that the good parts outweigh the bad parts and that they are necessary regardless of what we might think of the bad parts. So obviously properly implemented secure communication technologies will always be problems for law enforcement and intelligence agencies because they depend on sort of exclusive access to our data as part of their job. But ... I mean that's fine but there are other things at stake, right? There's individual liberty, there's freedom of speech, there's freedom of association, there's the ability to have secure technologies that will protect really important sensitive information, you know embarrassing stuff or your credit card number." (P8, Expert, Male)

There were some exceptions, however. One participant believed that privacy and national security are synergistic, and the protection of the rights of the people, including privacy, is itself a matter of national security.

"In my opinion, security is directly related to privacy and so is privacy to anonymity. I feel stronger tools are needed and are a benefit to society. Giving up any of the three (security, privacy, anonymity) means you can have none of the above. I understand the national security threat when the 'bad guys' use these tools, but they won't follow the rules anyway." (P6, Expert, Unspecified)

Another participant believed that Tor was detrimental to national security and should include a back door that allows access to the government.

"For national security reasons there is a need to have back hole [back door] access to certain things ... Tor is something that can be a very positive tool but at the same time it is used by a lot of illegal entities ... everything from child pornography to black market smuggling to terrorism, finances, planning, and coordination and so in that sense I think that there needs to be a certain degree of control from a government perspective." (P17, Non-expert, Male)

It must be noted that the views of the non-experts on this matter may have been influenced by some of the misunderstandings described in Sections 4.1 and 4.2. Specifically, a few non-experts believed that intelligence agencies, such as the Central Intelligence Agency (CIA) and the National Security Agency (NSA), are capable of defeating the protection Tor provides and have access to Tor network traffic. As mentioned earlier, one participant believed that Tor was capable of giving notices to the government if a Tor user is deemed suspicious by government standards, but would not do so because of the business implications of such an action. Yet, most non-experts believed that Tor helped foster important sociotechnical values, such as freedom of speech, uncensored information access, privacy, and personal security.

Category	Expert	Non-Expert
Mental Model	Complex network	On-demand service
Threat Model	Multiple threats	Specific (single) threat
Frequency of Use	Frequent	Mostly for specific uses
Discovery	Varied	Mostly through news
Morality of Tor	Good, positive	Varied, mostly positive

Table 2: Comparison of notable aspects of the understanding and the use of Tor across the experts and the non-experts.

5. DISCUSSION

Table 2 summarizes the notable aspects of our findings across the experts and the non-experts. In addition to the findings related to our research questions, we found that several experts and non-experts mentioned enhancing their anonymity and privacy by engaging in 'compartmentalization' via the use of a separate device for Tor use. Such a practice indicates greater attention to privacy and security among Tor users in comparison with non-users. While the adoption of Tor in the general population remains low, our sample shows that its user base is heterogenous and not composed only of domain experts with deep technical knowledge.

As expected, our findings confirm that the extent to which non-experts grasp the operational details of Tor differs substantially from the level of understanding of experts. Non-expert understanding of the operational details of Tor varied widely, possibly because of the differences in the frequencies and the motivations of use. Our findings shed light on the nature of these differences in terms of mental models and threat models. Regardless of the technical sophistication of these mental models, Tor, like any privacy enhancing technology, would benefit greatly from understanding and utilizing the mental models of its users [31]. For instance, the user interface as well as the documentation of Tor could draw upon the mental models to present the operational concepts more effectively.

The experts exhibited useful and complete knowledge of the Tor architecture and operation along with a nuanced understanding of its threat model. In contrast, the mental models of the non-experts were incomplete and overly abstract, leaving out or distorting important details that impact anonymity and privacy. For instance, bounding the entire Tor network within a single box may create a false sense of privacy and security by ignoring the potential at-

tacks by malicious exit nodes, such as capturing sensitive information passing through the node via insecure protocols [19]. Moreover, users operating under an assumption of anonymity may engage in behavior they might not want tied back to their identity. In contrast, viewing Tor as a centralized service could lead to the opposite effect. A belief that external parties, such as governments, law enforcement agencies, ISPs, and the Tor Project, can access decrypted Tor traffic has the potential to create a chilling effect, leading to self-censorship as well as unwillingness to use Tor. As mentioned earlier, even some experts exhibited gaps in understanding and engaged in behaviors that left them vulnerable to specific attacks, such as DNS leaks. This underscores that even the smallest of gaps in knowledge has the potential to defeat the anonymity protection a user seeks via Tor. Some of these issues, such as DNS leaks, can be addressed by the Tor software itself,[11] while others can be addressed by explicitly documenting the dangers of non-standard uses of Tor.

