

Individuality and Fairness in Public Health Surveillance Technology: A Survey of User Perceptions in Contact Tracing Apps

Ellen Hohma^{ID}, Ryan Burnell, Caitlin C. Corrigan^{ID}, and Christoph Luetge

Abstract—Machine learning algorithms are playing an increasingly important role in public health measures, accelerated by the Covid-19 pandemic. It is therefore vital that machine learning algorithms are applied in ways that are generally considered fair. However, the question of how to define fairness in a public health context is still an open one. In this study, we investigated people’s attitudes towards two ways of defining fairness in the context of Covid-19 contact tracing apps. In the first, ‘high-individuality’ approach, the likelihood of an algorithm asking a person to self-isolate would depend on the person’s individual characteristics, such as their risk of spreading the virus through regular contacts. In the second ‘low individuality’ approach, these individual characteristics would not be used to come to a decision. For each approach, participants rated its fairness, overall quality, and their privacy concerns, and answered questions about basic psychological need satisfaction. Participants rated the high-individuality approach as fairer and better overall compared to the low-individuality approach, despite having greater privacy concerns. Further, we found a strong correlation between the participants’ fairness perceptions and their overall impression of the tracking tool. Together, these findings suggest that people prefer individualised approaches in some contexts and perceive them as fairer. However, policy makers should consider the privacy trade-off of employing such measures.

Index Terms—Algorithmic decision-making, contact tracing, data privacy, fairness, individuality, machine learning, public health surveillance.

I. INTRODUCTION

THE OUTBREAK of the Covid-19 pandemic has highlighted the importance of continuous development and advancement of technologies to oppose health crises and to maintain public security and well-being. Multiple innovative approaches have quickly been developed in an effort to contain the spread of the virus, many equipped with Artificial Intelligence (AI) and Machine Learning (ML). In particular,

Manuscript received 31 January 2022; revised 17 May 2022 and 26 September 2022; accepted 26 September 2022. Date of publication 30 September 2022; date of current version 15 December 2022. This work was supported in part by the Technical University of Munich’s Institute for Ethics in Artificial Intelligence (IEAI) and in part by the Leverhulme Centre for the Future of Intelligence. (Corresponding author: Ellen Hohma.)

Ellen Hohma, Caitlin C. Corrigan, and Christoph Luetge are with the Institute for Ethics in Artificial Intelligence, School of Social Sciences and Technology, Technical University of Munich, 80333 Munich, Germany (e-mail: ellen.hohma@tum.de; c.corrigan@tum.de; luetge@tum.de).

Ryan Burnell was with the Dyson School of Design Engineering, Imperial College London, SW7 2BX London, U.K. He is now with the Leverhulme Centre for the Future of Intelligence, University of Cambridge, CB2 1TN Cambridge, U.K. (e-mail: rb967@cam.ac.uk).

Digital Object Identifier 10.1109/TTS.2022.3211073

the public health surveillance sector has seen a flood of tools that employ ML techniques to support public health monitoring, such as video surveillance for mask regulation compliance and fever or quarantine verification checks [1]. However, with the rapid development and deployment of these tools, ethical concerns about them have spread quickly, including concerns about whether they use individuals’ data in ways that are fair.

In addition to the obvious ethical issues surrounding algorithmic fairness, the extent to which ML algorithms are fair might have practical consequences if it affects people’s willingness to use AI-enabled tools or their acceptance of the outcomes recommended by the algorithms. For example, in the context of organisations, there is evidence for an interrelation between fairness concerns and overall satisfaction. Martinez-Tur et al. [2], for instance, showed that perceived distributive justice of gastronomy services (i.e., the perceived fairness of the outcome) was the primary determinant of customer satisfaction. Studying the effect of post-complaint behaviour, Blodgett et al. [3] concluded that although, in their case, distributive justice did not have an impact on complainant’s satisfaction, the way in which the outcome was communicated could compensate unfair treatments. Similar evidence was further found in the relationship between organisations and employees. Sudin [4], for example, observed that distributive justice has a significant impact on overall employee satisfaction when studying performance appraisal processes. In the case of public health measures, it is therefore vital to ensure that people view ML algorithms as fair to improve the overall acceptability of such measures.

However, there are many challenges to ensuring that ML algorithms are considered fair. For instance, as ML algorithms are designed to predict outcomes based on input data, any biases in the input will lead to biased outcomes [5], [6]. In addition, the algorithmic design itself can produce biases because the underlying model chosen for a ML based system is a crucial factor for determining the outputs [7]. A further problem stems from the difficulty in defining “fairness” in different contexts. ML researchers have proposed a variety of fairness definitions that could be used to guide the design and evaluation of algorithms. But deciding which definition is best is not always easy. In particular, one critical issue for determining the right notion of fairness is the intangibility of the concept itself, even in anthropological and psychological studies [8]. Especially problematic is the fact that people’s beliefs about what is fair differ depending on the context [9].

