

# Zooming Into Video Conferencing Privacy and Security Threats

Dima Kagan<sup>\*1</sup>, Galit Fuhrmann Alpert<sup>†1</sup>, and Michael Fire<sup>‡1</sup>

<sup>1</sup>Department of Software and Information Systems Engineering,  
Ben-Gurion University of the Negev, Israel

July 3, 2020

## Abstract

The COVID-19 pandemic outbreak, with its related social distancing and shelter-in-place measures, has dramatically affected ways in which people communicate with each other, forcing people to find new ways to collaborate, study, celebrate special occasions, and meet with family and friends. One of the most popular solutions that have emerged is the use of video conferencing applications to replace face-to-face meetings with virtual meetings. This resulted in unprecedented growth in the number of video conferencing users. In this study, we explored privacy issues that may be at risk by attending virtual meetings. We extracted private information from collage images of meeting participants that are publicly posted on the Web. We used image processing, text recognition tools, as well as social network analysis to explore our web crawling curated dataset of over 15,700 collage images, and over 142,000 face images of meeting participants. We demonstrate that video conference users are facing prevalent security and privacy threats. Our results indicate that it is relatively easy to collect thousands of publicly available images of video conference meetings and extract personal information about the participants, including their face images, age, gender, usernames, and sometimes even full names. This type of extracted data can vastly and easily jeopardize people's security and privacy both in the online and real-world, affecting not only adults but also more vulnerable segments of society, such as young children and older adults. Finally, we show that cross-referencing facial image data with social network data may put participants at additional privacy risks they may not be aware of and that it is possible to identify users that appear in several video conference meetings, thus providing a potential to maliciously aggregate different sources of information about a target individual.

**Keywords:** Video Conference, Video Conference Applications, Security and Privacy, Image Processing, Data Science, COVID-19

---

<sup>\*</sup>kagandi@bgu.ac.il

<sup>†</sup>fuhrmann@bgu.ac.il

<sup>‡</sup>mickyfi@bgu.ac.il

# 1 Introduction

The COVID-19 pandemic outbreak has had tremendous impacts on daily life in multiple aspects, ranging from basic hygiene, medicine, health science, economics, politics, and society. In particular, it dramatically affected the ways in which people communicate with their families, friends, and coworkers [1]. Social distancing resulted in the search for new ways to collaborate, study, celebrate birthdays, and even meet with parents and grandparents. One of the most popular solutions that have emerged is the use of video conferencing applications, such as Google Meet,<sup>1</sup> Microsoft Teams,<sup>2</sup> and Zoom.<sup>3</sup> During COVID-19 outbreak and accompanying stay-home quarantine strategies, people have started using video conferencing applications to replace face-to-face meetings in schools, workplaces, and social gatherings with virtual meetings. As a result, the number of video conferencing users along with the number of daily meetings surged sharply [2, 3]. Today, there are hundreds of millions of video conferencing meetings that take place daily, encompassing millions of users [3].

With the unprecedented growth in video conferencing usage, many security and privacy issues have been unraveled [4, 5]. These issues range from unencrypted communication for unpaid users [6] to vulnerabilities that allow malware execution on participants’ devices [5]. Moreover, the use of video conferencing applications has raised privacy concerns that may enable uninvited attendees to find ways to join meetings by guessing meetings IDs [7] or by simply searching for video conferencing meeting links that were made publicly available (e.g. meeting links published on social media websites [8]).

A malicious user that gains access to video conferencing meetings can collect sensitive and private data on users, such as their names, usernames, images of their faces, samples of their voice, and even exposure to personal data that has been shared as part of the conversations. Moreover, using accessible deepfake tools, a malicious user can attend video conferencing meetings under a false identity. For example, one can use realtime deepfake tools in order to join meetings using celebrity avatars that make him or her look like celebrities, such as Barack Obama or Elon Musk [9].<sup>4</sup> The participants’ personal data can later be used to jeopardize participants’ safety in both the virtual and the real-world [10]. For example, Acquisti et al. [11] demonstrated the threat of face recognition can be used to identify individuals both in the online and offline worlds. Namely, they used publicly available images from Facebook to identify students strolling through campus. They also illustrated that it is possible to predict personal and sensitive information from a face, such as the individuals interests, activities, and even his or her social security number.

In this study, we take an in-depth look into the world of video conferencing by analyzing images taken from thousands of video meetings, which were pub-

---

<sup>1</sup><https://meet.google.com>

<sup>2</sup><https://www.microsoft.com/en-us/microsoft-365/microsoft-teams/group-chat-software>

<sup>3</sup><https://zoom.us>

<sup>4</sup><https://github.com/alievk/avatarify>

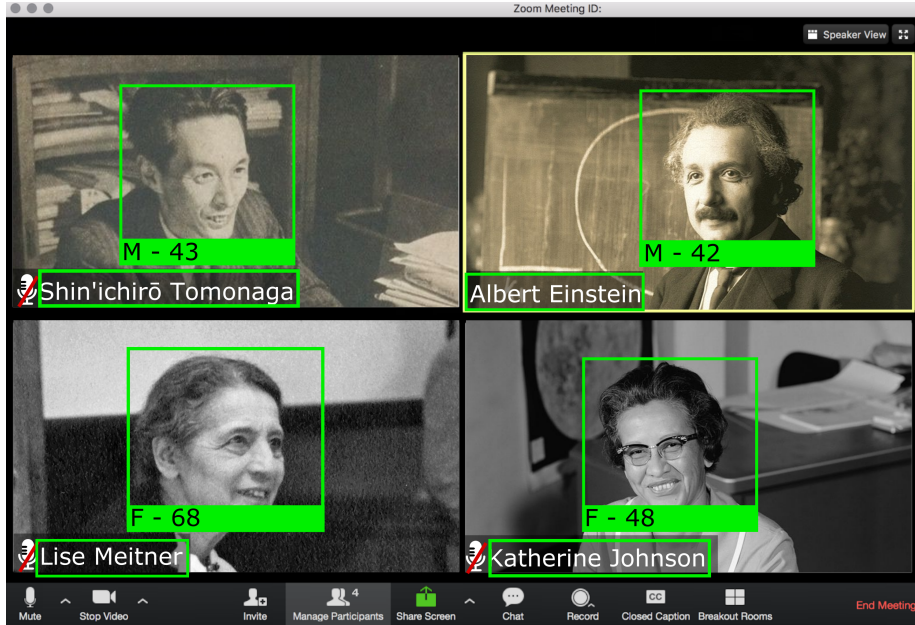


Figure 1: Zoom Image Collage with Detected Information, Along with Extracted Features of Gender, Age, Face, and Username.

licly published online. We curated an image dataset by collecting and analyzing ten-of-thousands images published on social media websites. We subsequently analyzed the dataset to extract information from each image regarding the meeting participants, such as their age, gender, username, as well as several other features.

