Distinguishing Group Privacy From Personal Privacy: The Effect of Group Inference Technologies on Privacy Perceptions and Behaviors*

JENNIFER JIYOUNG SUH, Department of Communication, UC Santa Barbara, USA
MIRIAM J. METZGER, Department of Communication, UC Santa Barbara, USA
SCOTT A. REID, Department of Communication, UC Santa Barbara, USA
AMR EL ABBADI, Department of Computer Science, UC Santa Barbara, USA

Machine learning and data mining threaten personal privacy, and many tools exist to help users protect their privacy (e.g., available privacy settings on Facebook, anonymization and encryption of personal data, etc.). But such technologies also pose threats to "group privacy," which is a concept scholars know relatively little about. Moreover, there are few tools to address the problem of protecting group privacy. This paper discusses an emerging class of software applications and services that pose new risks to group privacy by revealing group-level information based on individual information, such as social media postings or fitness app usage. The paper describes the results of two experiments that empirically establish the concept of group privacy and shows that it affects user perceptions of and interactions with information technology. The findings serve as a call to developers to design tools for group privacy protection.

Additional Key Words and Phrases: Privacy, social media, mobile apps, trending topics

ACM Reference Format:
Jennifer Jiyoung Suh, Miriam J. Metzger, Scott A. Reid, and Amr El Abbadi. 2018. Distinguishing Group Privacy From Personal Privacy: The Effect of Group Inference Technologies on Privacy Perceptions and Behaviors. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 168 (November 2018), 22 pages. https://doi.org/10.1145/3274437

1 INTRODUCTION

The uproar over the collection and use of Facebook data by Cambridge Analytica highlights a growing area of development in machine learning and data mining, namely the use of algorithms to profile individuals in ways that reveal information beyond what the person explicitly shared [20, 37]. Such algorithms have sparked new concerns about personal privacy in the age of big data. However, the risk to personal privacy is only one type of threat. Attention to safeguarding personal information overlooks problems caused by algorithms that extract group-level information [38]. Indeed, advances in data processing are increasingly focused on groups, classes, or subpopulations rather than individuals as the primary objects of value. Floridi offers this example [15] p. 98: *"Both*

*This work is partially supported by NSF Grant CNS-1649469.

Authors' addresses: Jennifer Jiyoung Suh, Department of Communication, UC Santa Barbara, Santa Barbara, CA, 93106, USA, suh@ucsb.edu; Miriam J. Metzger, Department of Communication, UC Santa Barbara, Santa Barbara, CA, 93106, USA, metzger@ucsb.edu; Scott A. Reid, Department of Communication, UC Santa Barbara, Santa Barbara, CA, 93106, USA, scottreid@ucsb.edu; Amr El Abbadi, Department of Computer Science, UC Santa Barbara, Santa Barbara, CA, 93106, USA, elabbadi@ucsb.edu.

2573-0142/2018/11-ART168 $15.00
https://doi.org/10.1145/3274437

friendly and hostile users of big data may not care about Alice at all, but only about the fact whether Alice, whoever she is, belongs to the group that regularly goes to the local church, or mosque, or synagogue, uses Grindr, or has gone to a hospital licensed to carry out abortions, or indeed shares [some other] feature of your choice. In military terminology, Alice is hardly ever a High Value Target, like a special and unique building. She is usually part of a High Pay-off Target, like a tank in a column of tanks. It is the column that matters." By collecting individual-level data, even anonymously, group-level information becomes available and carries with it novel threats to *group* privacy, as well as challenges to its protection.

Algorithmic profiling poses a range of potential privacy threats to groups and their members that could result in negative consequences for the group, including group infiltration, dissolution, discrimination, or annoyance for individuals whose group memberships are misclassified. Despite these threats, conceptualizations of privacy remain mostly at the individual level, as do legal and technical remedies to protect privacy [6, 14]. Scholars have only recently recognized the need for greater clarity concerning group privacy and its social, legal, and ethical implications [37]. This paper thus examines how emerging algorithmic inference technologies both create groups and threaten privacy, whether individuals recognize distinct personal and group privacy concerns, and whether these concerns affect user evaluations of and privacy-related behavior toward such technologies. Ultimately, this research can provide information that aids software engineers in the design of mechanisms for group privacy protection.

2 DIFFERENTIATING INDIVIDUAL AND GROUP PRIVACY

2.1 Existing Conceptualizations of Privacy and of Privacy Protection

Current conceptualizations of privacy and data protection are focused on the prevention of personal harm [15, 25, 37]. As a consequence, privacy is usually defined in terms of *individuals'* control over their information [42]. The majority of theory and research has focused on privacy risks to individuals and ways to help them manage those risks [6]. For example, individuals control their personal information by limiting their disclosure to online retailers and social network sites, or by using available privacy settings to adjust the size of the audience that can view their posts [23, 24, 26]. Prominent social science theories of privacy such as the privacy calculus model, behavioral economics, and most models based on social exchange theory examine how individuals weigh the benefits and risks of disclosure, and thus all strive to understand privacy exclusively at the individual level [4, 5, 30].

Legal and technological mechanisms to protect privacy in online contexts have similarly focused on the individual level, as illustrated by 'informed consent,' 'anonymization,' or 'differential privacy' approaches to privacy protection. Informed consent relies on individuals to protect their interests by legally mandating data gatherers to ask individuals to affirmatively cede control of their personal data prior to its collection. Anonymization of individuals' data records is a common technical means of privacy protection, as is allowing individuals to manage their own data privacy through making privacy settings available to users of apps and services. Differential privacy protection works through the injection of noise into large datasets that makes it impossible to trace any observed result back to a specific individual [13]. While these approaches can help protect individual privacy, none of them necessarily protect against violations to group privacy based on aggregation of individuals' information.

Nonetheless, some definitions have acknowledged group privacy [5, 41]. However, even conceptualizations that include the idea of group privacy define it in terms of individual privacy—i.e., as the collection of group members' individual privacy rights, rather than viewing group privacy as an autonomous entity [15, 25, 37]. This is flawed because data technologies that categorize people

into groups or monitor group activities without their knowledge "raise questions that go beyond the level of individual privacy harm" (p. 223) [39]. As such, strategies to protect individual privacy may not always be helpful in protecting group privacy.

Existing privacy protection remedies based on individual identifiability are indeed ineffective when the goal of an attacker is to identify or profile a group (e.g., an ethnic minority, a political network, or a group of people who consume the same product) rather than to identify individuals [21, 36]. Identifying individuals is not necessary for group profiling to occur, and people may be acted upon in harmful ways without their identity being revealed. For this reason, scholars are beginning to recognize that privacy problems raised by group-level inferences produced from algorithmic aggregation of individuals' data are different from those that arise when individuals are made personally identifiable [36]. They are calling for new privacy solutions that are not exclusively based on individual privacy rights [15], and to broaden conceptualizations of privacy to include groups' need to safeguard their collective privacy and data protection rights [38].

2.2 Algorithmically-Determined Groups and Group Privacy

Another important consideration in understanding group privacy is how emerging data technologies complicate the very notion of what is a "group." Traditional definitions conceive groups to be socially determined, meaning that members (and often outsiders too) are aware of the group's existence as well as their own membership status. But advances in data mining and machine learning are stretching such definitions, because groups can now be algorithmically determined. By extracting subsets from individual-level information or classes of similar individuals based on common habits and characteristics, technology can itself discover or "create" groups that may have consequences for members [21, 37]. But more than this, Floridi argues that technology can also "design" groups according to how the data gatherer specifies the algorithm. One tweak of the algorithm can produce a different group [15].

Algorithmically-determined groups thus tend to be more fluid, dynamic, and ad hoc, as opposed to socially-determined groups which tend to be more stable.[1] Also, while socially-determined groups are usually "self-proclaimed and self-aware" [21], meaning that the group is recognized by society (e.g., Rotary Club, Black Lives Matter, etc.) and that group members know they are members, this may not be true for algorithmically-determined groups. In such groups, people might or might not know that a group exists and/or that they have been classified as a member of the group. Kammourieh and colleagues describe this difference between groups that are determined socially versus algorithmically as "active" versus "passive" groups [21].

The question arises, then, as to whether algorithmically identified groups, including passive groups, have a legitimate claim to privacy. This is a matter of debate [37]. Some perspectives from social psychology suggest that "entitativity," which is a perception of the extent to which a collection of people is perceived as a group, is a necessary condition for groups to have attitudinal and behavioral significance for people at all [8]. This suggests that clusters of individuals identified via an algorithm may not generate group privacy concerns insofar that the individuals are not perceived as constituting a psychological group. However, other social psychological research from the minimal group paradigm shows that mere categorization on an ad hoc basis reliably produces group identification and discrimination against outgroup members [35]. This suggests that algorithmically-determined groups may indeed produce group privacy concerns for people, and thus warrant claims to group privacy rights. Kammourieh and colleagues moreover argue that just as active groups are impacted by the perception and treatment of the group by the rest of

[1]This does not imply that ad hoc groups are only limited to those that are determined algorithmically. For example, we might speak of the collection of people who are waiting at the same bus stop at the same time as a "group."

society, passive groups may also be treated as a group by society, making such groups "actionable" by opening them to targeting by third parties, and thus vulnerable to privacy attackers such as marketers or law enforcement agents. And this is true even in cases where group members are unaware that they are members of the group [21] (see also [15, 37]).[2]

2.3 Ways that Group Inference Technologies Threaten Privacy

The extent to which group privacy exists and is distinct from individual privacy are unsettled questions. One reason for this is that the concept of group privacy necessarily involves the privacy rights of individual group members and the privacy rights of the group as a whole. Taylor and colleagues refer to this as the difference between "their privacy" and "its privacy" [38]. This further implies that individual and group privacy are overlapping concepts. We argue, however, that they are separable as demonstrated by the fact that while anonymization of individual-level data may protect "their privacy," it does not protect "its privacy."