One key part of algorithmic design that is closely related to people's perceptions of fairness is how an algorithm draws on personal information about individuals. Essentially, there are two contrary extremes. In 'high-individuality' approaches, algorithms make use of personal information about people to treat them according to their personal needs. By contrast, 'low-individuality' approaches consider people as a homogeneous group, treating everyone in that group the same. This has led ML researchers to categorise previously developed fairness concepts into Individual Fairness models aiming for similar predictions for similar individuals, Group Fairness models treating different groups equally, and Subgroup Fairness models—a combination of the former two—that categorize individuals based on their personal features into subgroups and ensuring group fairness constraints for those subgroups [6]. In many situations individual fairness models seem the most appropriate and fair. The pandemic in particular has brought forward many examples of how people would like to see more actions tailored specifically to their circumstances, such as vaccination status for contact restrictions. At the same time, concerns have been raised as to how much data can be justifiably requested, particularly when it contains sensitive information (e.g., debates around compulsory vaccination at the workplace). This shows that balancing different ethical principles can be challenging, but it is vital because these factors are linked to the users' uptake and acceptance of tools. Of course, in order for 'high-individuality' algorithms to treat people according to their needs, they sometimes require more extensive data on users. Using more data in the decision-making process might be seen as an invasion of privacy in certain situations [10], [11]. Some researchers, e.g., [12], even argue that fairness always comes at the cost of privacy. This trade-off between individuality and data privacy, often referred to as the personalisation-privacy paradox, has been frequently identified and studied in literature, e.g., [13], [14], [15], [16]. Still, even holding data collection equal, we might expect that the ways in which data are used might affect perceptions of fairness.

The distinction between high and low individuality approaches is highly linked to the debate on equity vs. equality in distribution decisions. The concept of equity is based on the equity theory by Adams [17] and refers to treating individuals according to their needs in a way to ensure equal outcomes (e.g., using affirmative action to assist disadvantaged groups). Equality, by contrast, involves treating all individuals the same—for example by giving everyone equal amounts of resources, even if this ultimately leads to inequality of outcomes. Views on which approach is fairer differ and depend greatly on the context, e.g., [18], [19], [20]. It is therefore vital that we understand whether people think high- or low-individuality approaches are fairer across different contexts.

In the context of public health measures, there are some data that speak to this issue. For example, Srivastava et al. [21] found that demographic parity (i.e., having the same probability for a positive outcome regardless of an individual's group membership) was most appealing to participants, suggesting a preference for low-individuality approaches. On the other

hand, there is an emerging trend towards increasingly individualised medical treatments in the healthcare sector. Recent advancements in Individualised Medicine have made it possible to classify patients into subgroups based on their clinical characteristics instead of treating them as one homogenous group [22]. Taking this even further, Personalised Medicine—which involves practices such as analysing the patient's genome and resulting predictions on the patient's future health risks—has become a realistic possibility [23]. These approaches highlight the ongoing trend towards more individuality in data collection and processing in healthcare, as well as the resulting shift towards a rising focus on need-based treatment. In general, patients appear to support and value these personalised approaches, as they emphasize the uniqueness of each medical case and the corresponding individuality required to provide appropriate treatment [24].

But it remains unclear whether people feel the same when it comes to health control measures such as those being implemented in response to the Covid-19 pandemic. For example, many countries were using contact-tracing apps that rely on algorithms to decide who was required to self-isolate based on their contact with Covid-positive individuals. For such apps to work effectively, they need mass adoption among the population [25]. However, adherence to these apps tended to be low. Among others, Walrave et al. [26] have studied factors that influence people's adoption intention for using contact tracing apps. They found that fewer than 50% of their surveyed participants intended to use these apps, with app-related privacy concerns being one factor that negatively influenced users' intentions [26]. However, it remains unclear whether people consider it fair to use highly individualised approaches to decide who needs to isolate, or if they prefer low-individuality approaches that use fewer personal data and treat people more uniformly. We address this gap in the study reported here.