We compute statistics on video conferencing usage and demonstrate that video conferencing participants encompass a wide range of ages, from young children to older adults (see Figure 1). We apply deep-learning based image processing algorithms to demonstrate that it is possible to identify the same individual’s participation at different meetings by simply using either face recognition or other extracted user features. The extracted information about users has the potential to be harmfully used to uncover participants’ social networks and other privacy related factors (see Figure 1). Overall, our study emphasizes the privacy risks that video conferencing participants need to be made aware of. We demonstrate that there is danger not only from a malicious user but also from other users in the meeting that can innocently upload *a single photo*, which can be maliciously be used to affect their and their family’s privacy.

The remainder of the paper is organized as follows: In Section 2, we present an overview of relevant studies. In Section 3, we describe the datasets, methods, algorithms, and experiments used throughout this study. In Section 4, we present our results. In Section 5, we discuss the obtained results. In Section 6,

we explain the study’s limitations. In Section 7, we recommend mitigations. Lastly, in Section 8, we present our conclusions from this study.

## 2 Background

Over the years, many efforts were put into analyzing user security and privacy in a variety of online platforms, such as online virtual worlds [12], online trading systems [13], online social networks [10] and cryptocurrency platforms [14]. In this study, we focus on privacy and security issues in video conferencing applications, one of the most commonly used one known as Zoom. Such applications have rapidly emerged as cardinal ways of communication as a result of stay-home and social distancing orders due to the COVID-19 pandemic outbreak, and have been increasingly supporting, in addition to work-related and educational activities, also a wide range of social activities, including virtual happy hours, virtual yoga, blind dating, and worship services [15]. As a consequence, users of video conferencing platforms have been exposed to multiple security and privacy risks, similar to those related to online social networks.

Privacy Risks related to online social networks include the following:

- **Cyberbullying.** Cyberbullying (also referred to as Cyber Abuse) refers to bullying and harassment that take place within technological communication platforms, such as chats, mobile devices, and social networks [16]. Recently, video conferencing applications users have been exposed to a new type of online harassment, termed *Zoom Bombing*, where an uninvited person joins a video chat meeting and interrupts the meeting by sharing inappropriate content [8].
- **Information Leakage.** Information leakage refers to the detection and extraction of information that was unintentionally disclosed [17]. Visual data contains implicit information in the background that may reveal interesting, as well as possibly sensitive pieces of information. For instance, Weyand et al. [18] was able to detect a photo location exclusively based on image pixels. Reece and Danforth [19] demonstrated that it is possible to detect markers of depression in Instagram photos.
- **Information Linkage.** Information linkage occurs when an attacker is able to link several pieces of information from at least two separate data sources. For example, an attacker may match user’s publicly available details in one online account (e.g. username or profile picture) with other accounts and social network profiles of the same user, in order to uncover additional details about the user, such as full name, home location, and workplace.
- **Malware Attacks.** Malware refers to malicious software that is intentionally developed in order to disrupt a computer operation for the purpose of collecting credentials and gaining access to private information. Like



many other web platforms, video conferencing applications are also subject to malware attacks. For example, in parallel to Zoom’s dramatic rise in popularity, it has become a popular target for hackers [5]. Recently, it has been reported that Zero day vulnerability in Zoom allows for remote code execution and launching malware on target computers, [20].

- **Phishing Attacks.** Phishing attacks are a form of social engineering that is developed in order to steal user-sensitive and private information by masquerading as a trustworthy third party. Since COVID-19 outbreak, hackers have been increasingly impersonating video conferencing applications, including Zoom, Microsoft Teams, and Google Meet by sending phishing emails and registering over 2,449 Zoom related domains for phishing scams [21]. Moreover, a new type of Zoom related phishing campaign has recently been revealed in which the phishing mail contains a fake Zoom link and includes words like ‘termination’ to distract users from noticing that the link in the email is not linked to Zoom official website [22].
- **Face Recognition Attacks.** Face recognition algorithms are capable of identifying or verifying a person from a digital image or a video source. Identifying a person’s face from a video, and cross-referencing it with other datasets might be used to expose personal information about the individual [23]. Acquisti et al. [11] demonstrated that it is possible to identify strangers online (on a dating site where individuals protect their identities by using pseudonyms) and offline (in public space), based on photos that are publicly available on social network sites. Zoom, in particular, has enabled until very recently data mining capabilities to secretly display data related to meeting participants by linking participants’ names and email addresses with their LinkedIn profiles [24].

In addition, users of video conferencing applications are exposed to additional sets of risks:

- **Data Breach.** Some video conferencing applications provide recording and cloud storage capabilities for their users. Recently, the Washington Post [25] reported that thousands of private Zoom videos were publicly accessible online on either Zoom’s cloud storage, or on other external storage services. Washington Post reporters watched several of these videos and discovered that these videos contained sensitive and private information, including peoples names and phone numbers, private company and financial statements in small-business meetings, as well as elementary school classes, in which childrens faces, voices, and personal details were exposed. Even users that used a password to protect their videos on Zoom cloud were exposed to these same risk, as presented by a proof of concept example that used brute force to bypass the password protected videos.<sup>5</sup>
- **Fake Avatars.** Fake avatar is a new type of identity fraud, similar to fake profiles in online social networks [26]. Video conferencing application

---

<sup>5</sup><https://github.com/markbuffalo/zoombo>

users can be victims of this new type of identity fraud, in which other participants in a meeting can pretend to be a person who they are not. For example, the code project Avatarify<sup>6</sup> can be used to create photorealistic avatars for video-conferencing, that may mislead meeting participants to think they are meeting with a celebrity [27]. Similar to other deepfake based applications, this type of technology might be used for fraudulent activities. For example, in 2019, it is believed that fraudsters used deepfake voice software to mimic a CEOs voice in order to steal money from a company [28].