Tor is used by experts and non-experts across the world for a variety of purposes. Many of these purposes involve society's important values and causes, such as circumventing censorship, avoiding surveillance, sharing sensitive information of journalistic importance, communicating with informants, and so on. In addition, Tor serves sensitive and valuable personal purposes, such as protecting one's online activities from an abusive partner, avoiding targeted advertising, etc. In a large majority of these situations, the users involved are non-experts. In such circumstances, gaps and inaccuracies in the understanding of the operation and threat model of Tor that lead to deanonymization may hold serious repercussions, including account compromises, identity theft, financial losses (resulting from fraud), surveillance of communication and movements, civil or legal penalties, physical and/or psychological abuse, imprisonment, or, in extreme cases, death.

Interestingly, the responses of our non-experts show that they placed importance on the societal values that Tor aims to promote along with the corresponding usage scenarios tied to those values. Indeed, some of them seemed to be using Tor to make a value statement related to civil liberties and democratic principles, such as privacy, anonymity, freedom from surveillance, personal liberty, censorship circumvention, freedom of expression, etc. While the experts also recognized the connection of Tor to societal values, they preferred to describe Tor in terms of the architectural and engineering details of the software and the network. When considering whether Tor poses a problem for national security, participant opinions ranged from asserting that Tor acts as a force for freedom to believing that Tor is a tool for cybercriminals and terrorists.

Regardless of how they discovered Tor, the experts reported using Tor more frequently and for longer periods. In contrast, the interest of the non-experts tended to fade, with many claiming that they saw no need for the tool. While these usage differences have previously been observed in other privacy tools as well [25], they are especially crucial for an anonymity system such as Tor because the efficacy of its protection improves with an increase in the number of users. Given the awareness of Tor's value proposition exhibited by the non-experts, emphasizing that the use of Tor is a community and societal contribution could potentially boost its adoption.

A typical goal of Human Computer Interaction research is creating user experiences that facilitate effective use of a system without requiring deep knowledge of the underlying operation, thus making it easily accessible to non-experts. As discussed in Section 6, our findings can be applied to improve the Tor user experience for non-experts. However, a key aspect where Tor differs from typical systems is its use as a privacy and security tool, sometimes under circumstances of great importance as well as danger. As such, an incomplete or inaccurate understanding of its operational details has the potential for individual as well as societal harm. These risks lead to a tension between the need to promote technical understanding of the operational detail and the goal of making such knowledge unnecessary as a requirement for the correct use of the system. Addressing the issues uncovered by our findings could be a step in the direction of mitigating the potential risks and resolving the tension between the simultaneous needs for revealing as well as abstracting away the technical details of Tor operation.

6. IMPLICATIONS

Our findings can be applied to improve the Tor system in a variety of ways. These include refining the design of the user interface and the user experience of the Tor Browser Bundle, targeting specific operational aspects for enhancement and optimization, and facilitating learning, especially for non-experts. We discuss some of these below. In addition to these improvements, our data suggests that Tor users desire a reduction in latency.

6.1 Route Information

As discussed earlier, non-experts conceptualize Tor as a centralized service. A possible solution to avoid such a misunderstanding could be displaying information about the ownership of each Tor node in the current connection's route, when such information is available and verifiable. Such a feature would be an extension of the current Tor Browser Bundle functionality that allows clicking the onion logo to display route information, such as the node IP address and the country. It might also be useful to make such information readily available in the background without the need for explicit click-and-seek. The feature could be further expanded to indicate the encryption status of each link within the current route.

6.2 Safe Script Execution

Our findings suggest that a notable barrier to the adoption and the use of Tor is the demand and the need for using Web sites and services that utilize JavaScript. JavaScript is so ubiquitous that disabling it makes a large proportion of popular Web sites unusable [28]. As mentioned earlier, enabling JavaScript while using Tor may lead to deanonymization [29]. We advocate investigations of operational and architectural modifications that reduce the attack surfaces opened up by enabling JavaScript within the Tor Browser Bundle. Such technical improvements could facilitate a reasonable balance between preserving anonymity without

[11] For instance, DNS leaks can be addressed by raising a warning when the Tor network proxy receives a numeric IP address instead of a request for resolving a text based domain name.

overly compromising usability and utility.