Group inference technologies threaten both "their privacy" and "its privacy" in a number of ways. At the individual level, two kinds of threats arise from inferences made about membership in algorithmically-determined groups. The first threat derives from accurate inferences of individuals as group members and happens when a group member is publicly associated with a group against their will (i.e., *true categorization threat*). The second type of threat occurs as a result of inaccurate inferences, for example, when someone who is not part of a group is incorrectly associated with the group (i.e., *false categorization threat*). Both types of threat can have negative consequences for individuals. For example, an inferred group membership—even when correct—can be harmful to an individual if the group is the target of harassment (e.g., LGBT teenagers), and wrongful association with a group can have similar negative effects for individuals. In addition, individuals are disadvantaged in terms of protecting themselves against such threats if they are not aware that they have been algorithmically classified as a member of the group in the first place.

Group-level privacy threats may arise from things that expose the group as a whole to outside influence and possible harm. Legally-recognized group privacy interests include things like the right to keep secrets from outsiders (e.g., family or military secrets, lawyer-client or doctor-patient privilege) and freedom from discrimination based on group members' shared traits (e.g., race, religion). But anything that negatively impacts the group or that shifts control of the group away from group members can be considered a threat to group privacy. For example, third party access to information that can lead to group dissolution, unwanted discovery, infiltration, hostile takeover, or harassment of the group as a whole or of its members all pose group privacy threats. Group profiling through data mining and machine learning may facilitate such things, with harmful consequences for the group and its associated members.

These harms apply to both active and passive groups. In active (self-aware) groups, information may be known only to ingroup members, such as the group's existence, location, or identity of members, that are essential to their functioning (e.g., military groups) or identity as a group (e.g., a religious sect's secret rites). As such, revealing group information to outsiders may be harmful to the group. Even when some group information is not secret (e.g., the group's existence), revelation of other information inferred about it via technological means (its location) could be dangerous for

[2] Another objection to the idea of group privacy is that U.S. privacy law says that any claim to privacy is forfeited for behavior done 'in public' (i.e., is observable by others). Under this logic, all group behavior is public because it is observable at least by other group members. In the context of social media this means that as soon as information is posted publicly it cannot be considered private. Nissenbaum [27] argues that this perspective is anachronistic since the networked structure of social media can push public information to audiences beyond those it was intended to reach, thereby potentially creating a privacy breach. Because of this, she maintains that it is fair to discuss privacy and private behavior in the context of public information and she calls for updating privacy law to include the need to protect 'privacy in public.'

the group; a group of political dissidents is an example. Likewise, discovering and exposing the latent collective traits and characteristics of passive groups (their existence, demographic makeup, and/or location) can endanger the group by opening it to surveillance by police or corporations. It is easy to imagine people who engage in seditious speech anonymously online getting grouped and then targeted for government monitoring as a result of group-inference technologies, or a collection of people who share demographic characteristics receiving annoying marketing messages as a consequence of being algorithmically identified. And because members of passive groups are not aware that a group has been identified, they have no ability to protect the group from unwanted intrusions.

3 EXEMPLARS OF GROUP INFERENCE TECHNOLOGIES THAT THREATEN PRIVACY

This section describes two examples of specific group inference technologies that may threaten individual and group privacy. The first example pertains to active (self-aware) groups and the second to both active and passive (non self-aware) groups.

Example 1: Strava. The conceptual difference between individual and group privacy is perhaps best illuminated in instances where group privacy is violated while individual privacy is protected. Such an instance is illustrated by a recent news story about the fitness app Strava [1, 3]. Strava is a fitness app with social features that collects user data anonymously and thus protects individual users' privacy. However, Strava aggregates user data to produce and publish "heatmaps" that show where users exercise. Although these heatmaps cannot reveal information about an individual user, they can reveal popular exercise locations that are frequented by multiple users. Based on this, global heatmaps published by Strava were shown to expose the outlines of both known and secret military bases around the world in countries such as Afghanistan, Iraq, and Syria [3], with ominous national security implications [2]. Although Strava prevents heatmaps from being used to track the location of individuals by anonymizing user data, data from individuals can be used to track group movements in ways that threaten group privacy. This example shows that while aggregating anonymized individual data can protect individuals' information, such data still have privacy implications for groups that are identified or profiled by the technology.[3]

Example 2: Community-Aware Trending Topics Analysis. Another example where group and individual privacy may be threatened concerns group inference technologies that are applied to social media platforms. One instantiation of such technologies is based on "Trending Topics" analysis that is used on popular social media platforms. Trending topics analysis extracts the most popular topics that people are discussing on a social media platform and can be narrowed to what people in specific geographical locations are talking about. Recent work is extending trending topics analysis to identify a spectrum of social characteristics of discussants beyond location (e.g., age, gender, political affiliation, etc.) based on information that individuals in a social network share. For example, Community-Aware Trending Topics (CATT) works by correlating topics discussed in Twitter with the publicly-available personal data of individuals who post on a topic [16]. An important aspect of this is that the postings are made by independent individuals and not part of a collective or group effort, and yet posters are grouped together algorithmically. As a result, detailed

[3]Of course, individuals are also threatened in this case. If a military group is revealed, this puts individual members of that group at risk. So, by threatening group privacy, individual group members (e.g., soldiers stationed on a revealed military base) are threatened. In this way, Strava threatens both "its privacy" and "their privacy" rights. This example highlights flaws with current approaches to privacy protection that focus solely on the individual level. By only protecting personal privacy, we are not protecting group privacy, and by revealing group privacy we are compromising personal privacy. The implication is that we need to identify and protect both individual and group privacy in order to ultimately protect the individual.

aggregated group information, including the group's latent collective characteristics and traits, may become obtainable by third parties.

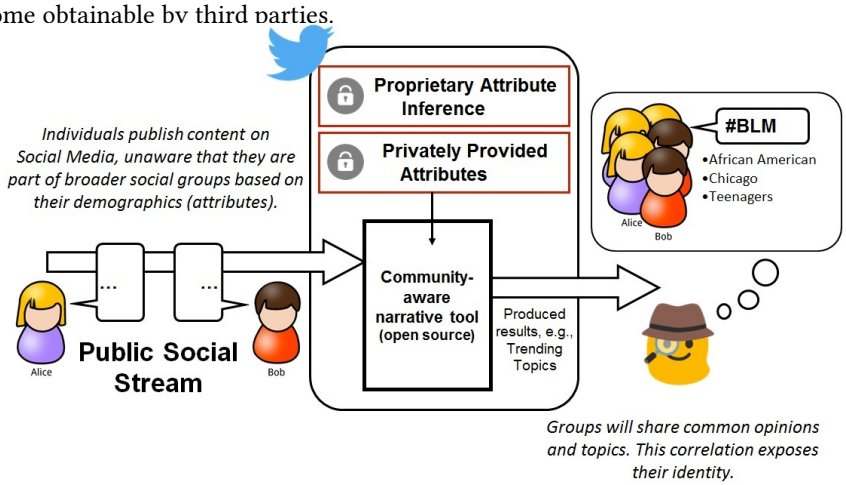

Fig. 1. Visualization of privacy attack model using CATT: an attacker can identify the demographics of individuals discussing certain trending topics, e.g., #BLM = Black Lives Matter

As illustrated in Figure 1, the public social media discussions of autonomous individuals can be combined with knowledge of their characteristics available from the public portion of their social media profiles to produce information that an attacker can reverse engineer to discover latent collective behavior and characteristics of groups. In this way, group privacy can be violated by the exposure of non-obvious collective behavior through the aggregation and mining of publicly visible actions of individuals in social media environments. Moreover, this information can allow attackers to draw inferences about group membership for people who share demographic and/or geographic characteristics with the discovered group, but who do not themselves necessarily post on the topic. This may create privacy problems at the individual level through either true or false categorization threats (e.g., correctly or incorrectly inferring that a person is homosexual).

The above scenarios apply to passive groups whose members lack awareness of being grouped at all, but even members of active groups are vulnerable to privacy violations stemming from group profiling technologies such as CATT (e.g., a political dissident group whose members are aware of each other but who post anonymously on social media to avoid government persecution). In sum, the Strava and CATT examples show how recent advances in data mining and machine learning threaten the privacy rights of individuals *and* groups, and thus underscore the need to study privacy at both levels, with the aim of developing mechanisms to protect not only individual but also group privacy.

4 RELATED WORK

A few researchers have examined group privacy (e.g., [6, 10, 31, 32, 42]) or group dynamics that challenge privacy management [9], but existing empirical research has been mostly qualitative studies with small samples. Theoretical work on group privacy is limited to how people in groups manage privacy rather than to understanding the concept of group privacy, if it differs from individual privacy, and how group privacy affects human behavior for both active as well as passive groups, which are the aims of the present study.

