

Users Can Deduce Sensitive Locations Protected by Privacy Zones on Fitness Tracking Apps

Jaron Mink
University of Illinois at
Urbana-Champaign
jaronmm2@illinois.edu

Amanda Rose Yuile
University of Illinois at
Urbana-Champaign
amandah3@illinois.edu

Uma Pal
University of Massachusetts Amherst
upal@cs.umass.edu

Adam J. Aviv
The George Washington University
aaviv@gwu.edu

Adam Bates
University of Illinois at
Urbana-Champaign
batesa@illinois.edu

ABSTRACT

Fitness tracking applications allow athletes to record and share their exercises online, including GPS routes of their activities. However, sharing mobility data potentially raises real-world privacy and safety risks. One strategy to mitigate that risk is a “Privacy Zone,” which conceals portions of the exercise routes that fall within a certain radius of a user-designated sensitive location. A pressing concern is whether privacy zones are an effective deterrent against common attackers, such as a bike thief that carefully scrutinizes online exercise activities in search of their next target. Further, little is known about user perceptions of privacy zones or how they fit into the broader landscape of available privacy precautions.

This work presents an online user study ($N=603$) that investigates the privacy concerns of fitness tracking users and evaluates the efficacy of privacy zones. Participants were first asked about their privacy behaviors with respect to fitness tracking applications. Next, participants completed an interactive task in which they attempted to deduce hidden locations protected by a privacy zone; we manipulated the number of displayed exercise activities that interacted with the privacy zone, as well as its size. Finally, participants were asked further questions about their impressions of privacy zones and use of other privacy precautions. We found that participants successfully inferred protected locations; for the most common privacy zone size, 68% of guesses fell within 50 meters of the hidden location when participants were shown *just 3 activities*. Further, we found that participants who viewed 3 activities were more confident about their success in the task compared to participants who viewed 1 activity. Combined, these results indicate that users’ privacy-sensitive locations are at risk even when using a privacy zone. We conclude by considering the implications of our findings on related privacy features and discuss recommendations to fitness tracking users and services to improve the privacy and safety of fitness trackers.

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 the author(s) 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.

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9157-3/22/04...\$15.00

<https://doi.org/10.1145/3491102.3502136>

CCS CONCEPTS

• Security and privacy → Human and societal aspects of security and privacy.

KEYWORDS

Fitness trackers, privacy, privacy zones, online survey, data sharing

ACM Reference Format:

Jaron Mink, Amanda Rose Yuile, Uma Pal, Adam J. Aviv, and Adam Bates. 2022. Users Can Deduce Sensitive Locations Protected by Privacy Zones on Fitness Tracking Apps. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3491102.3502136>

1 INTRODUCTION

Fitness tracking applications allow users to record their exercise activities in real time, post those activities online, and even compete with an online community of local athletes. By tracking their location, speed, and elevation changes using GPS, athletes can share detailed descriptions of their exercises on the fitness app [17] and other online social networks [37]. Fitness tracking applications have seen widespread success within the mobile application marketplace, boasting millions of users [11]. Spending on health and fitness applications tripled from 2016 to 2018 [3] and was valued at \$4.4 billion in 2020 [16]. That same year, Strava reported a user base of 55 million users, crediting the social and competitive elements of their services as a key factor in their growth [11].

Unfortunately, sharing GPS mobility data can lead to unintended privacy disclosures. Posting exercises activities online revealed the location and layout of secret military compounds [19, 22], led to the doxxing and stalking of application users [28], and has been linked to the theft of top-end bicycles and other exercise equipment [5, 39–41]. While fitness applications allow users to set posts to private or share them with only select followers, these options can conflict with the individual sharing preferences of users. For instance, while many users enjoy connecting with strangers from the athletic community rather than risk annoying their real-life friends, they still may be concerned about revealing their precise location to those strangers [2].

To offer athletes greater control over their online privacy, a specialized privacy control known as a “Privacy Zone” is offered by many fitness apps, including Strava [24] and Garmin [14]. Privacy zones grant users the ability to conceal portions of their exercise

route that occur near a sensitive location. Unfortunately, recent work raises concern as to the security offered by privacy zones [26]. Most notably, Hassan et al. present an algorithm that de-noises route data and solves a series of geometric equations to deduce protected locations of millions of athletes [18].

While prior work demonstrates that privacy zones leak information, it is unclear if this information can be weaponized by common attackers. Real-world thieves and stalkers may lack the necessary skills to deploy complex algorithms on millions of exercise routes scraped from the Internet [18]. Instead, they are more likely to abuse fitness apps by visually inspecting the exercise activities of individual targets. This more commonplace attacker, in spite of their technical disadvantage, may be also able to reason over exercise route data to de-anonymous users.

To gain a better understanding of privacy zones, and fitness privacy more broadly, in this work we consider the following questions:

RQ1 What are users' general perceptions and behaviors regarding privacy when using fitness apps?

RQ2 How effective are privacy zones at protecting users' sensitive locations, and what factors facilitate or impede their efficacy?

RQ3 How do users perceive the utility and effectiveness of privacy zones?

To answer these questions, we conducted an online user study on Amazon Mechanical Turk ($N=603$). Participants were first asked about concerns and management strategies regarding the privacy of their fitness activities. After completing a brief training on how privacy zones work, participants were then asked to complete an interactive task in which they attempted to infer locations protected by a privacy zone on a series of exercise activities. These exercise activities were drawn from unprotected (i.e., fully public) Strava posts that we retroactively applied privacy zones to based on the original Strava implementation. After the task, participants were asked to evaluate the utility and efficacy of privacy zones and indicate whether they would use privacy zones or other privacy precautions in the future.

We found participants to be largely successful at circumventing privacy zones. Smaller privacy zones ($\frac{1}{8}$ th vs. $\frac{5}{8}$ th mile) and more available routes (3 vs. 1) significantly increased the accuracy of participant guesses. Under the most common real-world scenario where the location was protected by the smallest privacy zone ($\frac{1}{8}$ mile) and multiple exercise activities (3) were displayed [18], 68% of participant guesses fell within roughly one house plot (50 meters, or 0.031 miles) of the true protected location. Participants in the 3-route condition were also more confident that they correctly inferred the sensitive locations as compared to the 1-route condition (63% vs. 43%). However, regardless of their condition assignment, participants expressed belief that privacy zones are effective. Furthermore, most participants reported that privacy zones did not meaningfully impact user experience.

Our results suggest that privacy zones, while flawed, are still viewed as an important tool for many users. In discussion, we make a series of recommendations to help users ensure their privacy, ranging from the use of the largest available privacy zones to out-of-app privacy precautions such as activating fitness trackers a few




Figure 1: An example of how privacy zones work. Lines represent individual exercise routes. On the left, 3 routes begin at a sensitive location (e.g., a user's home address). To protect this information, the user places a privacy zone at that address. Route information that falls within the privacy zone, and the privacy zone itself, will then be hidden from other users (shown on the right).

blocks away from the home. We also argue that it is in the best interest of fitness tracking companies to engage users in active and ongoing dialogue about how to manage their privacy and safety while using these apps. Finally, we discuss the implications of our findings on related fitness privacy mechanisms, including features just released by Strava in Summer 2021 [23, 38].

2 BACKGROUND AND RELATED WORK

Recognizing the potential privacy risks of posting exercise data online, fitness apps provide users with a variety of different options for managing their privacy. Features common to many online social networks, such as private accounts and hiding individual posts, are also available on fitness apps. However, these features can hamper users' enjoyment of the app by preventing them from connecting (and competing) with a community of athletes. To strike a better balance between privacy and utility, many fitness trackers offer the option to apply "privacy zones" to exercise activities, including Strava [24], Garmin [14], and others [4, 36]. Shown in Figure 1, privacy zones can be used to selectively hide the endpoints of an exercise route from other users. This allows athletes to share their performance without simultaneously publishing the private location where they started their exercise, such as their home or place of work. While not universally employed, privacy zones are an extremely popular feature. In 2018, they were observed to have been applied to 11% of exercises and were used twice as frequently as fully private activities [18].¹

Following growing recognition of the privacy concerns of athletes, the complexity and diversity of privacy zones mechanisms began to grow in 2018. Initially, all implementations behaved similarly to the example in Figure 1 – a privacy zone of a user-specified size was centered directly over a user-specified location. However, given that exercise endpoints marked the boundaries of the privacy zone and only a small number of possible radii were available (e.g., $\frac{1}{8}$, $\frac{2}{8}$, $\frac{3}{8}$, $\frac{4}{8}$, or $\frac{5}{8}$ of a mile), users began to observe that deducing the location hidden by a privacy zone was as simple as solving a basic geometry problem [29]. This intuition was confirmed by Hassan et al. [18], who leveraged computation methods to recover upwards of

¹To our knowledge, more recent usage data on privacy zones is not publicly available.

95% of hidden locations for users that had posted at least 3 exercises. In response, fitness apps have adopted a variety of approaches to further improve privacy, including randomizing the location of the hidden location within the circle [15] and “fuzzing” the intersection of the route with the privacy zone to make the center more difficult to guess [18]. As recently as August 2021, Strava released a variety of extended privacy zone options, including the ability hide up to 1 mile from the start/end points of routes on an activity-by-activity basis [23, 38].

While all these approaches to fitness privacy warrant consideration, this work focuses on the original privacy zone implementation depicted in Figure 1. By investigating an implementation that is known conceptually to leak information [18], we aim to investigate whether this information leakage is a real-world concern to a typical user, i.e., whether it would allow a common criminal to identify where a user lives. Further, this choice allows us to validate our experimental stimuli by testing against a known-vulnerable implementation, allowing our results and materials to serve as a baseline for the analysis of newer implementations in future work. We discuss the implications of our study on new privacy zones at length in Section 6.

Privacy Perceptions of Fitness Trackers. User concern and awareness for location privacy and sharing has been explored previously. Almuhammedi et al. investigated privacy notices with respect to location sharing in the context of mobile application permissions [1]. Motti and Caine investigated user privacy concern with respect to wearable technology [25], which are often used for sensor information in fitness tracking, showing that users have different privacy concerns depending on the type of wearable and the information being tracked.

Prior work has also explored these issues for fitness tracking specifically. Gabriele and Chiasson [13] survey fitness tracking users, observing that users have nuanced approaches to privacy but are often not fully aware of how this data is used. Alqhatani and Richter investigate sharing practices of wearable fitness data through semi-structured interviews in which participants discussed norms and self-presentation as issues with sharing [2]. Finally, Zimmer et al. apply communication privacy theory to how users manage fitness information through qualitative interviews, finding that participants exhibit low levels of privacy concern [43]. Our study differs from the above in our focus on mobility/GPS information, rather than general fitness/health data. We observe some of the same trends – for example, like Gabriele and Chiasson we observe that users develop their own ad hoc privacy precautions, and like Alqhatani and Richter we identify social pressures that influence sharing decisions. However, our work further differentiates itself through the introduction of a novel interactive task that allows participants to actively experiment with and evaluate privacy zones. This leads to participants that are more informed about the potential real-world issues, offering deeper insights into how users active fitness tracking users reason about privacy risks.

3 THREAT MODEL

In this section, we enumerate our assumptions about the knowledge and capabilities of an attacker that seeks to do harm to real users of fitness apps. Our study methodology casts participants in the

role of such an adversary, who possesses limited technical abilities. Specifically, our design is guided by the following assumptions about this attacker. We assume that the attacker is limited only to inferences that can be made through visual inspection of exercise posts while visiting a fitness tracking website. While these websites typically display each exercise post in a separate map, we assume that the attacker can composite multiple exercise activities onto a unified visualization. This could be done by drawing different exercise routes onto one physical map and thus does not require technical skills. Finally, we assume that this attacker has reviewed and understands basic information about how privacy zones work; this functionality is explained in accessible language by most fitness app websites (e.g., [4, 14, 24, 36]).

We argue that this threat model is more consistent with the concerns of typical users. Prior work on privacy zone efficacy assumed an extremely powerful attacker that analyzed tens of millions of exercise activities [18]. Instead, our study is designed with consideration for a less technically skilled attacker that is motivated by common crime such as theft [5, 39, 40] or spying on an intimate partner [9]. We can compare this adversary to Freed et al.’s “UI-bound” attacker in their study of technology abuse in intimate partner violence [12]. While both consider the threats posed by a technical layperson, whereas Freed et al.’s adversary has authenticated access to the mobile device of their victim, our attacker only has access on a social network website (e.g., can view public posts). The scenario we consider thus has an extremely low barrier to entry for the attacker.