In addition, the Tor Browser Bundle should be modified to warn users that it defaults to a *low* security level that has JavaScript enabled. In the current release, one must open a menu accessible by two clicks in order to discover this default setting. A prominent visual indicator of the current security level should be available at a glance in the Tor Browser Bundle interface.

6.3 Tor Friendly Web sites

In line with the empirical findings of Khattak et al. [14], most of our participants mentioned routinely having trouble due to the restrictions many Web sites place on Tor traffic. These sites typically restrict Tor traffic even under situations that pose minimal risk to the server, such as fetching static Web pages that do not involve user interaction or data input. We recommend that Web sites, especially those providing important information such as government Web sites, provide 'Tor friendly' versions of the pages that allow Tor users to at least fetch information, even if specific mechanisms, such as posting, are disallowed to protect against abuse. In addition, using CAPTCHAs to prevent abuse by Tor users should be limited only to submitting POST data, permitting GET requests without submitted parameters to proceed without such checks.

Moreover, site owners could consider serving an alternative Tor friendly version of the site for connections from Tor exit nodes. The process of creating such a Tor friendly version of a site could be made easier by promoting the creation of plugins for common Web development platforms, such as WordPress and Dreamweaver. Such plugins could ensure that Tor nodes are not blacklisted and automatically create versions of Web pages that reduce the amount of JavaScript to the bare minimum, or possibly none.

6.4 Compartmentalization

Our participants reported using the privacy enhancing strategy of compartmentalization by separating different tasks or personas through the use of separate computers and/or software. Yet, most current programs and operating systems make it challenging, if not impossible, to achieve meaningful compartmentalization of digital activities. We advocate explicit attention by system designers and Tor developers to the provision of compartmentalization functionalities as a privacy and security enhancing feature.

6.5 Maintaining Workflow

Although many participants in our study compartmentalized their Tor use, some participants indicated frustration at the burden of switching away from Tor in order to complete the tasks that could not be performed via Tor. These tasks included visiting Web sites that depend on flash or Java plugins or those that explicitly block Tor traffic. Currently, when a user wants to perform such tasks he or she must manually switch to another browser and copy/paste the site address. Subsequently, the user must remember to switch back to Tor once the task in the other browser is completed. While the task is ongoing, the user must switch back and forth between Tor and the other browser. The desire to minimize the disruption caused by the burden and frustration of managing the workflow and task switches can lead users to choose a non-Tor browser as the default. We suggest adding functionality within the Tor Browser Bundle that makes such task switches easier and faster whenever a task necessitates the use of a non-Tor browser. Such functionality has the potential to increase Tor adoption and usage by making it easier for users to stay within the Tor system as much as possible, switching away from Tor only when absolutely necessary for the task at hand. In addition to benefiting the individual user, increasing the time users spend using Tor would boost the overall utility of Tor by increasing the number of active users at any given time.

6.6 Contextual and Personalized Training

It would be beneficial to explore training and learning opportunities for non-experts in order to promote the development of useful conceptualization of the operation and threat model of Tor. Training has been shown to be effective in other cybersecurity domains, such as phishing [16]. In addition to an explicit focus on non-experts, training mechanisms could be customized to the person(s) and situation(s) at hand. For instance, different training modules could be developed for common use cases, such as circumventing censorship, avoiding surveillance, communicating securely and anonymously with a journalist, etc. Training activities could even be embedded into the user experience of the Tor Browser Bundle in a manner that utilizes learning theories and techniques, such as gradual knowledge building, periodic repetition, and effective assessment.

7. LIMITATIONS

A few limitations must be kept in mind when considering the generalizability of these findings. While we continued iterative coding of the interview responses until sufficient understanding emerged, we were unable to engage in purposeful additional sampling aimed at filling gaps. Although such a step is common in inductive qualitative analysis, the difficulties in finding and recruiting unbiased and unprimed Tor users limited our sampling efforts. Despite this limitation, we believe we reached reasonable saturation for the research questions at hand. In addition, the inherent difficulty in recruiting Tor users without bias or priming means that our study has a small sample size compared to other research on mental models. Further, advertising in Tor-specific groups such as Reddit's Tor community may have introduced bias in the sample. For privacy and anonymity, we did not collect demographic data beyond gender. As a result, we cannot account for cultural differences. Although we cannot be certain, advertising in online and offline communities in English leads us to believe that most of our participants were native residents of the US or Canada. Finally, we point out that our findings are derived from self-reports. Consequently, it is possible that the participants omitted, forgot, or misrepresented their understanding and behavior.

