



Safeguarding inclusion when using gestures in microtask crowdsourcing

Sanne Veringa
Supervisor(s): Garrett Allen, Ujwal Gadiraju
EEMCS, Delft University of Technology, The Netherlands

24-6-2022

**A Dissertation Submitted to EEMCS faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering**

Abstract

Microtask crowdsource workers are negatively influenced, mentally as well as physically, by the repetitive nature of the tasks they perform. Research is ongoing on whether using a gesture-based input technique could mitigate these negative effects. This paper identifies possible ways that using gestures as an alternative input modality could lead to exclusion by analysing survey responses, where $n=10$. While further research is necessary, there are indications that this could lead to cultural and physical exclusion of certain groups. This paper is not meant to discourage using gestures as an alternative method of input but is solely meant to bring attention to possible risks to take into account.

1 Introduction

Large tasks require lots of work. On crowdsourcing platforms, these large tasks can be broken down into many microtasks and distributed over a large set of people. These people, often referred to as the workers, get paid a small amount of money to finish one of these microtasks. Many different types of tasks can be considered, from image classification to sentiment analysis to surveys.

With a diverse group of workers on crowdsourcing platforms [1] it is vital to research the well-being of these workers and what work components impact their well-being. Repetitive tasks negatively influence both these forms of well-being in workers on crowdsourcing platforms. In an effort to reduce these negative effects, alternative input modalities could be considered. One of the alternatives offered is using gestures.

Instead of clicking a mouse or typing on a keyboard, one could imagine a broad range of gesture-based inputs. Most gestures are performed using hands or objects, but one could also think of using their head, fingers or even their whole body. [2]

As different solutions may impact people in different ways, research must be done to investigate the inclusivity of this form of alternative input modalities. To the best of the authors' knowledge, no research has been conducted on this specific problem and thus there is a knowledge gap.

It is of great importance that these platforms stay inclusive for multiple reasons. First, ensuring inclusivity in the workplace leads to greater job satisfaction and increased innovation [3]. Second, with crowdsourcing platforms often being used for research, and some platforms even specializing in providing a quick way to find research participants, such as Prolific Academic [4], it is important that these platforms can continue to offer a diverse participant pool.

The demographics of crowdsourcing workers are very diverse [5]. As has been revealed in the research of Gadiraju et al.[1], different workers perform the work in different environments and with different access to tools. From previous work in the crowdsourcing domain, we can see that it is a difficult task for platforms to cater to all people.

Turner et al. have investigated whether or not crowdsourcing platforms can be effectively used as a way to reach older

adult workers for research. They found this to be effective, however, they noted that respondents were not entirely representative of the larger older adult population. They noted, "At this point in time crowdsourcing does not appear to be an effective method for reaching older adults from underrepresented minority populations. Barriers to online access, historic mistrust of research and lack of outreach to minority groups may contribute to underrepresentation." [6, p.8]

In some cases, specialized crowdsourcing platforms have been created as existing platforms did not serve a wide enough audience. In the case of Natural Language Processing for Arab languages, the most used crowdsourcing platforms fell short as they were developed in a specific linguistic context. [7]

There are many ways crowdsourcing platforms could potentially miss out on being inclusive. In this research, we pursue to form a more general understanding of how using gestures as an alternative input modality for microtask crowdsourcing could lead to the exclusion of different groups. This research aims to answer this question by splitting it into four smaller parts.

- Could using gestures lead to physical exclusion of certain groups?
- Could using gestures lead to conceptual exclusion of certain groups?
- Could using gestures lead to economic exclusion of certain groups?
- Could using gestures lead to cultural exclusion of certain groups?

These questions will likely not cover all possible cases of exclusion, but must be seen as a starting point for more in-depth research. Answers to these questions bring a new angle to the discussion of the possibility of using gestures as an alternative input modality.

In section 2, related work will be analysed. In section 3, the methodology of the research will be laid out; in this section it is discussed how inclusion will be measured. Section 4 is about responsible research. In section 5, the results of the research are presented together with an analysis. This is followed by section 6, which is a discussion of the results. Finally, section 7 concludes with limitations of the research and ideas for future work.