- **Zoombombing Campaigns.** Zoombombing (also referred to as Zoom Raiding) is a practice that refers to unwelcome stranger participants joining a meeting, and disrupting video calls with offensive language and imagery. This type of behavior has recently become a dangerous method for coordinated harassment and hate speech [29]. For example, recent research performed by the New York Times uncovered that thousands of people gathered online to organize Zoom harassment campaigns, share meeting passwords and plan for sowing chaos in public and private meetings [29]. Also, in the UK, there is an ongoing investigation of 120 Zoombombing cases of Zoom video calls that were hijacked by people displaying images of child abuse [30].

### 3 Methods and Experiments

The primary goal of this study was to explore privacy aspects that may be at risk by attending virtual conference meetings, a practice that has become very common following the COVID-19 outbreak and related stay-home lockdown and social distancing measures.

To achieve this goal we performed the following steps:

#### 3.1 Curating an Image Dataset of Video Conferencing Collages

We curated an image dataset that contains images from thousands of video conference meetings by performing the following three steps (also illustrated in Figure 2):

First, to collect images from video conference meetings, we utilized web crawlers that collect data from Twitter and Instagram using online accessible Twitter and Instagram scraper tools.<sup>78</sup> We were specifically interested in collecting images that were likely to be taken in video conference meetings, thus we set the web crawlers to search for tweets that contain terms or hashtags from a

---

<sup>6</sup><https://github.com/alievk/avatarify>

<sup>7</sup><https://developer.twitter.com/en/docs>

<sup>8</sup><https://github.com/arc298/instagram-scraper>



Figure 2: Dataset Generation Process.

predefined set of meeting related target terms (e.g. “zoom school” and “#zoom-meeting”).<sup>9</sup> Using this method we were able to collect 89,305 Zoom related tweets and 90,395 Zoom related Instagram posts (a total of 179,700 meeting related posts).

Second, we filtered out tweets and Instagram posts that did not contain an image and extracted the images from all remaining tweets and posts (a total of 26,408 images from Twitter and 78,435 images from Instagram were collected). We then utilized Fastai framework [31] to construct an image classifier that identifies which of the images contain a Zoom video meeting collage of participants.

To train the classifier, we manually labeled 5,271 images, out of which 1,505 contained a collage of participants (positive examples), and 3,766 were labeled as other types of images (negative examples). Using the labeled images, we applied ResNet-50 [32] based transfer learning to train our model. The trained classifier achieved an accuracy of 0.969, a true-positive rate of 0.935, and a false-positive rate of 0.016 on a test set of 1,054 images. This process resulted in a dataset consisting of 16,133 Zoom collage images.

To remove duplicates and similar images (i.e same image with minor adjustments or crops), we utilized dhash [33] to compute the Hamming distance between all pairs of images. We removed all images with distance equal or smaller than an empirically set threshold 1.2.<sup>10</sup> Moreover, to remove additional similar images, we computed the Euclidean and cosine distances between all pairs of images using each image embedding generated by using the Zoom image classifier last layer representation. Then, we removed all the images that

<sup>9</sup>Our twitter crawler downloaded tweets that contained the following terms and hash-tags: “zoom birthday,” “zoom happyhour,” “zoom school,” “zoom party,” “#happyhour,” “#Zoom,” “#zoommeeting,” “#zoomus,” and “@Zoom.US.” The Instagram crawler downloaded posts containing the following hashtags: “#zoommeeting,” “#zoomparty,” and “#zoombirthday.”

<sup>10</sup>The selected threshold of 1.2 was chosen, by manually testing and evaluating several threshold values, searching for minimum false positives.

their distance was equal or lower of 0.0035 and 25 for cosine and euclidean distance respectively.<sup>11</sup> Finally, we remained with an image dataset that consists of 15,709 collage images automatically collected from Twitter and Instagram.

### 3.2 Feature Extraction by Image Processing of Collages

To extract structure data from collage image dataset, we performed the following steps on each image in the dataset:

1. **Faces Recognition.** Zoom collage images typically contain multiple faces. We utilized two face recognition tools to detect each of the faces in a collage. The first tool we used was a pre-trained model [34], based on MTCNN [35]. The second tool we utilized was Microsoft Azure Face API for face recognition [36]. Each of the two tools extract the bounding boxes of the faces in each collage image in our dataset. Then, the results of the two face recognition tools are combined in the following manner: We keep the faces that are uniquely identified by each tool, and in cases where bounding boxes of two faces intersect, we consider them as the same face (the bounding box from the first model is used). To estimate recall values, we manually evaluated faces detection recall rates in 100 random selected Zoom collage images. Our face detection approach achieved a recall of 80% (i.e detection of 80% faces of the available data).
2. **Face Embedding.** By using a dlib based tool [37, 38], each detected face in our dataset, we generated a 128-dimension numerical vector representation, which represents the features extracted from the face.
3. **Age Detection.** For each detected face, we estimated the age of the person, by applying two separate models for inferring the age of an individual based on his or her face image, a pre-trained model for age detection [39]<sup>12</sup> and the package of Microsoft Azure Face API [36]. We then define the estimated age as the average value of two predicted ages. We binned the average predicted ages into four categories [40]:

$$f(x) = \begin{cases} child, & \text{if } x \leq 12 \\ adolescent, & \text{if } x > 12 \text{ and } x \leq 17 \\ adult, & \text{if } x > 17 \text{ and } x < 65 \\ older adult, & \text{otherwise} \end{cases}$$

4. **Gender Detection.** For each face detected by Microsoft Azure Face API [36], we also utilized the API to detect the user’s gender. By manually inspecting 100 randomly selected Zoom collage images, we evaluated that Microsoft Azure Face API detects gender on our curated dataset with a recall of 98.9%.