4 METHODOLOGY

In this section, we describe the browser-based survey used in this study. First, participants were asked about their familiarity with fitness trackers and privacy-relevant behaviors, including privacy zone utilization. Participants were then asked to complete a task in which they attempted to identify protected sensitive locations using an interactive map of exercise activities. Finally, participants were surveyed about the utility of privacy zones and if they would use such a feature compared to, or alongside, other privacy precautions. In the rest of this section, we describe the survey design and procedure in detail.

4.1 Survey Procedure

The survey was comprised of three sections, a pre-task survey (RQ1), privacy zone inference task (RQ2), and a post-task survey (RQ3). The complete survey can be found in Appendix A.

Pre-Task Survey. After providing informed consent, participants were asked about their use of fitness apps and if they have any privacy strategies when using fitness apps. First, we defined fitness apps as below:

Fitness apps allow users to track their activities, such as runs or bike rides, using wearable devices or mobile phones. Examples of popular fitness apps include Strava, Garmin Connect, Endomondo, MapMyRun, and Nike+.

Participants were then asked whether they currently or previously used any fitness apps, which fitness apps they used, whether they shared activities using these fitness apps, their comfort level in sharing these activities, and whether they took any steps to protect




Figure 2: Interactive map interface used in the privacy zone inference task. The participant first views map that displays one or more routes that have been truncated by a privacy zone (see left). Next, the participants infer the size and placement of the applied privacy zone via a blue circle (see middle). Finally, the participant infers what the protected location is via a red pin (see right).

their privacy while using these apps. For those who do not use a fitness application, they were asked to report on apps they have heard of, if they would be comfortable sharing, and any steps they would take to protect their privacy when using fitness apps. These questions are marked Q1-Q7 in Appendix A.

Privacy Zone Inference Task. Following the pre-task survey, participants were briefed on what privacy zones are and how they function. All participants received the following definition of a privacy zone, as well as a visual explanation (see Appendix A.2).

To create a privacy zone, a user enters an address and then a virtual circle is placed on the map with the protected address at the center. Each user can still see the entire route of their exercise, but friends can only see a shortened version that appears to begin/end at the edge of the circle.

Next, participants completed a brief guided activity that served to familiarize them with the task interface and ensure that they possessed a baseline understanding of how privacy zones work. Shown in Figure 2, participants were presented with an interactive map interface that displayed an exercise route that had been protected by a privacy zone. The interface consisted of an OpenStreetMaps panel that could be zoomed or dragged, a toolbar that allowed participants to place a privacy zone circle, a pin representing their guess as to the exact location, and options to submit or reset their guess. In the familiarization activity, participants were presented with two privacy zone maps and were offered feedback if their guesses indicated a misunderstanding about how privacy zones work. The first map displayed one “out-and-back” exercise route that intersected the privacy zone twice, once at each endpoint (e.g., Fig. 2). Participants were prompted to select a privacy zone circle size, place the circle where it could intersect both endpoints, place a red pin where they thought the hidden location was in the circle, and finally submit their guess using the green checkmark button. If the circle touched both endpoints and the red pin was roughly in the center of the circle, participants were permitted to proceed; otherwise, they were prompted with feedback about their mistake and asked to try again, up to 10 times before proceeding. The second map proceeded similarly but depicted a “one-way” exercise route

where the start and end points were far apart, meaning that they could not *both* touch the privacy zone. One of these endpoints fell in a forested area while the other was in a populace area featuring many streets and buildings. Participants received feedback if they chose the forested area and were asked to try again, up to 10 times before proceeding.

After the familiarization activity, participants were presented with 12 more privacy zone maps and asked to use the same interface to identify the privacy zone size and hidden location. Across participants (between subjects), we varied the number of exercise routes (1 vs. 3) visible on the map. Across trials (within subjects), we varied the size of the privacy zone applied to the exercise routes ($\frac{1}{8}$, $\frac{3}{8}$, and $\frac{5}{8}$ mile radii) as well as the density of road intersections around the protected location (high, medium, or low). In these trials, participants were not given any feedback or asked to try again if their guess was incorrect. We describe how we selected the maps and exercises that were displayed in the inference task in Section 4.2.

Post-task Survey. After completing the inference task, participants were asked to self-report on the strategies they used, and how accurately they believe they were in identifying the privacy zone sizes and locations (Q10-Q15). We then asked participants to report on their perceptions of privacy zones as an effective mechanism for fitness app privacy, whether they would be willing to share their exercise on fitness apps, the impact of privacy zones on their experience in fitness apps, and any future privacy precautions they might use with fitness apps (Q16-Q22). Finally, participants were asked to provide demographic information (Q24-Q26).

4.2 Exercise Activity Selection

During the inference task, participants were presented with actual fitness activities from the Strava service that were originally collected by Hassan et al. [18]. The activities were drawn from fully public Strava posts recorded by 175,607 total athletes. Each activity in the dataset contained a GPS trace of an exercise and an associated athlete ID. Importantly, none of these activities had a privacy zone applied to them by the athletes. Instead, we retroactively applied a privacy zone to the GPS trace using the original privacy zone

mechanism described in Section 2 and depicted in Figure 1. Applying privacy zones to public exercise activities granted us ground truth as to the hidden location being protected by the privacy zone. It also ensured that we did not expose private athlete information to our participants; we explore the ethical considerations of our methodology further in Section 4.4.

To select a set of exercise activities for the inference task, we applied the following criteria to the dataset: First, following the procedure used by Hassan et al. [18], we removed athletes that had less than three total activities whose endpoints fell within 0.031 miles (50 meters) of one another. We required at least three activities so that the same maps could be used in both the one and three route conditions. We then generated modified versions of these activities using the three privacy zone sizes ($\frac{1}{8}$, $\frac{3}{8}$, and $\frac{5}{8}$ mile radii), with the centroid of the nearby endpoints serving as the protected location. These zone sizes are the smallest, medium, and largest of the privacy zone sizes available on Strava. We then removed athletes if the endpoints of the modified activities were not at least 0.031 miles *apart* from one another. This is because we wanted these endpoints to be visually distinct for participants in the inference task. From the remaining athletes, we then drew a random sample while controlling for geographic area, using road intersection density around the protected location ($\frac{5}{16}$ miles) as a proxy for urbanization. Finally, selected maps were visually inspected to ensure that the GPS trace was stable and consistent. If a large gap was detected, the athlete was randomly replaced by another with a similar road density.

Using these criteria, we selected 60 athletes with 3 exercise activities apiece to use in our inference task. We semi-randomly formed 5 ordered lists of 12 maps, each containing 4 maps with high, 4 with medium, and 4 with low intersection density. Each map set was rotated into three versions to control for ordering effects in the inference task. This resulted in a total of 15 map lists. For all lists, we also semi-randomly applied a different privacy zone size to each map such that a participant would see each privacy zone size an equal number of times and all maps were viewed under each privacy zone an equal number of times. This resulted in a total of 45 map lists (15 * 3 zone sizes). Finally, each map could also be displayed in one of four possible route configurations – three routes combined, or just one of the three individual routes. We chose to test each individual route separately to gain insight into whether certain features of the exercise route simplified the task. From the 60 initial athletes, this led to a total of 180 map lists (45 * 4 route configurations) and 720 unique map items (60 athletes * 3 zone sizes * 4 route configurations).

4.3 Limitations

There are a number of potential limitations with our survey methodology. First, because data collection took place on Amazon Mechanical Turk, our sample may not be fully representative of the entire population. As with many MTurk studies, participants tended to be younger and identify as male. In the context of fitness tracking, the over-representation of male respondents may obscure certain trends, e.g., women’s perception of the privacy and safety of fitness trackers. However, there is evidence that MTurk samples generalize with respect to reported security behavior [31].

Second, our sample may not be fully representative of the current population of fitness tracking users. As we discuss in Section 5, while most of our participants reported having used at least one fitness tracking application, many participants did not. Further, only some of these participants previously made use of exercise map functionality within the fitness app. To account for this limitation, where appropriate in our pre-task survey results (§5.1), we separately examine the responses of those participants that reported having used a fitness app (Q1). Further, our inference task includes a familiarization activity and does not depend on past exposure to fitness apps. The same is true of our post-task survey in which participants reflect on their experience in the inference task. Thus, no accommodations are needed when interpreting these results (§5.2-5.3).

Third, our inference task makes use of only three privacy zone sizes, as opposed to the Strava’s five ($\frac{1}{8}$ – $\frac{5}{8}$ miles) or Garmin Connect’s nine ($\frac{1}{10}$ – 1 miles) available options. In our experiments, we made use of the smallest, middle, and largest options on Strava, reasoning that any linear trends in performance would be captured using just these three options. We note also that $\frac{1}{8}$ mile privacy zone is by far the most popular [18] and is featured in our study. Finally, our survey may also suffer from response bias where participants overestimate their adherence to privacy preserving behaviors, particularly in the post-task survey after the participants were primed to think about privacy zones. We account for this possibility while interpreting our survey results.

4.4 Ethical Considerations

Our study was approved by the IRB at the authors’ institutions. The inference task asked participants to take on the role of an attacker attempting to deduce a sensitive location that was protected by a privacy zone. This methodology raises concerns that participants would gain access to private athlete data during the study. To account for this, the exercise activities used in our survey are all public posts to the Strava network that did not have an active privacy zone. This means that no actual protected locations were revealed during our study; instead, participants viewed maps that were already in the public domain on Strava’s website. To further reduce the risk to active Strava users, we also placed constraints on the map interface to reduce the likelihood that participants could deduce the displayed location. Specifically, we removed all street names, city names, and geographic markers using the OpenStreetMaps API. We also placed limits on how far out the participants could zoom, preventing them from easily identifying the country or region where the exercise took place.

A final concern raised by our methodology is that participants were being trained to violate user privacy on fitness tracker websites. To account for this in our design, we first avoided use attack-oriented terminology to prevent alarming participants or calling attention to how this task, if performed outside the study, could infringe on the privacy of others. For example, we use the terms “friend” or “others” when describing how privacy zones affect the visibility of routes, rather than speaking in the parlance of “attackers” and “victims”. Further, our study makes use of the original privacy zone implementation used by Strava, Garmin, and others

Table 1: Participant Background by Route Treatment. * *One Route* combines three treatment groups where participants viewed either the first, second, or third of the routes available on each map. This is why the counts are roughly three times as large as the *Three Routes* condition.

		One Route*	Three Routes	Total
Gender	Female	168	55	223
	Male	279	95	374
	Non-Binary	2	1	3
	Prefer not to say	2	1	3
Age	18-20	4	1	5
	20-24	29	5	34
	25-29	88	33	121
	30-34	122	37	159
	35-39	78	35	113
	40-44	51	16	67
	45-49	33	10	43
	50-54	20	6	26
	55-59	14	4	18
	60-64	3	5	8
	65-69	6	0	6
	70-74	1	0	1
	Prefer not to say	2	0	2
Location	Urban	146	60	206
	Suburban	230	73	303
	Rural	70	19	89
	Prefer not to say	5	0	5
Total		451	152	603

prior to late 2018. Newer implementations expose the same interface but apply subtle noise to the privacy zone’s placement such that the sensitive location is no longer reliably in the center of the circle. To our knowledge, no fitness tracking service still uses the exact implementation employed in our study, thus preventing participants from using the training they received to violate real users’ privacy. Of course, our use of the original privacy zone implementation is also a methodological limitation, albeit a necessary one. We consider how our findings impacts other privacy zone implementations in Section 6.

5 RESULTS

In this section we discuss the results of the online survey. First, we discuss participants’ familiarity with fitness apps, as well as their comfort with sharing activities with fitness apps, before being exposed to the inference task. Next, we analyze participants’ performance in inferring privacy zones from activity data, and the factors that influence that performance. Finally, we discuss whether exposure to attacks on location privacy affect participants’ comfort in sharing activities on fitness apps, and the strategies they employ to manage their privacy.

Participants. We recruited $N=603$ participants on Amazon Mechanical Turk between September and December 2021. which provided us with (at least) 9 completions for each of the 720 unique items used in the inference task. To account for low-effort or automated responses, which are a common problem in online surveys, 117 additional participants that completed the survey were




Figure 3: Prevalence of different fitness applications amongst participants who reported using a fitness tracker (Q2). Overall, 73.4% of participants were currently using or had previously used fitness apps (Q1).

excluded because they failed attention checks, provided irrelevant single word responses, or an author judged their responses to be copied from external online sources (e.g., responses were reused by multiple participants). The survey took a median of 31.34 minutes to complete and participants were compensated \$3.75 for their time.² Table 1 describes participant demographics: Participants ranged from 18-74 with the median age between 30-34; 62% of respondents identified as male and 37% as female.