2 Related work

Inclusion and diversity are terms that are often named together and know many different definitions [8]. Though often used interchangeably, the distinction between the two terms lies in that diversity defines a state, whereas inclusion is about an action [9].

Research into inclusion in businesses is often focused on employee sentiment with respect to other members in their organization [10]. As crowdworkers do not necessarily interact directly with other people that work on or for the platform, this hardly feels relevant in this context. A little closer to the problem at hand is the notion of inclusive technology. While the term was originally coined in the context of bridging the digital divide, we feel that certain elements of it can be placed

in a broader context. Barriers to inclusion mentioned in a paper by researchers at Concordia University include access to physical resources, digital information resources, human resources and social resources. In addition to that, they mention "... for a technology to be accepted by its intended users, it must be perceived as beneficial, easy to use, and socially endorsed, with an adequate infrastructure in place to support its use." [11, p.5].

It has also been stated that in Inclusive Design, while designers may provide accessibility options, the usability of these options is often not considered [12].

When combining the parts of these definitions relevant to our goal, we arrive at the following three conditions for inclusion: being accessible to all, being usable by all and being satisfactory for all. To reach the next layer, the layer above should be satisfied, meaning that we can not have something that is usable by all if it is not accessible to all. Likewise, we can not have something that is satisfactory for all if it is not usable or accessible by all. This definition gives us a way to reason about inclusion on different levels.

3 Methodology

3.1 Experimental setup

We measure inclusion through a survey sent out through Prolific Academic[4]. The answer options are a mix of open-ended response and response on a five-item Likert scale.

No prefiltering was done in selecting the participants. It was debated whether or not we should prefilter on people who would be likely to have experienced some form of exclusion at prior moments in life, as these people would know what types of issues to look out for. However, it was decided to do no prefiltering as the Prolific website only allowed for an 'and' type of filtering between filters.

Questions included in the survey can be split into three types. The first is questions about demographics. These questions are included to find trends in data among and within different groups. They can also be used to check the diversity of the respondent pool, which is important to ensure people across different circumstances are heard.

The following section of the survey consists of open questions aimed at finding different reasons for gestures not being inclusive. The questions consider inclusion across two dimensions. The first dimension focuses on the condition for inclusion that is limited, where the options are accessibility, usability and satisfaction. The second dimension focuses on the specific manner in which the conditions can be limited. Here it distinguishes between physical, economic, cultural and conceptual factors. Together, these two dimensions lead to a total of twelve open-ended questions.

The aim of dividing the questions across these two dimensions is both to help interpret the results and guide the respondents. For interpreting the results, it helps to more precisely spot where the main roadblocks on inclusion are. On the other side, the subdivision can guide respondents to recognise potential exclusion, which they may not have thought of with a more generic question. To indicate what type of responses we expect from participants, we included some examples for each factor. The examples used in the survey are:

- For **physical factors** you could think of not comfortably being able to move your arms up high or not being able to make specific gestures with your hands and fingers.
- For **economic factors** you could think of not having access to a webcam or mobile phone camera which would make it not accessible to you. Or you could think of not wanting to spend more time on tasks than you would when using mouse and keyboard.
- For **cultural factors** you could for example think of not wanting to use big arm gestures when working on these platforms in public which would make it less usable.
- For **conceptual factors** you could think of not understanding how to use gesture-based input and therefore making it less accessible or having difficulty with having to constantly think about it which could make it less usable.

Last, participants were asked about their current feelings on accessibility, usability and satisfaction on the platform and their predicted feelings on these topics if the platform were to switch to a gesture-based method of input. The responses to these questions were recorded on a five-item Likert scale. These questions were purposely put after the open-ended questions to ensure people thought the topic through before giving a judgement.

3.2 Biases

When dealing with self-reported data, one needs to be aware of the possibility of cognitive biases being present in the data. Cognitive biases are deviations in thinking, away from rationality. These types of biases are mostly present in environments where humans have limited time for decision making. As the nature of these crowdsourced microtasks is fast-paced, one should be mindful of these types of biases. Another factor that increases the possibility of cognitive biases being present in the data is when the information the worker is presented with is not (clear) enough to be able to accurately answer the question. Therefore, an effort was made to keep the questions posed as simple and clear as possible.

