

Towards Models for Quantifying the Known Adversary

Alaadin Addas
Ontario Tech University
alaadin.addas@uoit.ca

Julie Thorpe
Ontario Tech University
julie.thorpe@uoit.ca

Amirali Salehi-Abari
Ontario Tech University
abari@uoit.ca

ABSTRACT

The *known adversary threat model* has drawn growing attention of the security community. The known adversary is any individual with elevated first-hand knowledge of a potential victim and/or elevated access to a potential victim's devices. However, little attention is given on how to carefully recruit paired participants for user studies, who are qualified as legitimate known adversaries. Also, there is no formal framework for detecting and quantifying the known adversary. We develop three models, inspired by Social Psychology literature, to quantify the known adversary in paired user studies, and test them using a case study. Our results indicate that our proposed *adapted-relationship closeness inventory* and *known adversary inventory* models could accurately quantify and predict the known adversary. We subsequently discuss how social network analysis and artificial intelligence can automatically quantify the known adversary using publicly available data. We further discuss how these technologies can help the development of privacy assistants, which can automatically mitigate the risk of sharing sensitive information with potential known adversaries.

CCS CONCEPTS

• Security and privacy → Social aspects of security and privacy;

KEYWORDS

known adversary, insider threat, social psychology, autobiographical authentication, social closeness

ACM Reference Format:

Alaadin Addas, Julie Thorpe, and Amirali Salehi-Abari. 2019. Towards Models for Quantifying the Known Adversary. In *Proceedings of 2019 New Security Paradigms Workshop (NSPW'19)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/nmmnnnn.nnnnnnn>

1 INTRODUCTION

Authentication schemes are often tested under several different threat models for security assessment before wide-scale industry use. Recently, a special class of threat models, dubbed the *known adversary threat*, has called the attention of the security community [3, 21, 22, 29]. The known adversary is any individual with elevated first-hand knowledge of a potential victim and/or elevated access to

a potential victim's devices, that uses these privileges with malicious intent. The known adversary threat also accounts for the risk of unauthorized access to devices/accounts by social insiders [29].¹

Known adversaries often have physical access to the devices/first-hand knowledge of a potential victim, which gives them an advantage, enabling them to access accounts/devices without permission. These intrusions can be a major privacy breach to potential victims, and can cause social and financial harm (e.g., accessing a bank account and transferring funds without permission). As a result, a few authentication schemes have been tested for resilience to the known adversary threat [3, 17, 21, 22, 33].

A common practice for examining the resilience of authentication systems to the known adversary threat is to recruit study participants in pairs, who try to authenticate as each other [3, 21, 22]. The participating pairs are asked to self-declare their relationship status to each other (e.g., acquaintance, close friend, or spouse). Participants are then sorted into categories (e.g., strong or weak adversaries) based on their own relationship declarations. It is important to classify adversaries as weak or strong in order to gauge how knowledge of a potential victim can effect the security of an authentication system.

However, problems arise with this method of testing for the known adversary threat because users don't always have the most accurate reading of their social relationships. Another flaw with this method of testing for the known adversary threat is that it only considers certain social labels as strong adversaries (e.g., co-workers would be classified as weak acquaintances, unless indicated otherwise). Co-workers can have physical access to our devices, and due to their proximity or access to employee records, they might have elevated knowledge of a potential victim that can aid in bypassing security measures.

We propose a new process for testing the resilience of authentication systems to the known adversary. Our new process involves the utilization of questionnaires that paired participants must fill out at the beginning of a user study (in lieu of simply self-declaring their relationship status). We propose three models to achieve this: (i) the adapted-Relationship Closeness Inventory (RCI); (ii) the Known Adversary Inventory (KAI); and (iii) the Oneness Score. Each of these models has roots in the field of Social Psychology's methods to measure relationship closeness. We adapt these models to suit our need for a questionnaire that will assist researchers in quantifying the threat posed by the known adversary in a security/usability study.

We test the viability of our models for quantifying the known adversary through a case study. In our case study, participants were

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
NSPW'19, September 23–26, 2019, San Carlos, Costa Rica

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-7647-1...\$15.00
<https://doi.org/10.1145/nmmnnnn.nnnnnnn>

¹One can view the known adversary a more general form of unauthorized access by social insiders as it encompasses a wider variety of threat sources including ones that do not originate from socially close individuals. We also note that our known adversary definition is not necessarily covered insider threat commonly used in the field of business management; where it refers to the internal threat of employees to the organization [34].

recruited in pairs to test the resilience of a proposed authentication system to the known adversary threat. Instead of simply asking the participating pairs to self-declare their relationship status, we asked the participants to answer questionnaires composed of the adapted-RCI, KAI, and Oneness Score. Our results indicate that the adapted-RCI and KAI are far better quantifiers of the known adversary than simple relationship self-declaration. The Oneness Score failed to yield any substantive results in quantifying the known adversary.

The adapted-RCI, and the KAI, are valuable tools for assessing the resilience of an authentication system to the known adversary threat. However, this is a starting point for the development of a more rigorous and automated framework for detecting and quantifying the known adversary. While the questionnaires take less than 5 minutes to fill out, that is still considered a cognitive burden, and the validity of the answers provided in the questionnaires can be questioned due to social pressures and strategic adversaries.

Future work in this field should emphasize automatically detecting known adversaries through social network analysis, using metrics proven through our case study (e.g., frequent physical proximity, social media access etc.). Utilizing an automated tool for detection serves two purposes: (i) ensuring that the known adversary detection and measurement is not based on answers that may be influenced by social constraints (e.g., partners are pressured into answering that they are socially close), and (ii) decreasing the cognitive burden on participants in user studies to avoid fatigue. We envision utilizing social network analysis tools to automatically identify potential known adversaries, and identifying the extent to which a node in a social network is a potent known adversary. This can aid in the development of fine-grained privacy assistants that set targeted information sharing policies based on the potential for a security breach. For example, if a social acquaintance has been identified as a potentially potent known adversary, then we can stop the information flow towards that social acquaintance (limiting the viewer not the content). This is an interesting point of discussion because the extent to which we limit sharing can undermine the purpose of publicly sharing information on social media.

2 RELATED WORK

Authentication systems that have the potential to be compromised by first-hand or publicly available knowledge (e.g., from social media) of a potential victim are typically studied using paired adversary user studies or modelling [17, 21, 22, 22, 33]. In addition to autobiographical authentication systems, challenge questions are typically vulnerable to this type of attack and tested by a paired adversary user study [33]. In some research, participants are not always recruited in pairs but some models are used to simulate adversarial guessing [17, 24]. We briefly review of some of this research below.

Muslukhov *et al.* [29], investigates the prevalence of breaches by *insiders*, through the use of an online survey (n=724). They found that participants tended to be victims of insiders accessing their devices/accounts, with 12% of participants indicating that they were aware of an instance of unauthorized access by an insider, and 9% of participants admitting that they gained unauthorized access to a device/account. We propose an expansion of the definition of

the insider, to include adversaries that are not necessarily socially close to a potential victim, but are within physical proximity (e.g., co-workers). Hence, we define an insider as any individual with elevated first hand knowledge of a potential victim and/or elevated access to a potential victim's devices, that uses these privileges with malicious intent. We also refer to the insider threat as the known adversary threat, to avoid confusion with the common utilization of the term in the field of business management; where it refers to the internal threat of employees to the organization [34].

