

Fueling AI with Public Displays? A Feasibility Study of Collecting Biometrically Tagged Consensual Data on a University Campus

Simo Hosio
Andy Alorwu
first.last@oulu.fi
University of Oulu
Oulu, Finland

Niels van Berkel
n.vanberkel@student.unimelb.edu.au
The University of Melbourne
Melbourne, Australia

Miguel Bordallo López
Mahalakshmy Seetharaman
Jonas Oppenlaender
first.last@oulu.fi
University of Oulu
Oulu, Finland

Jorge Goncalves
jorge.goncalves@unimelb.edu.au
The University of Melbourne
Melbourne, Australia




Figure 1: The user interface of our data collection application. From left: the front-facing splash screen with a call to action and a mandatory consent request, a video capture window with a countdown to zero, and a metadata entry screen to tag the video with additional information.

ABSTRACT

Interactive public displays have matured into highly capable two-way interfaces. They can be used for efficiently delivering information to people as well as for collecting insights from their users. While displays have been used for harvesting opinions and other content from users, surprisingly little work has looked into exploiting such screens for the consensual collection of tagged data that might be useful beyond one application. We present a field study where we collected biometrically tagged data using public kiosk-sized interactive screens. During 61 days of deployment time, we collected 199 selfie videos, cost-efficiently and with consent to leverage the videos in any non-profit research. 78 of the videos also had metadata attached to them. Overall, our studies indicate that people are willing to donate even highly sensitive data about themselves in public but that, at the same time, the participants had specific ethical and privacy concerns over the future of their data. Our study paves the

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

PerDis '19, June 12–14, 2019, Palermo, Italy

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6751-6/19/06...\$15.00
<https://doi.org/10.1145/3321335.3324943>

way forward toward a future where volunteers can ethically help advance innovations in computer vision research across a variety of exciting application domains, such as health monitoring and care.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**.

KEYWORDS

public displays, computer vision, field study, ethics

ACM Reference Format:

Simo Hosio, Andy Alorwu, Niels van Berkel, Miguel Bordallo López, Mahalakshmy Seetharaman, Jonas Oppenlaender, and Jorge Goncalves. 2019. Fueling AI with Public Displays? A Feasibility Study of Collecting Biometrically Tagged Consensual Data on a University Campus. In *Proceedings of the 8th ACM International Symposium on Pervasive Displays (PerDis '19)*, June 12–14, 2019, Palermo, Italy. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3321335.3324943>

1 INTRODUCTION

Public display research has matured significantly in recent years, and interactive deployments in particular are increasingly exploited for a growing array of exciting and novel use cases. As installing public screens becomes easier in general, many such uses will eventually find their way out of the research laboratories and into the wild. Currently, touch-enabled public displays are already commonly used for

way-finding and navigation inside malls and larger campus areas. They offer various types of directory services, interactive advertising, and it is not uncommon to stumble upon an interactive camera-equipped display for capturing and sending a digital, holiday-themed snapshot from a hotel lobby in your favourite holiday destination.

A distinct characteristic of public displays is that they attract users through serendipity [9, 14, 20, 25]. This makes them excel in collecting data in an ad-hoc fashion from passersby who are willing to donate their time and effort, most likely because they simply have some free time on their hands. In this paper, we investigate the potential of interactive displays in collecting biometrically tagged media from passersby. While such data may have many uses, we are interested in advancing AI and, more specifically, computer vision research. Computer vision requires high-quality training data, and situated technologies, including but not limited to public displays, may offer an avenue worth exploring for the collection of ecologically valid media.

We created a public display deployment that collects ‘selfie videos’. Our application asks the users to donate a 15-second video selfie. Then, the users can supplement the video by adding various metadata. While such collection of biometrically tagged data with situated technologies for reuse by potentially unknown algorithms may sound eerily dystopian for some, there are interesting scenarios that could be enabled by this type of data collection, such as training algorithms to detect and monitor health conditions — indeed, up to 30 different symptoms and medical conditions can be detected from observing human faces with a camera [35]. The key contributions of our work are:

- A dynamic, easy-to install setup to collect media files that are tagged with biometric metadata.
- A feasibility study that analyses the collected material and highlights important contextual aspects that must be considered in future deployments.
- Commentary and analysis of perceived ethical issues and potential new consent models that may be necessary in the future digital research ecosystems that exploit public displays as citizen-facing data collection interfaces.

In practice, we deployed a kiosk-sized display setup on the corridor of our university campus for a total of 61 days. We collected 199 15-second video selfies, 78 of which had voluntarily participant-labelled metadata attached to them. We assessed the validity of the collected videos for computer vision research both by manual inspection by domain experts as well as algorithmically through frame analysis. To gain further insights into the acceptability and ethics of this data collection approach, we collected 22 online questionnaire responses among the study participants. Overall, our results speak for the validity of the approach but also highlight a plethora of issues that need to be considered in future implementations.