To be aware of possible cognitive biases in the research, a checklist proposed by Draws et al. was consulted[13]. Based on that and further research, four specific types of biases were identified and have been taken into account when developing the survey.

The first bias that was selected is the self-interest bias. As workers may think their answers to the questions can influence what the platform they use may look like in the future, they might change their answers to be more extreme than their actual feelings. Another bias that was considered is recency bias. People tend to put too much emphasis on events that are recent, and with that pay too little attention to things that happened longer ago. To measure this and to mitigate the effects, one of the questions asked to participants is for how long the person has been actively using the platform. A third bias to be wary of is the sunk cost fallacy. The more time a worker has invested in a task, the more they want to complete it, even if that means giving lower-quality responses. An attention check question was included to ensure the use of only

high-quality data. The last bias to be considered is loss aversion. Workers might think they may not get paid (fairly) after expressing their true feelings, as they might think it could be perceived as critique. Therefore it was clearly communicated that their answers would in no way affect (future) payouts.

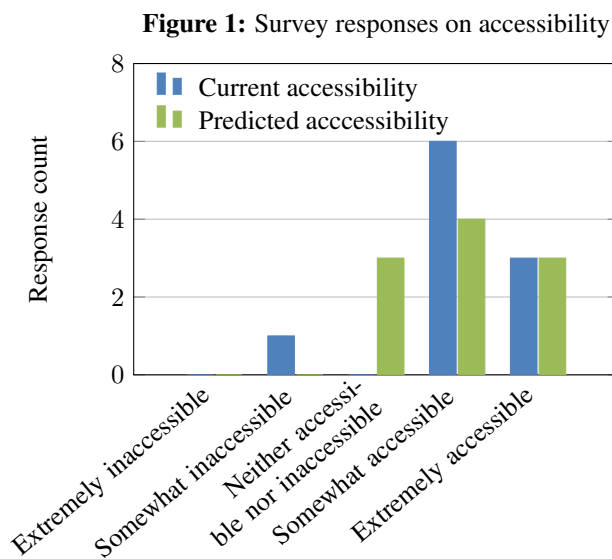
4 Responsible Research

In an effort to make this research as reproducible as possible, all parts of the experimental setup are reported. The full survey and the informed consent form that goes with it can be found in appendix A. For this research, approval was received from the Delft University of Technology Human Research Ethics committee, which ensures we only proceed with research of minimal risk. All personally identifiable data was anonymized soon after the data was gathered. The anonymized data was stored on a password-protected electronic device. All responses recorded are included in the analysis.

5 Results

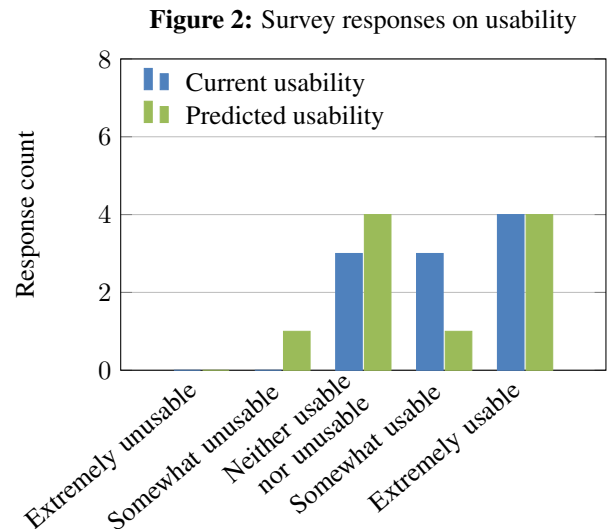
We were able to find 10 participants through Prolific. While that number is too low to perform a statistically powerful analysis, it serves as a good starting point for further research. Respondents' ages range from 18 to 44. Half of them identify as female and the other half as male. 9 out of 10 respondents spend less than 10 hours a week on the platform. Respondents reside in 7 different countries in Africa and Europe.