4.1 Previous Theoretical Work on Group Privacy

Based on Altman's conceptualization of privacy management as a bidirectional and social process [5], Petronio developed Communication Privacy Management (CPM) theory [30]. This theory says

that people manage their privacy through interpersonal boundary coordination and collaborative negotiation. A central argument of the theory is that once a person shares his or her information with others, they co-own that information and must collectively manage the privacy boundary around that information through agreeing on rules about information disclosure. This theory was developed for interpersonal communication contexts and thus assumes co-owners of the information are aware of each other's co-ownership. As such, this theory helps explain how people collaborate with others to manage information that individuals in active groups disclose, but not information about groups themselves (e.g., socio-demographic characteristics, group activities, etc.).

Nissenbaum's theory of contextual privacy also pushes privacy theorizing beyond the individual level by conceptualizing privacy as a collective property that is negotiated between people, in this case through social norms for information sharing that are situated in a specific context [28]. This theory draws attention to how online contexts require new ways of thinking about privacy protection that go beyond defining privacy as individuals' control over information about the self, or in terms of information that is secret (known only to the self). While this perspective is helpful in propelling scholarship towards the concept of group rather than individual privacy, the theory does not address group privacy directly.

4.2 Previous Empirical Studies of Group Privacy

Little empirical research has examined group privacy. The studies that come closest are based on theories that view privacy management as an interpersonal boundary management process, and accordingly examine how individuals engage in collective privacy management in the service of their groups. For example, Cho and Filippova found that people engage in collaborative strategies when sharing information pertaining to multiple individuals (e.g., passing the phone around to approve posting a group photo; see also [19]) [10]. However, the goal of using these strategies is successful privacy management of individuals in the group, not the group as a whole. These strategies help individual group members make disclosure decisions based on the potential harm to themselves using collective input from other group members.

Recently, a few studies have examined how people manage information about self-aware groups. A study by Lampinen shows how household members coordinate their interior and exterior privacy boundaries when hosting others in their homes, for example by delegating specific members to interact more with visitors [22]. In a study with members of Flemish youth organizations, DeWolf [11] found participants protected the privacy of group information by managing their Facebook content together (e.g., using Facebook groups to limit access to group information by outsiders, collectively choosing group photos to upload, deleting information that is harmful to their groups; see also [12]). While these strategies are geared towards managing group reputation or secrets (within their families or organizations), these studies do not discuss passive (non-self-aware) groups' privacy that is threatened by profiling technologies based on big data analytics.

Some prior work has also focused on designing technological mechanisms to support interpersonal and collaborative privacy management. Besmer and colleagues developed a tool to help users decide whether to share information based on the percentage of people in a user's own network who had shared similar types of information on their profiles [7]. Patil and Lai found that people prefer defining privacy permissions at the group level when using an awareness application at work [29]. Others have developed prototypes for tools that enable multiple people to control photos that include more than one person [18, 32, 33]. In sum, while our review of the literature uncovered some prior research relating to the collective management of private group information for active groups, there appears to be no existing studies of group privacy issues stemming from analyzing

individual-level private or public information mined from algorithmically-identified groups. Yet the Strava and CATT examples described earlier demonstrate that such research is needed.

5 PRESENT STUDY

Privacy research has yet to examine whether people are aware of privacy risks from group inference technologies. Awareness is important because it is a precondition for attempts to manage privacy risks with appropriate strategies or tools. It is also unknown whether people react differently to individual versus group privacy threats posed by group profiling and inference technologies, or whether those threats have differential effects on how people evaluate and use—or are willing to tolerate—such technologies. Unknown too is whether the kinds of threat engendered by apps like Strava that use individual-level data to reveal information about group location to third parties evoke similar personal and group privacy concerns as revealing the demographic characteristics of algorithmically-identified communities that discuss specific topics in social media which are produced by CATT.

Given these gaps in knowledge about group privacy and its effects on data subjects, two parallel experiments were designed that manipulated the salience of threats to personal (high/low salience) and to group privacy (high/low salience). Experiment 1 presented participants with an ostensible app based on Strava that we call "FitNow" that tracks personal workout data and produces geographically-localized heatmaps that could be exploited by third parties. Experiment 2 followed the same design but presented participants with information about a CATT-like fictional app we call "Trender" and the kinds of personal and group privacy threats that it portends. These experiments were conducted with the aim of answering the following research questions:

- *RQ1: Do people psychologically distinguish personal and group privacy threats and how does this affect their privacy concerns?*
- *RQ2: Do personal and group privacy threats affect user evaluations of group inference technologies?*
- *RQ3: Do personal and group privacy concerns predict different regulatory preferences for such technologies?*

Results of the experiments will help to better understand group (versus individual) privacy and its effects on users/data subjects, and ideally will increase awareness of group privacy threats among software engineers and prompt them to design tools to protect group privacy.

6 METHOD

6.1 Sample and Procedure

A total of 600 participants were recruited from Amazon's Mechanical Turk for the two experiments. After removing responses from participants who gave incorrect responses to the inserted quality check items (i.e., "Please answer 'strongly disagree' on this question"), spent too little time (< 5 minutes) on the survey, demonstrated clear response bias patterns (e.g., all 1's), had substantially incomplete surveys or duplicate IP addresses, or admitted inattention to the survey as measured by a self-report item, 507 participants remained. All participants lived in the U.S., and their ages ranged from 19-70 years. The sample was 53.6% male and 46.4% female. Across both experiments, the sample consisted of 77.4% Whites, 11.9% African-Americans, 11.1% Spanish, Hispanic or Latinos, 1.9% American Indians or Alaska Natives, 9.6% Asians, 0.2% Pacific Islanders or Hawaiian, and 1.8% indicated their race as "Other." On average 9.35% of the participants completed high school, 36 completed some college, 40.1% completed college, and 14.6% earned a post graduate degree.

Of the 260 participants who saw the descriptions of FitNow, 71.5% used fitness apps similar to FitNow (e.g., Fitbit, Strava, etc.). Among the 247 participants who read about Trender, 75.7% post

on Twitter and 93.3% post on other social media platforms (e.g., Facebook, Instagram, Snapchat, etc.). IRB approval was obtained from the authors' university prior to data collection.

Participants in each experiment were randomly assigned to one of four conditions in which everyone first read the same description of either FitNow or Trender. FitNow was described as "a new fitness app for runners and cyclists," and Trender was described as "a new app that tracks what topics people are talking about on social media." Each description also contained some information about the app's capabilities. Participants in Condition 1 (low individual, low group privacy threat) were then presented with a brief explanation of how the app works, but no privacy threats were described. Participants in Condition 2 (high individual, low group privacy threat) were presented with information about potential threats to individual privacy resulting from use of the app. No information about group-level privacy threats was provided in this condition. In Condition 3 (low individual, high group privacy threat), information was provided to participants describing potential threats to group privacy. No information about individual-level privacy threats was provided in this condition. Finally, participants in Condition 4 (high individual, high group privacy risk) saw the information about both individual and group privacy threats that were presented in Conditions 2 and 3 (Appendices A and B contain all stimuli materials used in the study).

6.2 Measures

After participants read the stimulus for their condition, they completed a questionnaire to measure their reactions to FitNow and Trender. All items for the variables in the study are shown in Table 1. *User evaluations* of FitNow (α = .94) and Trender (α = .94) were measured by calculating participants' mean response to seven items on 5-point scales (strongly disagree = 1; strongly agree = 5). *Privacy concerns* were measured at the individual and group levels (i.e., concern for personal privacy and concern for group privacy) with three items each, which were averaged to create scales. The scales for personal and group privacy concern were reliable for both the FitNow and Trender apps (Cronbach's α ranged from .82 - .87). The *cost/benefit trade-off* of using the FitNow and Trender apps was measured with one question, with scores anchored at 1 = more benefits, 4 = benefits and risks are equal, 7 = more risks. Next, participants' regulatory preferences for each app was measured in two ways. First, participants were given five *privacy settings* that they could choose to customize when using FitNow or Trender, with "yes"(1) versus "no" (0) response options. The sum of the items constituted a scale for the stringency of selected privacy settings. Scale reliability was found for both FitNow (α = .75) and Trender (α = .80). The second measure gauged participants' *endorsement of government oversight* of the technology and was measured as the average of 10 items. The scale was reliable for FitNow (α = .81) and Trender (α = .87).

Participants also indicated their demographics, including sex, age, education, race, and state of residence. Response options for the sex measure were "male" and "female." For age and state of residence, participants were presented with drop-down menus to select their age in years and their state respectively. For education, participants selected among: "Less than high school degree," "High school graduate (high school diploma or equivalent including GED)," "Some college but no degree," "Associate degree in college (2-year)," "Bachelor's degree in college (4-year)," "Master's degree," "Doctoral degree," "Professional degree (JD, MD)." For race, participants indicated if they were "White," "Black or African American," "Hispanic (Spanish, Hispanic, Latino)," "American Indian or Alaska Native," "Asian," "Native Hawaiian or Pacific Islander," or "Other."