2 RELATED WORK

2.1 Data Collection on Public Displays

Public displays have in recent years become popular means of conducting research, due to their low barrier of entry especially in situations where the tasks are simple and require little effort to complete. Their visibility makes them a good choice for creating awareness as well as inviting passersby for voluntary contributions to research [14, 25, 29].

To optimise the use of public displays, Alt et al. [3] suggest that public displays should be designed to fit both the social and the cultural context of the community where it will be deployed. Doing so ensures data quality and optimises for accuracy. Several types of data

can be collected using public displays. While it is common to use touch screens for input, also more novel interaction modalities such as gaze detection can be employed [19]. Public displays have been used to collect large-scale civic polling data through open-ended long format answers [15] or even classified ads directly using the screen [2]. Playfulness plays a crucial role in public display interactions and can be used in designing engaging applications [36]. For example, collecting situated snapshots by using the displays as camera devices have been a class of well-received applications [23]. The *mirror metaphor*, which entails participants seeing themselves or some parts of the screen responding to their body movement, seems to be a compelling playful design element [26].

2.2 Situated Crowdsourcing

Originally coined as a way to distribute large tasks as smaller chunks to the crowds online, crowdsourcing (CS) is now the primary means to collect high-quality data from human subjects at scale [30]. The online labour force is, however, highly self-selected and the markets are naturally by design limited to specific types of work [31]. A complement to online crowdsourcing is situated crowdsourcing, referring to purpose-built deployments and methods to elicit batches of contributions for highly specific tasks in typically geofenced contexts [8, 10, 16]. Public interactive displays excel in this type of work, due to their inherent characteristics of being location-bound and capable of self-advertising their contents to any potential curious passersby (e.g., [12, 14, 17, 25]).

Over the past decade, numerous crowdsourcing deployments have exploited public interactive screens. Just to provide one notable example, *Umati* is an augmented vending machine that was deployed on a university campus to crowdsource contributions from the student community [12]. The non-financial rewards (snacks from the vending machine itself) were found as sufficient to elicit high-quality contributions in expert work: grading exams.

2.3 Data Ethics

One visionary application domain of computer vision is health care and monitoring applications. Much progress has already been made in detecting health conditions based on video material only and more is expected to follow. Such future applications, however, are expected to rely not only on already available training data but increasingly on real-time, ubiquitous data collection methods and digital research ecosystems that are, as an unfortunate side effect, rapidly antiquating the concept of *informed consent* [6, 28] — the very ethical cornerstone of any research dealing with human contributions. As data collection is increasingly automated, human subjects may not be aware of data collection, how the data is processed, used as training data, or even monetised [6]. Concerning medical ethics in particular, participants may feel obliged to donate data under the current circumstances (e.g., enrolled in a study as a patient-participant) but reconsider afterwards [5]. How can one even give a truly informed consent, when algorithms mine data for presently unknown anomalies? All this calls for a change in data management to a more human-centric viewpoint [34] as well as thoroughly debating ethical issues with the data subjects themselves. Public displays are excellent citizen-facing access points to digital research ecosystems, as they allow for harvesting data from passersby fairly easily and cost-effectively. In our study, we set to collect data by asking an explicit consent first at the data collection point and then by initiating the discussion on ethics via an anonymous post-study questionnaire online.

3 SYSTEM DESIGN

In this section, we present the physical setup of our study and the videosourcing application that facilitated the collection of selfie videos and metadata.

3.1 The Physical Setup

We used a made-to-order desk with adjustable height and a circular wooden tabletop (60 cm in diameter) that hosts three Android tablet mounts. The tablet devices are enclosed in a metal casing (see Figure 2) and positioned at 120 degrees from each other, facing outward from the table. This setup makes it possible for 1–3 people to use the desk at the same time. Users cannot, however, easily see the screens of other users without consciously making an effort to peek by moving aside. We purchased a prepaid SIM card with unlimited data plan and used our own router, so that the deployment depended only on access to power and would not suffer from WiFi outages or poor connection quality. Similar setups have been used successfully for several situated studies, such as in the case of *UbiTable* to support easy access to extemporised face-to-face collaborations for small groups of people [32] or *TeamSourcer* in exploring team dynamics in a collocated crowdsourcing setup [17].