In addition, it is important to consider why people think particular application approaches are fairer than others. One factor that might play a role can be found in self-determination and its link to basic psychological need satisfaction. According to the Self-Determination Theory, an individual's motivation and engagement can be stimulated through three basic psychological needs: autonomy, competence and relatedness [27]. Autonomy reflects the extent to which individuals feel they are acting according to their own volition, willingness and choice. Competence reflects feelings of effectiveness and the capability of achieving important goals. Finally, Relatedness captures the feeling of being connected to and cared for by others. These basic needs have been associated with perceptions of fairness in organisational contexts. For instance, Olafsen et al. [28], found that employees' basic need satisfaction ratings were related to their judgments of the extent to which companies' payment distribution procedures are fair. Haar and Spell [29] found evidence for job autonomy to directly influence job satisfaction and, at the same time, moderate the relation between distributive justice and job satisfaction. Similar results were reported by Aryee et al. [30], who found a significant influence of justice on need satisfaction, which in turn was positively associated with intrinsic

motivation. These findings indicate that fairness can promote basic psychological need satisfaction. In the context of public health surveillance, satisfying the three basic needs, and hence, stimulating intrinsic motivation, could encourage uptake of such tools and, ultimately, increase their overall effectiveness. Therefore, we sought to determine the relationship between perceptions of fairness and psychological need satisfaction.

To do so, we focused on public health instruments that were developed during the Covid-19 crisis, with the specific use case of contact tracing apps. We investigated how incorporating more individuality to a decision-making process—in this case, the risk of spreading the virus a person poses to others—affects people’s perceptions of fairness and quality of such tools. We also investigated how these fairness perceptions related to basic psychological need satisfaction and frustration.

II. RESEARCH METHOD

An online vignette study using a within-subject design was conducted. To test our hypotheses on a prominent and real-world example from public health surveillance, contact tracing technology was chosen as a use case. To test the two extreme approaches, ‘high-individuality’ and ‘low-individuality’, we derived two policies on how contact tracing applications could use personalised data to determine who should be asked to self-isolate. Since contact-tracing apps were developed in various ways in different countries during the pandemic, people have divergent previous experiences with such tools according to their origin and place of residence. Because of this, we collected data from two countries: the U.K. and Germany. The two cases were selected based on the feasible access of the researchers to survey subjects in these countries, the similar design and policies surrounding the two countries’ national contact tracing apps.

This research received ethical approval from Imperial College London’s Research Governance and Integrity Team.

A. Participants

Participants were recruited through the Prolific survey platform—any participants who lived in the U.K. or Germany and who speak English were invited to participate. Not knowing how big the effect of our manipulation would be, we aimed to collect data from 150 participants from each country. Participants were only excluded if they failed one or both of the attention checks in the survey ($n = 31$). In total, 273 participants were considered in the analysis, 129 from the U.K. and 144 from Germany. The mean age of the participants was 28.01 ($SD = 8.43$). In terms of their highest level of education, 123 participants had finished secondary school, 92 held a Bachelor’s, 38 held a Master’s and 20 held a Doctoral degree.¹

B. Procedure

The survey was conducted in English for all participants. Participants read and rated, in a random order, two different approaches for how contact tracing apps could determine

¹We found no relationship between people’s age or education and their ratings of the approaches.




Fig. 1. Distribution of averaged fairness perception per participant for the low individuality approach (left) and high individuality approach (right).

who should be required to self-isolate following contact with infected individuals. In the “high-individuality” approach, the algorithm would consider the risk of a person spreading the virus to others in deciding whether to send that person to self-isolation. By contrast, in the “low-individuality” approach, the algorithm would not consider the risk of spreading the virus in its decision.

Because the primary purpose of this study was to investigate how the *usage* of data affects perceived fairness, we held constant the *amount* of data collected across conditions. Therefore, both approaches mentioned that the app collects on location and contacts with others—the only difference was whether these data would be used to make decisions about self-isolation. For each algorithmic approach, participants were asked to rate how fair they perceived it, to what extent they had privacy concerns about the approach, whether they would be in general satisfied with such an approach, and how much the policy supported or frustrated their feelings of autonomy, competence, and relatedness. Measurement items for perceived fairness in the algorithmic decision-making process were based on Wang et al. [31], perceived privacy on Roca et al. [32], and adapted to this study’s context. Need satisfaction and frustration was measured by adapting items from Peters et al. [33]. The full wording of the items is displayed in Table I. After completing these questions, participants were asked to provide demographic data, as well as answer some control variables as proposed by Wang et al. [31]. The survey ended with questions on the real Covid-19 contact tracing apps of the respective countries (NHS Covid-19 Tracing App for U.K., Corona-Warn-App for Germany), but because these data are not central to our research question, we do not report them here.