Three questions measured participant's past experience with privacy violations (see Table 1), and answer options these items included "yes" or "no." Participants in each experiment were asked about their use of similar apps. For those in the FitNow experiment, participants were asked if they currently use or have ever used a fitness app, with the answer options "yes" or "no." For Trender, participants were asked how frequently they use Twitter and other social media, and response

Variable	Items
Personal privacy concerns	If you were a FitNow/Twitter user, to what extent would you be concerned about FitNow/Trender compromising your personal privacy? I am worried about FitNow/Trender compromising my own personal privacy (if I used it). How much of a risk do you think FitNow/Trender poses to your own personal privacy (if you used FitNow/Twitter)?
Group privacy concerns	To what extent are you concerned about FitNow/Trender compromising group privacy (e.g., revealing the location or demographic characteristics of user groups? I am worried about FitNow/Trender compromising groups' privacy. How much of a risk do you think FitNow/Trender poses to groups' privacy (e.g., identifying military operations or groups based on race, gender, and/or location)?
User evaluations	FitNow/Trender sounds appealing. FitNow/Trender seems interesting to me. FitNow/Trender seems cool. I don't think anyone should use FitNow/Trender. (reverse coded) People who use fitness apps/Twitter should use FitNow/Trender. I like FitNow/Trender. If I wanted a fitness app I'd like to use FitNow./If I was a Twitter user, I'd like to use Trender.
Cost/benefit trade-off	On balance, do you think FitNow/Trender has more benefits or more risks?
Privacy settings	Below is a list of FitNow/Trender's privacy settings. Please customize the settings as you would if you wanted to use the app: Allow FitNow to create heat maps based on my data/ allow Trender to discover groups based on user data. Allow FitNow/Trender to collect my data without my name. Allow FitNow/Trender to share my data with other companies. Allow my data to be aggregated with other users. Allow FitNow/Trender to collect my data with my name. (reverse coded)
Endorse government oversight	Governmental resources should be dedicated to evaluating risks of these technologies. There should be government oversight of these technologies. The government should shut down apps like this. A governmental body should be created to regulate these technologies. I would sign a petition to block FitNow/Trender from the App Store and Google Play. It is important that FitNow/Trender allow users to opt-out of sharing their personal data. If I used this app, I would dedicate time and effort to learning how to use its privacy settings. Users of this app should be cautious about sharing any personal information with it. I would warn my friends about FitNow/Trender. I would use FitNow/Trender with a fake name if I wanted to use it.
Use of similar apps	Do you now use or have you ever used a fitness app? (used in FitNow experiment only) How often do you post information on Twitter? (used in Trender experiment only) How often do you post information online generally or on any other social media besides Twitter (e.g., Facebook, Instagram, Snapchat, etc.)? (used in Trender experiment only)
Demographics	What is your state of residence? What is your sex? What is your age? What is the highest level of school you have completed or the highest degree you have received? Choose one or more races that you consider yourself to be from the list below.
Attention check item	Thank you for completing our study. Please answer this last question honestly. No matter how you answer, you will still be paid. There is no right or wrong answer, we are only asking this question so that we can gauge the quality of the data we collect for this study: Overall, how much attention would you say you paid when reading and answering the questions on this survey?

Table 1. Scales and Items for All Variables

options included "never," "rarely," "occasionally," "often (a few times per week)," and "many times a day." Finally, quality check measures (i.e., an attention check item and distractor items) were included. Distractor items followed this example: "Please answer 'strongly disagree' for this question" and

were inserted into the instrument randomly in two places where respondents answered a block of questions whose answer options included "strongly disagree," "somewhat disagree," "neither agree nor disagree," "somewhat agree," and "strongly agree." The attention check item was placed at the end of the instrument (see Table 1 for exact wording). Response options included "I paid no attention at all," "I paid very minimal attention," "I paid a little attention," "I paid pretty close attention," and "I paid very close attention."

7 RESULTS

The results are organized around the three research questions described in Section 5. For each question, the FitNow and Trender data are presented serially.

7.1 Psychological Distinction Between Individual and Group Privacy Threats

The first research question asked: Do people psychologically distinguish personal and group privacy threats and how does this affect their privacy concerns? For both FitNow and Trender, personal privacy concern is highly correlated with group privacy concern ($r = .77$, $p < .001$, and $r = .82$, $p < .001$, respectively). These findings suggest that personal and group privacy concerns may be psychologically interchangeable. It is nonetheless possible that privacy concerns are affected differentially by the threats we presented in the experimental design. To find out we conducted between-subjects analyses of variance (ANOVA) on each privacy concern separately. If group- and individual-level threats influence privacy concern independently, we would have direct evidence that privacy concerns at different levels are psychologically distinct.

7.1.1 FitNow. A 2(Group Privacy threat: low/high) by 2(Personal Privacy threat: low/high) between-subjects ANOVA was conducted on personal privacy concern. The ANOVA showed a main effect of group privacy threat that approached but did not reach significance, $F(1,256) = 2.90$, $p = .09$, $\eta_p^2 = .01$. However, higher group privacy threat was associated with higher personal privacy concern ($M = 3.62$) than was lower group privacy threat ($M = 3.40$). The main effect for personal privacy threat was not significant, $F(1,256) = 1.12$, $p = .29$, $\eta_p^2 = .004$, as was the interaction, $F(1,256) = .96$, $p = .33$, $\eta_p^2 = .004$. For group privacy concern, the 2 x 2 ANOVA showed a significant main effect of group privacy threat, $F(1,256) = 23.73$, $p < .001$, $\eta_p^2 = .09$. Group privacy concern was greater when threats to group privacy were high ($M = 3.72$) rather than low ($M = 3.14$). The main effect for personal privacy threat was not significant, $F(1,256) = .41$, $p = .52$, $\eta_p^2 = .002$, as was the interaction, $F(1,256) = 1.05$, $p = .31$, $\eta_p^2 = .004$. Group privacy concern is greater when group privacy threats are highlighted, and this effect is independent of threats to personal privacy.

The results indicate that while concerns for the privacy of oneself and groups are highly positively correlated, these concerns are distinct, and that threats to individual and group privacy have independent effects on judgments of privacy concerns regarding FitNow.

7.1.2 Trender. A 2(Group Privacy Threat: low/high) by 2(Personal Privacy Threat: low/high) ANOVA showed a significant main effect of personal privacy threat, such that higher personal privacy threat led to higher personal privacy concern ($M = 3.99$) than the low personal privacy threat condition ($M = 3.69$), $F(1,243) = 5.59$, $p = .019$, $\eta_p^2 = .02$. There was no main effect for group privacy threat, $F(1,243) = .05$, $p = .82$, $\eta_p^2 < .001$, or interaction, $F(1,243) = .15$, $p = .70$, $\eta_p^2 = .001$. When group privacy concern was the dependent measure, the ANOVA showed a main effect of personal privacy threat, $F(1,243) = 5.43$, $p = .021$, $\eta_p^2 = .02$. Group privacy concern was greater when individual privacy threats were high ($M = 3.76$) rather than low ($M = 3.45$). There was no main effect for group privacy threat, $F(1,243) = .01$, $p = .91$, $\eta_p^2 < .001$, or interaction, $F(1,243) = .10$, $p = .75$, $\eta_p^2 < .001$. These findings provide no evidence that people distinguish privacy at the

personal and group levels when it comes to the Trender technology. Figure 2 illustrates the results for Sections 7.1.1 and 7.1.2.

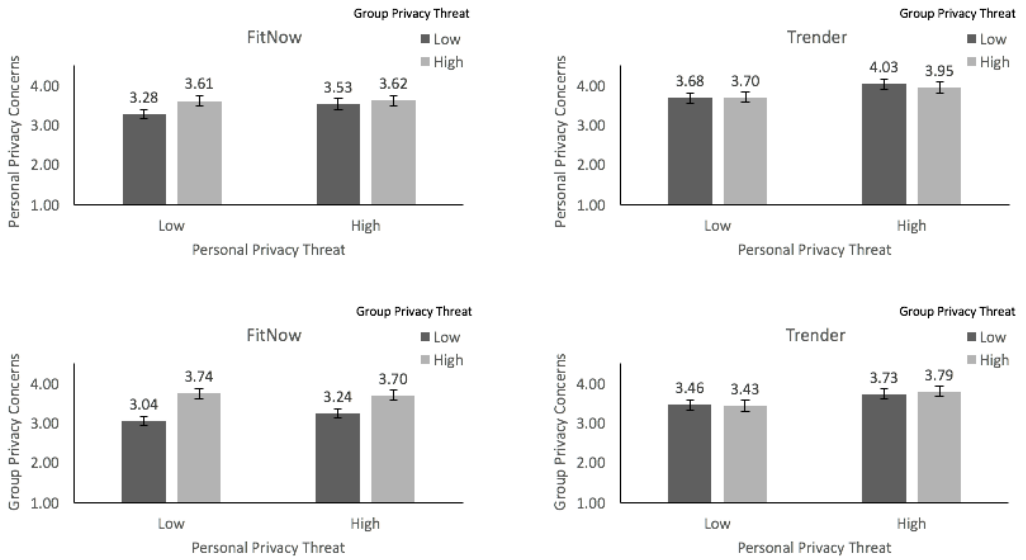

Fig. 2. ANOVA results for both experiments for personal concern (top) and group concern (bottom)

7.2 Effects of Individual and Group Privacy Threats on User Evaluations

7.2.1 FitNow. The second research question asked if individual and group privacy threats affect people's evaluations of group inference technologies. As before a 2 x 2 between-subjects ANOVA was conducted on user evaluations of FitNow. Results showed significant main effects for group, $F(1,256) = 5.66$, $p = .018$, $\eta_p^2 = .02$, and personal privacy threat, $F(1,256) = 7.51$, $p = .006$, $\eta_p^2 = .03$, but no evidence for an interaction, $F(1,256) = .15$, $p = .70$, $\eta_p^2 = .001$. The FitNow app was less appealing in the high (M = 3.19) than low group threat condition (M = 3.49), and less appealing in the high (M = 3.17) than low personal threat condition (M = 3.51). Personal and group privacy threats produce independent and additive effects on judgments of the appeal of the FitNow app, with the lowest appeal found when both personal and group privacy threats were highlighted, moderate levels of appeal in the low/high combinations, and the greatest appeal in the low/low combination.