Below the responses to the Likert scale questions are reported, split out by condition for inclusion and complemented by responses to the open-ended questions. An important thing to be noted is that from all open-ended questions (10 participants times 13 open-ended questions = 130 possibilities for open-ended response), we only got 16 answers, 5 of which gave insight into possible reasons for exclusion. In reporting the results below, we interpret the responses on the Likert scale as values from 1 to 5 and assume that the "distance" between each option is equal.



Accessibility

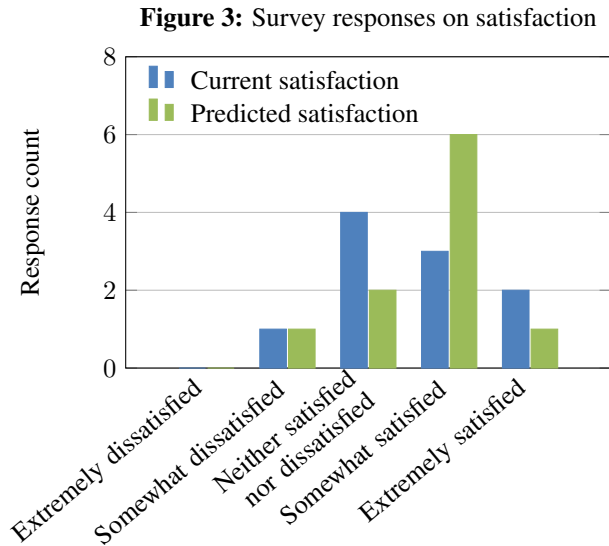
In figure 1, responses to the questions "How accessible do you feel this platform is for you at this moment?" and "How accessible do you feel this platform would be for you if it made use of a gesture-based input technique?" are laid out. Mean predicted accessibility decreases when using gestures as input modality compared to the current situation. Current mean accessibility is 4.10, which places it between somewhat accessible and extremely accessible. Predicted mean accessibility is 4.00, which corresponds to somewhat accessible. Because this research is focused on inclusion it makes sense to report on cases where people may not have felt included. In total, 3 out of 10 respondents foresee a decrease in accessibility. Among these three respondents, the average accessibility went down from 4.3, putting it somewhere between somewhat accessible and extremely accessible, to 3, putting it in the category neither accessible nor inaccessible. One reason given for a potential decrease in accessibility is privacy (GDPR) issues that result from filming people. If certain countries have laws in place that prevent recording of the workers, their citizens may have trouble participating in the microtasks. This indicates a form of cultural exclusion.



Usability

In figure 2, responses to the questions "How usable do you feel this platform is for you at this moment?" and "How usable do you feel this platform would be for you if it made use of a gesture-based input technique?" are laid out. Mean predicted usability decreases when using gestures as input modality compared to the current situation. Current mean usability is 4.10, which places it between somewhat usable and extremely usable. Mean predicted usability is 3.80, which places it between neither usable nor unusable and somewhat usable. In total, 3 out of 10 respondents foresee a decrease in usability. Among these three respondents, the average usability went down from 4.7, putting it between somewhat usable and extremely usable, to 2.7, between somewhat unusable and neither usable nor unusable. One respondent elaborates

that people may not feel comfortable being filmed. Another respondent notes that they may have trouble with gestures after playing sports, as their muscles would be sore. These comments suggest a form of physical and cultural exclusion.



Satisfaction

In figure 3, responses to the questions "How satisfied are you with this platform at this moment?" and "How satisfied do you think you would be using this platform if it made use of a gesture-based input technique?" are laid out. Mean predicted satisfaction increases from 3.6 to 3.7 when using gestures as input modality compared to the current situation, both means falling between neither satisfied nor dissatisfied. In total, 3 out of 10 respondents foresee a decrease in satisfaction. Among these three respondents, the average satisfaction went down from 4.3, falling between somewhat satisfied and extremely satisfied, to 3.0, corresponding to neither satisfied nor dissatisfied. Two respondents see a downside due to the impression that using gestures could have on other people. Specifically, they suggest that it could be seen as weird by nearby people when initially introduced, or that they do not like making big gestures due to cultural reasons. This indicates a form of cultural exclusion.