III. RESULTS

To test whether the level of individuality in decision-making affects fairness perceptions, we compared the mean fairness ratings for the two proposed app approaches. Fig. 1 shows the distribution of the averaged fairness ratings for each participant for the low individuality approach (left) and high individuality approach (right). We found that participants rated the high individuality approach as significantly slightly fairer than the

TABLE I
OVERVIEW OF THE ITEMS USED TO MEASURE STUDIED PERCEPTIONS

	Items
Fairness	In general, the above approach is fair
	The above approach is fair to people who are asked to self-isolate
	The above approach is fair to people who are at risk of serious consequences of COVID-19
	The above approach will lead the app to make fair decisions about who to send to self-isolation
Privacy	I am concerned that an app following the above approach will use my personal information for other purposes without my authorization
	I think that too much of my personal information will be collected by an app following the above approach
	I am concerned about the privacy of my personal information when using an app following the above approach
	I am concerned that personal information about me collected by an app following the above approach will be shared with other entities without my authorization
Autonomy	
	<i>Satisfaction</i>
<i>Satisfaction</i>	The above approach would give people a sense of choice and freedom in the things they undertake
	The above approach would allow people's decisions to reflect what they really want
<i>Frustration</i>	The above approach would get in the way of people's ability to make their own choices
	The above approach would be controlling
Competence	
	<i>Satisfaction</i>
<i>Satisfaction</i>	The above approach is easy to understand
	The above approach would be easy to implement
<i>Frustration</i>	The above approach would get in the way of people's ability to achieve their goals
	The above approach would reduce people's confidence in their abilities
Relatedness	
	<i>Satisfaction</i>
<i>Satisfaction</i>	The above approach would help people feel close to people they care about
	The above approach would help people feel connected to other people
<i>Frustration</i>	The above approach would make it harder for people to connect in a meaningful way
	The above approach would make people feel excluded from groups they want to belong to

low individuality approach, $M_{diff} = 0.27$, 95% CI [0.04, 0.51], which supports the hypothesis that increasing the degree of information individuality improves perceptions of fairness.

Moreover, individuality impacted participants' overall evaluations of the approaches—they rated the high individuality approach as significantly better overall than the low individuality approach, $M_{diff} = 0.23$, 95% CI [0.03, 0.43].

We also found that, within each condition, there was a strong correlation between participants' ratings of fairness and their overall evaluations of the approach, $r_{LowIndividuality(271)} = 0.74$, 95% CI [0.68, 0.79], $r_{HighIndividuality(271)} = 0.75$, 95% CI [0.69, 0.79], supporting our hypothesis that positive perceptions of overall fairness can increase satisfaction with public health technologies.

Next, we examined the impact of our manipulation on basic psychological need satisfaction and frustration.

TABLE II
DIFFERENCES IN NEED SATISFACTION AND FRUSTRATION BY CONDITION

		M_{diff}	95% CI
Autonomy	Satisfaction	0.13	[-0.05, 0.30]
	Frustration	-0.01	[-0.16, 0.15]
Competence	Satisfaction	-0.44***	[-0.27, -0.61]
	Frustration	-0.13*	[-0.01, -0.25]
Relatedness	Satisfaction	0.06	[-0.09, 0.20]
	Frustration	0.07	[-0.07, 0.21]

Note: *** indicates $p < 0.001$. * indicates $p < 0.01$. Positive mean differences represent higher ratings for the high individuality group.

TABLE III
BETA WEIGHTS FROM REGRESSIONS PREDICTING OUTCOMES

		Fairness Perception	Overall Impression
Autonomy	Satisfaction	0.18	0.2***
	Frustration	0.08	-0.01
Competence	Satisfaction	0.25***	0.30***
	Frustration	-0.22***	-0.22***
Relatedness	Satisfaction	0.18***	0.19***
	Frustration	0.03	-0.02
Privacy Concerns		-0.02	-0.10***

Note: *** indicates $p < 0.001$.

As Table II shows, we found no significant differences between the approaches in terms of autonomy satisfaction or frustration, nor in relatedness satisfaction or frustration. However, we found that the low individuality approach was rated significantly higher on both competence satisfaction and frustration.

Next, we conducted a linear regression with perceived fairness and overall impression as the dependent measures to examine their relationship with basic psychological needs. Table III presents the results of this regression. We found similar patterns for both dependant variables. For basic need satisfaction, competence satisfaction and relatedness satisfaction were both significantly, positively related to fairness and overall user impression. Autonomy satisfaction was only significantly, positively related to the user's overall impression. Investigating its counterpart basic need frustration revealed that competence frustration was negatively related to fairness as well as overall user impression. No evidence was found that autonomy frustration or relatedness frustration are related to perceptions of fairness or overall ratings. In other words, our results support the hypothesis that positive perceptions of fairness can stimulate basic psychological need satisfaction, although only for competence and relatedness.

Finally, we studied the effect of perceived data privacy on the participants' perceptions of fairness and overall satisfaction with the application. We found that people reported slightly greater privacy concerns for the high individuality approach compared to the low individuality approach, $M_{diff} = 0.15$, 95% CI [0.06, 0.24]. We also found a significant, negative effect of privacy concerns on the overall user impression in both conditions. However, we found no evidence that privacy concerns are related to perceptions of fairness in our study.