8. FUTURE WORK

To the best of our knowledge, we are the first to attempt to understand the mental models of users of any anonymity software. Our findings point to several opportunities for future sociotechnical research. In Section 6, we proposed potential solutions that involve changes to the Tor user interface and user experience. The effectiveness of these suggestions needs to be validated via empirical studies. In addition, we call for further design exploration in creating user experiences that balance the tension between revealing and abstracting the operational detail. Due to the qualitative

nature of our study, the findings are derived from a small sample. To validate generalizability, an online questionnaire could be formulated based on these findings and administered to a larger sample covering a broader population. We focused our investigation only on the Tor anonymity system. Further research is needed to examine whether these findings apply to other anonymity systems, such as the Invisible Internet Project (I2P) [34] and Freenet [2].

9. CONCLUSION

Anonymity systems, such as Tor, are an important tool for providing privacy and security in a landscape of growing online surveillance and censorship. In addition to enabling ordinary citizens to assert their civil liberties, Tor serves as a crucial anonymity and safety mechanism for society's important actors, such as journalists, political dissidents, whistle blowers, human rights activists, etc. A large majority of these actors are not technical domain experts. We found that non-experts conceptualize Tor via abstractions and metaphors that hide important operational aspects, thus potentially compromising the anonymity they seek. In contrast, experts understand the underlying technical operation and threat model and are highly likely to possess an accurate understanding of the level of privacy and security protection afforded by Tor. Fostering useful and complete understanding of the operation and threat model of Tor is a critical need to avoid deanonymizing vulnerable users as well as to promote adoption of Tor.

10. ACKNOWLEDGMENTS

We thank the participants of the study. We are grateful to Lesley Fosh, Jeffrey Ramdass, Martin Shelton, and Rick Wash for insightful feedback on draft versions of the paper. We acknowledge the anonymous reviewers for helpful comments. Thanks are due to Thomas Davis and Mihir Mahajan for editorial input. This work is supported by National Science Foundation (NSF) grant DGE-0966187.

11. REFERENCES

[1] J. Clark, P. C. van Oorschot, and C. Adams. Usability of Anonymous Web Browsing: An Examination of Tor Interfaces and Deployability. In *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS 2007)*, pages 41–51. ACM, 2007.

[2] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A Distributed Anonymous Information Storage and Retrieval System. In *Proceedings of Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, pages 46–66. Springer, 2001.

[3] R. Dingledine and N. Mathewson. Anonymity Loves Company: Usability and the Network Effect. In *Proceedings of Workshop on the Economics of Information Security (WEIS 2006)*, pages 547 – 559. Springer, 2006.

[4] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The Second-generation Onion Router. In *Proceedings of the 13th Conference on USENIX Security Symposium (USENIX Security 2004)*, pages 21–21. USENIX Association, 2004.

[5] R. Dingledine and S. J. Murdoch. Performance Improvements on Tor or, Why Tor is slow and what we're going to do about it. `https://www.torproject.org/press/presskit/2009-03-11-performance.pdf`. Accessed: 2017-06-15.

[6] B. Fabian, F. Goertz, S. Kunz, S. Müller, and M. Nitzsche. Privately Waiting – A Usability Analysis of the Tor Anonymity Network. In *Sustainable e-Business Management: Proceedings of the 16th Americas Conference on Information Systems (AMCIS 2010)*, pages 63–75. Springer, 2010.

[7] B. G. Glaser and A. L. Strauss. *The discovery of grounded theory: Strategies for qualitative research.* Transaction Publishers, 2009.

[8] D. M. Goldschlag, M. G. Reed, and P. F. Syverson. Hiding routing information. In *Proceedings of the International Workshop on Information Hiding*, pages 137–150. Springer, 1996.

[9] I. Ion, R. Reeder, and S. Consolvo. "... no one can hack my mind": Comparing Expert and Non-Expert Security Practices. In *Proceedings of Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 327–346. USENIX Association, 2015.

[10] A. D. Jaggard, A. Johnson, S. Cortes, P. Syverson, and J. Feigenbaum. 20,000 in league under the sea: Anonymous communication, trust, MLATs, and undersea cables. In *Proceedings on Privacy Enhancing Technologies (PoPETS 2015)*, pages 4–24. De Gruyter Open, 2015.