While there is no immediately apparent reason to feature three separate tablets in a study such as ours, and it would have been trivial to not include more than one tablet on the desk for the study, we opted to use three tablets for two reasons. First, our goal was to study feasibility from the viewpoint of media quality and not just to see if people donate any data. To this end, we hypothesised that having the displays — the front cameras of the mounted Android tablets — facing in different directions would help us explore the potential effect of different backgrounds that end up in the collected videos. Second, a desk with three tablets is a more disruptive element on the corridor, potentially evoking curiosity that typically attracts higher engagement [18]. In a way, this can be thought of as introducing a *novelty effect* on purpose, as our goal was to make people notice the deployment and our call to action.

3.2 VideoSourcing Application

We designed an Android application to facilitate the data collection: *VideoSourcing*. *VideoSourcing* is designed to be run on tablet devices that would later on act as our public kiosk-sized displays (see Figure 2).

3.2.1 Capturing Video Selfies. The user interface (UI) of *VideoSourcing* is straightforward, consisting of three stages, or activities, as depicted in Figure 1.

First, the splash screen features a large call-to-action as the heading with a smaller subheading that in a generic fashion simply invites users to ‘participate to win’ and help us in our research. Second, the splash screen contains two text areas that can be dynamically configured to provide task instructions. For instance, we can ask the users to make certain specific facial expressions. However, for this study the task was simply to ‘shoot a 15 second video selfie’ and ‘just be yourself,’ as can be seen in Figure 1. Most importantly, the splash page features a mandatory consent check box: any participant wishing to proceed forward from the splash page has to toggle a check box (unchecked by default), to give their consent for donating all of the subsequent data for scientific purposes. Finally, a link to a more nuanced data and privacy disclaimer (pop-up) as well as the email address of the responsible researcher was included in the footer of

the screen for participants to contact for questions or data removal requests.

After tapping the ‘start task’ button in the splash screen, a three second countdown timer is started. After the countdown is complete, the video recording begins. The user performs whatever task was assigned on the previous screen here (in this case, just shooting a 15-second selfie with instructions to ‘be yourself’). A new timer is then displayed at the top left corner of the screen informing the user of the time spent in performing the assigned task. After the duration set by the researcher has run out, the video recording stops (automatically) and the screen transitions to the final stage, which entails the collection of metadata. The video is first saved locally and then immediately uploaded to an Amazon S3 bucket, while the user is still in the metadata entry stage.

In the metadata entry screen, users are asked to provide basic demographic data (age, gender) and, for this study, we configured the screen to ask height, weight and a self-reported truthfulness value about the reported height and weight (using a scale from 1–7). We are aware that anyone can simply enter erroneous or mock data to all of this, but we were simply curious what would happen if we allow people to enter their personal data and then be conscious about its truthfulness, given how people tend to provide false data on their own height and weight online [27]. Further, using an open-ended text we asked the users to describe shortly, in one or two words, their current mood. Again such mood information and, to a degree, also the truthfulness information are interesting from the point of potentially teaching computers to detect these characteristics from faces in videos.

3.2.2 Server-Side Implementation Notes. *VideoSourcing* integrates the Amazon Web Services (AWS) SDK for Android to facilitate extremely robust as well as secure persistent storage by using a write-only API key. The metadata is stored in a MongoDB database and connected to the videos using a shared unique identifier that was linked to the videos through filenames of the videos. Notifications are managed using the OneSignal¹ SDK.

3.2.3 Dynamic Features. *VideoSourcing* was designed to be usable in different types of video collection studies far beyond this initial feasibility study. To this end, several parts of the UI are remotely configurable via push messages triggered through the OneSignal web dashboard or programmatically using the OneSignal API. The different configurable text fields are injected in the notifications as additional payload, and the moment that the device where *VideoSourcing* is installed receives the notification, i.e. the user does not have to even open the notification, *VideoSourcing* is configured with the new dynamic values. First, the task description (title + instructions) in the splash screen can be configured to encourage the user to perform a specific action or task while taking the video. Second, in the metadata entry screen both the slider as well as the open-ended items can be configured remotely. This way, the user can be asked one numerical value question (from 1–7) and one open-ended question. For instance, in this study we asked about the truthfulness of self-reported height and weight as well as a short open-ended description of the participant’s current mood.

¹<https://onesignal.com/>, accessed April 15, 2019



Figure 2: The *VideoSourcing* desk deployed at the corridor of our campus.

4 STUDY

4.1 Pilot Experiment

Before commencing a longer field deployment, we validated the technical solution and feasibility of data collection in two short studies: as an example of a technology probe during a lecture on Human-Computer Interaction and as a standalone installation in an open laboratory space within our own research premises. We used the deployment at this stage to undergo the local ethical review process, as we were not planning to change any major aspects of the study design itself. As a result, we included more comprehensive disclaimers to the application. During these studies, we discovered technical problems such as loss of data due to the devices not staying connected to the public WiFi network and minor usability problems with the screen flow of the application itself (*e.g.*, ambiguous wording in the final submit button).