Qualitative Response Coding. Our pre- and post-task surveys included a number of open-ended questions regarding user attitudes and behaviors surrounding fitness app privacy. To analyze open-ended responses we used thematic coding. For each open-ended response question, we assigned a primary coder from the research team (the first author) to code a random sub-sample ($n=300$)³ of the 603 responses for the purpose of developing a “codebook.” 300 responses was sufficient to reach code saturation for each question, which we visualize in Appendix B. A second coder used that codebook to independently code 20% of the previously coded responses. Inter-coder agreement was calculated using Cohens- κ . In cases when high agreement ($\kappa > 0.7$) was not reached, the two coders discussed their differences, resolved incorrect codings, and if a code was found to be inconsistently applied, updated the codebook and re-coded all responses accordingly. This process was repeated until high agreement was reached. Below, we report the results of the primary coder for each open-ended question; the full codebook can be found in Appendix B.

5.1 Pre-Task Survey (RQ1)

The pre-task survey was designed to measure users’ general perceptions and behaviors regarding privacy when using fitness apps.

Fitness App Sharing Attitudes and Behavior. Consistent with the large user bases reported by fitness tracking companies (e.g., [11]),

²This compensation amount was equivalent to prorated minimum wage for our anticipated completion time of 20 minutes, although participants took longer to complete the survey in the final study.

³In some cases responses were grouped based on a previous response. If less than 300 responses existed for a question, all responses were coded.




Figure 4: Proportion of participants that had posted (or had seen posts for) exercise maps using fitness trackers. (Q3)




Figure 5: Level of comfort reported in sharing maps of their exercise across all participants, including non-users. (Q4)

most participants (73%, $n=443$) were currently using or had previously used *at least* one fitness app. Participants reported using over 32 different fitness apps (Figure 3, Q2). Among participants that reported having used a fitness app, 39% reported sharing geolocation data within the app or another social network, as compared to 57% that had not (Figure 4, Q3). Among non-fitness app users ($n=160$), 49% reported having viewed an exercise map on a social network, while 47% had not viewed an exercise map on a social network. This indicates that exercise posts are highly visible online, even to those that are not fitness apps users.

Roughly half of our participants noted some level of unease when sharing exercise maps online (Figure 5, Q4). 53% of participants reported that they were *somewhat* or *very uncomfortable* sharing, while 45% were *somewhat* or *very comfortable*. When asked to explain the reason behind these preferences (Q5), 55% ($n=165$) of coded participant responses (see Qualitative Response Coding) expressed concerns that included, most commonly, the risk of their location being revealed. One participant (P139) noted that shared location data allows viewers to “know where I go, alone, and ...figure out when I’m not at my house.” Participants also mentioned peer judgement, being part of a vulnerable demographic group, and lack of trust in their friends on social networks as concerns. Conversely, 41% ($n=123$) of participants noted a lack of concern, often crediting their use of privacy precautions as the reason for this comfort. Additional reasons for comfort included the belief that their exercise routes did not contain sensitive information, practicing discretion

about what and when to share exercise data, and sharing amongst dedicated fitness-based social networks to minimize social awkwardness. Some participants (15%, $n=46$) mentioned the benefits of using fitness apps, most notably the social support that sharing an exercise provides. One participant (P112) said, “I want others to see my progress, keep me motivated in reaching my goals, and to continue to keep me accountable”. 10% ($n=29$) of participants mentioned that they saw no benefit to sharing exercises, and 2% ($n=7$) had no opinion.

Privacy Precautions. Among participants that had used a fitness app, 51% reported having taken active steps towards protecting their privacy in some form, 28% had not taken any privacy-oriented steps, and 21% were unsure (Q6). This rate of privacy precautions appears roughly consistent to those observed by Gabriele et al. [13], who found that 49% of fitness app users had actively set sharing preferences. There was no significant difference between those who had or had not taken active privacy precautions when it came to whether or not they had shared an exercise (Q3), Pearson chi-squared test: $\chi^2(1, 344) = 0.08, p = 0.784$.

When asked to describe the privacy precautions that they had employed (Q7), participants that had used fitness app mentioned using privacy features offered by the app (62%, $n=138$) or by the mobile device (27%, $n=60$). App-level defenses included setting their profile to private or only posting in a private group. While only a few participants specifically mentioned privacy zones ($n=3$), this usage rate appears to be consistent with Hassan et al.’s observation of 11% usage on Strava [18] when accounting for the fact that apps that offered privacy zones were less popular among our participants.⁴ Phone-level defenses included placing restrictions on the fitness app’s access to location data or the Internet. 17% ($n=37$) of participants limited or were faking the information they provided to the fitness app. Other participants took additional actions (10%, $n=23$) such as monitoring their social media for accidental disclosure of information or sharing via alternative methods such as screenshots and text messaging. 8% ($n=18$) also discussed ways they ensured protection through access data control mechanisms such as authentication or secure communication.

Key Takeaways We found that fitness app users are likely to share their activities with others, and that even non-users often view others’ exercise activities. Participants were roughly split on whether they are comfortable or uncomfortable sharing geolocation data on fitness apps, and only half of our participants reported employing any privacy defenses. These findings are consistent with prior reports on the proportion of users that actively manage their privacy [13] or use privacy zones [18]. We also observe that some fitness app users employ out-of-band privacy precautions such as starting and stopping the GPS tracking further away from their home. This suggests that some privacy behaviors may be difficult to observe and evaluate by focusing only on the fitness apps.

⁴Recall that not all fitness apps offer a privacy zone feature. The two most prominent fitness apps to offer privacy zones, as reported by our participants, are Strava and Garmin (see Figure 3).




Figure 6: Box plots of participant guess error by route assignment and privacy zone size. Middle lines indicate median performance, hinges mark the first and third quartiles, and whiskers extend up to 1.5 times the inter-quartile range.

5.2 Privacy Zone Inference Task (RQ2)

The inference task was designed to measure the efficacy of privacy zones at protecting users’ sensitive locations and determine the factors that contribute to an attacker’s success (or failure). Based on Hassan et al.’s findings [18], we predicted that more exercise routes and/or smaller privacy zones would improve accuracy in determining the hidden location. Reasoning that less populace areas would feature fewer buildings and thus offer a smaller anonymity set, we also predicted that participants would perform better when analyzing exercises performed in regions with fewer roads.

Accuracy in the inference task was assessed by measuring the distance between the true hidden location and the participant’s guess. As a reminder, the true hidden location lies in the center of the privacy zone circle. We call this distance “guess error.” Larger guess error indicates that the attempt is farther away from the true hidden location protected by the privacy zone.

Figure 6 depicts box plots summarizing the guess error by route assignment and privacy zone size. As can be seen, accuracy improved (smaller guess error) when participants received 3 exercise routes instead of just 1. In the 1-route condition, accuracy in determining the hidden location appears to vary as a function of privacy zone size, with $\frac{1}{8}$ mile privacy zones exhibiting a median error of just 0.08 miles (133 meters) while the median guess for $\frac{5}{8}$ mile zones fell 0.50 miles (805 meters) away. In contrast, median performance in the 3-route condition is roughly equal regardless of privacy zone size, with guess error for the smallest zone at just 0.01 miles (10 meters). We attribute this to an apparent floor effect in our design, where participants in the 3-Route condition perform well regardless of privacy zone size.

To quantify these impressions, we ran a linear mixed effects regression model using the `lmer` function from the `lme4` package [7] in R [30]. P-values were calculated using the `lmerTest` package [21]. **This model evaluates the independent effects of route assignment and privacy zone size (the main effects), while also measuring their interaction (i.e., whether the effect of route assignment on guess error changes with privacy zone size).**

Table 2: Linear mixed effects regression model. The unit for estimate and standard error is miles. Significance is denoted by * ($p < 0.001$).**

Variable	Estimate	Std. Err.	t
Intercept	0.411	0.070	5.896***
Route Assignment (Reference = 1 Route)			
3 Routes	-0.248	0.051	-4.841***
PZ Assignment (Reference = $\frac{5}{8}$ th mi)			
$\frac{1}{8}$ th mi	-0.163	0.036	-4.540***
$\frac{3}{8}$ th mi	-0.038	0.036	-1.053
Interaction (Reference = Route : PZ $\frac{5}{8}$ th mi)			
Route : PZ $\frac{1}{8}$ th mi	0.275	0.081	3.385***
Route : PZ $\frac{3}{8}$ th mi	-0.127	0.082	-1.553

Our model uses guess error as the outcome measure. Our fixed effects were route assignment (1 or 3 routes, between subjects), privacy zone size ($\frac{1}{8}$, $\frac{3}{8}$, $\frac{5}{8}$ miles, within subjects), and their interaction. We attempted to add, but subsequently removed, intersection density as a fixed effect because it did not account for any additional variance in the model and degraded the model fit. We used mean-centered effects coding, which adjusts for unequal amounts of information between groups, for our comparisons within the fixed effects. This was necessary because each of the routes in the 3-route condition were tested independently as 1-route prompts, such that the 1-route condition accounted for 74% of our data. Privacy zone Assignment was a three-level factor, allowing for two comparisons in our model. We chose to compare the small privacy zone to the largest ($\frac{1}{8}$ vs. $\frac{5}{8}$) as well as the medium to the large ($\frac{3}{8}$ vs. $\frac{5}{8}$). Finally, the model also includes the random slope for item by route assignment, random slope for item by privacy zone size, and random intercept for participant, which reflects the maximal random effects structure permitted by the data [6].

The model is summarized in Table 2, with negative values for estimate indicating a decrease in guess error. Participants’ guess error was significantly lower in the 3-route assignment than the 1-route assignment ($\beta = -0.248$, $SE = 0.051$, $t = -4.841$, $p < 0.001$). Participants’ accuracy was also influenced by privacy zone size: guess error was significantly lower when the true privacy zone size was $\frac{1}{8}$ rather than $\frac{5}{8}$ ($\beta = -0.163$, $SE = 0.036$, $t = -4.540$, $p < 0.001$). However, there was no significant difference in guess error between the $\frac{3}{8}$ and $\frac{5}{8}$ privacy zone sizes, which is likely due to the floor effect observable in the 3-Route condition. Finally, we observed a significant interaction of route assignment and the $\frac{1}{8}$ vs. $\frac{5}{8}$ privacy zone treatment ($\beta = 0.275$, $SE = 0.081$, $t = 3.385$, $p < 0.001$). This interaction can be seen in Figure 6. Here, the interaction’s positive coefficient effectively indicates that the 3-route condition receives little-to-no additional benefit in the case of a $\frac{1}{8}$ size privacy zone. In contrast, guess error is reduced for the 1-Route condition for the $\frac{1}{8}$ size privacy zone.

Comparison to Prior Work. Our analysis thus far has not yet provided a direct comparison between the performance of our participants and Hassan et al.’s automated circle finding algorithm [18]. In their work, accuracy was not a continuous variable but instead a binary determination of whether the algorithm’s guess fell within 0.031 miles (50 meters) of the true hidden location. According to

Table 3: Proportion of guesses where guess error was less than 50 meters from the hidden location.

Route Assignment	PZ Size	Prop. Guess Error ≤ 0.031 miles (50 meters)
1-Route	1/8th	38%
	3/8th	26%
	5/8th	11%
	Combined	25%
3-Route	1/8th	68%
	3/8th	63%
	5/8th	54%
	Combined	62%

the authors, this threshold was selected to represent the size of a typical suburban house plot. Table 3 reports on the proportion of participant guesses with guess error less than or equal to 0.031 miles. Under the most favorable conditions (3-route, $\frac{1}{8}$ mile privacy zone), the success rate was 68%. In contrast, the success rate reported by Hassan et al. under the same conditions was 95%. The lower success rate for human participants is likely due the computer algorithm’s ability to distinguish between nearby GPS coordinates with many decimal points of precision. In contrast, our participants were limited in their ability to differentiate endpoints that were nearby one another, preventing them from being able to visually trace the outline of the hidden circle.

Participant Strategies. We explored strategies participants used to successfully deduce protected locations in two ways. First, participants were asked to self-report on any strategies they employed during the task to help determine the privacy zone sizes and hidden locations. **For this qualitative question, we again coded a random subset of 300 responses using the procedure outlined above.** Participants most frequently mentioned using embedded topological information in the map to inform guessing the privacy zone size (60%, $n=179$) and pin location (69%, $n=203$). Participants noted the presence of residential neighborhoods, street paths, and physical boundaries (e.g., mountains, lakes) in deducing possible areas of interest. One participant (P118) noted making “sure the pin was not placed in a body of water, [and] placing it on or near a building”. Participants also used route features for determining circle size (50%, $n=149$) and location (16%, $n=47$). Many participants noted using all available endpoints and “tried to size the circle based on what ‘fit’” (P14). Some participants used multiple information sources in conjunction to make complex inferences, e.g., “I tried to predict based on the end point, what general direction would the person most likely have gone. From there, I chose one of the buildings along that path” (P6). When identifying the hidden location with the pin, most participants (53%, $n=158$) simply placed it near the center of an already identified circle. For determining the privacy zone size, a few participants simply recited that they followed the instructions, randomly guessed, had a preference of a particular circle size, or chose inferences based on how the map was shown in the survey (this was not a factor in the task).