Another 2 x 2 between-subjects ANOVA was conducted on evaluations of the costs versus benefits of FitNow. This ANOVA showed no significant main effects for group, $F(1,256) = 2.99$, $p = .085$, $\eta_p^2 = .01$, or personal privacy threat, $F(1,256) = 3.32$, $p = .07$, $\eta_p^2 = .01$, but a significant interaction, $F(1,256) = 4.03$, $p = .046$, $\eta_p^2 = .02$. The mid-point of the scale (4) was the point at which individuals judged benefits equal to risks for using the FitNow app. In the low group and low personal threat condition, participants judged the app as having relatively more benefits than risks (M = 3.77), whereas in all other conditions the perceptions were that FitNow produced more risks than benefits, and these risks were all at approximately the same level.

These findings show that people attend to the degree of personal and group privacy threats that FitNow presents and perceive the app as less appealing and costlier when there are more threats presented. Interestingly, the levels of perceived risk resulting from threats to individual and group privacy were of similar magnitude.

7.2.2 Trender. A 2(Group Privacy Threat: low/high) by 2(Personal Privacy Threat: low/high) between-subjects ANOVA was conducted on user evaluations of Trender. This ANOVA showed a significant main effect for personal privacy threat, $F(1,243) = 11.78$, $p = .001$, $\eta_p^2 = .05$. Trender was judged less appealing in the high ($M = 2.73$) than low ($M = 3.21$) personal threat condition. There was no main effect of group privacy threat, $F(1,243) = .02$, $p = .88$, $\eta_p^2 < .001$, or interaction, $F(1,243) = .21$, $p = .65$, $\eta_p^2 = .001$. Personal but not group privacy threats affected the perceived appeal of Trender.

A 2 x 2 between-subjects ANOVA conducted using the measure of the costs versus benefits of Trender showed a main effect for personal privacy threat only, $F(1,243) = 5.15$, $p = .024$, $\eta_p^2 = .02$. In the low personal threat condition, participants judged Trender as having relatively more risks than benefits ($M = 4.66$), and this was higher still in the high personal threat condition ($M = 5.12$).

There was no main effect of group privacy threat, $F(1,243) = .78$, $p = .38$, $\eta_p^2 = .003$, or interaction, $F(1,243) = .07$, $p = .79$, $\eta_p^2 < .001$. The findings for Trender show that people attend only to the degree of personal threats, but not group privacy threats when evaluating the appeal of the technology and perceptions of its cost benefit ratio. Figure 3 displays the results for Sections 7.2.1 and 7.2.2.

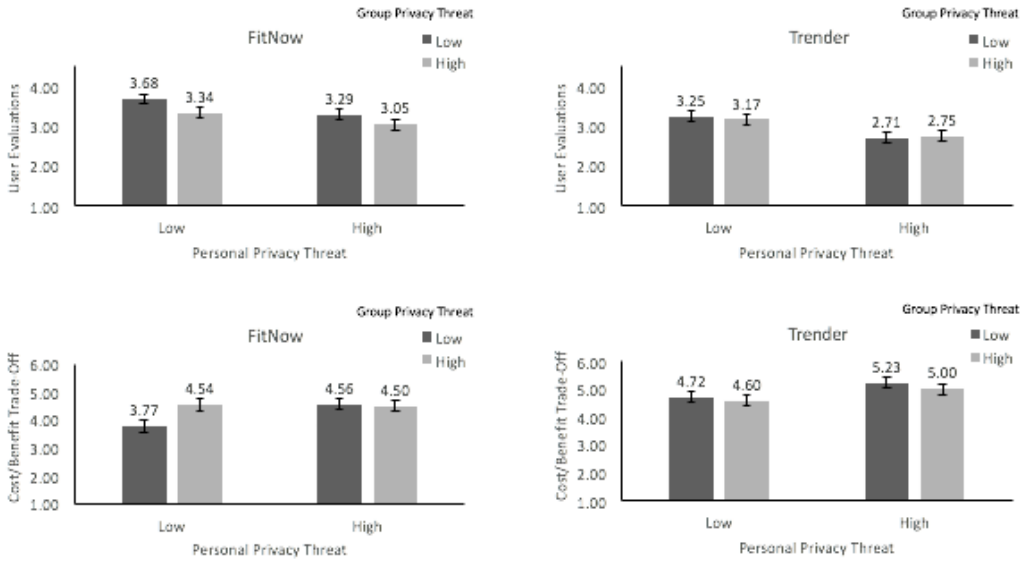

Fig. 3. ANOVA results for both experiments for user evaluations (top) and cost-benefit trade-off (bottom)

7.3 Effects of Individual and Group Privacy Threats on Regulatory Preferences

7.3.1 FitNow. Individuals perceive threats to group and personal privacy as correlated yet distinct, and those threats affect their evaluations of the appeal and costs of using the FitNow app. We can now ask whether individuals are also likely to change their regulatory preferences in terms of (a) the stringency of their privacy settings and (b) whether they would endorse government controls on the FitNow app when threats to their personal and/or group privacy are salient.

A 2(Group Privacy Threat: low/high) by 2(Personal Privacy Threat: low/high) between-subjects ANOVA was conducted on the stringency of privacy settings users chose. This analysis revealed significant main effects for group, $F(1,256) = 3.93$, $p = .048$, $\eta_p^2 = .02$, and personal privacy threats, $F(1,256) = 6.25$, $p = .013$, $\eta_p^2 = .02$, but no evidence for an interaction, $F(1,256) = .03$, $p = .87$, $\eta_p^2 < .001$. Participants reported invoking more stringent privacy settings in the high ($M = 8.26$) than low ($M = 7.87$) group privacy threat condition, as well as more stringent settings in the high ($M = 8.31$)

than low (M = 7.82) individual privacy threat condition. The effects of group and personal privacy threats on privacy setting stringency were independent and additive. A 2(Group Privacy Threat: low/high) by 2(Personal Privacy Threat: low/high) between-subjects ANOVA was also conducted on endorsement for government oversight of FitNow. This ANOVA showed a significant main effect for group, $F(1,256) = 12.94$, $p < .001$, $\eta_p^2 = .05$, but not personal privacy threats, $F(1,256) = 2.15$, $p = .14$, $\eta_p^2 = .008$, and no interaction, $F(1,256) = .20$, $p = .65$, $\eta_p^2 = .001$. Group privacy threats have unique effects on bolstering attitudes in favor of government regulation.

7.3.2 Trender. Participants indicated the privacy settings they would employ if Trender were in use. A 2(Group Privacy Threat: low/high) by 2(Personal Privacy Threat: low/high) between-subjects ANOVA was conducted on the stringency of privacy settings. This analysis revealed a main effect for personal privacy threats, $F(1,243) = 7.62$, $p = .006$, $\eta_p^2 = .03$. Participants reported invoking more stringent privacy settings in the high (M = 8.43) than low (M = 7.84) personal privacy threat condition. There was no evidence for a main effect of group privacy threat, $F(1,243) = .01$, $p = .95$, $\eta_p^2 < .001$, or interaction, $F(1,243) = 1.46$, $p = .23$, $\eta_p^2 = .006$.

Another 2 x 2 between-subjects ANOVA was conducted on endorsement for government oversight of Trender. This ANOVA showed no significant effects for personal privacy threat, $F(1,243) = 3.12$, $p = .079$, $\eta_p^2 = .01$, group privacy threat, $F(1,243) = .22$, $p = .64$, $\eta_p^2 = .001$, or their interaction, $F(1,243) = 2.41$, $p = .12$, $\eta_p^2 = .01$. Figure 4 illustrates the results for Sections 7.3.1 and 7.3.2.