[11] R. Jansen, F. Tschorsch, A. Johnson, and B. Scheuermann. The sniper attack: Anonymously deanonymizing and disabling the Tor network. In *Proceedings of the Network and Distributed System Security Symposium 2014 (NDSS 2014)*. Internet Society, 2014.

[12] R. Kang, L. Dabbish, N. Fruchter, and S. Kiesler. "My Data Just Goes Everywhere": User Mental Models of the Internet and Implications for Privacy and Security. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 39–52. USENIX Association, 2015.

[13] A. R. Kearney and S. Kaplan. Toward a methodology for the measurement of knowledge structures of ordinary people: The conceptual content cognitive map (3CM). *Environment and Behavior*, 29(5):579–617, 1997.

[14] S. Khattak, D. Fifield, S. Afroz, M. Javed, S. Sundaresan, V. Paxson, S. J. Murdoch, and D. McCoy. Do You See What I See? Differential Treatment of Anonymous Users. In *Proceedings of the Network and Distributed System Security Symposium 2016 (NDSS 2016)*. Internet Society, 2016.

[15] S. Köpsell. Low Latency Anonymous Communication – How Long Are Users Willing to Wait? In *Proceedings of Emerging Trends in Information and Communication Security (ETRICS 2006)*, pages 221–237. Springer, 2006.

[16] P. Kumaraguru, J. Cranshaw, A. Acquisti, L. Cranor, J. Hong, M. A. Blair, and T. Pham. School of phish: A real-world evaluation of anti-phishing training. In *Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS 2009)*, pages 3:1–3:12. ACM, 2009.

[17] L. Lee, D. Fifield, N. Malkin, G. Iyer, S. Egelman, and D. Wagner. A Usability Evaluation of Tor Launcher.

In *Proceedings on Privacy Enhancing Technologies (PoPETs 2017)*, pages 87–106. De Gruyter Open, 2017.

[18] P. Leon, B. Ur, R. Shay, Y. Wang, R. Balebako, and L. Cranor. Why Johnny Can't Opt out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2012)*, pages 589–598. ACM, 2012.

[19] D. McCoy, K. Bauer, D. Grunwald, T. Kohno, and D. Sicker. Shining Light in Dark Places: Understanding the Tor Network. In *Proceedings of the 8th International Privacy Enhancing Technologies Syposium (PETS 2008)*, pages 63–76. Springer, 2008.

[20] S. E. McGregor, P. Charters, T. Holliday, and F. Roesner. Investigating the Computer Security Practices and Needs of Journalists. In *Proceedings of the 24th USENIX Security Symposium (USENIX Security 2015)*, pages 399–414. USENIX Association, 2015.

[21] H. Nissenbaum. The meaning of anonymity in an information age. *The Information Society*, 15(2):141–144, 1999.

[22] G. Norcie, J. Blythe, K. Caine, and L. J. Camp. Why Johnny Can't Blow the Whistle: Identifying and Reducing Usability Issues in Anonymity Systems. In *Proceedings of the 2014 Workshop on Usable Security (USEC 2014)*. Internet Society, 2014.

[23] G. Norcie, K. Caine, and L. J. Camp. Eliminating Stop-Points in the Installation and Use of Anonymity Systems: A Usability Evaluation of the Tor Browser Bundle. In *5th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETS 2012)*. PETS Symposium, 2012.

[24] L. Rainie, M. Shelton, and M. Madden. Americans' privacy strategies post-Snowden. *Pew Research Center*, March 2015.

[25] K. Renaud, M. Volkamer, and A. Renkema-Padmos. Why Doesn't Jane Protect Her Privacy? In *Proceedings of Privacy Enhancing Technologies (PoPETs 2014)*, pages 244–262. Springer, 2014.

[26] B. Schneier. How to Remain Secure Against the NSA. `https://www.schneier.com/blog/archives/2013/09/how_to_remain_s.html`. Accessed: 2017-06-15.

[27] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal. RAPTOR: Routing Attacks on Privacy in Tor. In *Proceedings of the 24th USENIX Security Symposium (USENIX Security 2015)*, pages 271–286. USENIX Association, 2015.

[28] The Tor Project. FAQ. `https://www.torproject.org/docs/faq.html.en`. Accessed: 2017-06-15.

[29] The Tor Project. Tor security advisory: Old Tor Browser Bundles vulnerable. `https://blog.torproject.org/blog/tor-security-advisory-old-tor-browser-bundles-vulnerable`. Accessed: 2017-06-15.