Second, we manually reviewed the 10 best and worst performing items according to median guess error. In the 1-Route condition,




Figure 7: Participants’ evaluation of the utility impact of privacy zones. (Q18 & Q19)

we noticed all of the top 10 items contained “out-and-back” routes where both endpoints intersected with the privacy zone, while all of the bottom 10 items contained “one-way” routes that intersected the privacy zone just once. As we have shown in our primary analysis, access to additional routes (and, thus, endpoints) significantly reduces guess error, we speculate that out-and-back routes simplify the task by providing additional endpoints. In the 3-Route condition we notice a similar trend, with the top 10 items offering over two out-and-back routes on average (mean=2.27) and the bottom 10 items offering less than two (mean=1.99). This may suggest that guess error would continue to decrease if we had given participants more than 3 routes. Beyond the shape of the exercise routes, we were also curious about additional information sources that participants recruited to solve the task. To explore this, we inspected the 10 best performing items in the 1-Route condition *that featured one-way routes*, reasoning that these items must contain cartographic or environmental information sources that simplified the task. We notice that these top items tend to contain one endpoint in a small, urbanized area and another in an empty area without buildings. It may be that these “dead end” endpoints allowed participants to deduce which size of the route intersected the privacy zone, while the small, urbanized area ruled out the possibility of larger privacy zone sizes whose center would have overshoot the urban area. These observations are consistent with some of the self-reported participant strategies described above (e.g., P6). In contrast, the bottom items with one-way routes mostly fell in large city-grid that gave little indication as to which endpoint intersected the privacy zone or what its size might be.

Key Takeaways Our findings indicate that, when attempting to infer a sensitive location that is hidden by a privacy zone, attackers will have more success when given multiple routes and/or smaller privacy zones. What makes this result especially concerning is that our multi-route condition offered participants just three exercise activities; in contrast, regular users of Strava post many activities per month. Given this, we predict that the protection offered by privacy zones would further degrade as the observable exercise posts accumulate.

Under favorable conditions, participant guesses consistently fell within about one house plot of the true hidden location. As noted by Hassan et al. [18], this favorable condition is actually the most




Figure 8: Participant's evaluation of the efficacy of privacy zones in the inference task. From an "attack-oriented perspective," we asked participants to evaluate their own performance in the task (Q10 & Q12, 8a). From a "defense-oriented perspective," we asked participants their view of the effectiveness of and confidence in privacy zones (Q15 & Q20, 8b)

common one — about 55% of all observed privacy zones on Strava used the $\frac{1}{8}$ mile radius and the median activity count was 5 per month. While the observed success rate by our human participants pales in comparison to Hassan et al.'s reported 95%, we feel this result is nonetheless concerning given the lack of expertise needed to attempt this attack. While 3 activities were necessary and sufficient for an algorithm, success rates for humans would likely improve given the opportunity to view more exercises; we can reasonably expect a motivated attacker to have access to more exercises as their targets continue to post new activities. Furthermore, unlike Hassan et al., humans meaningfully utilize topological information such as neighborhoods, roads, and physical obstacles to inform their guesses about the size of the privacy zone and the protected location. Thus, our results likely underestimate the risks of this privacy zone mechanism.

5.3 Post-Task Survey (RQ3)

The post-task survey was designed to determine how users perceive the utility and effectiveness of privacy zones.

Perceptions of Privacy Zone Efficacy and Utility. We asked participants about the degree to which privacy zones affect the user experience from the perspective of both an activity poster (Q18) and an activity viewer (Q19). As shown in Figure 7, most participants reported that the use of privacy zones impacted the user experience *very little* or *not at all* for posting (73%) and viewing (69%) an activity. This indicates that most fitness app users find that privacy zones do not interfere with the utility of app, namely sharing GPS routes of activities.

Participants were also asked to provide their opinion on the efficacy of privacy zones via questions from both *attack-* and *defense-*oriented perspectives. Attack-oriented questions (Q10, Q12) asked the participant to self-evaluate their performance in the inference

task, while defense-oriented questions (Q15, Q20) asked participants about their level of confidence in privacy zones' ability to protect sensitive locations.

When asked about privacy zone effectiveness in an *attack*-oriented perspective (Q10, Q12), participants were more confident in their ability to identify the zone size than the sensitive location. Across all conditions, 68% of participants believed they found the zone size *somewhat* or *very often* and 48% believed they found the hidden location *somewhat* or *very often*. Later in the survey, when asked to reflect on their comfort in posting exercises online (Q17), many participants alluded to their perceived success in the task: participants that were confident in Q10/Q12 later noted that they were "surprised at how easy it was to pinpoint the privacy zone" (P9) and from there "it is pretty simple ... to guess someone's location based on the circle size" (P160), while less confident participants later remarked that "The privacy zones seem to do their job. I was guessing a lot" (P71) and "even with all of the information on how it worked, I struggled" (P123). When comparing across route assignment (Fig. 8a), we notice that participants that viewed three routes were more confident than those that viewed one route. The percent of participants that believed they identified the correct zone size *somewhat* or *very often* increased from 63% to 80% (1 vs 3 route), while belief in identifying the protected location *somewhat* or *very often* increased from 43% to 63% (1 vs. 3 route). These differences are significant under an independent measures Mann-Whitney U test: $U_{\text{home}}(443, 152) = 41357.5, p < 0.0001$, $U_{\text{pin}}(439, 149) = 39937, p < 0.0001$.

When asked about privacy zone effectiveness from a defense-oriented perspective (Q15, Q20), most participants found privacy zones to be effective. Across all conditions, 65% of participants believed privacy zones were *somewhat* or *very effective* at protecting sensitive locations, and 58% would be *slightly* or *very confident* in




Figure 9: Participants’ change in comfort level following the inference task. Pre-task comfort is organized by row, while post-task comfort is denoted in colored bars. (Q4 & Q16)

privacy zones’ ability to protect a personal location. Surprisingly, we found no difference in perceived effectiveness and confidence when splitting responses by route assignment (Fig. 8b). The percentage of participants believing privacy zones were *somewhat* or *very effective* decreased from 66% to 63% (1 route vs. 3 routes) and the percentage who were *slightly* or *very confident* in privacy zones’ ability to protect a personal location decreased from 59% to 57% (1 route vs. 3 route). Neither condition difference was not found to be significant under an independent measures Mann-Whitney U test: $U_{\text{effectiveness}}(426, 150) = 30596, p = 0.4093$; $U_{\text{confidence}}(451, 152) = 33407.5, p = 0.6243$.

Change in Privacy Perceptions. Participants were asked about their change in comfort after the task via Likert response (Q16), then asked to explain their response. Over all conditions, 20% became less comfortable, 42% of participants experienced no change of comfort and 39% of participants experienced an increase in comfort. When viewed by route condition, the differences in comfort change do not appear to follow an obvious trend: the percentage of those less comfortable increases from 18% to 27% (1 route vs. 3 route) while the percentage of those who experience no comfort change or increased comfort decreases from 43% to 35% (1 route vs. 3 route) and 39% to 38% (1 route vs. 3 route) respectively. We did not find evidence of the number of presented routes significantly influencing comfort change via an independent measures Mann-Whitney U test: $U(443, 150) = 31916, p = 0.49497$.

To better understand the factors that affect comfort change, we thematically coded a random sample of 300 participants’ comfort change explanations (Q17). A large proportion (52%, $n=156$) mentioned the efficacy of privacy zones as a primary reason, with 28% ($n=84$) having increased comfort and 20% ($n=62$) having decreased comfort. A further 4% ($n=12$) note that the provided privacy of privacy zones to be highly dependent on the region of the activity, e.g., “it does a decent job... with lots of residential zones around. But if I were in the middle of the country, I still wouldn’t feel adequate” (P35). 20% ($n=61$) expressed that their prior beliefs played a major role and thus the inference task did not greatly affect them: “I am already uncomfortable posting my things online so this didn’t change very much for me” (P31). 14% ($n=42$) of participants believed that the defense did not address a concern of theirs, such as someone knowing where they run (9%, $n=28$) regardless of a sensitive




Figure 10: Privacy precautions participants said they would apply in the future when sharing an exercise (Q22).

location: “even a ‘privacy zone’ doesn’t help since someone can still track you outside of it” (P83). 10% ($n=30$) believed this defense was unnecessary given the availability of private groups (3%, $n=8$) or the fact that they would never use a fitness app anyhow (6%, $n=19$).

To determine if this relative change in comfort was affected by the initial pre-task comfort level (Fig 9), a non-parametric Spearman rank-order test was ran and a statistically significant moderate, positive correlation was found between prior discomfort (Med:4, IQR:2)⁵ and change of discomfort (Med:3, IQR:1)⁶: $r_s(593) = 0.3145, p < 0.0001$. This indicates that participants that reported being comfortable sharing before the task reported being more comfortable sharing after the task, while participants that reported being uncomfortable sharing before the task reported being less comfortable sharing after the task. While this could be interpreted as evidence that participants became more entrenched in their initial beliefs over the course of the study, we note a number of alternative explanations. First, we did not design our experiment as a belief change task; we deliberately abstained from providing participants feedback on their performance because it ran contrary to our goal of isolating the factors that facilitate or impede privacy zone efficacy. As a consequence, however, this means that participants were not given an obvious reason to re-evaluate their comfort level. Second, because the pre-experiment question (Q4) assessed absolute comfort level while the post-experiment question (Q16) assessed comfort change, we are reluctant to interpret this result.

Comparison to Other Privacy Precautions. We asked participants which privacy mechanisms they would use if they were to share GPS activities in the future, recording both the individual counts (Fig 10, Q22) and the reasons behind them (Q23). 67% ($n=407$) participants reported that they would use a privacy zone for future sharing. Of these participants, only 10% ($n=39$) would use privacy zones alone, while 90% ($n=364$) would use some combination of traditional privacy mechanisms (e.g., private profile, private post, block user) and a privacy zone. 28% ($n=166$) of participants opted to only utilize generic privacy mechanisms (e.g., private profile, private post, block user).

To better understand how privacy zones relate to other privacy precautions, we analyzed the participant responses that specifically mentioned the privacy zone mechanism in the follow-up explanation (Q23) by matching the words “privacy zone” and “zone”. We

⁵From “Very Comfortable”(1) to “Very Uncomfortable”(5), 3 is unused

⁶From “More Comfortable”(1) to “Less Comfortable”(5)

then subdivided the responses based on whether using a privacy zone was selected/deselected in (Q22). This resulted in $n=117$ responses that selected privacy zone and mentioned it specifically, and an additional $n=13$ responses that did not select a privacy zone and mentioned it specifically.

Of the participants that mentioned privacy zones in their explanation and indicated they would use it ($n=117$), most (64%, $n=75$) described that they would intend to employ multiple protections simultaneously, believing that a privacy zone could provide defense in depth ($n=17$) and “adds a few extra layers of security” (P502). A few ($n=8$) were more focused on preventing leakage of sensitive locations outside a trusted group, e.g., in the event “someone took a screenshot and shared it publicly” (P92). Conversely, a few participants (3%, $n=3$) do not trust those *within* their social network group and thus wished to use a privacy zone. As P471 noted, “There are many varying degrees of trust between myself and various people in my social network. Some I would have no issue with them knowing where I live or frequent, some I wouldn’t want knowing that information”. A few participants mentioned that they would not typically use a privacy zone, but if they were to share more publicly then they believed it would be a good option (4%, $n=5$). Other participants even went so far to note that they believed privacy zones as brittle and ineffective but still intended to utilize them (9%, $n=11$). This was justified in a few ways. First, a few felt that any protection, albeit imperfect, was better than nothing: “any attempt to obfuscate the data is better than none at all” (P330) ($n=6$). Another participant relied on the fact that it required effort to break the privacy zone: “someone would have to be specifically determined to figure out my location in order for the zone to be ineffective and honestly if someone is that determined, they can probably figure out where I live or work by other means”. Several participants felt that there was little harm in using a privacy zone (5%, $n=6$) and that they tended to feel safer as a result (5%, $n=6$).