In these pre-studies, we collected 11 videos and most likely lost several more due to the aforementioned technical challenges. The 11 participants (10 male, 1 female) also responded to a brief questionnaire online about the deployment. In the invitation, we clarified that the responses were not going to be connected with any of the data they donated earlier in the app itself and that the purpose of the study was to inform our longer future field study. From these results, we note two interesting findings: 1) the considered monetary value of their data (video and metadata combined) to range between 0-5 euros and 2) at least one of the respondents was willing to donate all of their data, including information on matters considered highly sensitive such as STDs and even serious health conditions. Insights like these were useful in designing the final questionnaire for the field deployment as well as in ensuring that the concept is feasible in general.

4.2 Field Deployment

Shortly after the pilot study and making small modifications based on the early findings to the setup itself and to the online questionnaire, we launched a field study on our University campus. To encourage passersby to interact with the setup, we provided printed A4-posters about a raffle that would take place after the study is over. We did not specify any guaranteed reward for everyone, nor did we specify

exactly how many lunch vouchers would we raffle. The desk was positioned at a somewhat quiet spot along a corridor so that it would not disturb the passersby too much by *e.g.* blocking a passageway. To ensure that the passersby could not tamper with the system and exit the *VideoSourcing* application, we used the *SureLock*² kiosk software for Android.

4.2.1 Additional Data Collection. For the participants who provided their email addresses in the metadata entry screen, we emailed a separate invite for the final online questionnaire. The invite was automatically sent via a dedicated Gmail account for this study and by using a *Zapier*³ automation flow (when a new document was added to the MongoDB, *Zapier* would extract it and trigger an outgoing email in Gmail). The questionnaire was hosted in Google Forms where users could contribute fully anonymously to avoid any ethical issues and contained items as follows.

- Basic demographic details (gender, age).
 - Perceived monetary value of the contributed data, separated by purpose of use: non-profit academic use or any use, including commercial purposes.
 - Choosing among a list of different types of data types that the participant would feel comfortable to disclose: age, gender, origin country, height, weight, sexual orientation, mood (happy, sad, neutral, etc.), non-serious medical conditions (flu symptoms, fever, etc.), serious health conditions (cancer, severe depression, diabetes, etc.), family details (number of children, married status, etc.), nothing or none of the aforementioned, and an open-ended option ‘Other...’.
 - Elaborating on the previous item: why and what?
 - Open-ended text fields specifically asking the participant’s opinions on any potential ethical issues of the setup as well as any perceived dangers or opportunities afforded by the data collection.
- Further, we probed for which purposes would the user be willing to donate data with a similar setup in the future.
- Strictly academic, non-profit research.
 - For commercial research (companies train algorithms that make money).
 - As part of a public free online dataset that contains everything – the video itself and all the data you provided us (age, gender, height, weight, ...).
 - Basically for anything, but I want to be compensated somehow (*e.g.* entering a prize draw, money, other rewards).
 - Basically for anything, no compensation needed.
 - Nothing, or none of the above.
 - Other...

5 RESULTS

Over the course of 61 days (4 days of pilot study + 57-day field study), we collected a total of 199 selfie videos, corresponding to 3 videos per day. This slightly exceeded our expectations, but at the same time we must note that our setup does not allow for calculating things like conversion ratio (of passersby or even from people starting to interact with the application) and we will leave this as future work. Further, we received 78 metadata submissions to supplement the videos. 63 of those left their email addresses, and of those 22 proceeded to provide online questionnaire responses (a 35% conversion ratio).

²<https://www.42gears.com/products/surelock/>, accessed April 15, 2019

³<https://zapier.com/>, accessed April 15, 2019

5.1 Metadata Analysis

Concerning the 78 metadata entries (48 male, 20 female, 5 other, 2 not disclosed; mean age 28.6, SD 12.2), we note that 62 people in total left diverse and rich open-ended mood information, such as “*super hungry*”, “*normal, little tired*”, “*fine, a bit confused*”, “*stressful because of the thesis*” or “*A little stressed, assignment due at midnight, slightly amused by the weight truth question :-)*”. 72 people provided information on their height and 71 on their weight. The average self-reported truthfulness, with no observed statistically significant differences between genders of age groups, was 5.4 (SD: 0.8).