These results suggest that perceived data privacy is related to evaluations of the proposed tools.

We found no differences in perceptions of fairness, overall rating, or privacy concerns between the U.K. and Germany ($p > .41$). Therefore, we combined the data from these two countries for the main analyses.

IV. DISCUSSION

The study's aim was to examine the extent to which people think two different approaches to ML-based public health surveillance technologies are fair. In our study, we found evidence that participants preferred high-individuality approaches to contact-tracing—participants rated the high-individuality approach as both fairer and better overall. Moreover, we found a strong correlation between participants' fairness perceptions and their overall impression, suggesting that perceptions of fairness are tightly linked with people's evaluation of public health tools. However, we did not find evidence that need satisfaction can explain these effects.

A. Ethical Implications

Issues of justice and fairness have been emphasized repeatedly in ethical frameworks for healthcare and AI-based tools [34]. Accelerated by the spread of the Covid-19 pandemic, recent literature has identified numerous instances of bias and unfairness in public health surveillance due to the fact that these technologies collect and analyse large amounts of data, often including socio-economic information such as race, ethnicity, gender or political affiliation [35]. For example, biased data collection strategies can result in subgroups not being visible or being stigmatised as they lack the needed technical devices [36], mobile communication or Internet access [37], [38]. The inevitable trade-off between individual freedom and civil obligations necessitates a delicate balance between collecting all the information needed to best protect public benefits while avoiding discrimination of certain populations [35]. Essentially, this leads to the dilemma of how far we can limit personal freedoms for the public benefit that has driven many controversies during the Covid-19 pandemic.

Although we did not explicitly study individuality with regard to demographic characteristics, our findings suggest that users, at least in the U.K. and Germany, value a more personal treatment based on their individual characteristics in health surveillance applications. While people feel discriminated when judgement is based on their demographic attributes, they seem to likewise feel treated in an unfair way if they are regarded as a fully anonymous, homogeneous group. Clearly, more work is therefore needed to determine the individualized uses of data that people see as discriminatory and those that contribute to positive perceptions of fairness. While in general, it is likely that individual and unchangeable traits, such as gender or race, might be counted among the discriminatory ones, our findings suggest that personal parameters resulting from an individual's actions instead of traits—such as in the context of this study the risk of spreading the virus that a person poses to others—might be among those features where disclosure is accepted to enable a fairer decision.

However, the further and more concrete identification of such a distinction of features can foster the pursuit of a solution to the question of how to balance personal rights against the well-being of the broader society.

In this study, we found that people preferred the highly individualized approach despite reporting greater privacy concerns regarding the use of data with that approach. This finding indicates that, in certain contexts, people might consider some invasions of privacy or limitations of freedoms fair and acceptable, at least to the extent that they are important for public health. Of course, in this regard, context is crucial. Nissenbaum [39] argues that contextual integrity is the benchmark of privacy, and consent to the use of data is only given in relation to its respective circumstances. Empirical field studies and scenario-based surveys, such as [40] or [41], support this notion. Perhaps, then, our participants were willing to sacrifice some data privacy because they viewed the high-individualization approach as fairer. The circumstances under which privacy is seen as an acceptable trade-off for fairness is worth of further investigation.

It is also worth thinking about how individual differences might affect people's preferences and perceptions of fairness. For example, individualistic persons or cultures put a higher focus on personal autonomy and self-fulfilment and base identity on themselves as well as their personal achievements [42]. By contrast, collectivistic persons or cultures value group belonging and loyalty and derive beliefs from group decisions and the social system [42]. Studies that measured public acceptance of digital contact tracing applications during Covid-19 have found the acceptance rate to be nearly twice as high in collectivist countries such as China than in individualistic countries such as Germany [43]. We might also expect cultural and social norms to affect people's evaluations of fairness and preferences for individualized approaches. In particular, people who value individualism might be more likely than users who value collectivism to prefer high-individuality approaches to satisfy their 'personalization' demand.

B. Practical Implications

Trying to solve the issue on how to incorporate fairness in ML algorithms, researchers have already gathered and developed numerous definitions of fairness, e.g., [6], [44], [45] and translated them into several distinct fairness models, e.g., [5], [46], [47], [48], [49]. The ultimate goal is to translate intangible notions such as fairness to statistically measurable features and probability equations. To achieve this goal, we need theoretical and empirical work that investigates what people consider fair in different contexts. More broadly, we need methods to concretely define, optimise, and evaluate fairness algorithms. In an effort to ease the model selection, researchers have categorised the identified definitions along their degree of personalisation, into individual, subgroup and group fairness models, e.g., [6], [7], [45]. Individual fairness models compare features of individuals under investigation to ensure that individuals with similar feature scores obtain similar predictions, whereas group fairness models cluster individuals into groups and ensure certain statistical paradigms between the groups.