---

<sup>11</sup>The selected threshold values of 0.0035 and 25 were empirically chosen, by testing and manually evaluating several threshold values.

<sup>12</sup>The code was refactored into a library instead of a CLI tool. The modified code is available on GitHub.

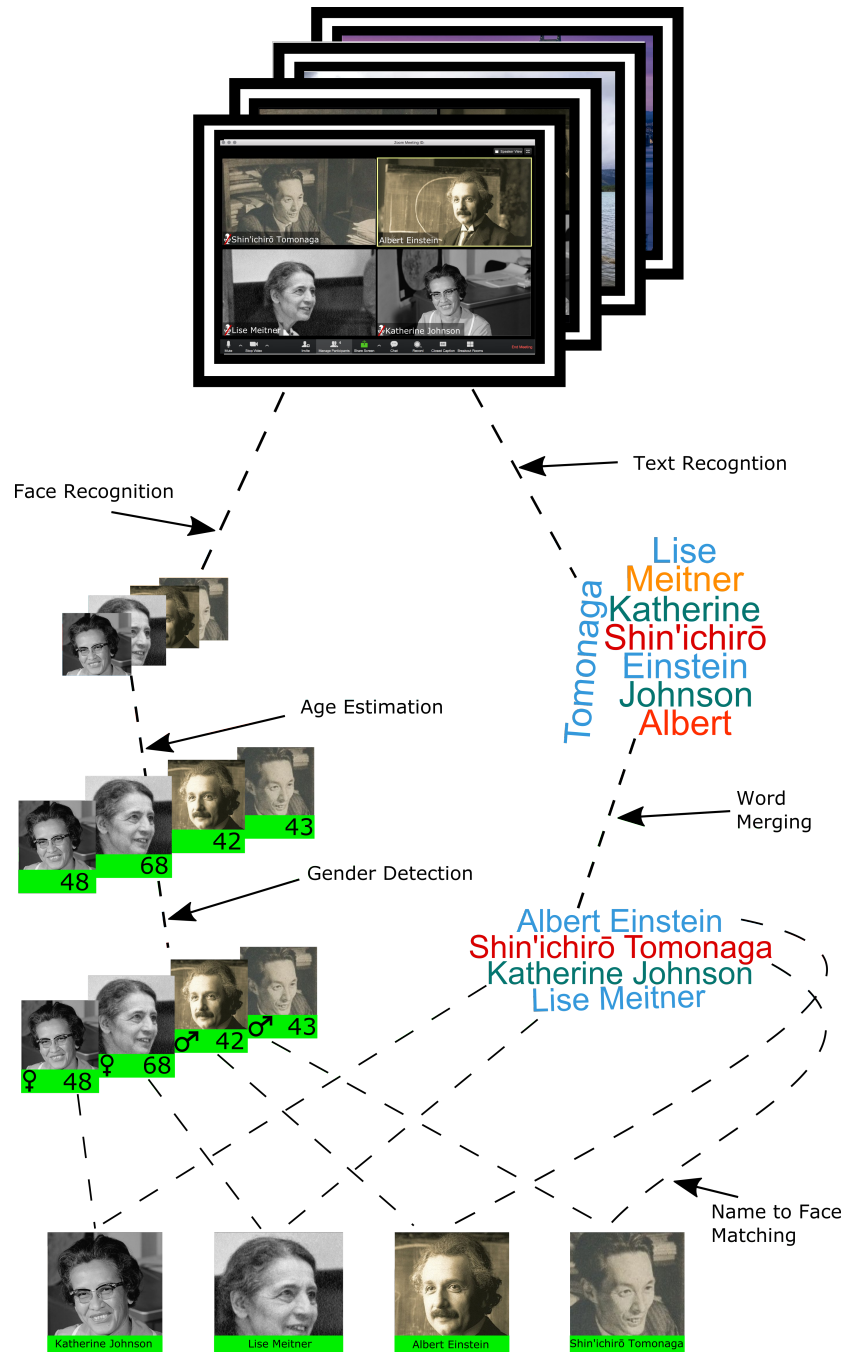


Figure 3: Data Extraction Process

5. **Username Recognition.** To identify participants’ usernames, we utilized a publicly available scene text recognition library.<sup>13</sup> The library uses a text detection approach that is based on EAST [41] for detecting text location and MORAN [42] for subsequently recognizing the composing characters. To filter out words that are not username candidates, we used a list of Zoom meeting-related words.<sup>14</sup> We used spaCy lemmatization [43] to filter out all dictionary words, aiming at removing words that are not names. One of the issues with the text recognition model is that it detects only single words. However, in many cases usernames are a combination of multiple words, such as first name, middle name, and last name. To combine single detected words into the associated full usernames we applied a custom heuristic approach based on the unique image structure of Zoom collages. We noticed that it is seldom to find two words located very close one to another that are not a part of the same name. Based on this observation, we developed a simple heuristic to join recognized words that are closely-located in image space into a single username. Namely, for each identified word in each image collage, we searched for another word that is the nearest to its top right point. If the nearest word was in a euclidean distance smaller or equal to 10,<sup>15</sup> we combined both words. To construct usernames of more than two words, we repeated this process repeatedly until no words were further combined. We evaluated this heuristic on 100 randomly selected Zoom collage images and found that it combines 97.6% of the usernames longer than a single word into the correct usernames. To evaluate the username recognition (scene text detection that is followed by word merging) we compared the actual username with the detected username in 100 randomly selected Zoom collage images. The method was able to detect 63.4% of the usernames correctly.<sup>16</sup>

### 3.3 Linking Personal User Data to Social Network Data

We explored whether video conference participants are vulnerable to information leakage attacks, such as linkage to social networks. To demonstrate that this type of attack is indeed possible using video conference data, we reconstructed participants’ links and social networks by matching users that participated in multiple meetings using the following methods:

(a) *Cross Referencing Usernames:* We match users across video conference meetings by explicitly using participants’ usernames. This method considerably

<sup>13</sup>[https://github.com/gtsoukas/scene\\_text](https://github.com/gtsoukas/scene_text)

<sup>14</sup>The list consists of the following words: “help,” “view,” “meeting,” “edit,” “participant,” “speaker,” “security,” “screen,” “record,” “video,” “stop,” “share,” “manage,” “view,” “exit,” “full,” “chat,” “end,” “reactions,” “recording,” “mute,” “zoom,” “participants,” “from,” “recording,” “window,” “search,” “invite,” “leave,” “unmute,” “option,” “delete,” “raise,” “hand,” “new,” “type,” “message,” and “here.”