To summarise, the survey respondents point to various reasons why they might experience using gestures as less inclusive. The majority of these reasons can be categorised as cultural, related to privacy and the impression it has on surrounding people. It also shows that it could lead to physical exclusion to people who are already physically active in other parts of their lives, such as sports.

6 Discussion

This research aims to identify ways that using gestures as an alternative input modality for microtask crowdsourcing could lead to the exclusion of different groups, either physically, conceptually, economically or culturally.

We find that for all layers of inclusion there are people who foresee a decline. An interesting thing to note is that three

out of ten people reported a decrease in either accessibility, usability, satisfaction or a combination of those, without providing a reason to do so in the open-ended questions. This could indicate that either the categories physical, economical, cultural or conceptual were too limiting, or that the conditions for inclusion, accessibility, usability and satisfaction were too limiting.

From the data, no indication can be seen that using gestures as alternative input could lead to economic or conceptual exclusion. A potential reason for a lack of data pointing towards economic exclusion could be that the respondents generally came from high GDP per capita countries. Future research could be conducted on how the findings translate to lower-income people. A potential reason for lack of data pointing towards conceptual exclusion is the relatively low age of people in our dataset.

Two respondents mentioned cultural reasons for exclusion when using gestures. One of them mentioned the possibility of it being perceived as weird by people nearby. Research on workers on the platform Amazon Mechanical Turk in India found that while the work is online, workers often meet in physical places, such as internet cafes. This research also mentions that some families pressured workers to find more traditional forms of employment [14]. A hypothesis is that because using gestures as an alternative input modality is farther away from traditional work, these pressures will increase. Further research should be conducted on whether or not these are reasons for workers to feel excluded when using gestures as an alternative input modality. Another cultural reason for exclusion a respondent saw was privacy issues connected to using filming methods for gesture recognition. One respondent mentioned physical reasons for exclusion, namely that they may have trouble with gestures after playing sports, as their muscles would be sore. Future research could be done on how to mitigate these potential forms of exclusion.

A remarkable response is one where a respondent clearly sees an improvement in using gestures as an alternative input modality. This respondent increased their feelings of accessibility, usability and satisfaction all with three points on the Likert scale. In this data point, we have to be wary of the self-interest bias.

With only a few data points, we cannot make any strong statistical statements based on the data. We can however theorize about why responses are the way they are. The first thing that stands out when looking at the data is the lack of elaborate responses. This could indicate that the amount of money offered per respondent to complete the survey was insufficient and would lead to a low hourly rate, which could lead respondents to hastily go through the survey. This can however be debated by looking at the recorded response time, where the average recorded time was a little less than eight minutes, which is below the anticipated time to spend on the survey. Furthermore, we saw more elaborate responses in participants that had far below average response times. One explanation could be that the more elaborate answers came from people with a better grasp of the concept of inclusiveness, who were thus able to answer the questions at a higher speed. Another explanation for the lack of response to open-ended questions could be an unwillingness to talk about private matters.

7 Conclusion

In this research, we investigate how using gesture-based input for microtask crowdsourcing could lead to exclusion of workers. We find various sources of exclusion, generally centered around physical and cultural exclusion.

The research is based on survey responses on the platform Prolific Academic. Due to the low number of respondents, the statistical significance of the results in this paper is limited. Furthermore, the data collected on expectations regarding different forms of exclusion when dealing with gesture-based inputs rely heavily on the predictive power of the participants in the research. Because surveys are prone to certain biases, this brings another limitation to the research. Although an active effort was made to mitigate most cognitive biases, some might still be present in the data. Furthermore, statistical bias, contextual bias and prejudices can affect the data.

In future work, we could consider posing a similar survey with less limiting answer options by not splitting the open-ended questions into the four categories. Another option to consider is to switch to or supplement with focus groups. Focus groups lend themselves to this type of research by eliciting responses and getting people to bounce off each other's ideas; sometimes, something someone else has said can be the trigger for a new idea. A difficulty with this can be that we would ideally like a diverse group. Since crowdworkers are located all over the world, this might be difficult to arrange. One option could be to carry out the focus group online, but this could take away from the safe, calm environment that encourages people to speak up in focus groups.