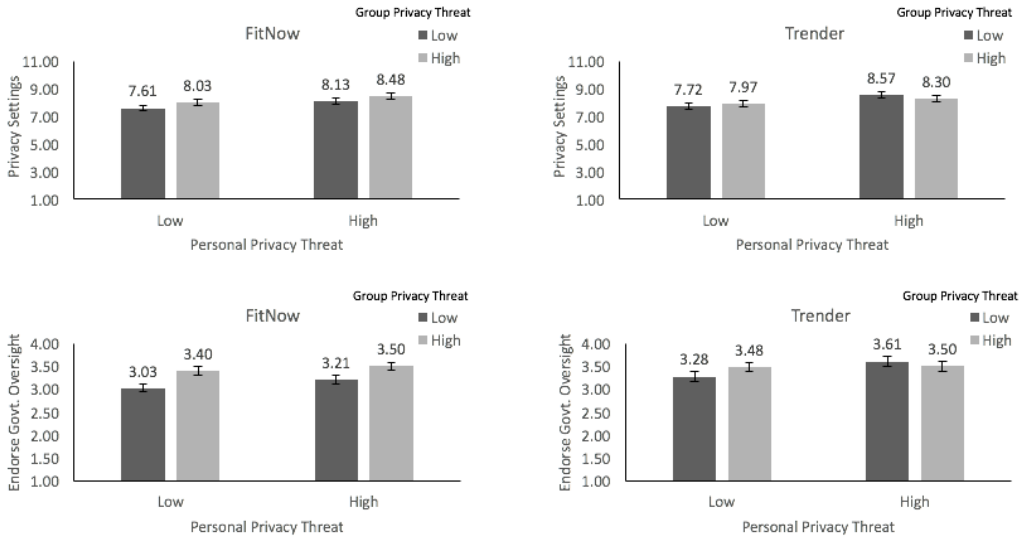

Fig. 4. ANOVA results for both experiments for stringency of privacy settings (top) and endorsement of government oversight (bottom)

7.3.3 Follow-up Mediational Analysis. We further investigated how threats to individual and group privacy affect attitudes about technology regulation by exploring which type of privacy concern explains the causal relationship linking privacy threats and regulatory preferences. We tested mediation models for both personal and group privacy concern (IVs) on stringency of privacy settings and support for government oversight (DVs) for each app, using Hayes' PROCESS module. For FitNow, only group privacy concern mediated the relationship between group privacy threat salience and stringency of privacy setting choices 95% CI (.03, .39). Personal privacy was not found to be a significant mediator. When the models were re-run with endorsement of government regulation as the dependent variable, again the only significant mediator was group privacy concern,

95% CI (.04, .30). These analyses show that group privacy concern was greater when group privacy threats were salient, and group privacy concern explains participants' privacy setting choices and support for government oversight of FitNow, despite a strong positive correlation existing between personal and group privacy concerns. Figures 5 and 6 present these results visually.

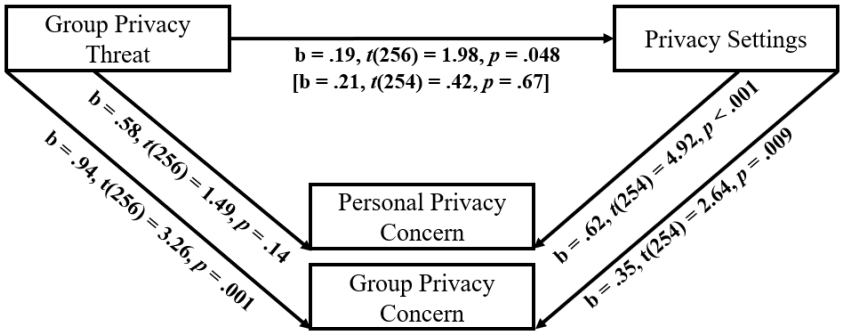

Fig. 5. Experiment 1, FitNow. Higher group privacy threat leads users to select more stringent privacy settings. This effect was mediated by group privacy concerns, but not personal privacy concerns

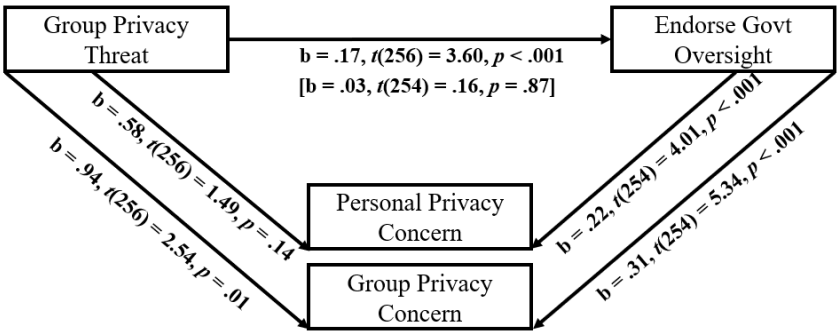

Fig. 6. Experiment 1, FitNow. Higher group privacy threat leads to stronger attitudes in favor of government oversight. This effect was mediated by group privacy concerns, but not personal privacy concerns.

By contrast, there was no evidence of mediation through either type of privacy concern for Trender. However, the findings showed positive relationships between personal privacy concern, b = .56, t(241) = 3.45, p < .001, and (marginally) group privacy concern, b = .28, t(241) = 1.81, p = .07, with the stringency of selected privacy settings. These analyses control for the privacy threats introduced in the experimental conditions, and therefore show that, on average, people who have greater personal and group privacy concerns are more likely to select restrictive privacy settings, independent of experimentally-induced privacy threats. Repeating this model with group privacy threats as the independent variable also found no mediation, but did show that people with greater personal, b = .25, t(241) = 3.92, p < .001, and group privacy concerns, b = .30, t(241) = 4.93, p < .001, were more in favor of government oversight of Trender. Again, these relationships held while controlling for the experimentally induced levels of personal and group privacy threat.

8 DISCUSSION

To our knowledge, no prior privacy research has considered threats to the privacy of algorithmically-determined groups or the privacy implications of group inference technologies that are based on public data supplied by individuals but have privacy implications that go beyond those actual individuals. The experiments presented in this paper thus help to shed new light on both individual and group privacy.

The first research question asked whether people psychologically distinguish personal and group privacy threats, and if this affects their privacy concerns. Results showed that while personal and group privacy concerns were highly positively correlated, they were distinct in that each type of privacy threat increased group and individual privacy concerns independently. However, this was only true for the FitNow app. The second research question examined whether personal and group privacy threats affect users' evaluations of group inference technologies. Results showed that personal and group privacy threats independently affected users' judgements of the apps, with increased threat lowering users' perceptions of the appeal and degree of risk posed by the technology. Again, however, the results varied by app. For FitNow, both personal and group privacy threats similarly affected users' evaluations whereas for Trender only personal threats did so.

These findings have important theoretical and practical implications. First, they suggest theoretical approaches such as CPM [30] that focus on how individuals negotiate their privacy boundaries should be expanded to consider group-level negotiations, not only between group members but also with companies that use group-inference technologies. This could serve as a means to avoid boundary turbulence. For example, if risks to group privacy had been made clearer to users at sign-up, Strava might have been able to avoid its recent negative news coverage. Second, the results of our studies showing group privacy concern manifest for active (i.e., FitNow) but not passive (i.e., Trender) groups supports Nissenbaum's contextual approach to privacy [28] and psychological theories described earlier (e.g., [8]) that claim group entitativity (self-awareness) is a precondition for group privacy concerns to arise. That people appear less concerned about the privacy of passive groups is not surprising given the psychologically abstract nature of these groups. People may not easily perceive algorithmically-defined groups as groups at all. However, this result should not be interpreted by theorists or software designers as an argument against protecting the privacy of passive groups, as privacy violations based on group inference algorithms can be harmful to people even when they are unaware of being grouped at all [21].

The last research question investigated whether personal and group privacy concerns predict different regulatory preferences for group inference technologies. For FitNow, making either personal or group privacy threats salient caused additive effects on the selection of more stringent privacy settings, although when it came to endorsement of government regulation of the app, it was only group privacy risks that were associated with a preference for greater oversight. Group privacy concern also mediated the relationships between group privacy threats and both the stringency of privacy settings selected and attitudes toward government regulation. Hence, for the FitNow app group privacy concerns were not just distinct from personal privacy concerns, as group privacy concerns also uniquely motivated attitudinal and behavioral reactions to the technology. Results for Trender showed a different pattern, as increases in personal but not group privacy threats affected privacy settings (endorsement of government oversight approached significance) and there was no evidence of any mediation. However, the data did show independent effects of personal and group privacy concerns on both outcomes such that people who had more personal and group privacy concerns about Trender selected more stringent privacy settings and were more favorable towards government oversight of the app.

Recent scandals involving the privacy of data subjects such as with Cambridge Analytica are re-focusing policy makers' attention on regulation. The findings pertaining to RQ3 indicate user support for policy to protect not only personal privacy, but group privacy as well, and such support is likely to intensify as the number of applications that incorporate group-inference technologies increases. Policy makers need to be aware of threats to privacy from algorithmically-determined groups to protect users appropriately, and we hope our study is a first step toward that end. Here again though we urge caution in interpreting our Trender results: Despite seemingly less concern for privacy issues concerning passive (compared to active) groups found in our study, this is perhaps

where policy is most needed, as users are unable to protect themselves when they cannot detect their own group memberships. Because the threat to group privacy is more amorphous and distal for data subjects in passive groups to perceive, companies can use this to their advantage at the expense of group privacy, which increases the urgency for government oversight to protect group privacy rights (see [21]).

In a similar vein, Kammourieh and colleagues argue for new legal privacy protections that go beyond the collection and transfer of individuals' data to include any data processing or algorithms that may be applied to the data, and for requirements on companies to report their data processing methods and the potential risks of those methods to the public [21]. They also suggest self-regulation on the part of data gatherers as a way forward by increasing transparency about their classification and grouping algorithms, adopting policies that make clear to the public when and how group-related information is used, and providing redress for violations to these policies, or by limiting companies' use of potentially sensitive information pertaining to groups.