5.2 Media Analysis

In order to statistically assess the quality of the faces collected in the videos and their usability for face biometrics and computer vision in general, we analysed them using a state-of-the-art face detector, based on the SSD-framework [21] and ResNet [11] as implemented in OpenCV [4]. The results of the automatic analysis show that:

- 179 videos (90% of the total) show a detected face during at least one full second, and are thus considered useful for several machine learning tasks as training data
- 113 videos contain a detected face during 100% of the duration of the video, 145 over 90% and 155 over 80%
- 20 videos do not contain a single detected face and could be discarded from a possible face database build from our results

While these can be considered remarkably good results and speak for the feasibility of the overall approach, an issue with some of the videos is that users not always show the whole face for the camera. Sometimes the users were short and either did not realise that the tablet stand can be tilted or simply did not bother to tilt it so that the full face would be visible. To assess the impact of this issue, we computed the detected facial sizes and their locations in the screen. In this context we considered faces to be too small and prone to limited usefulness if their size is below 130x130 pixels, and too big and prone to have occlusions or bad camera angles if their size is over 280x280 pixels. We also flagged the videos where the detected face is too close to the border since they are very likely to present occlusions or not full faces. This analysis shows that:

- 32 videos are perfect in terms of face size, screen location and detection percentage.
- 39 additional videos have faces with perfect face size and location over 70% of the time.
- 56 additional videos have perfect segments during at least five seconds.
- 57 videos have faces positioned too close to the bottom, 22 faces too close to the top, 15 have faces too small (11 of them with good location) and 68 have faces too big (22 with good location).
- Up to 100 videos could have benefited from better positioning of the user in respect to the camera.

By visually inspecting the videos, we noticed that sub-optimal illumination was a problem in up to 60 of the videos. Of them, 43 videos show different types of spotlights close to the facial area and thus have a limited dynamic range, while 17 additional videos present different problems such as back-light illumination or being captured in the dark. Also, hand movements around face and head caused issues (partially hidden facial features during such movements).

Finally, our design allows the same user to record several videos even in a short time span. We have manually identified 47 repeated users in the total 179 videos (132 unique users, 73.7%).

5.3 Insights from the Online Questionnaire

First, it was clear that most people answering the items about monetary value of the data had no real credible insights to offer, as the answers ranged anywhere from zero to hundreds of thousands of dollars, with no convergence points. This leaves us with an excellent chance to design a future study around the monetary valuation of selfie videos and their metadata, e.g., by using a reverse second-prize auction (also known as a *reverse Vickrey auction*) [22] that is ideal for extracting the real valuation in this type of scenarios of auctioning intangible goods [13, 33].

As for what type of information people would feel comfortable disclosing, we noticed significant differences. Of the 22 respondents, 21 would be willing to reveal their gender and 19 their age. This can be contrasted to e.g. height and weight (15 and 12, respectively), or to sexual orientation and non-serious health conditions (seven respondents for both). Five participants indicated to be willing to disclose serious health conditions and family details. Further, 19 participants were willing to donate data for strictly academic, non-profit research, whereas only five (5) were happy to donate data for commercial purposes. Zero participants chose to agree to “for practically anything, for no compensation whatsoever,” indicating that people seem to perceived at least some value to their data, regardless of the inability to clearly articulate such value in numbers.

5.3.1 Qualitative Results. An in-depth thematic analysis is out of the scope of this feasibility study but we provide a brief account of the key aspects and thoughts raised by the participants. First, the respondents provided interesting insights when elaborating on why they would donate or not donate data, and how they could be persuaded to disclose more about themselves. Naturally, many participants hesitated donating data if they do not fully trust the organization managing the collection or for concerns about the future use of the data, as exemplified by comments such as: “*I would tell more about myself if I knew exactly how that info will be used.*”, “*If I trusted the researcher and the organization that is dealing with my data, I would be willing to tell more.*”, or “*I am just scared information will go too public and will be bullied.*”

Some participants also took a far more nuanced approach in their responses: “*Age, gender and mood are very general. Also my height and weight aren’t that secure info for me. To make me give more info the use of it has to actually help someone in non monetary way.*” or “*I would be willing to disclose pretty much anything that doesn’t get me into trouble (e.g. drug use, or other criminal activities that I might be involved in). I wouldn’t want to disclose information that might be used to develop better mass-manipulation method; but if the motives are pure, e.g. to better understand the human condition, then why not. I could be further persuaded by informing me that my disclosures wouldn’t come back to bite me.*” and “*Yes [I do see a problem]. They might get lost or stolen and misused. Also the laws can change in 20 years and my data might be used in ways I didn’t agree. I can also become a public person and this data could be used against me.*”

Overall, helping research was seen as a good motive for donating data, however: “*If it’s used for non-profit research it’s fine, as long as people whose data is gathered are fine with it. If it’s for commercial purpose, I think some kind of compensation for those people is necessary (they helped you to make money after all).*” and “*If I was told that it was safe to do so and that I wouldn’t be identified I would share pretty much anything for research purposes. For other purposes I would be more careful and would demand more money.*” Further, one participant came up with a suggestion to not only do research but exploit the deployment as a civic feedback medium: “*This could be a great*

way for people to express their ideas about issues and not just be about personal data collection.”