Hang *et al.* [22] studies the security of a location-based geographical authentication system by recruiting paired participants and testing against the known adversary threat model under three classes of self-declared adversaries: (i) socially close adversaries with no access to the internet; (ii) socially close adversaries with access to the internet; and (iii) stranger adversaries with access to the internet. The authentication system was resilient, and no adversaries were able to guess a single location. Out of 15 participants (some of which acted both as an adversary and as main participant), 4 relationships did not match (e.g., one participant described the pair as a good/best friend while the pair described the participant as a acquaintance/good friend).

Hang *et al.* [21] performed a user study on another authentication scheme that relies on autobiographical data (e.g., incoming/outgoing calls) and tested the systems' resilience to the known adversary threat. In the user study (n=11), the participants were asked to bring along two adversaries; one close adversary, and one acquainted adversary. In this user study no contradictions in the self-declared relationships were observed. Generally, the close adversaries were far better at guessing authentication credentials than acquainted adversaries.

Albayram *et al.* [3] also performed a user study on an authentication scheme that relies on autobiographical data from everyday activities. To test the authentication system against the known adversary threat model, the researchers recruited participants in pairs (n=12 pairs). The participants were asked to self-declare their relationship status on a five-point Likert scale. No discrepancies in the self-reported relationship characterizations were reported. However, participants were asked to bring along socially close individuals (e.g., spouse, or close friend), 4 pairs in the study were spouses/significant others and 8 pairs were close friends. The researchers modelled the weak (naive) adversary by randomly pairing. Generally, the strong adversaries performed much better at guessing authentication credentials than weak adversaries in this work.

The general trend in this related work is that there is a reliance on self-declaring relationship status. However, there is no framework to determine if the recruited participants were a good sample of known adversaries. Our proposed models, in this paper, can benefit future research with paired adversarial models. Using our models, one can identify relationship characteristics (e.g., constant physical presence) of pairs rather than relying on a broad relationship label. Researchers using our models can know more precisely what social factors result in a high adversarial capability.

It is important to note that our definition of the known adversary does not supersede more established variants of the definitions of the *insider threat* [10–12, 14]. Previous attempts at formally defining the insider threat have focused on the insider threat to organizations, our definition of the known adversary is more broad

with a focus on social insiders. Organizations typically have security policies and defined physical perimeters, therefore an insider can be defined based on the breach of a security policy or a defined physical perimeter. However, individual users do not have a security policy, therefore it is difficult to define the insider threat to individuals based on the breach of a security policy.

Brackney and Anderson [14] define an insider as any individual with privilege, access, or knowledge of information systems and services [14]. Bishop [11] takes a different approach to the definition of the insider, where the insider is a trusted entity with the power to violate a security policy [11]. The insider becomes a threat if the security policy is violated [11]. Bishop and Gates [12] later refined this definition to differentiate insiders based on their access. Insiders that have legitimate access but violate a security policy are categorized differently than insiders that do not have legitimate access but somehow obtain it.

3 MODELS FOR QUANTIFYING THE KNOWN ADVERSARY

To quantify whether a pair is indeed a known adversary, we must first be able to quantify the relationship between the adversary and the participant in security research. In the quest to quantify relationship closeness, we discuss several different relationship closeness indices/scales proposed or used in this paper, and detail the design, scoring and motivation for each of our models. Our models for quantifying the known adversary have roots in tools from Social Psychology for measuring closeness in social relationships [4, 5, 9, 15, 19].

3.1 Adapted-Relationship Closeness Inventory (RCI)

The Relationship Closeness Inventory (RCI) [9] quantifies the closeness of two people (see Appendix A). However, it is not specifically tailored towards security research and quantifying the known adversary. For example, the RCI does not contain any questions regarding pairs accessing each other's devices. The RCI was created during that period of time accounts/devices were not a large part of everyday life. This motivated us to adapt the RCI to reflect modern life with questions regarding social media and device/account access. Furthermore, we also identified frequent physical proximity to the pair as a security compromise to our authentication system. This drove us to adapt the RCI to ask several questions regarding physical proximity. Additionally, we recognized that locations can be shared through social media, so we included questions regarding location posts. We stress that our adapted-RCI and KAI models are tailored towards location based autobiographical questions.

Before adding any components, we begin by removing the frequency and diversity sub-scales of the RCI. The frequency sub-scale includes questions regarding the frequency of time spent alone with a partner/spouse (e.g., during the past week, what is the average amount of time per day that you spent alone with your spouse/partner in the morning?). The diversity sub-scale focuses on the breadth of activities completed together with a spouse/partner (e.g., during the past week, did you do laundry alone with your spouse/partner). We removed the frequency sub-scale and the diversity sub-scale because the RCI places a great emphasis on time spent

alone/activities completed alone with a partner/spouse rather than general physical proximity and activities experienced/completed with groups. We then include relevant questions that account for access to devices, access to social media accounts, and physical proximity to devices (see Appendix B). We refer to this modified version of RCI as the *adapted-RCI*. In this work, we study the correlation between the number of successful guesses by pairs in our case study, and the adapted-RCI score using a linear regression analysis. This helps our understanding of the effectiveness of the adapted-RCI for quantifying the known adversary.

The scoring scheme of the adapted-RCI is based on the answers given by a participant. The highest possible score in the adapted-RCI is 141 (higher scores should indicate a more potent or effective known adversary). By utilizing the adapted-RCI, we measure not only the relationship strength, but also proximity to a potential victim (e.g., co-worker or classmate). These two are important factors given the known adversary definition. The questions are weighed based on the answers, and the scoring system is designed based on an analysis of the relevance of the question. For example, do you and your pair follow each other on social media? That question is weighted at 10 points for "yes" and zero points for "no." This is supported by high tendency of users in sharing sensitive information on social media that can lead to a security compromise for accounts/devices [32]. After an in-depth analysis of our case study data, the adapted-RCI is the model that best quantifies the known adversary (see Section 4.1 for details).

3.2 Known Adversary Inventory (KAI)

The Known Adversary Inventory (KAI) is the collection of the questions that we added to the RCI (to create the adapted RCI), without any original components from the RCI. The KAI is a subset of the adapted-RCI. These questions reflect more modern measures of closeness, and encompass physical proximity, device access, and social media access (see Appendix C). The scoring scheme of each shared question in adapted-RCI and KAI is the same for consistency. The maximum score in the KAI is 92. We use KAI in our case study to evaluate whether or not our proposed modern social closeness measures alone are effective for quantifying the known adversary. After an in-depth analysis of our case study data, the KAI is the second-best model for quantifying the known adversary. The performance of the KAI and adapted-RCI are very close for quantifying the known adversary (see Section 4.2 for details).

3.3 Oneness Score

Another tool for the measurement of social closeness is the Oneness Score [15]. The Oneness Score is a compact relationship measurement tool composed of the *Inclusion of Other in the Self Scale* (IoS Scale) [4] and the *We Scale* [15].² The first component of the Oneness Score (i.e., the IoS Scale) is a simple pictorial tool that allows the participant to pick the image that most closely resembles the relationship to the pair (see Fig. 1). The We Scale component [15] presents a Likert Scale question with seven levels. The question asked is "Please, select the appropriate number below to indicate to what extent you would use the term "WE" to characterize you and

²By compactness, we mean that this model has considerably fewer questions compared to others.

this individual.” The lowest score is 1 = “not at all” and the highest one is 7 = “very much so”. This scale is a compact way to measure closeness. The average of the IoS Scale and the We Scale is known as the Oneness Score [15].