Our findings should also prompt software engineers to develop technological means to protect group privacy. This will require a shift because, as mentioned earlier, the current focus of privacy research in computer science is premised on protecting individual privacy. Traditional approaches are based on the idea of hiding the individual in a "crowd," and aggregating data to protect information about an individual. These techniques derive from early notions of k-anonymity where a set of data records is obfuscated in a way that hides the record of any individual record among at least k other records [34].[4] The differential privacy approach [13] obfuscates data by adding enough noise so that any query on the data will return the same answer irrespective of whether the individual record is in the data set or not. But these means to protect privacy are so focused on personal privacy that they ignore or even may unintentionally reveal group privacy. K-anonymity, almost by design, reveals group characteristics, and thus may compromise group privacy. In k-anonymity, the algorithm attempts to hide the individual in a group of size at least k. This approach inherently identifies and hence reveals a group. This is exactly the case in the Strava example, where the privacy of the individual was hidden completely, but as a result, a group was identified and group privacy was compromised. Differential privacy, on the other hand, adds enough noise to a data set so that the record corresponding to an individual has negligible effect on the answer to queries on the data set, but again the focus is on personal privacy, not on group privacy.

We believe that our findings demonstrate the need for new approaches within computer science to understand how to preserve group privacy. This is challenging, as these approaches must protect privacy while not compromising the utility of the obfuscated data. While traditional approaches critically depend on hiding the individual in a larger group, to preserve group privacy, the "group" needs to be obfuscated. This challenge may lead to innovative privacy preserving technologies that do not use notions of generalization for obfuscation. Also, the focus on the individual records in differential privacy will need to be modified to hide group identities, where some groups might not even be pre-defined and known a-priori (i.e., passive groups). However, the cost of differential privacy is typically bounded by the amount of noise inserted in the data, and as the dataset increases in size, so does the noise, which reduces utility. This demands an in-depth understanding of group privacy by computer scientists and has the potential for interesting and innovative technical discoveries for protecting both group and individual privacy.

From a more practical standpoint, another challenge when developing software tools for privacy protection is how to balance users' privacy, be it individual or group, with the self-interests of

[4]Obfuscation can be achieved in different ways, for example, by generalizing the address of an individual to the county or state where they live, thus hiding the individual among all those residing in the county or state, as long as there are at least k such individuals in the data set.

companies whose revenue stream often derives from analyzing group characteristics of users for targeted advertisements. We argue that group privacy provides a unique opportunity, which might not even be an option for individual privacy: Namely, an individual's membership in a group is usually defined by a collection of attributes. An individual may want to hide (or even generalize) some specific attributes and not others—for example by revealing their ethnicity and age group but not their sexual orientation, or a user may be willing to reveal their gender and age group, but want to generalize their location from say San Diego to California. Companies may cooperate with users to tailor which attributes are strictly private while others may be revealed or generalized. This is a novel and interesting aspect of privacy protection that group privacy provides and should be further explored.

Finally, new information literacy approaches are needed as well. The current educational approach assumes people will manage their own data privacy through available settings once they are informed about them. But this is both unrealistic and flawed, as a lot of people don't manage their settings well even when aware of the options, and because available privacy settings don't protect from many types of privacy attacks anyway [37]. Future information literacy efforts should strive to increase public awareness of new threats to privacy from group inference technologies as a starting point for empowering users and data subjects alike. One potential avenue for increasing awareness is a software daemon that users could install on their devices that gives advice when posting or tweeting. This advice can help protect a user's group privacy when a post may compromise their group privacy. This approach has been proposed by Zakhary et al. in the form of a cyborg which resides on a user's computer and protects the user from privacy violations [17, 43]. It works similarly to "privacy nudge" which provide timely, soft paternalistic messages that guide users towards more thoughtful and informed privacy-related decisions (e.g., what information to disclose). Such nudges have been shown to reduce users' privacy-threatening behavior in social media contexts (e.g., [40]). Nudges might also be useful for increasing awareness among software developers and big data analysts of the social and ethical consequences of their work in terms of group privacy. Nudges could be designed to inform developers of apps such as Strava and Trender of potential dangers to group privacy during the app design process.

8.1 Limitations

Our research is not without limitations. For example, we studied only two instances of group inference technologies (FitNow and Trender) and participants in our experiments were not representative of all users of data technology, as our sample included only people residing in the U.S. at the time of data collection. But users of information technology, including fitness apps and social media, come from all over the globe, and attitudes toward both privacy and government regulation vary significantly across cultures. This indicates that future research that examines a greater range of group-inference technologies, and includes users from a wider variety of locations to examine cultural variation in privacy concerns is warranted before firm conclusions about group privacy can be reached.

Also, our experimental design may have primed participants to be more privacy concerned than they would be naturally. To guard against this problem, questions to participants were ordered such that after reading the description of FitNow or Trender, they were first asked to provide their positive and negative evaluations of the app, then asked about their perceptions of risks and benefits, and immediately after that to customize the app privacy settings. This ordering was intended to prime participants to think about *the positives and negatives* of the app prior to asking about their desired privacy settings. Of course, this question ordering strategy may not have completely mitigated the possibility of priming affecting our results, and so research methodologies using unobtrusive behavioral data (e.g., where people choose app privacy settings without prior

exposure to any information about privacy risks) will be important to establish the ecological validity of our findings. At the same time, our findings indicate that risks from group inference technologies should be more clearly highlighted to users when signing up for services as a means to help users minimize risks to personal or group privacy violations.

Despite these limitations, our hope is that this research inspires further investigations of group privacy generally as well as in the specific but important context of group inference technologies. Such research is exciting because we believe it affords opportunities for significant and novel discoveries in both computational and social sciences.

REFERENCES

[1] 2018. Fitness app Strava lights up staff at military bases. https://www.bbc.com/news/technology-42853072/. (2018).
[2] 2018. The latest data privacy debacle. https://www.nytimes.com/2018/01/30/opinion/strava-privacy.html/. (2018).
[3] 2018. Strava fitness app can reveal military sites, analysts say. https://www.nytimes.com/2018/01/29/world/middleeast/strava-heat-map.html/. (2018).
[4] Alessandro Acquisti and Jens Grossklags. 2007. What can behavioral economics teach us about privacy. *Digital privacy: theory, technologies and practices* 18 (2007), 363–377.
[5] Irwin Altman. 1975. The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding. (1975).
[6] France Bélanger and Robert E Crossler. 2011. Privacy in the digital age: a review of information privacy research in information systems. *MIS quarterly* 35, 4 (2011), 1017–1042.
[7] Andrew Besmer, Heather Richter Lipford, Mohamed Shehab, and Gorrell Cheek. 2009. Social applications: exploring a more secure framework. In *Proceedings of the 5th Symposium on Usable Privacy and Security*. ACM, 2.
[8] Donald T Campbell. 1958. Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Behavioral science* 3, 1 (1958), 14–25.
[9] Yunan Chen and Heng Xu. 2013. Privacy management in dynamic groups: understanding information privacy in medical practices. In *Proceedings of the 2013 conference on Computer supported cooperative work*. ACM, 541–552.
[10] Hichang Cho and Anna Filippova. 2016. Networked privacy management in facebook: A mixed-methods and multinational study. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, 503–514.
[11] Ralf De Wolf. 2016. Group privacy management strategies and challenges in Facebook: A focus group study among Flemish youth organizations. *Cyberpsychology-Journal of Psychosocial Research on Cyberspace* 10, 1 (2016).
[12] Ralf De Wolf, Koen Willaert, and Jo Pierson. 2014. Managing privacy boundaries together: Exploring individual and group privacy management strategies in Facebook. *Computers in Human Behavior* 35 (2014), 444–454.
[13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.
[14] Luciano Floridi. 2014. Open data, data protection, and group privacy. *Philosophy & Technology* 27, 1 (2014), 1–3.
[15] Luciano Floridi. 2017. Group privacy: A defence and interpretation. In *Group Privacy*. Springer, 83–100.
[16] Theodore Georgiou, Amr El Abbadi, and Xifeng Yan. 2017. Extracting topics with focused communities for social content recommendation. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*.
[17] Theodore Georgiou, Amr El Abbadi, and Xifeng Yan. 2017. Privacy Cyborg: Towards Protecting the Privacy of Social Media Users. In *Data Engineering (ICDE), 2017 IEEE 33rd International Conference on*. IEEE, 1395–1396.
[18] Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2011. Detecting and resolving privacy conflicts for collaborative data sharing in online social networks. In *Proceedings of the 27th Annual Computer Security Applications Conference*. ACM, 103–112.
[19] Haiyan Jia and Heng Xu. 2016. Autonomous and interdependent: Collaborative privacy management on social networking sites. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4286–4297.
[20] Isaac Johnson and Brent Hecht. 2017. Inferred Profiles: Examining How People Understand and Control What Algorithms Infer about Them. In *CSCW Workshop on Privacy Ethics*. ACM.
[21] Lanah Kammourieh, Thomas Baar, Jos Berens, Emmanuel Letouzé, Julia Manske, John Palmer, David Sangokoya, and Patrick Vinck. 2017. Group Privacy in the Age of Big Data. In *Group Privacy*. Springer, 37–66.
[22] Airi Lampinen. 2016. Hosting together via Couchsurfing: Privacy management in the context of network hospitality. *International Journal of Communication* 10 (2016), 20.
[23] Kevin Lewis, Jason Kaufman, and Nicholas Christakis. 2008. The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication* 14, 1 (2008), 79–100.