Finally, only six of the 22 participants found potential issues or raised related additional questions concerning ethical aspects, such as: “However, am I allowed to have access to my data at any point in time should I request it?”, “People might end up sharing more than they are actually willing to and the information might end up somewhere they don’t want to.”, or “I would be willing to disclose my age and gender in a non profit or industrial research but I won’t like to disclose my nationality because of the racial reasons where it drags down to deal the data differently.”

6 DISCUSSION

Overall, our results speak favourably for the feasibility of the described data collection concept and the subsequent exploitation of the collected video and metadata in e.g. computer vision research. The setup functions well, technically speaking, and people do not hesitate to donate data with it. Further, people are willing to contribute — even when the rewards are not guaranteed: we employed the prize draw model to encourage participation without specifying how many prizes were available.

A common criticism toward the legal disclaimers protecting practically all of the contemporary data-intensive platforms online is their poor readability: users do not read them, nor do they fully understand them. Who is then to blame when things go wrong with personal data? Prior research has identified a clear dichotomy in the sense of responsibility when things backfire: people blame themselves for trusting any of those platforms in the first place and the platforms for poor conduct [7].

We deployed on a university campus and most likely benefited from the prevailing sense of trust in the academic context. Further, the users of our deployment could not even proceed to the video capture stage before agreeing to our simple but, as far as research goes, all-encompassing disclaimer that we may use all of their data for our own non-profit purposes. What was interesting, however, was that so many respondents expressed the willingness to provide much more detailed and sensitive information about themselves as long as they could trust the organization behind the data collection, the intention of data use, and the security of the platform. Nevertheless, we must remember that there is a mismatch between what people say and what they do in relation to data privacy [1].

Further studies (preferably outside the university environment) need to be conducted to ascertain whether our experiences obtained within the academic community apply outside the campus as well. The needed human-centric viewpoint [34] to ethical data collection has yet to be exhaustively discussed within the scope of public displays. We hope that our study is a *call to action* in itself concerning these aspects; How can we start leveraging public displays ethically for what they really excel in — collecting data serendipitously and largely on autopilot — for purposes beyond one single application at a time? Indeed, the whole prospect of data collection with ubiquitous technologies for common good via open science is an interesting proposition. The big corporations have all the data, and such data can be turned into profit. Is the academic line of work falling behind the curve? Is there a way to collaborate?

6.1 Design Implications

The user concerns about our intentions with the data as well as ethics concerns in general alert us to rethink the user-facing introduction

to the concept, the invitation to donate data. Indeed, “*why*” is more important than “*what*” [7] in explaining the intricacies of data collection. While a short disclaimer such as ours is perhaps technically all-encompassing, such deployments should in the very beginning of the user interaction flow include a mandatory disclaimer that proactively addresses the *why*-question. Why exactly are the data collected? And why is all this useful for research? This would in all likelihood increase participation, reduce concerns about data use and make the collection procedure more ethical and less suspicious in general.

Second, the design itself needs to optimise for the quality of the captured media. To encourage users to take videos where the head is positioned in the middle of the frame and correctly sized, a ‘target’ icon or simply drawing a square on the screen as an overlay would most likely *nudge* [24] participants to place the face in the centre. The square with a textual hint would help users adjust their distance from the camera and optimise the size of the face visible in the video.

Finally, while we explicitly set out to place our deployment in a well-lit corridor, poor illumination was still a problem in some instances. More specifically, the ceiling lights were often visible in the background, which may degrade the performance of computer vision algorithms. Therefore, similar deployments should be placed in physical contexts where there is ample ambient light available and that encourage users to stare horizontally at the camera to avoid sub-optimal camera angles.

6.2 Future Work

We plan to extend the deployment first by ensuring granular data use rights, by letting the users indicate their preferences prior to taking a video and donating metadata. Second, we plan to crowdsource additional labels by automating the labelling process via Amazon’s Mechanical Turk. How other humans perceive certain attributes (e.g., age, gender, origin, mood) is a fascinating challenge in the field of computer vision. The actual *chef d’oeuvre* of this work in our vision is a longitudinal and distributed across several campuses deployment that provisions an open database online for researchers to use. People already upload selfies online daily, for no apparent scientific benefit but rather for the profit of commercial platforms, i.e. social networking services. This type of non-profit, academic endeavour would certainly be something we are interested in collaborating and hope to find willing partners to work with in PERDIS ’19.