Gachter *et al.* [19] performed a study (n= 772) on Amazon’s Mechanical Turk with the aim of establishing a correlation between components of the Oneness Score, and several other relationship closeness scales/inventories. The result showed a high correlation between the IoS Scale (a component of the Oneness Score) and the RCI. This motivated us to investigate the Oneness Score as a tool to quantify the known adversary because of its compactness.

We attempt to establish a correlation between the number of successful guesses in our case study and the Oneness Score, to validate the use of the Oneness Score for quantifying the known adversary. Due to the compactness of Oneness, it is more favourable than others by reducing cognitive burden and duration of user studies (see Section 4.3 for details).




Figure 1: Remake of the original IoS Scale [4]

3.4 Further Discussion on Models

There are a wide variety of tools for measuring relationship closeness in Social Psychology [4–6, 9, 19]; however, we have selected and designed our models while considering their effectiveness for our goal and ethical concerns. For example, the Personal Acquaintance Measure (PAM) is a tool for measuring relationship closeness on a wide spectrum (from acquaintance to close friend/spouse). The PAM utilizes six dimensions to measure relationship closeness: (i) duration of acquaintance, (ii) knowledge of goals, (iii) frequency of interaction, (iv) social network familiarity, (v) level of self-disclosure, (vi) degree of physical intimacy. Each dimension contains three Likert scale questions. Results from the PAM are also correlated to the RCI and the Oneness Score [19]. Due to this strong correlation, and ethical concerns regarding the degree of physical intimacy dimension in the PAM, we chose the RCI and the Oneness Score for our study.

We also note that the accuracy of our models is sensitive to the choice of scoring. One interesting direction is to learn scores of each question in the model from data. This learning problem can be viewed as a linear regression problem where the goal is to find the appropriate scores (or weights) to minimize the fitting error. Of course, the validation of these weights would need some testing data. Due to our limited number of samples, we could not split up our data to perform this test.

4 RESULTS

Our known adversary measurement models are tested using GeoSQ authentication system [1, 2]. GeoSQ is an authentication system that logs autobiographical location data and queries the user about their previously visited locations. Our user study which spanned two sessions, was approved by our university’s Research Ethics Board. We controlled for risks such as a data leakage (which could have potentially damaging social consequences) by anonymizing the data and ensuring that it is transmitted and stored in a secure manner. In Session 1, participants (n=19 pairs) were asked to install the GeoSQ application, answer the questionnaires presented (adapted-RCI, KAI, and Oneness Score questionnaires), and self-declare their relationship status with their pair.³ In Session 2, participants (n=18 pairs) were asked to authenticate themselves by answering 10 location questions. Followed by that, pairs would attempt to guess each others’ location questions (both pairs performed all parts of the study, we did not have an adversary pair and a main participant). See [1, 2] for more details regarding the user study.

Our user study was designed without a main participant (i.e. victim) and an adversary participant so that we could analyze this two sided data. The symmetric assumption is a current shortfall for paired adversary studies in the literature when they don’t swap attacker/victim positions [3, 21, 22]. Related work utilizes labels, such as acquaintance or close friend, to classify the relationship of adversaries (as strong or weak), who test the strength of an authentication system. In this literature, it is not clear if the solicited labels from victims, attackers, or both take into account for classification of adversary relationship to the victim. We discuss and comment further on the asymmetry of social relationships and implications in Section 5.

We test the efficacy of the models for quantifying the known adversary by performing a linear regression analysis on the number of questions guessed correctly, and the scores of each of the models. We refer to the number of questions guessed correctly by pairs as a guessing score. We then compare the correlation results of our models to self-declared relationship status for further comparison and validation.

Participants in our user study performed both roles of the victim and the attacker. Upon conducting our analysis, we noted that many relationships were often asymmetric: significantly different relationship scores were obtained for a pair of individuals for a given known adversary measurement model (e.g., RCI-adapted, KAI, etc.). We then conducted the analysis from two different perspectives. Our first approach was to correlate the attacker’s relationship score (about the victim) with the guessing score that the attacker achieved. This analyses exhibits statistically significant results reported below. Our second approach was to correlate the potential victim’s relationship scores (about the attacker) with the attacker’s guessing scores. This second approach yielded insignificant results. Thus, we report on our first approach (i.e., correlating the attacker’s relationship score with the guessing score that the attacker achieved), and skip our insignificant results. Our significant results are corrected

³The order of various questionnaires (e.g., adapted-RCI, KAI, and Oneness Score questionnaires) was not randomized in our user studies. Due to this weakness of our study design, we can’t ensure that the order is not inadvertently affecting the way our participants answer the questions.

using a Bonferroni correction given all tests that we have run on our data including those yielded insignificant results and are not reported in this paper (i.e., the total of 9 tests).

4.1 Adapted-Relationship Closeness Inventory Results

We performed a linear regression analysis on the guessing score of each participant and the corresponding adapted-RCI score. The adapted-RCI score was the predictor (x) and the guessing score was (y). We attained a correlation of $r=0.772$ ($p=3.47 \times 10^{-8}$) which is significant at $\alpha = 0.005$ after multi-test correction. Figure 2 shows a scatter plot illustrating the linear relationship between the guessing score and the adapted-RCI score hence the positive correlation of $r=0.772$ attained. Figure 3 features two box-plots that display the quartiles and the mean of the GeoSQ guessing scores, and the adapted-RCI scores. The average adapted RCI score according to our box-plots is 60/141, while the average guessing score is roughly 4/10. The adapted-RCI box-plot shows us that we have a pseudo-normal distribution.

Furthermore, Figure 4 shows a density plot for the adapted-RCI scores and the GeoSQ guessing scores. We use the density plots in order to test for normality and the validity of the analysis.

4.2 Known Adversary Inventory (KAI) Results

For the KAI, we obtained a high correlation value of $r = 0.728$ ($p = 4.703 \times 10^{-8}$) which is significant at $\alpha = 0.005$ after multi-test correction. The KAI is a subset of the adapted-RCI (without the questions from the original RCI). The removal of the questions from the original RCI (which are included in the adapted-RCI but not the KAI) did not have a major difference on the correlation between relationship scores and guessing scores (KAI: $r=0.728$, adapted-RCI: $r=0.772$). Figure 5 shows the scatter plot of the guessing scores against the results of the KAI. Figure 6 features two box-plots that display the quartiles and the mean of the guessing score and the KAI score. While Figure 7 showcases the density plot, which shows a similar distribution to the adapted-RCI scores. These numbers suggest that even without the RCI component, the questions we added




Figure 2: Scatter plot of guessing scores/adapted-RCI scores.

to the adapted-RCI (i.e., the KAI model) are also good predictors for a participant’s guessing ability.

4.3 Oneness Score Results

We performed the same analysis for the Oneness Score model. The correlation coefficient was weak at $r=0.096$ ($p=0.563$). Figure 8 shows the scatter plot of the Oneness scores and the GeoSQ guessing scores, and shows no clear correlation between the guessing scores and the Oneness Scores.