[30] The Tor Project. Understanding and Using Tor - An Introduction for the Lay(wo)man. `https://trac.torproject.org/projects/tor/wiki/doc/TorALaymansGuide`. Accessed: 2017-06-15.

[31] R. Wash. Folk Models of Home Computer Security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security (SOUPS 2010)*, pages 11:1–11:16. ACM, 2010.

[32] R. Wash and E. Rader. Too Much Knowledge? Security Beliefs and Protective Behaviors Among United States Internet Users. In *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 309–325. USENIX Association, 2015.

[33] P. Winter and S. Lindskog. How the Great Firewall of China is Blocking Tor. In *Proceedings of the 2nd USENIX Workshop on Free and Open Communications on the Internet*. USENIX Association, 2012.

[34] B. Zantout and R. Haraty. I2P data communication system. In *Proceedings of the 10th International Conference on Networks (ICN 2011)*, pages 401–409. ACM, 2011.

[35] M. E. Zurko and R. T. Simon. User-centered Security. In *Proceedings of the 1996 Workshop on New Security Paradigms (NSPW 1996)*, pages 27–33. ACM, 1996.

APPENDIX

A. SCREENING QUESTIONNAIRE

We invite you to participate in our study. Your participation will benefit science and help us understand user perceptions of software.

Our study involves a one-on-one interview. You may participate in-person or remotely via telephone or Voice-over-IP solutions, such as Skype. The interview will take a maximum of 45 minutes. Each participant will be compensated with a $20 Starbucks gift card.

To register, please answer the brief questionnaire below. We will contact you if we have an available position. Slots are limited, so if you wish to participate, please sign up as soon as possible.

If you have any questions, please contact us via email.

1) Age

- 18-24
- 25-34
- 35-44
- 45-54
- 55+
- Prefer not to say

2) Gender

- Male
- Female
- Other
- Prefer not to say

3) Email
Please enter your email address. This is the email address we will use to contact you.

4) Which of the devices below do you own and use? (Check all that apply.)

- Desktop Computer
- Laptop Computer
- Smartphone
- Tablet
- Other

5) Which of the technologies and services below have you ever used? (Check all that apply.)
[NOTE: Options were presented in random order.]

- Social Networking (Facebook, Twitter, LinkedIn, etc.)
- Online Audio and Video Conferencing (Skype, Facetime, etc.)
- Anonymization Software (Tor, etc.)
- Office Software (Word, Excel, Powerpoint, etc.)
- Online Music, TV, and Media
- Version Control Software (Git, Subversion, etc.)
- Online File Sharing (Dropbox, OneDrive, etc.)
- Mobile Messaging (Kik, Telegram, Snapchat, etc.)
- Online Banking
- Encryption Software
- Online Communities (Reddit, etc.)
- Computer Programming
- Online Shopping
- Blogging

B. INTERVIEW PROTOCOL

Thank you for taking the time to participate in this interview. The purpose of this interview is to discover your views, opinions, and understanding regarding how Tor works. Many people use Tor everyday for many reasons, from reading their email to accessing blocked Web sites. Please keep in mind that there is no single correct answer to these questions. Please answer the questions based on your own knowledge and experiences.

1. What do you do for a living? What does that entail?
2. What kind of computer(s) or mobile device(s) do you use? What are the differences (if any) in what these device(s) can do and how you use them?
3. On which of these device(s) do you use Tor?
4. When did you start using Tor? Why did you start using Tor?
5. How did you discover Tor?
6. Why do you currently use Tor?
7. This is a drawing exercise. Keeping background processes in mind, please draw what happens when you use the Tor Browser Bundle. Also note of who can access information about you. Please think aloud and explain your thought process while you are drawing.
8. How often do you use Tor?
9. What other browsers do you use?
10. Under which circumstances do you use the Tor Browser Bundle instead of another browser or vice versa?
11. Describe your feelings regarding the advantages and disadvantages of using Tor.
12. In what ways, if any, do you use Tor differently on your mobile device(s) than your computer(s)? (If applicable.)
13. Please fill out the given table of tasks and various entities involved in those tasks. For each of the tasks, mark the entities that you believe can access information about you when you perform the task using Tor. Please also mention what information you believe they can access. (See Table 1.)
14. Currently, a debate is going on about the role of privacy tools in matters pertaining to national security. Some people claim that strong privacy tools like Tor are good, while others claim they are bad. This is a part of a larger discussion about the trade-off between privacy and national security concerns. What is your opinion on this matter?
15. Is there anything else you would like to tell us? Is there anything that we should have asked?