7 CONCLUSION

We investigated exploiting interactive kiosk-sized displays for collecting short selfie videos with metadata. We conclude the approach as feasible since passersby started donating videos. Further, a remarkably high percentage of the videos contained the participants’ faces and are therefore usable for training algorithms that, for instance, focus on learning about the aspects collected among the metadata. Concerning ethics and data collection in general, our participants did express concerns, but as is also evident from the number of videos we collected, they were perhaps even surprisingly willing to contribute data, as long as the data was used for non-profit academic research and guaranteed to not leak to commercial use.

ACKNOWLEDGMENTS

This research was funded partially by the Australian Government through the Australian Research Council (project code DP190102627) and The Melbourne School of Engineering’s Visiting Fellows scheme.

REFERENCES

- [1] Alessandro Acquisti and Ralph Gross. 2006. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *International workshop on privacy enhancing technologies*. Springer, 36–58.
- [2] Florian Alt, Thomas Kubitzka, Dominik Bial, Firas Zaidan, Markus Ortel, Björn Zuurma, Tim Lewen, Alireza Sahami Shirazi, and Albrecht Schmidt. 2011. Digified: insights into deploying digital public notice areas in the wild. In *Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia*. ACM, 165–174.
- [3] Florian Alt, Nemanja Memarovic, Ivan Elhart, Dominik Bial, Albrecht Schmidt, Marc Langheinrich, Gunnar Harboe, Elaine Huang, and Marcello P Scipioni. 2011. Designing shared public display networks—implications from today’s paper-based notice areas. In *International Conference on Pervasive Computing*. Springer, 258–275.
- [4] G. Bradski. 2000. The OpenCV Library. *Dr. Dobbs’s Journal of Software Tools* (2000).
- [5] Lawrence H Brenner, Alison Tytell Brenner, and Daniel Horowitz. 2009. Beyond informed consent: educating the patient. *Clinical orthopaedics and related research* 467, 2 (2009), 348–351.
- [6] Markus Christen, Josep Domingo-Ferrer, Dominik Herrmann, and Jeroen van den Hoven. 2017. Beyond Informed Consent—Investigating Ethical Justifications for Disclosing, Donating or Sharing Personal Data in Research. In *Philosophy and computing*. Springer, 193–207.
- [7] Casey Fiesler and Blake Hallinan. 2018. We Are the Product: Public Reactions to Online Data Sharing and Privacy Controversies in the Media. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 53.
- [8] Jorge Goncalves, Denzil Ferreira, Simo Hosio, Yong Liu, Jakob Rogstadius, Hannu Kukka, and Vassilis Kostakos. 2013. Crowdsourcing on the Spot: Altruistic Use of Public Displays, Feasibility, Performance, and Behaviours. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp ’13)*. ACM, New York, NY, USA, 753–762.
- [9] Jorge Goncalves, Simo Hosio, Yong Liu, and Vassilis Kostakos. 2014. Eliciting situated feedback: A comparison of paper, web forms and public displays. *Displays* 35, 1 (2014), 27–37.
- [10] Jorge Goncalves, Simo Hosio, Niels van Berkel, Furqan Ahmed, and Vassilis Kostakos. 2017. CrowdPickUp: Crowdsourcing Task Pickup in the Wild. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 51 (Sept. 2017), 22 pages.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [12] Kurtis Heimerl, Brian Gawalt, Kuang Chen, Tapan Parikh, and Björn Hartmann. 2012. CommunitySourcing: engaging local crowds to perform expert work via physical kiosks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1539–1548.
- [13] Simo Hosio, Denzil Ferreira, Jorge Goncalves, Niels van Berkel, Chu Luo, Muzamil Ahmed, Huber Flores, and Vassilis Kostakos. 2016. Monetary Assessment of Battery Life on Smartphones. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI ’16)*. ACM, New York, NY, USA, 1869–1880. <https://doi.org/10.1145/2858036.2858285>
- [14] Simo Hosio, Jorge Goncalves, and Vassilis Kostakos. 2013. Application discoverability on multipurpose public displays: popularity comes at a price. In *Proceedings of the 2nd ACM International Symposium on Pervasive Displays*. ACM, 31–36.
- [15] Simo Hosio, Jorge Goncalves, Vassilis Kostakos, and Jukka Riekk. 2015. Crowdsourcing Public Opinion Using Urban Pervasive Technologies: Lessons From Real-Life Experiments in Oulu. *Policy & Internet* 7, 2 (2015), 203–222.
- [16] Simo Hosio, Jorge Goncalves, Vili Lehdonvirta, Denzil Ferreira, and Vassilis Kostakos. 2014. Situated Crowdsourcing Using a Market Model. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST ’14)*. ACM, New York, NY, USA, 55–64.