4.4 Self-Declared Relationships

To show that utilizing our models for measuring the known adversary is indeed a more accurate measure than study participants self-declaring their relationship status, we perform the same analysis on the self-reported relationships and GeoSQ guessing scores. We obtained a negative correlation coefficient $r = -0.068$ ($p= 0.692$), which is not significant.

The self-reported relationship characterization was split across four levels that participants could choose from. The first level was no relationship, the second level was not very close, the third level




Figure 3: Box-plots for the adapted-RCI and GeoSQ guessing scores.

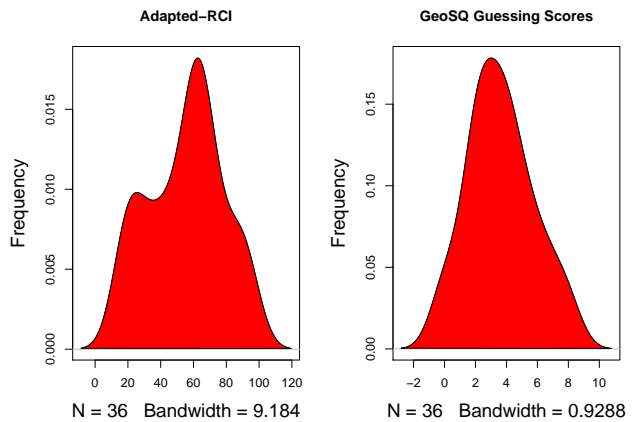


Figure 4: Density plot for the adapted-RCI and GeoSQ guessing scores.




Figure 5: Scatter plot of guessing scores/KAI scores.

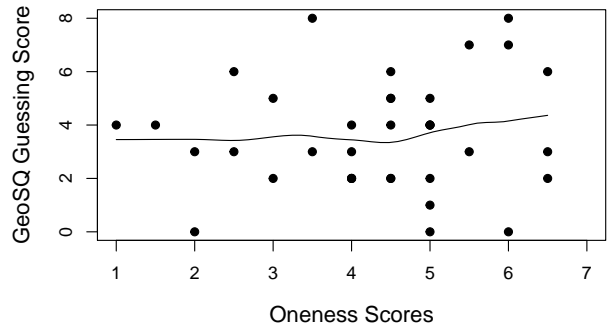


Figure 8: Scatter plot of Oneness and GeoSQ guessing scores.

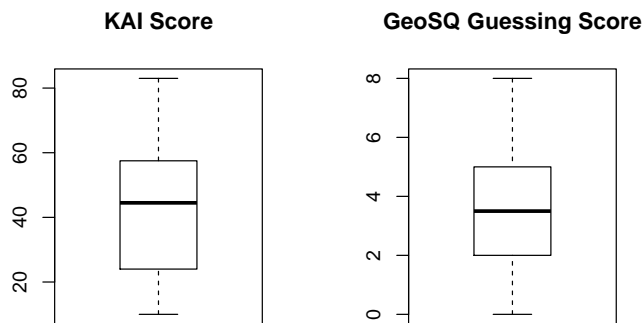


Figure 6: Box-plots for the KAI and GeoSQ guessing scores.

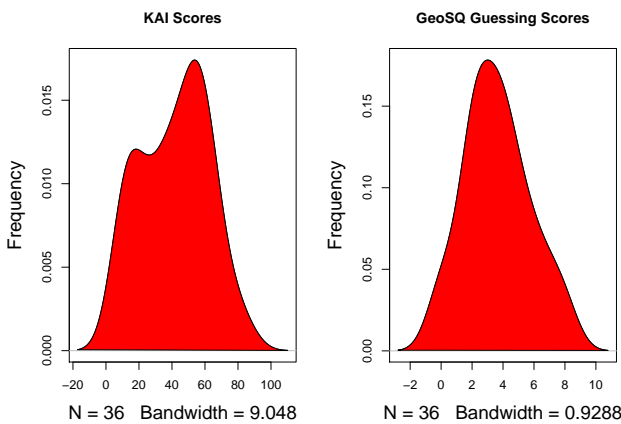


Figure 7: Density plot for the KAI and GeoSQ guessing scores.

was close, and the fourth level was very close. Of the 18 pairs that successfully completed all parts of Session 1 and Session 2, there were 7 instances in which the participant pairs self-reported their

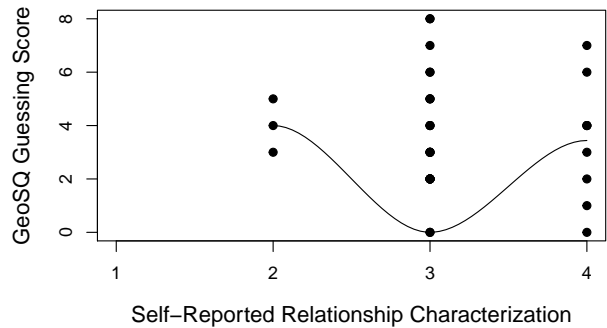


Figure 9: Scatter plot of guessing scores and self-reported relationships. 7 pairs labelled their relationship differently.

relationship characterization differently. All the differences were one level apart, meaning that it was typically a participant who picked a similar but different value from their pair.

Figure 9 shows the relationship of the self-reported relationships with the GeoSQ scores.

5 DISCUSSION AND FUTURE WORK

Utilizing our frameworks provides researchers with a better measure of the potency of a known adversary. Utilizing a label can lead to inaccuracies in the classification of the known adversaries. For example, a weak adversary may be labelled as a strong adversary (based on a social label). This can have a significant effect on the conclusion the researchers will reach when testing an authentication system for resilience to the known adversary. The most adverse consequence is deploying an authentication system where the testing methodology is not sound because it can lead to various security breaches.

While many tools exist for measuring relationship closeness we utilized the RCI due to its focus on physical presence between

pairs. This physical presence focus was critical for our use case of studying the known adversary for autobiographical authentication system. The other social measurement scales that we considered (e.g., PAM) do not focus on the frequency of physical presence with the pair, so we decided against using them in our study. This is in addition to our concern with the physical intimacy dimensions of the PAM [6], which could cause discomfort to the participants.

The adapted-RCI and the KAI were the two best models in terms of quantifying the known adversary, while the Oneness Score and self-declared relationships proved to be highly ineffective in quantifying the potency of the known adversary. We reiterate that we are measuring the potency of the known adversary by the number of successful guesses. The adapted-RCI and the KAI are highly similar in terms of model scores correlating to higher guessing score. Since the KAI is a subset of the adapted-RCI, we perform an analysis on the components of the adapted-RCI that were extracted from the original RCI. The result shows a weak correlation ($r=0.25$, $p=0.127$), thus suggesting that the components of the original RCI add very little value in terms of accuracy of quantifying the known adversary threat.