There were $n=13$ participants that specifically mentioned privacy zones in their explanation but did not select privacy zones as a mechanism they would use. The most common reason ($n=7$) was that when sharing in private groups consisting of close friends and family, they were not worried about revealing sensitive locations that were already known to this group: “My friends know me, and I trust them” (P570). Some noted that privacy zones are simply ineffective ($n=6$) and so they should not be used: “privacy zones are too predictable to be of substantial benefit” (P234). Two participants were dissuaded by the relative complexity of privacy zones: “it is more difficult to set-up and to effectively communicate the concept of a privacy zone so I’m more inclined to use the easier and simpler options” (P382).

Key Takeaways Our findings indicate that participants were confident in their ability to identify and defeat privacy zones, and that this confidence grows as they are provided with more of the information needed to succeed in this task. This suggests that would-be attackers are likely to be aware when they have collected sufficient information to circumvent a privacy zone. In spite of this, however, most participants still considered privacy zones to be effective. It may be that participants were comparing privacy zones to no defense at all, found them to be useful as a complement to other precautions, or were evaluating them relative to the minimal impact

on user experience. This persistent belief in privacy zone efficacy may be an artifact of our survey design; this could be explored in future studies by employing a belief change paradigm in which participants were given explicit feedback about their performance.

When considering which privacy mechanisms to use for sharing exercise activities, many participants indicated that they would use a privacy zone. However, most would use them in combination with another mechanisms such as a private group. Surprisingly, some participants were willing to use privacy zones despite recognizing their flaws. Those that indicated they would not use a privacy zone specifically noted that private groups and selective sharing is probably enough, and that a privacy zone increases the complexity and may not actually offer additional benefits.

Several participants also highlighted concerns that were not addressed by privacy zones. Some participants were uncomfortable with the idea of others viewing any part of their exercise activity, not just the endpoints that privacy zones protect. Of particular concern were scenarios where the activity had occurred in a secluded area. Some participants also noted discomfort with sharing any identifiable fitness information. These concerns inform our discussion of the design space of additional privacy mechanisms in Section 6.2.

6 DISCUSSION & CONCLUSION

With regards to our qualitative findings (RQ1 & RQ3), we observed that while fitness app usage is widespread, many users and prospective users remain concerned about sharing exercise route information online. This finding is consistent with prior work on user perceptions of fitness [2, 13, 27, 33, 42, 43] and location sharing [8, 20, 32, 34, 35] which demonstrates that while users tend to better understand the implications and risks of sharing location better than other data types, they may still underestimate the degree of leakage. Indeed, we find that in cases when users feel comfortable, they do so because they employ the use of a privacy mechanism. In contrast to prior work, however, our study centers the efficacy of these privacy mechanisms by asking participants to reflect on the unique challenges posed by sharing location information on fitness apps. We find that participants’ privacy and safety concerns extend beyond the immediate area of their home, making privacy zones alone insufficient to protect users’ sensitive information. Participants describe protecting themselves with a combination of in-app and out-of-app behaviors to rectify these gaps in privacy protections. To the best of our knowledge, this is the first systematic investigation of real-world users’ perceptions of fitness app privacy mechanisms and their efficacy. These observations provide useful insights for the design of future privacy mechanisms.

The inference task presented in our work validates participants’ privacy and safety concerns — given just a few reference activities, everyday people can reliably deduce the locations hidden by privacy zones at a rate of up to 68%. Further, participants that were placed in a position to succeed in the task (i.e., with three routes) were confident in their ability to identify the hidden sensitive locations. This finding not only confirms the conceptual vulnerability identified in prior work [18], but also demonstrates that attackers with limited technical ability can confidently reason about and circumvent the privacy zones of real users. Participants expressed confidence in the

efficacy of privacy zones even after completing the inference task, in many cases indicating an increase in their comfort level. At the same time, though, participants noted that they were more likely to use privacy zones in combination with other privacy mechanism rather than relying on privacy zones alone, indicating that they are a “net positive” in spite of their limitations.

6.1 Implications for Current Fitness App Privacy Features

Our inference task was based on an implementation of privacy zones that was widely used until 2018. Following widespread reports that they leaked information [18], many companies changed their implementation to increase the difficulty of de-anonymizing user locations. For example, Strava introduced notions of spatial cloaking into their privacy zones, such that the hidden location was equally likely to appear anywhere within a hidden inner area of the circle [15]. Garmin chose a different approach, adding noise to the endpoints of the route that intersected with the privacy zone so that it would be more difficult to identify the exact boundary of the circle [18]. Privacy zones are continuing to evolve – this past summer (2021), Strava created a new feature that allowed participants to apply privacy zones to the start/end points of routes on an activity-by-activity basis [23, 38]. Our preliminary experiments with this new feature indicate that this new feature is based at least partially on their spatial cloaking approach.

In this study, we chose to make use of the original privacy zone implementation, such that many of our findings are not directly applicable to today’s privacy mechanisms. We chose to study the original implementation so that we could better answer the question of whether information leakage vulnerabilities in privacy zones could actually be harnessed by everyday people. Common criminals that are of far greater concern to the typical user, such as a burglar or stalker, are likely to be non-technical laypersons. While prior work has demonstrated conceptual privacy zone vulnerabilities, it has not demonstrated that these vulnerabilities could account for any of the actual crimes that have been linked to fitness app usage (e.g., [5, 39–41]). To the best of our knowledge, our work is the first to demonstrate that privacy zone information leakage could pose a real-world threat. Further, choosing the original implementation was also necessary to validate our experimental design. Information leakage vulnerabilities have not been demonstrated in the new privacy zones; had we used these mechanisms, a negative result in our inference task would have been inconclusive – was the mechanism secure, or was there simply a fault in our design? Having now validated our stimuli and study materials, we are interested in evaluating these new privacy zone models in future work.

In spite of not having directly analyzed these new privacy zones, some of our findings still carry significant implications for their use. To our knowledge, all fitness apps that offer privacy zones also attempt to hide the fact that a user has enabled them. This is because the presence of a privacy zone itself leaks information about the user’s fitness habits, while uncertainty about a privacy zone’s presence may even provide some level of protection for fully public posts. However, our results indicate that typical users can reliably determine the placement and size of privacy zones. This finding is directly applicable to the new implementations, suggesting that

fitness apps’ privacy model should not depend on the assumption that the existence of privacy zones can be hidden.

Second, our results indicate that users are able to identify and extract subtle cartographic information when attempting to identify protection locations. While we did not quantitatively evaluate this effect, our participants reported using a diverse array of strategies and information sources when attempting to circumvent privacy zones. These included taking note of residential areas, the directions that roads and paths were headed in, and physical boundaries like park space or bodies of water. Our post-hoc analysis of the best performing items in the 1-Route condition suggests that, even when participants were not put in a position to succeed in the task, they were still sometimes able to identify the hidden location by leveraging these information sources. Thus, even though newer privacy zone implementations incorporate notions of information theoretic privacy, the practical protection offered by these features may be greatly diminished because they do not account for the presence of additional cartographic information. A likely possibility is that privacy zones would offer different levels of protection for different individuals based on topology and urbanization levels is concerning. Ultimately, further work is required to determine the effects of topology and the level of protection offered by these features; the methods and stimuli presented in this work can be used to such ends.

6.2 Recommendations to Fitness Tracking Services

In addition to our quantitative measurement of privacy zone revealing conditions with low efficacy, our qualitative results indicate that even perfectly functioning privacy zones are not fully meeting the privacy needs of fitness app users. This implies that users may be interested in alternate privacy mechanisms. Inspired by participant-provided concerns, we consider the space of alternate privacy mechanisms that offer different trade-offs between sharing and safety.

First, several participants raised concerns that privacy zones do not provide safety or privacy during the majority of the exercise route. To avoid this, many participants reported selectively sharing routes they felt were sufficiently public and safe. Unfortunately, in today’s fitness apps this means that users would often be unable to take advantage of the community features such as leaderboards. Many fitness apps already include a “segment leaderboard” in which athletes race on known segments, but do not allow users to *only* share segments, or only share with a set of segments they deem safe. One possible solution would be to allow users to exclusively share route portions that are associated with a segment leaderboard. These segments generally fall on popular exercise paths and are also often brief, such that a user could still participate in the community while protecting the majority of their movements. By creating a safe sharing feature for leaderboard participation, users would have the option to exclusively share mobility data from public, populous spaces, leaking minimal personal information.

Second, several participants noted that *any* location information would make them uncomfortable. To accommodate such users, some fitness trackers already offer the ability to exclusively share aggregate exercise statistics, such as the total number of miles

completed. However, one can also imagine other precautions where users are able to share fine-grained properties of their activity without revealing the location. For example, fitness apps could allow users to place their exercise activity in a completely different area with similar properties. Similar to virtual rides on stationary exercise equipment, fitness apps could also offer users the option to resituate their exercise activity into a virtual world or exotic locale. Because such features would still share certain properties of the true exercise route, further investigation would be required to ensure they do not leak information.

Third, out of an abundance of caution for their privacy and worries of peer judgement, several users did not want to show *any* identifiable exercise information, including basic performance statistics. At present, the only privacy option for such users is to make all of their exercise posts fully private. In spite of their reluctance to share, however, participants still expressed belief that competition was beneficial. One possible way of reconciling these seemingly contradictory beliefs is to incorporate notions of differential privacy into fitness apps' community features. For example, via a percentile-based query system [10], users would be able to contribute to and discover their relative ranking on a segment leaderboard while formally bounding the possible leakage of personal information. This would allow for both complete anonymity and the benefits of fitness sharing, at least for the most popular and heavily trafficked segments.

More broadly, our results underscore the complexity and context-sensitivity of privacy risks when using fitness apps. To our knowledge, these nuanced concerns are not being communicated to users by fitness tracking companies. Beyond new privacy mechanisms, we believe these issues could be better communicated to users by fitness tracking companies. Our more privacy-savvy participants reported using their own ad-hoc solutions like waiting to turn their trackers on, providing fake information to the app, or sharing screenshots of activities via private message rather than posting them. We argue that it is in the long-term interest of fitness tracking services to actively communicate with users about risks to their privacy and safety, and also suggest the use of such possible management strategies. This messaging should be integrated into various aspects of app functionality. During normal interactions like posting a new exercise activity, fitness apps could check with users to make sure that the activity took place in a safe and populated area. While waiting for posts to upload, apps could make additional suggestions like exercising with friends, starting the tracker a few blocks away from home, or using larger privacy zones. In their privacy settings, apps should be forthright about the potential limitations of privacy zones. Explicit privacy features like privacy zones and private posts should also feature prominently in the user interface. In the past several years, fitness trackers have repeatedly been portrayed negatively in the media for enabling privacy disclosure [19, 22], harassment [28], and theft [5, 39–41]; engaging users in an active and ongoing dialogue about risks could help to restore and maintain trust in these services.

6.3 Recommendations to Fitness Tracking Users

While privacy zones are far from a perfect defense, our findings indicate that under certain circumstances they can provide added protection for users. Contrary to prior work [18], performance in our inference task indicates that large privacy zones provide considerably more security. When a $\frac{5}{8}$ mile privacy zone was applied, only 11% and 54% of guesses fell within 50 meters for the 1-Route and 3-Route conditions, respectively, as compared to 38% and 68% for the smallest privacy zone of $\frac{1}{8}$ miles. These reductions in confidence may be enough to deter a would-be attacker. We strongly recommend that users make use of large privacy zones, particularly in concert with additional layers of privacy precautions, such as sharing in a private group. For example, if a user is sharing with a private group of joggers that are known in their community, the risk of burglary may be low. That said, the user may still find comfort in the knowledge that a privacy zone keeps their home address from being broadcast to the entire private group. Users should also consider that the most effective privacy precautions may present themselves outside of the app itself. By bicycling a few blocks before turning on the fitness tracker, users can make it that much more difficult for others to learn their exact address. We observed participants noting this very tactic when using fitness tracking apps.

6.4 Conclusions

Fitness tracking apps allow users to connect with a broader community of athletes, but the very nature of sharing health and mobility data online raises serious privacy and safety concerns. In this work, we design and administer a novel interactive task that allows us to reason about the practical security offered by fitness app privacy features in real-world settings. Our results indicate that, indeed, it is possible for technically unskilled attackers to confidently identify and circumvent privacy zones, particularly for the original privacy zone implementation. Many of the strategies leveraged by participants, such as referencing cartographic information, suggest that even state-of-the-art privacy zone techniques may continue to leak hidden locations. In spite of this, participant responses suggest that privacy zones should continue to play an important role in helping to manage the risks of fitness tracker usage, especially as a complement to additional precautions. These findings illuminate promising future directions for the design and evaluation of fitness tracking privacy mechanisms.