Subgroup fairness models form a combination of the former two categories, categorizing individuals based on their personal features into subgroups and ensuring group fairness constraints for those subgroups [6]. Taking this classification as a basis, the degree of individuality is crucial for examining the various notions of fairness and needs to be weighed to determine which fairness model should be chosen for a specific AI-enabled technology. However, deciding which fairness model category to draw from almost always require an in-depth understanding of the specific context. In the context of public-health surveillance, our data suggest that people indeed valued some degree of individuality in the decision even at the expense of data privacy. This means that, in the field of public health surveillance and the context of our study design, our findings suggest that some individuality is desired over complete homogeneity, pointing towards models like “Fairness through Awareness” [50] that allow for a greater consideration of individual personal characteristics.

Of course, this leads to the question of how such a balance between data privacy and fairness perceptions can be ensured. We suggest that it can be targeted with a clearer classification of attributes into those that users consider as purely privacy-intrusive or those that are perceived to contribute to enhancing fairness. For this study, we chose to examine a scenario in which it was not clear which approach people would view as fairer. Preferences regarding individualised decisions would probably look different if we selected inherent traits as personalisation factors, such as gender or social status. The fact that the chosen attributes are derived from people’s actions or decisions (i.e., characteristics that can be more consciously and more easily influenced) might make these more acceptable factors for individualization. Therefore, when separating parameters into those that are discriminatory and those that are acceptable for algorithmic decision-making, it is important to consider the type of individualisation and the attribute’s specific nature.

C. Limitations and Future Research

Although we found empirical evidence for a preference towards more individuality in public health surveillance tools, there was considerable overlap in the distributions of people’s responses across the two approaches. One explanation for this small effect is that the distinction between the individuality approaches was not stressed precisely enough in our experiment. Another possibility is that the selected feature, a person’s risk of spreading the virus, was not perceived as sufficiently individual to substantially impact fairness perceptions. Future work should examine how other factors affect fairness perceptions in public health contexts.

Although the user’s perceived data privacy did not predict perceived fairness, we found evidence that it might still affect the user’s satisfaction with the application. While we interpreted this as indication that the way data are used in decision-making can be important for perceptions of data privacy, it is possible participants did not fully understand that the data collected was the same across the two approaches. Future research should take this into account when developing

similar experiments, as studying the relative and cumulative effects of data collection and data use could help inform policy decisions.

Furthermore, widening the focus of this study in future research to include people from a broader range of cultural backgrounds and to examine other public health measures could complement the picture to a more holistic overview.

V. CONCLUSION

In this study, we investigated the relation between the individuality of a ML-based public health surveillance method and the perceived fairness as well as overall impression of that tool on the example of contact tracing applications.

Our findings suggest that users (in the U.K. and Germany) value higher degrees of individuality in health surveillance related decisions and perceive ‘high-individuality’ contact tracing app versions as fairer and more satisfactory overall. This pattern held despite the fact that people viewed higher levels of individuality as more privacy intrusive. Moreover, our findings suggest that perceptions of fairness are important for people’s evaluations of public health surveillance tools and could affect people’s adoption and acceptance of those applications.

Our results support the general trend towards more personalisation in healthcare also in health surveillance technologies and inform the design of future ML-enabled public health surveillance tools. While more individuality seemed more appealing for participants in our study, the nature of attributes that are used within a decision seems to be crucial for fairness perceptions, pointing towards a greater need for research to distinguish the parameters considered as fair or discriminating.

REFERENCES

- [1] F. Piccialli, V. S. di Cola, F. Giampaolo, and S. Cuomo, “The role of artificial intelligence in fighting the COVID-19 pandemic,” *Inf. Syst. Front.*, vol. 23, pp. 1467–1497, Apr. 2021.
- [2] V. Martinez-Tur, J. M. Peiro, J. Ramos, and C. Moliner, “Justice perceptions as predictors of customer satisfaction: The impact of distributive, procedural, and interactional justice,” *J. Appl. Soc. Psychol.*, vol. 36, no. 1, pp. 100–119, 2006.
- [3] J. G. Blodgett, D. J. Hill, and S. S. Tax, “The effects of distributive, procedural, and interactional justice on postcomplaint behavior,” *J. Retailing*, vol. 73, no. 2, pp. 185–210, 1997.
- [4] S. Sudin, “Fairness of and satisfaction with performance appraisal process,” *J. Global Manag.*, vol. 2, no. 1, pp. 66–83, 2011.
- [5] A. Chouldechova and A. Roth, “The frontiers of fairness in machine learning,” 2018, *arXiv:1810.08810*.
- [6] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, “A survey on bias and fairness in machine learning,” 2019, *arXiv:1908.09635*.
- [7] S. Feuerriegel, M. Dolata, and G. Schwabe, “Fair AI: Challenges and opportunities,” *Bus. Inf. Syst. Eng.*, vol. 62, no. 4, pp. 379–384, 2020.
- [8] N. Grgic-Hlaca, E. M. Redmiles, K. P. Gumjadi, and A. Weller, “Human perceptions of fairness in algorithmic decision making: A case study of criminal risk prediction,” in *Proc. World Wide Web Conf.*, 2018, pp. 903–912.
- [9] D. Shin, “User perceptions of algorithmic decisions in the personalized AI system: Perceptual evaluation of fairness, accountability, transparency, and explainability,” *J. Broadcast. Electron. Media*, vol. 64, no. 4, pp. 541–565, 2020.
- [10] C. Top and B. J. Ali, “Customer satisfaction in online meeting platforms: Impact of efficiency, fulfillment, system availability, and privacy,” *Amazonia Investiga*, vol. 10, no. 38, pp. 70–81, 2021.
- [11] N. D. Nayeri and M. Aghajani, “Patients’ privacy and satisfaction in the emergency department: A descriptive analytical study,” *Nursing Ethics*, vol. 17, no. 2, pp. 167–177, 2010.