The KAI model (and Adapted-RCI) rely on three factors: (i) physical proximity, (ii) device access, (iii) social media access (through following or friendship on different social media mediums) to a great extent. The aforementioned factors seem to be the best predictor of the potency of the known adversary in guessing authentication credentials. GeoSQ (the application which we were testing) relies on autobiographical location data to authenticate a user, and autobiographical location data is highly vulnerable to the known adversary threat [3, 17]. The adapted-RCI and the KAI models require further verification by testing on different types of authentication systems. Our case study only suggests that the models work for autobiographical based authentication systems. We predict that most types of fallback authentication (e.g., security questions, email resets, and SMS resets) are vulnerable to physical proximity, device access, and social media access. The adapted-RCI and the KAI should be altered to match the factors that compromise an authentication system. In our case with location based autobiographical authentication, physical proximity plays a large role, while in other authentication systems (e.g., email resets or SMS resets), device access even for a brief period of time.

Furthermore, the Oneness Score is a compact tool to measure relationship closeness that also has a high correlation with the RCI [19]. We tested it because of its compactness (composed of only two questions). This compactness is an important factor for reducing a participant's cognitive burden in user studies. We reiterate that the adapted-RCI and the KAI are not currently generalizable to other authentication systems when it comes to measuring the potency of a known adversary; the relationship closeness tools must be adapted in order to suit the authentication system under investigation.⁴ However, the general framework of adapting tools for measuring

closeness for measuring known adversary threats in a user study is our recommendation.

For other authentication systems, one can develop tailored scales to identify known adversary. We propose a plan for creating similar scales following psychometric practices: (i) identify the most pertinent factors to an authentication system's security. For example, SMS resets are especially vulnerable to a known adversary with fettered or unfettered access to a potential victim's device; (ii) identify the related social traits that may have a negative impact on an authentication system based on the pertinent factors identified. For example, working closely together in an employment setting could lead to unauthorized device access hence putting the security of SMS resets at risk; (iii) construct the scale that is composed of questions that can query participants about the prevalence of social traits that have been identified as a threat to the security of an authentication system. At this stage, security researchers can search for scales from various fields that query users about a specific social trait. The adaption process can begin from that point. It is important to note that external factors can greatly effect the results obtained (e.g., the order that questions are presented in or social/peer pressure). Researchers must normalize for such factors whenever possible; and (iv) investigate the optimal weights for each question in the scale.

Labelling social relationships is limiting in paired user studies with an adversary-guessing component. One cannot extract definitive relationship characteristics from a label, and labels can be very broad and subject to interpretation. For example, an adversary with the label of acquaintance is very subjective for both researchers, and users, and does not imply any informative reason why an acquaintance should or should not be able to guess an authentication credential. Furthermore, it is difficult to find a list of labels that has full coverage of all the possible known adversaries with various potential. Therefore, the strongest known adversaries might not be identified because they do not fall into a pre-specified set of labels. Finding alternate means of measuring relationship closeness is important to help gauge how potent of a known adversary a paired participant will be, by identifying the factors that will affect the security of a new authentication system (e.g., in our case, frequency of physical presence was an important factor).

Furthermore, self-declared pair relationships can have discrepancies. In our user study, 7 pairs did not label the relationship in the same way, and 10 participant pair's had relationship closeness scores that were more than 30 points apart. The occurrence of discrepancies can cause consistency issues especially in authentication user studies where the pair's relationship is thought of as a symmetric relationship. For example, if an individual declares that their pair is an acquaintance while the pair states that they are friend, one cannot classify the pair as a weak or strong adversary if symmetric assumptions for relationships are in place.

One interesting phenomenon that arose as a result of the lack of symmetry in relationship scores and relationship labels (pairs often don't see the relationship in the exact same way) was that the correlation between relationship scores and GeoSQ guessing scores changed depending on the point of view of the analysis. In our analysis (which yielded a high correlation, see Section 4) between the adapted-RCI scores/KAI scores and the GeoSQ guessing scores we uniformly utilized the relationship score of the adversary and

⁴Adapting our models to other forms of authentication systems might be an outstanding problem depending on operational knowledge of the practitioner. We believe also, there is a need general framework to build on. This will aid in the assignment of points and the normalization of the process so authentication systems can be compared even when utilizing a different type of underlying data (e.g., autobiographical location data or text data for passwords).

correlated that score to the GeoSQ guessing score. To investigate further, we changed the point of view of the analysis and used the relationship score given by the potential victim and correlated it to the guessing score achieved by the attacker. In that analysis the correlation was very weak between relationship scores and GeoSQ guessing scores. This phenomenon suggests that pairs not only view their relationships differently, but also that pairs might have varying degrees of knowledge regarding each other's activities.

Related work [3, 21, 22] conducted analysis on one-sided data (i.e., the participants were either potential victims or adversaries). In our study, the participants switch roles which means we have relationship closeness scores from our models and guessing scores from all our participants. Through our analyses (see Section 4), we found that the attacker's closeness scores were highly correlated to the guessing score of the attacker, rather than those of victims. This indicates that we should utilize an attacker's relationship closeness scores because the adversary is ultimately the best judge of the relationship.

Our approach of utilizing models in the form of questionnaires to quantify the known adversary might be limited because we are relying on input from a user study participant and assuming that this information is robust and reliable. However, we have no way of verifying whether or not this information is true. Study participants can easily manipulate the results by inputting incorrect answers to questions, and study participants might be inclined to answer a certain way due to social pressure. Hence, a more robust method of detecting the known adversary is necessary.

In our analysis, we assumed that all of our adversaries were non-strategic. Assuming that an adversary is non-strategic might be ecologically valid in cases like our user study where the participants were not told that they would be guessing each other's location questions before answering relationship closeness questions. However, when the goal is to automate the detection of known adversary in practice (e.g., privacy assistants discussed below), we cannot rely on the adversary to be truthful. This is an issue because our results indicate that the potential victim's scores are not predictive of whether or not the adversary will be highly successful in guessing authentication credentials or not.

Peer/social pressure to answer a certain way may have affected the results we obtained. Related work on the effects of peer pressure [13] has previously found a correlation in behaviour because of peer pressure. This is also applicable in our user study. To decrease the effects of social/peer pressure, we took the following precautions: (i) we assured participants several times throughout the user study that the relationship closeness information will never be disclosed, (ii) During the user study, participants were not allowed to sit next to each other while answering questions about relationship closeness, (iii) to avoid peer pressure, we also assured the participants that all the data was anonymized both verbally and in a consent form that they were required to read. Despite all these measures, peer pressure could have been a factor in the way participants answered relationship closeness questions; however, we have no way of measuring the extent to which peer or social pressure affected the way participants answered questions.

Strategic adversaries can also have adverse effects on user study results. Strategic adversaries can affect the results by purposefully

answering relationship questions or labelling the relationship incorrectly. This has the potential to skew the security results of an authentications system significantly. The effect of strategic adversaries can be minimized in user studies by utilizing the relationship scores from the potential victim. Our results indicate that utilizing the attacker's relationship scores is a more accurate indication than relationship scores from the potential victim. Future work should address how to limit strategic adversaries from effecting the results, especially in cases where both the potential victim and the adversary are being strategic (untruthful).

Thus, for strategic environments and applications, one should consider redesigning our models so potential victims' scores are more predictive. How that could be achieved is still an open question. Automatically detecting a potentially strong known adversary based on some factors (e.g., local social network structure, social activity patterns, etc.) is a possible direction. Another alternative is to design scoring systems, which help potential victims assess the knowledge of potential attackers about themselves, rather than them assessing their own knowledge of the potential attackers.