<sup>15</sup>The selected threshold of 10 was chosen, by manually testing and evaluating several threshold values.

<sup>16</sup>We used a strict measurement where even if only one character is different between the actual username and detect one it will not be considered as correct detection.




Figure 4: Zoom users gender and age distribution.

reduces the search space of Zoom participant identities. Moreover, in some cases, participants have unique usernames, which greatly assists in finding these same individuals in multiple meetings

(b) *Cross Referencing using Face Embedding*: We match users across different meetings by comparing the detected participants’ faces images. Namely, we utilized the extracted face embeddings (see Section 3.2), and calculated the euclidean distance between all pairs of detected faces. We consider two face images to be images of the same user if the euclidean distance between their vector representation is smaller than an empirically chosen threshold.<sup>17</sup>

Using the cross-referenced data, we constructed a social graph  $G := \langle V, E \rangle$  of the participants, where  $V$  is the set of video conference users, and  $E$  is a set of links between users that according to the collage image dataset participants in a meeting together.

In order to demonstrate that video conference meeting participants are susceptible to information leakage attacks, we tested whether by using the participants’ usernames and facial images, it is possible to match video conference users with their online social networks’ profiles. To achieve this goal, we manually searched for usernames on social media sites and compared them to a video conference username with the same face. In addition, after matching a video conference user with his or her online social network profile, we explore whether it is possible to gather additional data about his or her friends to reverse search and match the identified friends profiles with their video conference profiles.

<sup>17</sup>To determine that two face images are of the same person, we empirically chose a distance threshold between the vector representation of two faces to be less or equal 0.3, which is a lower than the default distance value of 0.6 used by the *Face Recognition* python library Geitgey [38]. We chose a relatively strict threshold since many of the faces extracted from Zoom collage images suffer from low resolution, thus resulting in many false matches of pairs of faces.

## 4 Results

To analyze privacy threats in video conference collage images (see Section 3), we collected 15,783 Zoom collage images that were publicly posted on Twitter and Instagram. We explored the dataset of collected Zoom collages and found an average of 9.04 participants in a video conference meeting collage. From the crawled Zoom collage images, we extracted a total of 142,001 face images for analysis. From an age perspective (see Figure 4), the average estimated age of the participants is 29.23 and the median is 28.56. By examining age distributions, we observed that 87.49% of the participants are adults, 6.18% children, 6.23% adolescents, and only 0.09% of participants are older adults. Additionally, the gender prediction algorithms (see Section 3) were able to identify 29,048, and 50,221 males and females respectively.

We were interested in how discrete users are about their real world identities, as reflected in the choice of their usernames. We thus extracted usernames from the Zoom image collages, and found 85,616 distinct usernames, out of which 48,818 (57% of the usernames) consist of more than a single word (see Figure 5), suggesting possibly more disclosed information about the user. By manually inspecting these 48,818 multi-word usernames we observed that many of these words represent distinct names that can be utilized to match a user’s social network profile. In fact, 2,522 of these usernames appeared in several collage images from different meetings. Moreover, some users use the name of their phone model as a username, thus not disclosing personal identity by the choice of username. In fact, we found that in our dataset the username “iPhone” was the most popular single word username (395 appearances).

By analysing the participants facial images, we identified 1,153 faces that likely appeared in several different Zoom meeting images. Using the cross-referenced images we constructed a large-scale social network of Zoom users with 16,842 nodes (participants) and 197,765 edges (participation in joint meetings). Each node in the network represents a single meeting participant, and each edge between a pair of participants in the network represents that they jointly participated in at least in one video conference meeting. Users may jointly participate in meetings from different worlds of content. The network consists of distinct 345 connected components. On average, each separate component consisted of 48.8 participant nodes and 573.2 joint meeting edges respectively. The largest component consisted of 3,066 nodes and 55,035 edges (see Figure 6). By inspection of randomly selected participants, we were able to manually locate their personal social network profiles. We also observed networks where all




Figure 5: Distribution of the Number of Words Composing a Zoom Username






Figure 6: The Largest Component Observed in the Social Network of Zoom Users in Our Dataset. The network was constructed by matching meeting participants' faces and usernames across meetings as explained in Methods (see Section 3). Each node in the network represents a Zoom meeting participant, and each edge between users in the network represents that both users participated in at least in one video conference meeting.

participants were co-workers. Finally, we found that some users are more privacy aware than others, using smiley emojis above their faces to protect their privacy.

## 5 Discussion

Our results point to several important findings:

First, we demonstrate how easy it is to collect personal data from an abundance of Zoom collage images that are publicly posted on the Web. Using Instagram and Twitter crawlers, we were able to collect over 15,000 collage images (see Section 3.1), all taken during Zoom video conference meetings. From those Zoom collages, using face recognition algorithms, we were able to extract a dataset of over 140,000 faces and over 85,000 distinct usernames. These results indicate that many images taken from video conference meetings are already publicly available online, and relatively easy to collect. Moreover, the faces and usernames collected in this type of process can be used to construct a facial image dataset, which contains personal details about meeting participants, including facial characteristics, age, gender, usernames, and sometimes even full names. This type of facial image dataset can vastly and easily jeopardize people’s security and privacy both in the online and real-world. Furthermore, in this study, in order to demonstrate this point, we generated our dataset by collecting publicly available data only. However, malicious users, such as zoom-bombers, can potentially actively harvest data from video conference meetings and construct more precise facial datasets of video conference participants’ data without their consent and knowledge.