ACKNOWLEDGMENTS

We would like to thank our reviewers for their valuable feedback, Wajih Ul Hassan for sharing and supporting our use of his Strava datasets, and Daniel Johnston, Dawei Wang, and Klaus Zou for their assistance in programming our survey software. This material is based upon work supported by the National Science Foundation under Grant Nos. CNS - 1951852 and 1955228, as well as the Graduate Research Fellowship Program under Grand No. DGE - 1746047. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of their employers or the sponsors.

AVAILABILITY

Our materials, data, and analysis scripts are available at
<https://bitbucket.org/sts-lab/epz-game-chi22>.

REFERENCES

- [1] Hazim Almuhiemedi, Florian Schaub, Norman Sadeh, Idris Adjerd, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 787–796. <https://doi.org/10.1145/2702123.2702210>
- [2] Abdulmajeed Alqhatani and Heather Richter Lipford. 2019. "There is Nothing That i Need to Keep Secret": Sharing practices and concerns of wearable fitness data. In *Proceedings of the fifteenth USENIX conference on usable privacy and security (SOUPS'19)*. USENIX Association, Santa Clara, CA, USA, 421–434.
- [3] App Annie. 2019. State of Mobile 2019 Report. <https://www.appannie.com/en/go/state-of-mobile-2019/>
- [4] Athletic.net. 2020. GPS Privacy Zones. <https://support.athletic.net/article/ghuixgcw6-gps-privacy-zones>
- [5] [Mrs Balls]. 2017. Strava and stolen bikes. <https://community.bikehub.co.za/topic/166972-strava-and-stolen-bikes/>
- [6] Dale J. Barr, Roger Levy, Christoph Scheepers, and Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68, 3 (2013), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- [7] Douglas Bates, Martin Maechler, Ben Bolker, Steven Walker, Rune Haubo Bojesen Christensen, Henrik Singmann, Bin Dai, and Fabian Scheipl. 2012. Package 'lme4'. CRAN. R Foundation for Statistical Computing, Vienna, Austria (2012).
- [8] Igor Bilogrevic and Martin Ortlieb. 2016. "If You Put All The Pieces Together...": Attitudes Towards Data Combination and Sharing Across Services and Companies. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 5215–5227. <https://doi.org/10.1145/2858036.2858432>
- [9] R. Chatterjee, P. Doerfler, H. Orgad, S. Havron, J. Palmer, D. Freed, K. Levy, N. Dell, D. McCoy, and T. Ristenpart. 2018. The Spyware Used in Intimate Partner Violence. In *2018 IEEE Symposium on Security and Privacy (SP)*. 441–458. <https://doi.org/10.1109/SP.2018.00061>
- [10] Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. 2013. Broadening the scope of differential privacy using metrics. In *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 82–102.
- [11] David Curry. 2020. Strava Revenue and Usage Statistics (2020). <https://www.businessofapps.com/data/strava-statistics/>
- [12] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. 2018. "A Stalker's Paradise": How Intimate Partner Abusers Exploit Technology. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174241>
- [13] Sandra Gabriele and Sonia Chiasson. 2020. Understanding Fitness Tracker Users' Security and Privacy Knowledge, Attitudes and Behaviours. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376651>
- [14] Garmin. 2021. Understanding Privacy Zones in Garmin Connect. <https://support.garmin.com/en-US/?faq=B9dIXYxQlr97DQwho5TBR7>
- [15] Matthew Gordon. 2018. Update to Privacy Zones Functionality. <https://medium.com/strava-engineering/update-to-privacy-zones-functionality-98a570f6ebb>
- [16] Grand View Research. 2021. Fitness App Market Size, Share & Trends Analysis Report By Type (Exercise & Weight Loss, Activity Tracking), By Platform (Android, iOS), By Device (Smartphones, Wearable Devices), And Segment Forecasts, 2021 - 2028. <https://www.grandviewresearch.com/industry-analysis/fitness-app-market>
- [17] Xinning Gui, Yu Chen, Clara Caldeira, Dan Xiao, and Yunan Chen. 2017. When Fitness Meets Social Networks: Investigating Fitness Tracking and Social Practices on WeRun. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (CHI '17). Association for Computing Machinery, New York, NY, USA, 1647–1659. <https://doi.org/10.1145/3025453.3025654>
- [18] Wajih Ul Hassan, Saad Hussain, and Adam Bates. 2018. Analysis of privacy protections in fitness tracking social networks: You can run, but can you hide?. In *Proceedings of the 27th USENIX conference on security symposium (SEC'18)*. USENIX Association, Baltimore, MD, USA, 497–512.
- [19] Alex Hern. 2018. Fitness tracking app Strava gives away location of secret US army bases. *The Guardian* (Jan. 2018). <https://www.theguardian.com/world/2018/jan/28/fitness-tracking-app-gives-away-location-of-secret-us-army-bases>
- [20] Predrag Klasnja, Sunny Consolvo, Tanzeem Choudhury, Richard Beckwith, and Jeffrey Hightower. 2009. Exploring privacy concerns about personal sensing. In *International Conference on Pervasive Computing*. Springer, 176–183.
- [21] Alexandra Kuznetsova, Per Bruun Brockhoff, Rune Haubo Bojesen Christensen, et al. 2015. Package 'Imertest'. *R package version 2*, 0 (2015).
- [22] Steve Loughran. 2018. Advanced Deanonymization through Strava. <http://steveloughran.blogspot.com/2018/01/advanced-deanonymization-through-strava.html>
- [23] Ray Maker. 2021. Strava Adds Major New Privacy Zone Features, Plus More Privacy & Map Options. <https://www.drainmaker.com/2021/08/strava-new-privacy-zone-features-options.html>
- [24] Megan McFadden. 2021. Privacy Zones. <https://support.strava.com/hc/en-us/articles/115000173384-Privacy-Zones>
- [25] Vivian Genaro Motti and Kelly Caine. 2015. Users' privacy concerns about wearables. In *International Conference on Financial Cryptography and Data Security*. Springer, 231–244.
- [26] Nick Mueller. 2018. Strava's privacy zone feature is almost worthless. <https://medium.com/@nickolaus/strasvas-privacy-zone-feature-is-almost-worthless-f47c5cebfece>
- [27] Mark W. Newman, Debra Lauterbach, Sean A. Munson, Paul Resnick, and Margaret E. Morris. 2011. It's Not That i Don't Have Problems, i'm Just Not Putting Them on Facebook: Challenges and Opportunities in Using On-line Social Networks for Health. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (Hangzhou, China) (CSCW '11)*. Association for Computing Machinery, New York, NY, USA, 341–350. <https://doi.org/10.1145/1958824.1958876>
- [28] Olivia Nuzzi. 2020. What It's Like to Get Doxed for Taking a Bike Ride. <https://nymag.com/intelligencer/2020/06/what-its-like-to-get-doxed-for-taking-a-bike-ride.html>
- [29] Liarna La Porta. 2018. Strava's privacy zones have been revealing users' hidden locations. <https://www.wandera.com/strasvas-privacy-zone/>
- [30] R Core Team. 2013. R: A language and environment for statistical computing. (2013).
- [31] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2019. How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples. In *2019 IEEE Symposium on Security and Privacy (SP)*. 1326–1343. <https://doi.org/10.1109/SP.2019.00014> ISSN: 2375-1207.
- [32] Stefan Schneegass, Romina Poguntke, and Tonja Machulla. 2019. *Understanding the Impact of Information Representation on Willingness to Share Information*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3290605.3300753>
- [33] Jeroen Stragier, Tom Evens, and Peter Mechant. 2015. Broadcast yourself: an exploratory study of sharing physical activity on social networking sites. *Media International Australia* 155, 1 (2015), 120–129.
- [34] Karen P. Tang, Jason I. Hong, and Daniel P. Siewiorek. 2011. Understanding How Visual Representations of Location Feeds Affect End-User Privacy Concerns. In *Proceedings of the 13th International Conference on Ubiquitous Computing* (Beijing, China) (UbiComp '11). Association for Computing Machinery, New York, NY, USA, 207–216. <https://doi.org/10.1145/2030112.2030141>
- [35] Karen P. Tang, Jialiu Lin, Jason I. Hong, Daniel P. Siewiorek, and Norman Sadeh. 2010. Rethinking Location Sharing: Exploring the Implications of Social-Driven vs. Purpose-Driven Location Sharing. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing* (Copenhagen, Denmark) (UbiComp '10). Association for Computing Machinery, New York, NY, USA, 85–94. <https://doi.org/10.1145/1864349.1864363>
- [36] Nick Tatt. 2014. Hide sensitive locations with privacy zones. <https://www.mapmytracks.com/blog/entry/hidden-sensitive-locations-with-privacy-zones>
- [37] Rannie Teodoro and Mor Naaman. 2013. Fitter with Twitter: Understanding Personal Health and Fitness Activity in Social Media. In *ICWSM*, Vol. 7.
- [38] The Strava Club. 2021. More Privacy, More Control and More Fun. <https://www.strava.com/clubs/231407/posts/17432797>
- [39] Alan Thompson. 2012. Lock your GPS data against bicycle thieves. <https://www.scarletfire.co.uk/lock-your-gps-data-against-bicycle-thieves/> Section: Cycling.
- [40] Alan Thompson. 2014. Strava Security: Why bike thieves are laughing at Strava users. <https://www.scarletfire.co.uk/bike-thieves-strava-security/> Section: Cycling.
- [41] Philippe Tremblay. 2018. Thieves allegedly use Strava to identify and steal cyclist's \$21,000 bike collection. <https://cyclingmagazine.ca/sections/news/thieves-allegedly-use-strava-to-help-steal-cyclists-21000-bike-collection/>
- [42] Jessica Vitak, Yuting Liao, Priya Kumar, Michael Zimmer, and Katherine Kritikos. 2018. Privacy attitudes and data valuation among fitness tracker users. In *International Conference on Information*. Springer, 229–239.
- [43] Michael Zimmer, Priya Kumar, Jessica Vitak, Yuting Liao, and Katie Chamberlain Kritikos. 2020. "There's nothing really they can do with this information": Unpacking how users manage privacy boundaries for personal fitness information. *Information, Communication & Society* 23, 7 (2020), 1020–1037.

A SURVEY MATERIALS

A.1 Pre-Task Survey

What are Fitness Applications?

Fitness apps allow users to track their activities, such as runs or bike rides, using wearable devices or mobile phones. Examples of popular fitness apps include Strava, Garmin Connect, Endomondo, MapMyRun, and Nike+. When a user is exercising, these fitness apps keep track of statistics such as their distance, pace, speed, elevation, and a map of their exercise route. When the user finishes, they can share their activity with friends or other athletes using the fitness app. If desired, they can even share the exercise to social networks like Facebook.

For example, Sam uses the Strava fitness app to track his exercise. When he wants to go on a run, he browses the public activities posted to Strava by other users and finds a route he's never been on before. He then tells Strava to start recording his run and leaves the house. When Sam gets back to the house, he checks to see how he did by comparing his performance to a public leaderboard. Finally, Sam posts the map of his exercise to Strava. Because Sam has connected his Facebook account to Strava, the run can also be seen by Sam's friends on Facebook.

Q1 Based on this description, have you ever used or currently use a fitness app to track your exercise?

- ☐ Yes ☐ No

Your Experience with Fitness Apps

Q2 If Q1 == Yes Which fitness apps have you used before? Select all that apply:

If Q1 == No Which fitness applications have you heard of before, if any? Select all that apply:

- ☐ Strava ☐ Endomondo
☐ Garmin Connect ☐ Map My Run
☐ Runtastic ☐ Nike+
☐ RunKeeper ☐ Other

Q3 If Q1 == Yes Have you ever shared a map of your exercise on a fitness app or social network?

If Q1 == No Have you ever seen someone share a map of their exercise on a fitness app or social network?

- ☐ Yes ☐ No ☐ Unsure

Q4 If Q1 == Yes How comfortable are you sharing a map of your exercise?

If Q1 == No How comfortable would you be sharing a map of your exercise?

- ☐ Very Comfortable
☐ Somewhat Comfortable
☐ Somewhat Uncomfortable
☐ Very Uncomfortable
☐ Cannot Say

Q5 Please explain your answer to the previous question:

Q6 If Q1 == Yes Do you take any steps to protect your privacy when using a fitness app?

If Q1 == No Are you aware of any ways someone could protect their privacy while using a fitness app?

- ☐ Yes ☐ No ☐ Unsure

Q7 If Q1 == Yes && Q6 == Yes Please describe the steps you have taken to protect your privacy when using a fitness app.

If Q1 == Yes && Q6 == No Please describe the steps to protect privacy you have heard about, but choose to not take when using a fitness app.