- [12] H. Chang and R. Shokri, "On the privacy risks of algorithmic fairness," in *Proc. IEEE Eur. Symp. Security Privacy (EuroSP)*, 2021, pp. 292–303.
- [13] S. Kokolakis, "Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon," *Comput. Security*, vol. 64, pp. 122–134, Jan. 2017.
- [14] N. Gerber, P. Gerber, and M. Volkamer, "Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior," *Comput. Security*, vol. 77, pp. 226–261, Aug. 2018.
- [15] S. Barth and M. D. T. De Jong, "The privacy paradox—Investigating discrepancies between expressed privacy concerns and actual online behavior—A systematic literature review," *Telematics Inform.*, vol. 34, no. 7, pp. 1038–1058, 2017.
- [16] E. Aguirre, A. L. Roggeveen, D. Grewal, and M. Wetzels, "The personalization-privacy paradox: Implications for new media," *J. Consum. Market.*, vol. 33, no. 2, pp. 98–110, 2016.
- [17] J. S. Adams, "Inequity in social exchange," in *Advances in Experimental Social Psychology*, vol. 2. Amsterdam, The Netherlands: Elsevier, 1965, pp. 267–299.
- [18] K. S. Cook and K. A. Hegtvedt, "Distributive justice, equity, and equality," *Annu. Rev. Sociol.*, vol. 9, no. 1, pp. 217–241, 1983.
- [19] B. Kabanoff, "Equity, equality, power, and conflict," *Acad. Manag. Rev.*, vol. 16, no. 2, pp. 416–441, 1991.
- [20] M. Deutsch, "Equity, equality, and need: What determines which value will be used as the basis of distributive justice?" *J. Soc. Issues*, vol. 31, no. 3, pp. 137–149, 1975.
- [21] M. Srivastava, H. Heidari, and A. Krause, "Mathematical notions vs. human perception of fairness: A descriptive approach to fairness for machine learning," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, 2019, pp. 2459–2468.
- [22] B. Hüsing, "Individualisierte Medizin—Potenziale und Handlungsbedarf," *Zeitschrift für Evidenz, Fortbildung und Qualität im Gesundheitswesen*, vol. 104, no. 10, pp. 727–731, 2010.
- [23] A. Agusti, "The path to personalised medicine in COPD," *Thorax*, vol. 69, no. 9, pp. 857–864, 2014.
- [24] J. P. Richardson *et al.*, "Patient apprehensions about the use of artificial intelligence in healthcare," *NPJ Digit. Med.*, vol. 4, no. 1, pp. 1–6, 2021.
- [25] S. Trang, M. Trenz, W. H. Weiger, M. Tarafdar, and C. M. K. Cheung, "One app to trace them all? Examining app specifications for mass acceptance of contact-tracing apps," *Eur. J. Inf. Syst.*, vol. 29, no. 4, pp. 415–428, 2020.
- [26] M. Walrave, C. Waeterloos, and K. Ponnet, "Ready or not for contact tracing? Investigating the adoption intention of COVID-19 contact-tracing technology using an extended unified theory of acceptance and use of technology model," *Cyberpsychol. Behav. Soc. Netw.*, vol. 24, no. 6, pp. 377–383, 2021.
- [27] E. Deci and R. Ryan, *Intrinsic Motivation and Self Determination in Human Behavior*. New York, NY, USA: Plenum, 1985.
- [28] A. H. Olafsen, H. Halvari, J. Forest, and E. L. Deci, "Show them the money? The role of pay, managerial need support, and justice in a self-determination theory model of intrinsic work motivation," *Scandinavian J. Psychol.*, vol. 56, no. 4, pp. 447–457, 2015.
- [29] J. M. Haar and C. S. Spell, "How does distributive justice affect work attitudes? The moderating effects of autonomy," *Int. J. Human Resour. Manag.*, vol. 20, no. 8, pp. 1827–1842, 2009.
- [30] S. Aryee, F. O. Walumbwa, R. Mondejar, and C. W. L. Chu, "Accounting for the influence of overall justice on job performance: Integrating self-determination and social exchange theories," *J. Manag. Stud.*, vol. 52, no. 2, pp. 231–252, 2015.