Second, the collected dataset provides us with a demographic glimpse into the world of video conference meetings. By analyzing age and gender features of over 140,000 participants (Figure 4), we observed that today’s video conferencing is widely used by various segments of the population. While the vast majority (87.49%) of meeting participants were adults, as might be expected considering that due to COVID-19 many companies have been encouraging employees to work from home [44], the age span is in fact from children to older adults. This result highlights the risk that privacy and security issues may affect not only adults but also more vulnerable segments of the population, such as young children and older adults.

Third, we showed that cross-referencing facial image data with social network data may put participants at additional privacy risks they may not be aware of. In fact, as one example, we explored a meeting that contained a group of adults. By performing an in-depth analysis of these images and comparing the participants’ usernames and facial images with data available from online social networks, we were able to conclude that the meeting participants worked in the same company and that we could additionally infer social links among the meeting participants. This illustrates that not only individuals’ privacy is at risk from data exposed on video conference meetings, but also the privacy and security of organizations. For example, by inspecting the social network of participants in an organization, it possible to expose a variety of private details

regarding the organization itself [45]. Thus, an attacker can potentially utilize the organization’s social network in order to attack the organization [46].

Fourth, we demonstrated that by using meeting participants’ facial characteristics and usernames, it is possible to identify users that appear in several video conference meetings, thus to collect and aggregate their data. For example, as depicted in Figure 6, it is possible to construct an individual’s social network by collecting data from several meetings he or she participated in. Moreover, we demonstrate that it is possible to use data collected from video conference meetings along with linked data collected in other video meetings with other groups, such as online social networks, in order to perform a linkage attack on target individuals. This can result in jeopardizing the target individual’s privacy by using different meetings to discover different types of connections. For instance, one may aggregate information regarding job ,personal, and political related social connections.

Finally, during this study, we observed that some meeting participants were more successful than others in protecting their privacy. Such privacy protection practices included the use of generic usernames, replacing their faces with smiley images, etc. Interestingly, even for such privacy aware individuals, we realized that it is possible to identify the same individuals across different meetings by using other means, such as identifying unique backgrounds, another potential privacy risk that users may be unaware of.

## 6 Research Limitations

It is worth noting that the methods used in this study are prone to several limitations: First, we collected only publicly available images from Twitter and Instagram postings. Therefore, our collage image dataset is partial and does not contain data from meetings that were not publicly published online. Additionally, we analyzed the images using state-of-the-art facial recognition tools. Nevertheless, these tools and algorithms have been reported to be biased in some cases [47]. To overcome such types of face recognition issues we combined the results of two facial recognition models (see Section 3.2). Also, to diversify the datasets, we collected data from multiple sources, Twitter, and Instagram. Moreover, the quality of the Zoom collage images in our dataset is not consistent, affecting our extracted data. Some images are screenshots and others taken by smartphones. Additionally, there are differences in the image resolution, which directly affects the number of pixels available for analysis of each face and username. The lower the image quality the harder it to extract accurate data. We also observed that some actions can improve the method general accuracy, including image alignment, face verification, and training a model for splitting a collage into rectangles may improve performance of name to face matching as well as face detection. Lastly, retraining the models on Zoom data may significantly improve the accuracy of the models used in Section 3.2.

## 7 Recommendations

As we have demonstrated throughout this study, video conference users are facing prevalent and varied security and privacy threats. In this section, we provide several easy-to-apply methods which can assist video conference participants to improve their security and privacy:

- **Avoid Video Streaming Whenever Possible.** As we have shown in this study, images taken from video conference meetings can jeopardize individual security and privacy in multiple ways. Even in cases of private meetings, it is sufficient for a single collage image to be taken by one of the participants in order to jeopardize the privacy of all associated meeting participants. Therefore, whenever possible, we strongly recommend video conference users to not share videos of meetings.
- **Avoid Uploading Virtual Meeting Photos and Videos onto Social Media.** Avoid uploading videos and collage photos of video conference meetings onto social media.
- **Use Generic Pseudonyms for Video Conferencing.** As demonstrated in this study, using a full name or a unique username for video conferencing can easily expose an individual to information linkage attacks. Moreover, using a generic name, such as “iPad” or “iPhone” makes it more challenging to cross-reference a user’s identity with other datasets. Therefore, we strongly recommend to video conference participants not to use their real name and advise using a generic pseudonym instead.
- **Use Generic Background Images as Virtual Backgrounds.** While video streaming, using a virtual background is highly recommended. Using a real background can compromise a user’s security and privacy. For example, items in the background of an individual’s room can expose his or her name and even geographic location. Additionally, not using a virtual background or using a unique background image may help malicious users in fingerprinting a user account across several meetings.
- **Use Anti-Facial Recognition Accessories.** In recent years, some companies have started selling anti-facial recognition cloth and makeup to protect the privacy of users Holmes [48]. Some accessories can be homemade. For example, Sharif et al. [49] presented a method that is using a home printer to create skin for glasses that disrupt facial recognition algorithms. Moreover, we observed that in many cases, we were not able to detect the faces of people wearing even a simple party mask.
- **Organizations Should Inform Employees on Video Conference Privacy Risks.** Many organizations, especially during times of pandemic, suggest that their employees work from home and conduct meetings using video conference applications. This results in a new set of security and privacy threats to organizations that not all organizations are aware of.

In this study, we demonstrated that sharing data from video conference meetings can compromise organizations’ security and privacy. Therefore, we recommend organizations be vigilant to the new risks, which are raised with the increased usage of this technology. Moreover, we recommend organizations to update their policies on what employees can share and not share while using video conference applications.