Our work is a stepping stone for a more developed automated framework for the detection of the known adversary. Our case study has provided valuable data regarding what are the important metrics when it comes to identifying the known adversary. The important metrics that we have identified, can potentially guide us in the process of mining publicly available data (e.g., social media data, forum data, etc.). This data combined with artificial intelligence techniques can be used for analysis and eventual automated detection of the known adversary.

A promising direction for the detection of the known adversary is social network analysis. Social network data can be extracted from public sources (e.g., social media, forums, public domain data, etc.) and analyzed using concepts in network science [8, 30] and graph theory such as: centrality metrics, degree distribution, ego networks, link prediction, dispersion, and node proximity [7, 18, 23, 25–27]. Link prediction is the process of inferring observed or hidden interactions based on social network structures [26]. One of the most common approaches utilized for link prediction is analyzing the proximity of nodes in a network [26]. Link prediction can be useful in discovering connections that are not observed, and have a potential to be known adversary.

Furthermore, dispersion might be an important metric for automated detection of known adversary. Dispersion is the measure of mutual friend connectivity [7]. Romantic relationships have been successfully identified in the past using the dispersion metric [7]. That means if two individuals' mutual friends are sparsely connected, those two individuals have high dispersion, and more likely have been involved in romantic relationships. At the other end, low dispersion indicates that two individuals' mutual friends are well-connected. Dispersion might be effective for automatically detecting potential known adversaries as in one end romantic partners have high dispersion, and in another end, co-workers usually have low dispersion (i.e., mutual friends are well connected as they are possibly co-worker of each others).

Identifying common traits and patterns on social networks using social media has become possible using data shared on social media. Homophily is the tendency of individuals with similar traits to be friends [28]. This phenomenon could be used for detecting known

adversaries through publicly available data and matching between individuals in a social network based on shared traits. Research on recommender systems has proven a correlation between the individuals we interact with, and their preferences (e.g., product preferences, political party preferences) [31]. We can utilize the same metrics to discover and quantify relationships in social networks, for detecting a potential known adversary. This approach has its own challenges: identifying traits and common patterns on social media could be difficult because this data may not be publicly available from the individuals in question.

We can also view the detection of a potential known adversary as a node labelling problem (or node classification problem) [35] in a social network; where we are trying to label potential known adversary nodes, for a given target node. We can utilize social network metrics such as dispersion, proximity of nodes, degrees of homophily as features for this node classification problem. The challenge is finding the correct metrics and being able to test it based on some labeled data with ground truth. Extending this approach, we can detect to what extent a node in a social network is a potential known adversary. In other words, based on what this node in your social network knows about you, here is how potent of a potential known adversary he/she would be.

An automated method for detecting known adversaries could be more robust than current methods of validation against the known adversary threat. When we have metrics that we can extract automatically from social networks, and determine to what extent an individual would be a potent known adversary, then we can rely less on a participants' answers. At least, the automated detection of potential known adversaries can help more effective recruiting procedures for user studies. As one can recruit those pairs of individuals for a user study such that in each pair, at least one can be a strong potential attacker to another one. Detecting potential known adversaries can also assist in the development of privacy assistants [16], by providing privacy setting recommendations based on the extent to which a person in a social network is a potential known adversary.⁵ This extension could potentially lead to a more fine-grained privacy assistant for detecting the potency of a potential known adversary and restricting the flow of information to that user via social mediums such as Twitter, Facebook, Instagram, or Snapchat.

We envision the automated detection of potential known adversaries serves privacy assistants to set or recommend more restrictive information sharing policies for potential known adversaries. Based on the recommendation of the privacy assistant, the user can choose whether or not to limit the flow of information to a particular individual. Our proposed privacy assistant can be utilized by everyday users to enhance their privacy. This tool would be powered by social network analysis metrics and can also help known adversary testing by providing the researchers with relevant information regarding the potency of an adversarial pair (e.g., to what extent is this adversary a strong known adversary?). This is an approach that can rid researchers of having to account for

social pressures or strategic adversaries (to some extent). Instead of asking the participants for this information they can now mine the information from publicly available sources.

The creation of this automated tool for detection of known adversaries faces many challenges. The proposed automated tool would use social network data, required to be mined from social media platforms. The most pressing challenge is collecting this data in light of the download or privacy restrictions many social media platforms would impose on a third party tool. For example, one would have to collect information from the point of view of the potential victim, and the point of view of the potential known adversary/attacker. Depending on the privacy settings of the potential known adversary/attacker on a social media platform, this may not be possible. If we utilize one sided data (e.g., data from the point of view of the potential victim), we may face some relationship asymmetry issues discussed earlier.

Additionally, identifiability of social media is another aspect to discuss [20]. This could be useful for identifying adversaries that are not necessarily part of our social network. The availability of data to the public on social media platforms with identifiable traits (e.g., usernames, profile names, profile pictures, and locations) can make certain authentication systems highly vulnerable. This is an interesting avenue for discussion and future work.

6 CONCLUSION

We utilized two models inspired from the field of Social Psychology (the adapted-RCI, and the Oneness Score) [4, 9, 15], and created a model (the KAI) to attempt to quantify the known adversary empirically without relying on relationship labels. After testing our proposed models using a case study that involved testing an authentication system (GeoSQ) using paired participants [2], we concluded that social closeness as declared by participants is not the defining factor in terms of quantifying the known adversary threat that an individual poses.

Our two most accurate models for quantifying the known adversary according to our case study are the KAI and adapted-RCI. However, several limitations such as relying on answers that might not be truthful due to social pressures from the participants still remain. We propose an extension to this work that involves the automatic detection of the known adversary based on social network analysis [7, 25–27, 31]. Known adversary detection is in its essence a node labelling problem, where we have a social network and based on proven networks we can determine to what extent a node (i.e., an individual) in this network is a potent known adversary. We can also utilize link prediction [26] to discover hidden or non-obvious relationships (e.g., relationships that are kept secret for a certain purpose). This extension has many implications, including creating a more fine grained privacy assistant, which adjusts information sharing on various mediums based on the potential threat (e.g., not sharing location information with co-workers).

7 ACKNOWLEDGMENT

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), funding reference numbers 402500-2013 and RGPIN-2018-05903.

⁵Privacy assistants would help the user identify information sharing risks in their networks. Of course, it would be up to the user to decide which type of information flows to which individuals, even if the privacy assistant declares several individuals with a high risk. Note that the privacy assistant would restrict some viewers' access to content, not the content itself.