In the long term, randomised control trials where workers actually use gestures as input could be used as the next step in analysing the impact on inclusion. This would also reduce the reliance on the predictive power of survey respondents.

A Survey

Below you find the informed consent form and the survey questions as used in the research.

A.1 Informed consent

In this research we want to investigate the possibility of using gestures as an alternative way of input for microtask crowdsourcing. For example, consider pointing a finger in a certain direction to indicate picking an option instead of using a mouse to click on the option. For the extent of this questionnaire we ask you to consider for the gestures things like waving your hand, raising your arm or using fingers to create for example an "OK" sign. In the questions below you will be asked about your expectations on accessibility, usability and satisfaction regarding gesture based input. Participation in this survey is voluntary and you can withdraw at any moment. It is important that you know that your answers to these questions will not reach Prolific and will in no way have any influence on your (future) payouts. All collected data will be anonymized and stored password protected. This survey consists of 26 questions, of which 14 are open ended. By pressing 'I agree' below, you indicate that have read and understood this consent form and that you agree to participate in this research

A.2 Survey questions

Demographic questions

1. May we know your age?
 - 18-24
 - 25-34
 - 35-44
 - 45-54
 - 55-64
 - 65+
 - Prefer not to say
2. May we know your gender?
 - Male
 - Female
 - Non-binary/third gender
 - Other (Respondents had an option to fill in their preference)
 - Prefer not to say
3. In which country do you currently reside?
 - Respondents were able to chose from a dropdown menu of all countries
4. For how long have you been actively using Prolific or other crowdsourcing platforms?
 - Less than half a year
 - Half a year to a year
 - More than a year
 - More than two years
5. How many hours a week on average do you spend on Prolific or other crowdsourcing platforms?
 - Less than 2 hours a week
 - 2 to 10 hours a week
 - 10 to 20 hours a week
 - 20 to 30 hours a week
 - More than 30 hours a week

Open ended questions

For the following questions respondents were able to fill in an open-ended response.

These next 12 questions will be about accessibility, usability and satisfaction when using gestures as an alternative way of input. Feel free to answer with a short 'No' if this is not applicable for you.

For the physical factors in next three questions you could think of things like not comfortably being able to move your arms up high or not being able to make specific gestures with your hands and fingers.

6. Are there any physical factors that would make using gestures as an alternative way of input not accessible for you?
7. Are there any physical factors that would make using gestures as an alternative way of input this not usable for you?

- Are there any physical factors that would make you dissatisfied with using gestures as an alternative way of input?

For economic factors in the next three questions you could think of not having access to a webcam or mobile phone camera which would make it not accessible to you. Or you could think of not wanting to spend more time on tasks then you would when using mouse and keyboard.

- Are there any economic factors that would make using gestures as an alternative way of input not accessible for you?
- Are there any economic factors that would make using gestures as an alternative way of input this not usable for you?
- Are there any economic factors that would make you dissatisfied with using gestures as an alternative way of input?

For cultural factors in the next three questions you could for example think of not wanting to use big arm gestures when working on these platforms in public which would make it less usable.

- Are there any cultural factors that would make using gestures as an alternative way of input not accessible for you?
- Are there any cultural factors that would make using gestures as an alternative way of input this not usable for you?
- Are there any cultural factors that would make you dissatisfied with using gestures as an alternative way of input?

For conceptual factors in the next three questions you could think of not understanding how to use gesture based input and therefore making it less accessible or having difficulty with having to constantly think about it which could make it less usable.

- Are there any conceptual factors that would make using gestures as an alternative way of input not accessible for you?
- Are there any conceptual factors that would make using gestures as an alternative way of input this not usable for you?
- Are there any conceptual factors that would make you dissatisfied with using gestures as an alternative way of input?