- **Monitor Childrens Video Conferencing Activity.** As we observed in our study, children are active participants in video conference meetings. While video conference may seem as a safe and harmless place to be, they are at risk of numerous threats ranging from malware to zoombombing as described in Section 2. Video conference is not different from regular online activity and parents should be actively involved in monitoring their children’s activity the same way. To reduce risks, parents should explain to their kids not to accept a request to join a meeting from someone they do not know in real life, and also remove personal details from the child’s account and adjust the privacy settings.
- **Video Conference Operators Should Add and Support Privacy Mode.** In the past couple of years, researchers presented methods that can disrupt facial recognition. For instance, Wilber et al. [50] showed that adding Gaussian noise to an image can disrupt facial recognition while keeping the face still recognizable for humans. Even without the help of video conference operators in some applications, users may opt to use filters in order to hide their appearance and to avoid automatic facial recognition.

## 8 Conclusions

In 2020, the COVID-19 pandemic drastically and globally changed the way people communicate, immensely increasing communication via video conferencing. As a result, novel privacy and security issues have emerged, endangering hundreds of millions of video conference participants. In this study, we dive into the privacy issue of video conference applications by analysing over 15,000 video conference images of over 140,000 meeting participants. From these images, we extracted multiple private information features including gender, age, and real name. We demonstrate that this information can be utilized to uncover additional details about video conference participants.

We discovered that cross-referencing participants in multiple conference meetings can uncover the participant’s social network and subsequently may lead to compromising his/her or other participants’ privacy and social accounts. That said, we offer several recommendations in which users can protect themselves, for instance by using generic usernames and background, wearing small masks, or using a filter. We observed that some users pasted emojis on their faces to protect their privacy, and we found this helpful for protecting their privacy. In

the current global reality of social distancing, we must be sensitive to online privacy issues that accompany changes in our lifestyle as society is pushed towards a more virtual world.

## 9 Data availability

Due to the private nature of the data, an anonymized version of the constructed social network is available only upon request from the corresponding author.

## References

- [1] Darcy Palder, Amy Mackinnon, and Kelly Kimball. How the coronavirus pandemic has changed daily life around the world. <https://foreignpolicy.com/2020/05/07/lockdown-covid-19-changed-lives-around-the-world/>, 5 2020. (Accessed on 06/15/2020).
- [2] Covid-19 outbreak: Video conferencing demand rises due to social distancing - researchandmarkets.com. <https://www.businesswire.com/news/home/20200507005631/en/COVID-19-Outbreak-Video-Conferencing-Demand-Rises-due>, 5 2020. (Accessed on 06/14/2020).
- [3] Tom Warren. Zoom admits it doesnt have 300 million users, corrects misleading claims - the verge. <https://www.theverge.com/2020/4/30/21242421/zoom-300-million-users-incorrect-meeting-participants-statement>, 4 2020. (Accessed on 06/15/2020).
- [4] Chaim Gartenberg. Google meet, microsoft teams, and webex are collecting more customer data than they appear to be - the verge. <https://www.theverge.com/2020/5/1/21244058/google-meet-microsoft-teams-webex-personal-data-collection-privacy-policy-concerns>, 5 2020. (Accessed on 06/15/2020).
- [5] Paul Wagneseil. Zoom security issues: Here’s everything that’s gone wrong (so far). <https://www.tomsguide.com/news/zoom-security-privacy-woes>, 6 2020. (Accessed on 06/14/2020).
- [6] Adi Robertson. Zoom says free users wont get end-to-end encryption so fbi and police can access calls. <https://www.theverge.com/2020/6/3/21279355/zoom-end-encryption-calls-fbi-police-free-users>, 6 2020. (Accessed on 06/14/2020).
- [7] Jay Peters. Automated tool can find 100 zoom meeting ids per hour. <https://www.theverge.com/2020/4/2/21206061/zoom-meeting-id-zwardial-automated-tool>, 4 2020. (Accessed on 06/15/2020).
- [8] Taylor Lorenz. zoombombing: When video conferences go wrong. <https://www.nytimes.com/2020/03/20/style/zoombombing-zoom-trolling.html>, 3 2020. (Accessed on 06/15/2020).

- [9] Samantha Cole. This open-source program deepfakes you during zoom meetings, in real time - vice. [https://www.vice.com/en\\_us/article/g5xagy/this-open-source-program-deepfakes-you-during-zoom-meetings-in-real-time](https://www.vice.com/en_us/article/g5xagy/this-open-source-program-deepfakes-you-during-zoom-meetings-in-real-time), 4 2020. (Accessed on 06/15/2020).
- [10] Michael Fire, Roy Goldschmidt, and Yuval Elovici. Online social networks: threats and solutions. *IEEE Communications Surveys & Tutorials*, 16(4): 2019–2036, 2014.
- [11] Alessandro Acquisti, Ralph Gross, and Frederic D Stutzman. Face recognition and privacy in the age of augmented reality. *Journal of Privacy and Confidentiality*, 6(2):1, 2014.
- [12] Tal Zarsky. Privacy and data collection in virtual worlds. 2006.
- [13] Sidney Tsang, Yun Sing Koh, Gillian Dobbie, and Shafiq Alam. Detecting online auction shilling frauds using supervised learning. *Expert systems with applications*, 41(6):3027–3040, 2014.
- [14] Mauro Conti, E Sandeep Kumar, Chhagan Lal, and Sushmita Ruj. A survey on security and privacy issues of bitcoin. *IEEE Communications Surveys & Tutorials*, 20(4):3416–3452, 2018.
- [15] Taylor Blose, Prasanna Umar, Anna Squicciarini, and Sarah Rajtmajer. Privacy in crisis: A study of self-disclosure during the coronavirus pandemic. *arXiv preprint arXiv:2004.09717*, 2020.
- [16] Robert Slonje and Peter K Smith. Cyberbullying: Another main type of bullying? *Scandinavian journal of psychology*, 49(2):147–154, 2008.
- [17] Information leakage. <https://www.whitehatsec.com/glossary/content/information-leakage>. (Accessed on 06/16/2020).
- [18] Tobias Weyand, Ilya Kostrikov, and James Philbin. Planet-photo geolocation with convolutional neural networks. In *European Conference on Computer Vision*, pages 37–55. Springer, 2016.
- [19] Andrew G Reece and Christopher M Danforth. Instagram photos reveal predictive markers of depression. *EPJ Data Science*, 6(1):1–12, 2017.
- [20] Zero day vulnerability in zoom allows remote code execution in windows & malware attacks. <https://www.securitynewspaper.com/2020/04/01/zero-day-vulnerability-in-zoom-allows-remote-code-execution-in-windows-malware-attacks/>, 4 2020. (Accessed on 06/16/2020).
- [21] Jay Peters. Hackers are impersonating zoom, microsoft teams, and google meet for phishing scams. <https://www.theverge.com/2020/5/12/21254921/hacker-domains-impersonating-zoom-microsoft-teams-google-meet-phishing-covid-19>, 5 2020. (Accessed on 06/15/2020).