If Q1 == Yes && Q6 == Unsure Please describe the steps that you think people might take to protect their privacy when using a fitness app.

If Q1 == No && Q6 == Yes Please describe the ways one could protect their privacy while using a fitness app.

If Q1 == No && Q6 == No Please describe the ways you think might be possible for someone to protect their privacy while using a fitness app.

If Q1 == No && Q6 == Unsure Please describe the ways you think might be possible for someone to protect their privacy while using a fitness app.

A.2 Privacy Zone Inference Task


What are Privacy Zones?

Many fitness app users start and/or finish exercising at their home or place of work, which they may not wish to share online. Some fitness apps allow users to create a "Privacy Zone" that protects these sensitive locations.


To create a Privacy Zone, a user enters an address and then a virtual circle is placed on the map with the protected address at the center. Each user can still see the entire route of their exercise, but friends can only see a shortened version that appears to begin/end at the edge of the circle. These friends can't tell whether or not the Privacy Zone exists or which parts of the route are being concealed. Users can also choose the size of the Privacy Zone's circle depending on how much space they would like to protect around the address.

How Privacy Zones Work

The portion of an exercise that starts or stops inside of a Privacy Zone will be hidden.




If the user passes through a Privacy Zone during the middle of an exercise, this portion will **not** be hidden.



How to perform your task

You will be shown a map of exercises for a user that has set up a Privacy Zone around an address. The hidden Privacy Zone circle will be touching either the start of the route (the green circle), the end of the route (the red circle), or both. The protected address often exists near the center of the privacy zone. Using the exercise routes as a guide, your goal is to identify the Privacy Zone and protected address.

To learn how to complete your task, press the *begin tutorial* button.



Interactive Tutorial via Intro.js hints

- (1) During this section of the survey, you will be shown different maps such as the one below.

- (2) The maps will depict a user's exercise. Each map contains a hidden Privacy Zone whose boundary touches the start of the route, the end of the route, or both.
- (3) The maps will depict a user's exercise. Each map contains a hidden Privacy Zone whose boundary touches the start of the route, the end of the route, or both.
- (4) You can zoom in and out of the map via the "+ -" controls or with your mouse wheel.
- (5) You can look around by clicking on the map and dragging.
- (6) To place a Privacy Zone, select a circle size from the toolbar and click on the map. To resize, select a different circle.
- (7) To move a Privacy Zone, click on its center and drag it into the desired location on the map.
- (8) When you think you've found the Privacy Zone, click on the Red Pin icon to add it to the map.
- (9) Click and drag the Red Pin to the address you believe is being protected at the center of the Privacy Zone.
- (10) You can reset your selection by clicking the refresh icon.
- (11) Once you are happy with the placement of the circle and pin, press the green check mark to submit your guess.
- (12) You're now ready to practice with two guided exercises. Feel free to experiment with some guesses. Pop-up hints will guide you towards the right answer.
- (13) When you're ready to begin, click continue.


Guided Practice (Performed twice with two different maps)

Try to find the Privacy Zone and protected address! Here are some helpful tips:

- At least one of the visible endpoints touches the hidden privacy zone.
- Typically the protected location is in a developed area (by a building/street), not in a field or body of water.
- Use the street layout to see where the athlete might have been headed as they entered the Privacy Zone.
- Once the Privacy Zone is found, look at buildings near the center of the circle to determine a likely protected address.

If first map Hint: For this particular map, all endpoints touch the privacy zone!

If second map Hint: for this particular map, not all endpoints touch the privacy zone!



You Are About to Begin Your Task

Identifying the protected locations may be difficult, but you should always use your best judgment based on the context available to you.

In total, you will be presented with 12 maps. The yellow progress bar at the top indicates your progress on completing maps. *If you experience any issues, you can reload the page without any penalties.*

Technical Issues (After completing the interactive task)

Q8 Did you experience any technical issues while completing this task?

- ☐ Yes ☐ No ☐ Unsure

Q9 If so, Please Describe... (Write n/a if none)

A.3 Post-Task Survey

Q10 How often do you think you were correct when selecting the circle size for the Privacy Zone?

- ☐ Very Often
☐ Somewhat Often
☐ Somewhat Infrequently
☐ Very Infrequently

- ☐ Cannot Say

Q11 Please describe the strategies you used when selecting the circle size of the Privacy Zone:

Q12 How often do you think that you were correct when identifying the location with the pin?

- ☐ Very Often
☐ Somewhat Often
☐ Somewhat Infrequently
☐ Very Infrequently
☐ Cannot Say

Q13 Please describe the strategies you used to identify the pin's location:

Q14 What is the shape of a red ball?

- ☐ Blue ☐ Round
☐ Red ☐ Square

Q15 Based on your experiences with this task, how effective are Privacy Zones at protecting sensitive locations?

- ☐ Very Effective
☐ Somewhat Effective
☐ Somewhat Ineffective
☐ Very Ineffective
☐ Unsure

Q16 Based on your experience with this task, would you be more or less comfortable posting your exercises online?

- ☐ More Comfortable
☐ Slightly More Comfortable
☐ No Change In Comfort
☐ Slightly Less Comfortable
☐ Less Comfortable

Q17 Please explain why your comfort has or has not changed:

Q18 Do you think Privacy Zones negatively impact user experience when posting an activity?

- ☐ To A Great Extent
☐ Somewhat
☐ Very Little
☐ Not At All

Q19 Do you think Privacy Zones negatively impact user experience when viewing others' activity?

- ☐ To A Great Extent
☐ Somewhat
☐ Very Little
☐ Not At All

Q20 If you were to setup a Privacy Zone around an address, how confident would you be in its ability to protect that location?

- ☐ Very confident
☐ Slightly confident
☐ Slightly not confident
☐ Not confident

Q21 What is the color of a red ball?

- ☐ Blue ☐ Round
☐ Red ☐ Square

Q22 If you were to share an exercise using a fitness application, would you take any of the following steps to protect your privacy? (select all that apply)

- ☐ Set my profile to private (only friends can see my profile and activities).
☐ Set an individual post to private (only friends can see that individual post).
☐ Block specific users.
☐ Use a Privacy Zone to protect my home address or other privacy-sensitive location.
☐ I am unlikely to use any privacy feature
☐ Other

Q23 Please explain your choices to the previous question:

Please enter some demographic information

Q24 What is your age?

- ☐ 18-20 ☐ 40-44 ☐ 65-69
☐ 20-24 ☐ 45-49 ☐ 70-74
☐ 25-29 ☐ 50-54 ☐ 50-54
☐ 30-34 ☐ 55-59 ☐ Prefer not to disclose
☐ 35-39 ☐ 60-64 ☐ 75+

Q25 What is your gender?

- ☐ Female ☐ Prefer to Self-Describe
☐ Male ☐ Prefer not to disclose
☐ Non-binary

Q26 Where you live is best described as

- ☐ Urban ☐ Rural
☐ Suburban ☐ Prefer not to disclose

One More Thing

Q27 Please indicate if you've honestly participated in this survey and followed instructions completely. You will not be penalized/rejected for indicating 'no' but your data may not be included in the analysis:

- ☐ Yes
☐ No

Thank you for taking the survey

Return to MTurk and enter the following finish code to prove you've completed the work: XXXXXXXX.

B CODEBOOKS AND SATURATION GRAPHS




Figure 11: Saturation of unique codes for reasons for sharing comfort/discomfort *before* task [Q5] ($n=300$)




Figure 12: Saturation of unique codes for privacy measures taken by participants that use fitness trackers [Q7] ($n=227$, fully-coded)




Figure 13: Saturation of unique codes for strategies report by participants for zone size inference [Q11] ($n=300$)




Figure 14: Saturation of unique codes for strategies report by participants for location inference. [Q13] ($n=300$)




Figure 15: Saturation of unique codes for participants' change in comfort level following the inference task [Q17] ($n=300$)




Figure 16: Saturation of unique codes for reasons for privacy zone with traditional mechanisms use for participants' future exercise sharing [Q23] ($n=117$, fully-coded)

Table 4: Qualitative codebook: Reasons for sharing comfort/discomfort *before* task [Q5]. ($n = 300$)

Primary Code	Prim. Freq.	Secondary Code	Freq.	Description	Participant Sample
CONCERN	165	REVEAL LOCATION	99	Related to the location of the user	"I don't like the idea of people knowing exactly where I've been and therefore might be in the future."
		GENERAL PRIVACY	31	Related to any revealed data. Often centered around principals	"I value my privacy therefore I believe it's not anybodies business where I'm running or what I'm doing"
		SOCIAL ISSUES	25	Fear of judgement/annoyance from viewers	"I'm not very athletic so I would feel like I'd be judged for not being able to go very far"
		ONLINE	24	Related to public online postage	"I am very private about what information I share online"
		DISTRUST COMPANY	9	Distrust of company and their uses of data	"Companies like Facebook already surveil us too much. Why help them out?"
		UNKNOWN HARMS	5	Fear over what <i>might</i> be able to be learned	"I'm wary about giving away too much private information as it's hard to know how it can be used to harm me"
		VULNERABLE DEMOGRAPHIC	4	Concern related to the gender and/or race of the user	"As a woman I would feel hesitant to use these apps"
		HEALTH SENSITIVE	3	Related to health implications of revealed data	"I want to keep my health information confidential"
		LACKS SECURITY	2	Distrust of company's security statue	"I don't want a map of where I've been on social media, given their track record of data protection"
		DISTRUST SOCIAL NETWORK FRIENDS	2	Distrust of users in private network	"I'm very uncomfortable ... because some of the 'friends' I have on there I am not very close to and dont want them to know my whereabouts"
NO CONCERN	123	ACCESS CONTROL	41	Comfort given controlled sharing with known friends	"I'm sharing with my friends who already know my location so I'm fine with that"
		NOTHING TO HIDE	19	Did not believe exercise data was sensitive	"I don't do anything especially secretive. My life is wonderfully boring and I don't mind sharing what I do."
		SELECTIVE SHARING	13	Only shares based on behavioral and app usage alterations	"As long as it wasn't from my front door, it's probably fine"
		SAFE EXERCISE AREA	7	Only share/exercise in safe public spaces	"I think it would be okay because it would probably be ... a common place for people to run."
		FITNESS-BASED SOCIAL NETWORK	4	Comfortable sharing among other enthusiasts	"I don't usually share that sort of thing on other social media networks. I'm comfortable with people on strava seeing my strava map"
		SAFE DEMOGRAPHIC	1	Comfort related to the gender and/or race of the user	"I would not be bothered by privacy issues created by sharing a map based on my location and demographics"
		PRIVACY ZONE	1	Comfort given use of privacy zone	"If I share exercise from my home. I have the exact location blocked"
BENEFITS	46	CORPORATE USAGE ONLY	1	Unconcerned with corporate data usage	"I don't care about the privacy aspect of it though in terms of corporate data collection as that information is already accessible"
		SOCIAL SUPPORT	29	Motivation and support from other users	"I want others to see my progress, keep me motivated in reaching my goals, and to continue to keep me accountable"
		FUN	11	Joy in act of sharing	"I felt accomplished and happy when sharing a map of my exercise online with friends."
		TRACK FITNESS	6	Utility in the stored data	"Sharing a map of my exercise help me to track my exercise data."
		SHARE GOOD ROUTES	2	Benefits from communal sharing of good paths	"I would be surprised if anyone cared about my exercise other than sharing good hiking places with people."
		FITNESS BENEFITS	2	General fitness benefits	"I like the fitness apps, the fitness apps and exercises give me a tricks for fit body"
NO BENEFITS	29	–	–	No found benefits to sharing exercise data	"I just don't care to share this about myself"
NO OPINION	7	NOT APPLICABLE	5	Unfit to comment	"I don't use fitness apps, and I don't normally exercise outside of a gym"
		UNSURE	2	Generally unsure of their stance	"I only used it once, but I never shared a map of my exercise, so I don't know if I would be comfortable sharing my map"

Table 5: Qualitative codebook: User Privacy-Defensive Measures [Q7]. ($n = 227$)