REFERENCES

- [1] Alaadin Addas, Amirali Salehi-Abari, and Julie Thorpe. 2019. Geographical Security Questions for Fallback Authentication. In *Proceedings of the 17th Annual Conference on Privacy, Security, and Trust (PST'19)*. 1–7.
- [2] Alaadin Addas, Julie Thorpe, and Amirali Salehi-Abari. 2019. Geographical Security Questions for Fallback Authentication. *CoRR arXiv:1907.00998v1* (2019), 1–18.
- [3] Yusuf Albayram and Mohammad Maifi Hasan Khan. 2016. Evaluating Smartphone-Based Dynamic Security Questions for Fallback Authentication: A Field Study. *Human-Centric Computing and Information Sciences* 6 (2016), 16. Issue 1.
- [4] Arthur Aron, Elaine N Aron, and Danny Smollan. 1992. Inclusion of Other in the Self Scale and the Structure of Interpersonal Closeness. *Journal of Personality and Social Psychology* 63, 4 (1992), 596–612.
- [5] Arthur Aron, Elaine N Aron, Michael Tudor, and Greg Nelson. 1991. Close Relationships as Including Other in the Self. *Journal of Personality and Social Psychology* 60 (1991), 241–253. Issue 2.
- [6] Katherine B Starzyk, Ronald Holden, Leandre Fabrigar, and Tara Macdonald. 2006. The Personal Acquaintance Measure: A Tool for Appraising One's Acquaintance With Any Person. *Journal of Personality and Social Psychology* 90 (2006), 833–47.
- [7] Lars Backstrom and Jon Kleinberg. 2014. Romantic Partnerships and the Dispersion of Social Ties: A Network Analysis of Relationship Status on Facebook. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computer (CSCW'14)*. 831–841.
- [8] Albert-László Barabási et al. 2016. *Network Science*. Cambridge University Press.
- [9] Ellen Berscheid, Mark Snyder, and Allen M Omoto. 1989. The relationship Closeness Inventory: Assessing the Closeness of Interpersonal Relationships. *Journal of Personality and Social Psychology* 57 (1989), 792–807. Issue 5.
- [10] Matt Bishop. 2005. The Insider Problem Revisited. In *Proceedings of the 2005 Workshop on New security Paradigms (NSPW'05)*. 75–76.
- [11] Matt Bishop. 2005. Position: Insider is Relative. In *Proceedings of the 2005 Workshop on New security Paradigms (NSPW'05)*. 77–78.
- [12] Matt Bishop and Carrie Gates. 2008. Defining the Insider Threat. In *Proceedings of the 4th Annual Workshop on Cyber Security and Information Intelligence Research: Developing Strategies to Meet the Cyber Security and Information Intelligence Challenges Ahead (CSIRW'08)*.
- [13] Rainer Böhme and Stefanie Pötzsch. 2011. Collective Exposure: Peer Effects in Voluntary Disclosure of Personal Data. In *Proceedings of the International Conference on Financial Cryptography and Data Security (FC'11)*. 1–15.
- [14] Richard C Brackney and Robert H Anderson. 2004. *Understanding the Insider Threat*. *Proceedings of a March 2004 Workshop*. Technical Report. RAND CORP, Santa Monica, California, USA.
- [15] Robert B Cialdini, Stephanie L Brown, Brian P Lewis, Carol Luce, and Steven L Neuberg. 1997. Reinterpreting the Empathy-Aaltruism Relationship: When One Into One Equals Oneness. *Journal of Personality and Social Psychology* 73 (1997), 481–494. Issue 3.
- [16] Anupam Das, Martin Degling, Daniel Smullen, and Normman Sadeh. 2018. Personalized Privacy Assistants for the Internet of Things: Providing Users with Notice and Choice. *IEEE Pervasive Computing* 17 (2018), 35–46. Issue 3.
- [17] Sauvik Das, Eiji Hayashi, and Jason I Hong. 2013. Exploring Capturable Everyday Memory for Autobiographical Authentication. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'13)*. 211–220.
- [18] Sergey N Dorogovtsev, Jos FF Mendes, and Alexander N Samukhin. 2001. Size-Dependent Degree Distribution of A Scale-Free Growing Network. *Physical Review E* 63 (2001), 062101. Issue 6.
- [19] Simon Gächter, Chris Starmer, and Fabio Tufano. 2015. Measuring the Closeness of Relationships: a Comprehensive Evaluation of the Inclusion of the Other in the Self Scale. *Public Library of Science* 10 (2015). Issue 6.
- [20] Ralph Gross and Alessandro Acquisti. 2005. Information Revelation and Privacy in Online Social Networks. In *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society (WPES'05)*. 71–80.
- [21] Alina Hang, Alexander De Luca, and Heinrich Hussmann. 2015. I Know What You Did Last Week! Do You?: Dynamic Security Questions for Fallback Authentication On Smartphones. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI'15)*. 1383–1392.
- [22] Alina Hang, Alexander De Luca, Matthew Smith, Michael Richter, and Heinrich Hussmann. 2015. Where Have You Been? Using Location-based Security Questions for Fallback Authentication. In *Proceedings of the 11th Symposium On Usable Privacy and Security (SOUPS'15)*. 169–183.
- [23] Salman Jamali and Huzefa Rangwala. 2009. Digging Digg: Comment Mining, Popularity Prediction, and Social Network Analysis. In *Proceedings of the 2009 International Conference on Web Information Systems and Mining (WSIM'09)*. 32–38.
- [24] Patrick Gage Kelley, Saranga Komanduri, Michelle L Mazurek, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, and Julio Lopez. 2012. Guess Again (and Again and Again): Measuring Password Strength by Simulating Password-Cracking Algorithms. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy (IEEE S&P'12)*. 523–537.
- [25] Longbo Kong, Zhi Liu, and Yan Huang. 2014. Spot: Locating Social Media Users Based On Social Network Context. *VLDB Endowment* 7 (2014), 1681–1684. Issue 13.
- [26] David Liben-Nowell and Jon Kleinberg. 2007. The Link-Prediction Problem for Social Networks. *Journal of the American Society for Information Science and Technology* 58 (2007), 1019–1031. Issue 7.
- [27] Julian McAuley and Jure Leskovec. 2014. Discovering Social Circles in Ego Networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 8 (2014), 4. Issue 1.
- [28] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology* 27 (2001), 415–444. Issue 1.
- [29] Ildar Muslukhov, Yazan Boshmaf, Cynthia Kuo, Jonathan Lester, and Konstantin Beznosov. 2013. Know Your Enemy: The Risk of Unauthorized Access in Smartphones by Insiders. In *Proceedings of the 15th international Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI'15)*. 271–280.
- [30] Mark Newman. 2010. *Networks: An Introduction*. Oxford University Press.
- [31] Amirali Salehi-Abari and Craig Boutilier. 2015. Preference-oriented Social Networks: Group Recommendation and Inference. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys'15)*. 35–42.
- [32] David Tayouri. 2015. The Human Factor in the Social Media Security—Combining Education and Technology to Reduce Social Engineering Risks and Damages. *Procedia Manufacturing* 3 (2015), 1096–1100.
- [33] Abrar Ullah, Hannan Xiao, Trevor Barker, and Mariana Lilley. 2014. Evaluating Security and Usability of Profile Based Challenge Questions Authentication in Online Examinations. *Journal of Internet Services and Applications* 5, 1 (2014), 2.
- [34] Merrill Warkentin and Robert Willison. 2009. Behavioral and Policy Issues in Information Systems Security: The Insider Threat. *European Journal of Information Systems* 18 (2009), 101–105. Issue 2.
- [35] Elena Zheleva and Lise Getoor. 2009. To Join or Not to Join: the Illusion of Privacy in Social Networks with Mixed Public and Private User Profiles. In *Proceedings of the 18th international conference on World Wide Web (WWW'09)*. 531–540.