- [22] Lee Mathews. New phishing attacks prey on job loss fears with fake zoom meeting invites. <https://www.forbes.com/sites/leemathews/2020/04/28/new-phishing-attacks-prey-on-job-loss-fears-with-fake-zoom-meeting-invites/#562bd08d4602>, 4 2020. (Accessed on 06/16/2020).
- [23] James Richard Ortega et al. *HIDING IN PLAIN SIGHT? THE IMPACT OF FACE RECOGNITION SERVICES ON PRIVACY*. PhD thesis, 2019.
- [24] Aaron Krolik and Natasha Singer. A feature on zoom secretly displayed data from peoples linkedin profiles. <https://www.nytimes.com/2020/04/02/technology/zoom-linkedin-data.html>, 4 2020. (Accessed on 06/16/2020).
- [25] Drew Harwell. Thousands of zoom video calls left exposed on open web. <https://www.washingtonpost.com/technology/2020/04/03/thousands-zoom-video-calls-left-exposed-open-web/>, 4 2020. (Accessed on 06/16/2020).
- [26] Michael Fire, Dima Kagan, Aviad Elyashar, and Yuval Elovici. Friend or foe? fake profile identification in online social networks. *Social Network Analysis and Mining*, 4(1):194, 2014.
- [27] Samantha Cole. This open-source program deepfakes you during zoom meetings, in real time. [https://www.vice.com/en\\_us/article/g5xagy/this-open-source-program-deepfakes-you-during-zoom-meetings-in-real-time](https://www.vice.com/en_us/article/g5xagy/this-open-source-program-deepfakes-you-during-zoom-meetings-in-real-time), 4 2020. (Accessed on 06/16/2020).
- [28] Catherine Stupp. Fraudsters used ai to mimic ceos voice in unusual cyber-crime case. *The Wall Street Journal*, 30, 2019.
- [29] Taylor Lorenz, , and Davey Alba. zoombombing becomes a dangerous organized effort. <https://www.nytimes.com/2020/04/03/technology/zoom-harassment-abuse-racism-fbi-warning.html>, 4 2020. (Accessed on 06/15/2020).
- [30] More than 120 cases of child abuse zoombombing in uk being investigated — express & star. <https://www.expressandstar.com/news/uk-news/2020/05/18/more-than-120-cases-of-child-abuse-zoombombing-in-uk-being-investigated/>, 5 2020. (Accessed on 06/16/2020).
- [31] Jeremy Howard et al. fastai. <https://github.com/fastai/fastai>, 2018.
- [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [33] Neal Krawetz. Kind of like that - the hacker factor blog. <http://www.hackerfactor.com/blog/index.php/?archives/529-Kind-of-Like-That.html>, 1 2013. (Accessed on 06/17/2020).



- [34] ipazc/mtcnn: Mtcnn face detection implementation for tensorflow, as a pip package. <https://github.com/ipazc/mtcnn#id1>. (Accessed on 06/22/2020).
- [35] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016.
- [36] Facial recognition — microsoft azure. <https://azure.microsoft.com/en-us/services/cognitive-services/face/>. (Accessed on 06/17/2020).
- [37] Davis E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [38] Adam Geitgey. ageitgey/face\_recognition: The world’s simplest facial recognition api for python and the command line. [https://github.com/ageitgey/face\\_recognition](https://github.com/ageitgey/face_recognition). (Accessed on 06/17/2020).
- [39] yu4u/age-gender-estimation: Keras implementation of a cnn network for age and gender estimation. <https://github.com/yu4u/age-gender-estimation>. (Accessed on 06/17/2020).
- [40] Age. <https://apastyle.apa.org/style-grammar-guidelines/bias-free-language/age>. (Accessed on 06/17/2020).
- [41] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: an efficient and accurate scene text detector. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 5551–5560, 2017.
- [42] Canjie Luo, Lianwen Jin, and Zenghui Sun. A multi-object rectified attention network for scene text recognition. *arXiv preprint arXiv:1901.03003*, 2019.
- [43] Matthew Honnibal and Ines Montani. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear, 2017.
- [44] Loten Angus. For many, remote work is becoming permanent in wake of coronavirus. *The Wall Street Journal*, 05 2020.
- [45] Michael Fire and Rami Puzis. Organization mining using online social networks. *Networks and Spatial Economics*, 16(2):545–578, 2016.
- [46] Abigail Paradise, Rami Puzis, and Asaf Shabtai. Anti-reconnaissance tools: Detecting targeted socialbots. *IEEE Internet Computing*, 18(5):11–19, 2014.
- [47] Clare Garvie and Jonathan Frankle. Facial-recognition software might have a racial bias problem. *The Atlantic*, 7, 2016.

- [48] Aaron Holmes. These clothes and accessories outsmart facial recognition tech - business insider. <https://www.businessinsider.com/clothes-accessories-that-outsmart-facial-recognition-tech-2019-10#a-japanese-college-professor-designed-goggles-fitted-with-leds-that-thwart-facial-recognition-5>, 6 2020. (Accessed on 07/01/2020).
- [49] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 acm sigsac conference on computer and communications security*, pages 1528–1540, 2016.
- [50] Michael J Wilber, Vitaly Shmatikov, and Serge Belongie. Can we still avoid automatic face detection? In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016.