[24] Eden Litt. 2013. Understanding social network site users' privacy tool use. *Computers in Human Behavior* 29, 4 (2013), 1649–1656.
[25] Alessandro Mantelero. 2017. From group privacy to collective privacy: towards a new dimension of privacy and data protection in the big data era. In *Group Privacy*. Springer, 139–158.
[26] Miriam J Metzger. 2007. Communication privacy management in electronic commerce. *Journal of Computer-Mediated Communication* 12, 2 (2007), 335–361.
[27] Helen Nissenbaum. 1998. Protecting privacy in an information age: The problem of privacy in public. *Law and philosophy* 17, 5-6 (1998), 559–596.
[28] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
[29] Sameer Patil and Jennifer Lai. 2005. Who gets to know what when: configuring privacy permissions in an awareness application. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 101–110.
[30] Sandra Petronio. 2012. *Boundaries of privacy: Dialectics of disclosure*. Suny Press.
[31] H Jeff Smith, Tamara Dinev, and Heng Xu. 2011. Information privacy research: an interdisciplinary review. *MIS quarterly* 35, 4 (2011), 989–1016.
[32] Anna Cinzia Squicciarini, Mohamed Shehab, and Federica Paci. 2009. Collective privacy management in social networks. In *Proceedings of the 18th international conference on World wide web*. ACM, 521–530.
[33] Anna C Squicciarini, Heng Xu, and Xiaolong Zhang. 2011. CoPE: Enabling collaborative privacy management in online social networks. *Journal of the American Society for Information Science and Technology* 62, 3 (2011), 521–534.
[34] Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 05 (2002), 557–570.
[35] Henri Tajfel. 1970. Experiments in intergroup discrimination. *Scientific American* 223, 5 (1970), 96–103.
[36] Linnet Taylor. 2017. Safety in numbers? Group privacy and big data analytics in the developing world. In *Group Privacy*. Springer, 37–66.
[37] Linnet Taylor, Luciano Floridi, and Bart Van der Sloot. 2016. *Group privacy: New challenges of data technologies*. Vol. 126. Springer.
[38] Linnet Taylor, Luciano Floridi, and Bart van der Sloot. 2017. Introduction: A new perspective on privacy. In *Group Privacy*. Springer, 1–12.
[39] Bart van der Sloot. 2017. Do Groups Have a Right to Protect Their Group Interest in Privacy and Should They? Peeling the Onion of Rights and Interests Protected Under Article 8 ECHR. In *Group Privacy*. Springer, 197–224.
[40] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2367–2376.
[41] Alan F Westin and Oscar M Ruebhausen. 1967. *Privacy and freedom*. Vol. 1. Atheneum New York.
[42] Heng Xu. 2011. Reframing privacy 2.0 in online social network. *U. Pa. J. Const. L.* 14 (2011), 1077.
[43] Victor Zakhary, Cetin Sahin, Theodore Georgiou, and Amr El Abbadi. 2017. Locborg: Hiding social media user location while maintaining online persona. In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 12.

APPENDIX A FITNOW

Stimuli for the FitNow study are pictured below. Participants saw two pages: The first page described the app and the second page differed in each of the experimental conditions as shown in the four lower panels.

FitNow Stimulus Page 1 (this page was the same for all participants in each FitNow condition)

Below is a description of an app. Please read the description:

FitNow is a new fitness app for runners and cyclists who want to record their activities, track performance, and share photos, stories, and highlights of their workouts with friends and other users. FitNow can also show where users like to work out, which can help people find the best places to be active and help advocacy groups and planners improve bike and pedestrian infrastructure.

Condition 1 Page 2 Stimulus (low personal, low group)

FitNow works by collecting personal data on users' workout routes and times anonymously and aggregates data across users to produce "heat maps" that show where people work out most frequently (an example heat map is pictured below).

Condition 3 Page 2 Stimulus (high personal/low group)

FitNow works by collecting personal data on users' workout routes and times anonymously and aggregates data across users to produce "heat maps" that show where people work out most frequently (an example heat map is pictured below).

However, there are potential risks associated with recording this data. For example, revealing the location or identity of FitNow users could enable stalkers, and result in identity theft or unwanted target marketing ads aimed at individual users.

Condition 2 Page 2 Stimulus (low personal/high group)

FitNow works by collecting personal data on users' workout routes and times anonymously and aggregates data across users to produce "heat maps" that show where people work out most frequently (an example heat map is pictured below).

However, while the identity of individual FitNow users is protected, there are potential risks associated with publishing this data in the form of heat maps.

For example, if soldiers on a secret military base use the app, FitNow's heat maps might reveal the location of their base and troop movements, and this information could be used by enemies to target U.S. military operations. And the same problem could apply to other sensitive groups as well.

Condition 4 Page 2 Stimulus (high personal/high group)

FitNow works by collecting personal data on users' workout routes and times anonymously and aggregates data across users to produce "heat maps" that show where people work out most frequently (an example heat map is pictured below).

However, there are potential risks associated with recording and publishing these data. For example, revealing the location or identity of FitNow users could enable stalkers, and result in identity theft or unwanted target marketing ads aimed at individual users.

Further, if soldiers on a secret military base use the app, FitNow's heat maps might reveal the location of their base and troop movements, and this information could be used by enemies to target U.S. military operations. And the same problem could apply to other sensitive groups as well.

APPENDIX B TRENDER

Stimuli for the Trender study are pictured below. Participants saw two pages: The first page described the app and the second page differed in each of the experimental conditions as shown in the four lower panels.

Trender Page 1 Stimulus (this page was the same for all participants in each Trender condition)

Below is a description of an app. Please read the description:

Trender is a new app that tracks what topics people are talking about on social media such as Twitter. Trender can help people find others to talk about specific topics, learn about the community of people who share their interests, and help companies discover audiences for their messages. For example, Trender can report that "18-25 year old women in Orange, CA are discussing #StrangerThingsNetflix"

Condition 1 Page 2 Stimulus (low personal, low group)

Trender works by correlating popular topics on Twitter with demographic information (e.g., age, gender, location) of the people who post about these topics, and so can show what kinds of people are discussing what topics on Twitter (the table below shows some examples).

Topic	Frequency	Characteristics	Size
#NFL	1534	Gender: Male, Location: USA	1212
#FreeJustina	54	Gender: Female, Location: Miami, Politics: Democrats	51
#RedSox	528	Gender: Male, Age: 19-11, Location: Boston	411
#ObamaCare	5090	Gender: Male, Location: USA, Politics: Republicans	4818

Condition 3 Page 2 Stimulus (high personal/low group)

Trender works by correlating popular topics on Twitter with demographic information (e.g., age, gender, location) of the people who post about these topics, and so can show what kinds of people are discussing what topics on Twitter (the table below shows some examples).

Topic	Frequency	Characteristics	Size
#NFL	1534	Gender: Male, Location: USA	1212
#FreeJustina	54	Gender: Female, Location: Miami, Politics: Democrats	51
#RedSox	528	Gender: Male, Age: 19-11, Location: Boston	411
#ObamaCare	5090	Gender: Male, Location: USA, Politics: Republicans	4818

However, there are potential risks associated with recording this data because Trender puts information about users collected from their social media profile together with what kinds of topics they talk about, and then uses that information to reveal things about their personality and lifestyle, including their interests, tastes, political preferences, and location. This type of revealed information could result in unwanted target marketing ads aimed at individual users, or if a controversial topic is discussed, individuals might be subject to surveillance by law enforcement.

Condition 2 Page 2 Stimulus (low personal/high group)

Trender works by correlating popular topics on Twitter with demographic information (e.g., age, gender, location) of the people who post about these topics, and so can show what kinds of people are discussing what topics on Twitter (the table below shows some examples).

Topic	Frequency	Characteristics	Size
#NFL	1534	Gender: Male, Location: USA	1212
#FreeJustina	54	Gender: Female, Location: Miami, Politics: Democrats	51
#RedSox	528	Gender: Male, Age: 19-11, Location: Boston	411
#ObamaCare	5090	Gender: Male, Location: USA, Politics: Republicans	4818

However, while Trender does not reveal any information about users beyond what they post publicly on social media, there are potential risks associated with Trender aggregating information across users.

For example, Trender can discover and profile groups of people based on the topics they discuss, and connect this information with location, age, and race. This information could then be exploited by law enforcement, enable government surveillance, and target marketing campaigns.

Condition 4 Page 2 Stimulus (high personal/high group)

Trender works by correlating popular topics on Twitter with demographic information (e.g., age, gender, location) of the people who post about these topics, and so can show what kinds of people are discussing what topics on Twitter (the table below shows some examples).

Topic	Frequency	Characteristics	Size
#NFL	1534	Gender: Male, Location: USA	1212
#FreeJustina	54	Gender: Female, Location: Miami, Politics: Democrats	51
#RedSox	528	Gender: Male, Age: 19-11, Location: Boston	411
#ObamaCare	5090	Gender: Male, Location: USA, Politics: Republicans	4818

However, there are potential risks associated with recording this data because Trender puts information about users collected from their social media profile together with what kinds of topics they talk about, and then uses that information to reveal things about their personality and lifestyle, including their interests, tastes, political preferences, and location.

Further, Trender can discover and profile groups of people based on the topics they discuss, and connect this information with location, age, and race. This information could be exploited by law enforcement, enable government surveillance, and target marketing campaigns.

Received April 2018; revised July 2018; accepted September 2018