A ORIGINAL RCI

The original Relationship Closeness Inventory (RCI) by Berscheid *et al.* [9], is split into three sub scales: (i) frequency, (ii) diversity, and (iii) strength. The frequency sub scale measures the frequency of activities being performed together by the pair. The diversity sub scale measures the diverse breadth of activities that pairs undertake together. The strength sub scale measures the amount of influence an individual has over his/her pair. The sum of all three sub scales make up the RCI score which quantifies how close two people are to each other. While the adapted-RCI is mostly inspired from the original RCI, we only directly utilize a portion of the original RCI from the strength sub-scale, in the adapted-RCI (shown below verbatim).

The following questions concern the amount of influence X has on your thoughts, feelings, and behaviour. Using the 7-point scale below, please indicate the extent to which you agree or disagree by writing the appropriate number in the space corresponding to each item.

	1	2	3	4	5	6	7
1-My vacation plans.	○	○	○	○	○	○	○
2-My plans to have children.	○	○	○	○	○	○	○
3-My plans to make <u>major</u> investments (house, car, etc...).	○	○	○	○	○	○	○
4-My plans to join a <u>club</u> , social organization, church, etc...	○	○	○	○	○	○	○
5-My school-related plans.	○	○	○	○	○	○	○
6-My plans for achieving a particular financial standard for living	○	○	○	○	○	○	○

B ADAPTED-RCI

The adapted-RCI is presented below. Questions extracted from the original RCI are clearly marked with an asterisk. Each question is paired with possible answers and corresponding scores. While question 1 and question 6 might appear to be identical questions, question 1 refers to time spent sleeping in proximity (indicating a closer relationship or being roommates). Question 6 specifies being awake together because that signifies frequency of activity, this is a subtlety adapted from the original RCI.

- (1) What is the approximate number of hours in physical proximity to your pair per week? This refers to the total number of hours you and your pair are in the same house, in the same room, same cafeteria, or same work-space etc...

Scoring Range:

Answer	0-1	2-7	8-10	10-15	16-20	21-25	26-30	31-40	41-45	45+
Score	1	2	3	4	5	6	7	8	9	10

- (2) Approximately how many meals do you share with your pair per week?

Scoring Range:

Answer	0	1-3	4-7	8-11	12-15	16-19	21-24	24+
Score	0	1	2	3	4	5	6	7

- (3) Does your pair have access do any of your social media accounts? This could mean a social media account password.

Scoring:

Yes = 10 points

No = 0 points

- (4) Does your pair have restricted physical access to your devices/accounts? (i.e. , does your pair have access to your locked device?)

Scoring:

Yes = 10 points

No = 0 points

- (5) Does your pair have unrestricted access to your devices/accounts? (i.e. , do they know one or more of your passwords?)

Scoring:

Yes = 15 points

No = 0 points

- (6) In the past week, approximately how many hours have you spent with your pair awake? (i.e. , at home together, in the same work space, or in the same social gathering).

Scoring Range:

Answer	0	1-10	11-20	21-30	31-40	41-50	51-60	61-70	70+
Score	0	1	2	3	4	5	6	7	8

- (7) Does your pair affect your vacation plans? *

Likert Scale question with seven levels. The score is the level chosen (i.e., 1 = pair has no power, 7 = pair has significant power to affect your vacation plans).

- (8) Does your pair have any affect/influence on your marriage plans?*

Likert Scale question with seven levels. The score is the level chosen (i.e., 1 = pair has no power, 7 = pair has significant power to affect your marriage plans).

- (9) Does your pair have any affect/influence on your plans to have children? *

Likert Scale question with seven levels. The score is the level chosen (i.e., 1 = pair has no power, 7 = pair has significant power to affect your plans to have children).

- (10) Does your pair have any affect/influence on your plans to make major investments? *

Likert Scale question with seven levels. The score is the level chosen (i.e., 1 = pair has no power, 7 = pair has significant power to affect your major investment plans).

- (11) Does your pair have any affect/influence on your plans to join a club/social organization? *

Likert Scale question with seven levels. The score is the level chosen (i.e., 1 = pair has no power, 7 = pair has significant power to affect your plans to join a club/social organization).

- (12) Does your pair have any affect/influence on your school related plans? *

Likert Scale question with seven levels. The score is the level chosen (i.e., 1 = pair has no power, 7 = pair has significant power to affect your school related plans).

- (13) Does your pair have any affect/influence on your financial standing/wellbeing? *
- Likert Scale question with seven levels. The score is the level chosen (i.e., 1 = pair has no power, 7 = pair has significant power to affect your financial standing/wellbeing).

- (14) Do you and your pair follow each other on social media?

Scoring:

Yes = 10 points

No = 0 points

If the answer to question 14 was no, questions 15-17 were not asked.

- (15) On how many different mediums do you follow your pair on social media?

Scoring Range:

Answer	0	1	2	3	4	5	6	7	8+
Score	0	1	2	3	4	5	6	7	8

- (16) How often do you see your pair's social media posts per week?

Scoring Range:

Answer	0	1	2	3	4	5	6	7+
Score	0	1	2	3	4	5	6	7

- (17) How many times per week do see your pair's social media posts relating to his/her location? (This could be a check in , a Snapchat post indicative of location etc...).

Scoring Range:

Answer	0	1	2	3	4	5	6	7+
Score	0	1	2	3	4	5	6	7

Adapted-RCI maximum score: 141

C KAI

The KAI is presented below. Each question is paired with possible answers and corresponding scores.

- (1) What is the approximate number of hours in physical proximity to your pair per week? This refers to the total number of hours you and your pair are in the same house, in the same room, same cafeteria, or same work-space etc...

Scoring Range:

Answer	0-1	2-7	8-10	10-15	16-20	21-25	26-30	31-40	41-45	45+
Score	1	2	3	4	5	6	7	8	9	10

- (2) Approximately how many meals do you share with your pair per week?

Scoring Range:

Answer	0	1-3	4-7	8-11	12-15	16-19	21-24	24+
Score	0	1	2	3	4	5	6	7

- (3) Does your pair have access do any of your social media accounts? This could mean a social media account password.

Scoring:

Yes = 10 points

No = 0 points

- (4) Does your pair have restricted physical access to your devices/accounts? (i.e. , does your pair have access to your locked device?)

Scoring:

Yes = 10 points

No = 0 points

- (5) Does your pair have unrestricted access to your devices/accounts? (i.e. , do they know one or more of your passwords?)

Scoring:

Yes = 15 points

No = 0 points

- (6) In the past week, approximately how many hours have you spent with your pair awake? (i.e. , at home together, in the same work space, or in the same social gathering).

Scoring Range:

Answer	0	1-10	11-20	21-30	31-40	41-50	51-60	61-70	70+
Score	0	1	2	3	4	5	6	7	8

- (7) Do you and your pair follow each other on social media?

Scoring:

Yes = 10 points

No = 0 points

If the answer to question 7 was no, the following questions were not asked.

- (8) On how many different mediums do you follow your pair on social media?

Scoring Range:

Answer	0	1	2	3	4	5	6	7	8+
Score	0	1	2	3	4	5	6	7	8

- (9) How often do you see your pair's social media posts per week?

Scoring Range:

Answer	0	1	2	3	4	5	6	7+
Score	0	1	2	3	4	5	6	7

- (10) How many times per week do see your pair's social media posts relating to his/her location? (This could be a check in , a Snapchat post indicative of location etc...).

Scoring Range:

Answer	0	1	2	3	4	5	6	7+
Score	0	1	2	3	4	5	6	7

KAI maximum score: 92