- [17] Simo Hosio, Jorge Goncalves, Niels van Berkel, Simon Klakegg, Shin’ichi Konomi, and Vassilis Kostakos. 2018. Facilitating Collocated Crowdsourcing on Situated Displays. *Human-Computer Interaction* 33, 5–6 (2018), 335–371.
- [18] Steven Houben and Christian Weichel. 2013. Overcoming Interaction Blindness Through Curiosity Objects. In *CHI ’13 Extended Abstracts on Human Factors in Computing Systems (CHI EA ’13)*. ACM, New York, NY, USA, 1539–1544. <https://doi.org/10.1145/2468356.2468631>
- [19] Mohamed Khamis, Ludwig Trotter, Markus Tessmann, Christina Dannhart, Andreas Bulling, and Florian Alt. 2016. EyeVote in the wild: do users bother correcting system errors on public displays?. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia*. ACM, 57–62.
- [20] Vassilis Kostakos, Hannu Kukka, Jorge Goncalves, Nikolaos Tselios, and Timo Ojala. 2013. Multipurpose Public Displays: How Shortcut Menus Affect Usage. *IEEE Computer Graphics and Applications* 33, 2 (2013), 56–63.
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*. Springer, 21–37.
- [22] R Preston McAfee and John McMillan. 1987. Auctions and bidding. *Journal of economic literature* 25, 2 (1987), 699–738.
- [23] Nemanja Memarovic, Ava Fatah gen. Schieck, Holger Schnädelbach, Efstathia Kostopoulou, Steve North, and Lei Ye. 2016. Longitudinal, Cross-site and ‘in the Wild’: A Study of Public Displays User Communities’ Situated Snapshots. In *Proceedings of the 3rd Conference on Media Architecture Biennale (MAB)*. ACM, New York, NY, USA, Article 1, 10 pages. <https://doi.org/10.1145/2946803.2946804>
- [24] Tobias Mirsch, Christiane Lehrer, and Reinhard Jung. 2017. Digital nudging: Altering user behavior in digital environments. *Proceedings der 13. Internationalen Tagung Wirtschaftsinformatik (WI 2017)*(2017), 634–648.
- [25] Jörg Müller, Florian Alt, Daniel Michelis, and Albrecht Schmidt. 2010. Requirements and design space for interactive public displays. In *Proceedings of the 18th ACM international conference on Multimedia*. ACM, 1285–1294.
- [26] Jörg Müller, Robert Walter, Gilles Bailly, Michael Nischt, and Florian Alt. 2012. Looking glass: a field study on noticing interactivity of a shop window. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 297–306.
- [27] Isabelle Niedhammer, I Bugel, Sébastien Bonenfant, M Goldberg, and Annette Leclerc. 2000. Validity of self-reported weight and height in the French GAZEL cohort. *International journal of obesity* 24, 9 (2000), 1111.
- [28] Onora O’Neill. 2003. Some limits of informed consent. *Journal of medical ethics* 29, 1 (2003), 4–7.
- [29] Jonas Oppenlaender and Simo Hosio. 2019. Experizone: Integrating Situated Scientific Experimentation with Teaching of the Scientific Method. In *Adjunct Proceedings of the 2019 ACM CHI Conference on Human Factors in Computing Systems (CHI ’19)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3290607.3313043>
- [30] Aditya Parameswaran, Akash Das Sarma, and Vipul Venkataraman. 2016. Optimizing Open-Ended Crowdsourcing: The Next Frontier in Crowdsourced Data Management. *Bulletin of the Technical Committee on Data Engineering* 39, 4 (2016), 26.
- [31] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* 70 (2017), 153–163.
- [32] Chia Shen, Katherine Everitt, and Kathleen Ryall. 2003. UbiTable: Impromptu face-to-face collaboration on horizontal interactive surfaces. In *International Conference on Ubiquitous Computing*. Springer, 281–288.
- [33] Jacopo Staiano, Nuria Oliver, Bruno Lepri, Rodrigo de Oliveira, Michele Caraviello, and Nicu Sebe. 2014. Money walks: a human-centric study on the economics of personal mobile data. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 583–594.
- [34] Melanie Swan. 2012. Health 2050: The realization of personalized medicine through crowdsourcing, the quantified self, and the participatory biocitizen. *Journal of personalized medicine* 2, 3 (2012), 93–118.
- [35] Jérôme Thevenot, Miguel Bordallo López, and Abdenour Hadid. 2018. A survey on computer vision for assistive medical diagnosis from faces. *IEEE journal of biomedical and health informatics* 22, 5 (2018), 1497–1511.
- [36] Martin Tomitsch, Christopher Ackad, Oliver Dawson, Luke Hespanhol, and Judy Kay. 2014. Who cares about the content? An analysis of playful behaviour at a public display. In *Proceedings of The International Symposium on Pervasive Displays*. ACM, 160.