Likert scale questions

For the following questions respondents were able to pick a response on a five-item Likert scale

Please answer the following questions about accessibility. *Extremely inaccessible - Somewhat inaccessible - Neither accessible nor inaccessible - Somewhat accessible - Extremely accessible*

- How accessible do you feel this platform is for you at this moment?

- How accessible do you feel this platform would be for you if it made use of a gesture based input technique?

Please answer the following questions about usability. *Extremely unusable - Somewhat unusable - Neither usable nor unusable - Somewhat usable - Extremely usable*

- How usable do you feel this platform is for you at this moment?
- This is an attention check question, please respond with 'Somewhat usable'.
- How usable do you feel this platform would be for you if it made use of a gesture based input technique?

Please answer the following questions about satisfaction. *Extremely unsatisfied - Somewhat unsatisfied - Neither satisfied nor unsatisfied - Somewhat satisfied - Extremely satisfied*

- How satisfied are you with this platform at this moment?
- How satisfied do you think you would be using this platform if it made use of a gesture based input technique?

Ending question

For the final question respondents were able to fill in an open-ended response.

- Is there anything else you would like to mention?

References

- U. Gadiraju, A. Checco, N. Gupta, and G. Demartini, "Modus operandi of crowd workers: The invisible role of microtask work environments," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, pp. 1–29, 09 2017.
- M. Karam and m. Schraefel, "A taxonomy of gestures in human computer interactions," *Electronics and Computer Science*, 01 2005.
- K. Brimhall and M. Mor Barak, "The critical role of workplace inclusion in fostering innovation, job satisfaction, and quality of care in a diverse human service organization," *Human Service Organizations: Management, Leadership Governance*, vol. 42, pp. 474–492, 10 2018.
- "Prolific academic." Available from: <https://prolific.ac/>. Accessed: 2022.
- G. Kazai, J. Kamps, and N. Milic-Frayling, "The face of quality in crowdsourcing relevance labels: demographics, personality and labeling accuracy," pp. 2583–2586, 10 2012.
- A. Turner, T. Engelsma, J. Taylor, R. Sharma, and G. Demiris, "Recruiting older adult participants through crowdsourcing platforms: Mechanical turk versus prolific academic," *AMIA ... Annual Symposium proceedings. AMIA Symposium*, vol. 2020, pp. 1230–1238, 01 2021.
- D. Alahmadi, A. Babour, K. Saeedi, and A. Visvizi, "Ensuring inclusion and diversity in research and research output: A case for a language-sensitive nlp crowdsourcing platform," *Applied Sciences*, vol. 10, no. 18, 2020.

- [8] S. Thompson, "Defining and measuring 'inclusion' within an organisation," 06 2017.
- [9] M.-F. Winters, *From Diversity to Inclusion: An Inclusion Equation*, ch. 7, pp. 205–228. John Wiley Sons, Ltd, 2013.
- [10] L. Romansky, M. Garrod, K. Brown, and K. Deo, "How to measure inclusion in the workplace," *Harvard Business Review*, vol. 27, 2021.
- [11] K. Pitula, "A framework and process for designing inclusive technology," pp. 64–64, 09 2007.
- [12] A. Newell, P. Gregor, M. Morgan, G. Pullin, and C. Macaulay, "User-sensitive inclusive design," *Universal Access in the Information Society*, vol. 10, pp. 235–243, 08 2011.
- [13] T. Draws, A. Rieger, O. Inel, U. Gadiraju, and N. Tintarev, "A checklist to combat cognitive biases in crowdsourcing," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* (E. Kamar and K. Luther, eds.), vol. 9, pp. 48–59, Association for the Advancement of Artificial Intelligence (AAAI), 2021. Green Open Access added to TU Delft Institutional Repository 'You share, we take care!' - Taverne project <https://www.openaccess.nl/en/you-share-we-take-care> Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.; The Ninth AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2021 ; Conference date: 14-11-2021 Through 18-11-2021.
- [14] N. Gupta, D. Martin, B. V. Hanrahan, and J. O'Neill, "Turk-life in india," in *Proceedings of the 18th International Conference on Supporting Group Work, GROUP '14*, (New York, NY, USA), p. 1–11, Association for Computing Machinery, 2014.