- [31] R. Wang, F. M. Harper, and H. Zhu, "Factors influencing perceived fairness in algorithmic decision-making: Algorithm outcomes, development procedures, and individual differences," in *Proc. CHI Conf. Human Factors Comput. Syst.*, 2020, pp. 1–14.
- [32] J. C. Roca, J. J. García, and J. J. De La Vega, "The importance of perceived trust, security and privacy in online trading systems," *Inf. Manag. Comput. Security*, vol. 17, no. 2, pp. 96–113, 2009.
- [33] D. Peters, R. A. Calvo, and R. M. Ryan, "Designing for motivation, engagement and wellbeing in digital experience," *Front. Psychol.*, vol. 9, p. 797, May 2018.
- [34] L. Floridi *et al.*, "AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations," *Minds Mach.*, vol. 28, no. 4, pp. 689–707, 2018.
- [35] U. Gasser, M. Ienca, J. Scheibner, J. Sleight, and E. Vayena, "Digital tools against COVID-19: Taxonomy, ethical challenges, and navigation aid," *Lancet Digit. Health*, vol. 2, no. 8, pp. E425–E434, 2020.
- [36] M. M. Mello and C. J. Wang, "Ethics and governance for digital disease surveillance," *Science*, vol. 368, no. 6494, pp. 951–954, 2020.
- [37] C. Klingler, D. S. Silva, C. Schuermann, A. A. Reis, A. Saxena, and D. Streh, "Ethical issues in public health surveillance: A systematic qualitative review," *BMC Public Health*, vol. 17, no. 1, p. 295, 2017.
- [38] J. Budd *et al.*, "Digital technologies in the public-health response to COVID-19," *Nat. Med.*, vol. 26, pp. 1183–1192, Aug. 2020.
- [39] H. Nissenbaum, "A contextual approach to privacy online," *Daedalus*, vol. 140, no. 4, pp. 32–48, 2011.
- [40] P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. Wagner, and K. Beznosov, "Android permissions remystified: A field study on contextual integrity," in *Proc. 24th USENIX Security Symp. (USENIX Security)*, 2015, pp. 499–514.
- [41] Y. Wang, H. Xia, and Y. Huang, "Examining American and Chinese Internet users' contextual privacy preferences of behavioral advertising," in *Proc. 19th ACM Conf. Comput. Supported Cooperative Work Soc. Comput.*, 2016, pp. 539–552.
- [42] G. Hofstede, "Motivation, leadership, and organization: Do American theories apply abroad?" *Org. Dyn.*, vol. 9, no. 1, pp. 42–63, 1980.
- [43] M. V. Zetterholm, Y. Lin, and P. Jokela, "Digital contact tracing applications during COVID-19: A scoping review about public acceptance," *Informatics*, vol. 8, no. 3, p. 48, 2021.
- [44] R. Binns, "Fairness in machine learning: Lessons from political philosophy," in *Proc. Conf. Fairness Accountability Transparency*, 2018, pp. 149–159.
- [45] B. Hutchinson and M. Mitchell, "50 years of test (Un)fairness: Lessons for machine learning," in *Proc. Conf. Fairness Accountability Transparency*, 2019, pp. 49–58.
- [46] S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq, "Algorithmic decision making and the cost of fairness," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, 2017, pp. 797–806.
- [47] S. Mitchell, E. Potash, S. Barocas, A. D'Amour, and K. Lum, "Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions," 2018, *arXiv:1811.07867*.
- [48] N. A. Saxena, K. Huang, E. DeFilippis, G. Radanovic, D. C. Parkes, and Y. Liu, "How do fairness definitions fare?: Examining public attitudes towards algorithmic definitions of fairness," in *Proc. AAAI/ACM Conf. AI Ethics Soc.*, 2019, pp. 99–106.
- [49] S. Verma and J. Rubin, "Fairness definitions explained," in *Proc. IEEE/ACM Int. Workshop Softw. Fairness (FairWare)*, 2018, pp. 1–7.
- [50] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *Proc. 3rd Innov. Theor. Comput. Sci. Conf.*, 2012, pp. 214–226.