Primary Code	Prim. Freq.	Secondary Code	Freq.	Description	Participant Sample
APP LEVEL	138	PRIVATE PROFILE	66	No posting/sharing of any data	"I don't share my location or post where I have been or where I'm going."
		PRIVATE GROUP	43	Posting/sharing to a private group of users	"I only share to a very limited audience. I try to set privacy options to restrict visibility of my information."
		SELECTIVE SHARING	11	Posting/sharing of data after manual inspection for sensitive info	"I make sure to review what it is that I am sharing and who I am sharing it with."
		OPT OUT	8	Opt out of sharing diagnostics and additional data to company	"When using a fitness app, I make sure to not have diagnostic data shared"
		NO AUTO SHARE	8	Disable autossharing data to social media profiles after exercise	"I make sure that it doesn't share my information with out it telling me."
		HIDE ROUTE	6	Post exercise but not geolocation data	"Anymore, I don't post locations of my runs but may post distances and times."
		HIDE TIME	4	Post exercise but not time of day	"I usually change a few details like what day or time I was there before I post."
		PRIVACY ZONE	3	Post exercise with privacy zone enabled	"I use privacy zones in Strava to block my home and other home addresses."
		DELETE DATA	1	Delete data after use	"After I finish the run I delete the data"
		STRICT APP PERMISSION	31	Enabling strict generic permissions on app	"Try to allow as few permissions as possible, idk whatever else I can do too"
PHONE LEVEL	60	LIMIT LOCATION PERMISSION	29	Disabling location permission permanently or intermittently	"I also turn off my location as soon as I'm done using the fitness app so as not to be continually tracked."
		LIMIT NETWORK PERMISSION	6	Disabling internet permission permanently or intermittently	"I turn off Internet. "
		BLOCK 3RD PART APPS	2	Only run verified applications	"I do not permit third party access to the Health App"
		CLEAR CACHE	2	Clear cache storage off app on phone	"I clear caches of that app regularly and hide my some personal activities by switching it off."
LIMIT SUPPLIED DATA	37	DON'T CONNECT SOCIAL MEDIA	16	Don't connect fitness account to social media	"I will switch off the external link activity like connecting with Social Profiles"
		ONLY SUPPLY BASIC INFO	14	Only supply required personally identifying info.	"I try as much as possible not to share any personal information like my place of residence."
BEHAVIOR CHANGE	23	SUPPLY FAKE INFO	13	Supply falsified personally identifying info.	"I use only my initials or a fake name on my profile."
		CHANGE RUNNING HABITS	9	Alter the running activity posted to social media	"I wouldn't map a run or walk if I was starting at my house because it can show where you live."
		READ POLICIES	6	Read privacy policy	"I also try to scan an app's privacy policy to see how the app uses and shares my data."
		READ REVIEWS	3	Read application reviews	"I research the apps before using them regarding privacy."
		OUT-OF-BAND SHARING	2	Share data with friends without using social media	"If I want to share it I will screenshot it and share it that way"
		STOP USING APP	2	Stop utilizing application	"I actually have stopped using fitness apps when I've found out their data has been hacked "
		SURVEIL SOCIAL MEDIA POSTS	1	Check social media to ensure nothing sensitive was posted	"I also randomly check my different social media accounts to make sure they aren't picking up anything from the app either."
		DON'T STORE SENSITIVE DATA	1	Don't store sensitive info on phone	"I keep privacy-sensitive things off of my phone usually."
		PHONE/APP AUTHENTICATION	10	Log into the phone/app via authentication (pin, pass, etc)	"I have secured my app using PIN so it can't be accessed without my authorizati"
		2FA	5	Log in via two factor authentication to app	"I have enabled two-factor authentication or 2FA"
SECURE STORAGE/TRANSMISSION	18	SECURE NETWORK CONNECTION	3	Only share/download data over trusted wifi/internet protocol	"ensur[e] that apps use HTTPS"
		ENCRYPT DATA	2	Encrypt data at rest	"and I use 2fa with iOS to encrypt iCloud backups of Health data"
		CHANGE PASSWORD OFTEN	1	Change password to phone/app often	"I also change my passwords often."

Table 6: Qualitative codebook: Features used in zone size [Q11] and location [Q13] inference strategies. ($n = 300$, $n = 300$)

Primary Code	Prim. Freq. (Q11,Q13)	Secondary Code	Freq. (Q11,Q13)	Description	Participant Sample
MAP	179, 203	RESIDENTIAL	162, 188	Use of buildings or building layouts	"I just looked for the populated areas and used that as a big reference point, working out from there."
		ROAD	37, 40	Use of shown paths/roads	"I also tried to look at where the roads/paths went and pick something that made sense with the route..."
		PHYSICAL BOUNDARY	7,3	Exclusion of areas based on open water, green space, or other impediments	"I tried to fit the circle between boundaries of highways, parks, open land, and water."
ROUTE	149, 47	ENDPOINTS	130, 26	Use of endpoints	"Usually I look at the gap between the beginning and end points and see if it looks like a circle can fit between them."
		PREDICTION	27, 25	Attempt to predict a continuation of the route	"I tried to imagine the route the runner would take"
		LENGTH	2, 1	Use of route length	"I did account for the distance run, that is I was more willing to look at the larger circles on bigger maps"
		AVOID OVERLAP	4, -	Ensure circle is not placed over visible route	"Also I would make sure the circle didn't go over the outlined path because then it shouldn't be there."
CIRCLE SIZE PREF.	17, -	SMALL	13, -	Preferred the smallest viable circle	"I tried to use the smallest circle that would still touch both endpoints."
		MEDIUM	1, -	Preferred the middle circle	"if [the endpoints] are at the same place, [I] picked the middle one"
		LARGE	3, -	Preferred the largest viable circle	"I grabbed the largest circle and it went into some non-residential area or into the water, I went down in size and adjusted as needed."
CENTER OF CIRCLE	-, 157	-	-	Guessed location was in the center of the circle	"I tried to line the circle up so there were neighborhoods at its center or homes at its center."
MAP ZOOM	4, -	-	-	Used the zoom of the map to determine circle size	"Based on how zoomed out from the map I was."
INSTRUCTIONS	20, 7	-	-	Only recalled following tutorial / instructions	"I tried to go off of the original instructions as close as possible."
RANDOM/INTUITION	6, 12	-	-	No relayed strategy, random guessing and hunches	"totally random. could not ascertain a strategy"

Table 7: Qualitative codebook: Reasons for sharing comfort change *after* task [Q17]. ($n = 300$)

Primary Code	Prim. Freq.	Secondary Code	Freq.	Description	Participant Sample
PRIVACY ZONE EFFICACY	156	ROBUST	84	Viewed privacy zone obfuscation as robust	"I would be comfortable as I know that I can hide my address from those online."
		WEAK	62	Viewed privacy zone obfuscation as weak	"They can find where I live within very close proximity. If they know my face, it would be very easy to find which is scary."
		VARIABLE	12	Viewed privacy zone obfuscation as dependent on the topological context	"If you live somewhere remote like I do, where developments are more sparse, that makes it easier to figure out where someone is."
PRIOR PRIVACY BELIEF	61	HIGH PRIVACY CONCERN	39	Strong prior privacy-conscious beliefs carried throughout	"I already did not want to post my location online and the level of privacy from the zones didn't change that."
		LOW PRIVACY CONCERN	22	Strong prior insensitivity to privacy concerns carried throughout	"I'm just not all that concerned about privacy and this task didn't change anything"
OTHER CONCERN	42	EXPOSED ROUTE PORTIONS	28	Exposed routes could lead to harm	"Someone could still use common routes you take to do something to you. They don't need your address to stalk you, or set up a place to rob you "
		OTHER METHODS OF ATTACK	8	Other easier ways to find home location	"If someone really wanted to find me they probably could do it with our without me posting my exercise route."
		DISTRUST COMPANY	4	Data still available to company	"I would be more concerned about the app maker ... There is nothing presented here that suggests they would not have access to location information when using privacy zones."
		SOCIAL ANXIETY	2	Social concerns over sharing still exist	"I wouldn't have been comfortable posting my exercises online before doing this experiment, and that was for social reasons, unrelated to privacy."
		USAGE IMPLICATIONS	1	Implication of wrong-doings by using privacy mechanisms	"This means I'm up to something. Why would I make a privacy zone. Wife would think I'm cheating."
PRIVACY ZONE UNNECESSARY	30	NOT UTILIZE APP	19	Would not utilize app regardless	"Posting my activities isn't a privacy matter... I don't post because it's not that interesting to me."
		PRIVATE GROUP IS ENOUGH	8	Private groups and limiting access to known people provide needed privacy	"I limit my social media to close friends and family anyways. I would feel comfortable with them knowing where I live."
		OUT-OF-BAND DEFENSE	2	Uses other method to provide needed privacy/safety	"There would be safer ways like working out with friends or going to a gym."
		SELECTIVE SHARING	1	Ad-hoc verification of privacy before posting provides needed privacy	"I think I'm already pretty good at protecting my privacy before I upload anything to the Internet"
UNSURE	2	WANT TO LEARN MORE	1	Need better understanding before judgement	"I would have to study more to be comfortable"

Table 8: Qualitative codebook: Reasons for using/not using privacy zones [Q23]. ($n = 117$, $n = 13$, $n = 34$)

Primary Code	Prim. Freq. (PZ, NO PZ, NONE)	Secondary Code	Freq. (PZ, NO PZ, NONE)	Description	Participant Sample
IN COMBINATION WITH OTHER DEFENSES	75, 0	PROVIDES EXTRA PROTECTION	17, 0, 0	Additional protection when applied with other defenses	"I would still use the privacy zone because it would be like double protection."
		PROTECTION AGAINST LEAKAGE	8, 0, 0	Protects from against accidentally disclosure outside of private group	"I would still use privacy zone just for added protection in case my friends share my information/maps with others I don't know."
		DISTRUST FRIENDS	3, 0, 0	Protects against lesser trusted social media friends	"I wouldn't want everyone (even if we're friends on social media) to know where I live"
EFFECTIVE	47, 0, 0	PROTECT HOME	24, 0, 0	Effective in protecting a home address	"Yes, I would use a privacy zone to help me to protect my home address."
		INCREASES EFFORT	1, 0, 0	Increasing effort is an effective deterrent	"I'd still use a privacy zone, because someone would have to be specifically determined to figure out my location in order for the zone to be ineffective"
INEFFECTIVE	10, 6, 0	BETTER THAN NOTHING	5, 0, 0	Imperfect, but better than an unprotected route	"Although privacy zones aren't perfect ... I do think that an additional level of security is never a bad thing."
		RECORD AWAY FROM HOME	1, 0, 0	Would not typically endanger home address with routes	"this particular feature is not very useful to me because I usually run at a location away from my home"
		DON'T CENTER ON HOME	1, 0, 0	Requires shifting the center to to use properly	"If I use privacy zone, I should ... set up privacy zone in odd shape or not to focus middle of circle as my house"
PRIVATE GROUP IS ENOUGH	1, 7, 1	-	-	Private posts within a group is enough protection	"If you are a friend or someone I know already, then I am already comfortable with you knowing where I live"
NO HARM IN USE	6, 0, 0	WOULD TRY	3, 0, 0	Would attempt to use out of curiosity	"I have been convinced of using the privacy zone, or at least trying it"
		NO REASON TO NOT USE	1, 0, 0	No disadvantages to use	"I figure why not use it? If you're going to go out of your way to post your route you might as well try and protect yourself to some degree."
		NOT CONCERNED ENOUGH FOR PRIVATE	1, 0, 0	Private settings are too extreme for concern	"I would put up privacy zones only. I don't really care about others looking at my route."
		PRIVATE AFFECTS USER EXPERIENCE	1, 0, 0	Private settings detract from apps purpose	"privacy zones are the perfect midway point for me to feel comfortable while still being able to use social media the way that is even worth using in the first place"
FOR PUBLIC POST	5, 0, 0	-	-	Would use for non-private posts	"If I did want to share with others than close friends and family, I would definitely use the privacy zone to safeguard my location."
FEEL SAFE	5, 0, 0	-	-	Provides a feeling of comfort	"I will feel safer with the use of privacy zone"
NO WORRIES	0, 0, 24	-	-	No concerns over public posting of data	""
HARMFUL	0, 2, 3	COMPLEX/TOO MUCH WORK	0, 2, 2	Too confusing or too much effort to use	"I just think that a lot of times privacy features are more hassle than they are worth"
		ENJOY MEETING PEOPLE	0, 0, 2	Want to interact and meet strangers	"I think it might be fun to interact with others, Maybe I will meet a new friend or something."
		PZ TRUNCATES ROUTE	0, 0, 1	Unsatisfied with obfuscation shortening route	"I don't like that it makes my route look shorter, sometimes much shorter. I want a tracking of all I did."
FUTILE	0, 0, 1	-	-	Futile to try to protect one's privacy	"It really makes little difference we can be tracked through our phones"
WOULD NOT POST	0, 0, 4	-	-	Would not post to begin with	"I wouldn't want to share my exercises online ever."
PREFERENCE	0, 0, 1	-	-	Preference to use/not use as dominant reason	"I am unlikely to use privacy zone"