Finding Waldo and whom he is talking to.
A rotational approach to find social interaction groups.

Ardy Zwanenburg Supervisor(s): Stephanie Tan, Jose Vargas Quiros, Hayley Hung EEMCS, Delft University of Technology, The Netherlands
22 June 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

In our daily life people encounter many social interactions, for example in the supermarket, at work and in schools. Currently the most reliable way to find social interactions in groups, is to manually annotate the data. Manual annotation takes a lot of time and human resources and as the information stream goes faster and faster, the manual labour cannot keep up. Therefore an automated program to replace this is desirable. One method to automate this, is by looking at the proximity of people, which has been done by Dikker [1]. The results look promising, however they are sensitive to errors. The F1-score (equation 3) is only 0.625 and the precision (equation 5) is 0.696. This means that there are still a lot of false positives. With orientation data one can determine if the people in question are facing each other before they are assigned to be in social interaction. With a view frustum (cone shape) it can be determined who are facing at each other by checking if the cones overlap. This cone can take different sizes, making the view area bigger or smaller. Because the absolute rotation were only used, it was used to find out who cannot face each other. Using the proximity results as a baseline, the impossible orientations were removed. With the goal to reduce the amount of false positives. From the results it can be concluded that a small cone will perform worse compared to the baseline. The bigger the angle of the cone becomes, the closer it gets to the results from Dikker [1]. However, in none of the results it turns out that using absolute orientations improves the current work significantly.

1 Introduction

In our daily life, people encounter many social interactions; for example in the supermarket, at work and in schools. Finding these social interactions in groups can help in many ways. One example is to find social interaction and warn the people within this group if something has happened with one of the group members. This became very relevant during the COVID-19 pandemic in 2020. Social interaction had a massive influence on how fast the virus spread [2]. To prevent further spread of the virus, different apps were developed (e.g. DP3T [3]). If a user tested positive for corona, all other users who had been in close contact would be notified by the app to warn them. Another example where finding social interactions is useful, are conferences. If group forming can be detected, it could help participants to find a good moment to nudge themselves into the group. This maximises the network potential within the group or even within the event.

Before finding social interactions, a definition of what a social interaction group is, is needed. One interpretation of such a group is called an F-formation. An F-formation is described as a group of people who gather together with the goal to conversing and exchanging information with each other [4]. For this research the definition of an F-formation is used to describe a social interaction group.

Currently, the most reliable way of finding F-formations is by manually annotating the data with the given definition of an F-formation. This works well for training data for models, but not for real time applications. The manual annotation of data is time inefficient, has high costs due to manual labor that has to be paid, and is error sensitive. An annotator can annotate something as F-formation if it is an F-formation in their opinion, however, another annotator can interpret it differently.

In the last years, research has been performed to do the annotations digitally using different types of data. Earlier this year Dikker [1] researched the performance of Received Signal Strength Indicator (RSSI) of Bluetooth Low Energy (BLE). The results look promising, but can still be improved. For the research Dikker used data from a smart sensor, the Midge [5]. A Midge is a smart sensor made by the Socially Perceptive Computing Lab (SPCL). It was used to record different types of data during a social interaction event. Among this data is audio, acceleration, gyroscope, magnitude, proximity and rotation data. This data together is called the ConfLab data set of the SPCL at the TU Delft. The dataset used for this research will be the same as the one used in Dikker [1] his research.

Currently, the RSSI approach simply looks at who is close to each other in a given time interval. Sometimes this approach can cause a false positive. People who are close to each other will be identified as being in a group together, while they might not have a social interaction. This can happen for example if there is a wall between them. Adding rotational data gives a possibility to remove false positives. By setting a lower affinity when people are not facing each other.

The orientation data from the Midges are absolute. Because the orientations are absolute, the location of the Midges compared to each other cannot be found with this alone. Therefore the method needs to take this into account.

Orientations then can indicate who is facing each other and who not. Which gives a new way of measuring social interaction. The probability of social interaction is higher when a group is facing each other, than when they are facing away from each other. With these insights orientation data became more interesting, especially as an addition to the work from Dikker [1].

The goal of this report is to research the performance of including orientation data for finding F-formations in the social interaction data from the Conflab data set and compare the results with the proximity-based data as the baseline.

2 Methodology

This research project is directly related to Dikker [1], since it uses the same methods. However, this project is an extension as it uses the orientation data which originally was not used. The affinity matrix created is used as input for the orientation algorithm. The orientation algorithm processes the data, compares it to the the current affinity matrix and updates it where it found that the affinity of the orientations is too low. The resulting affinity matrix is then used as an input for the dominant set algorithm. The dominant set algorithm was written in such a way that it already can be compared to the ground-truth using an evaluation method.

2.1 Conflab

The Conference Living Lab (ConfLab) is a collection of data about real life in-the-wild free-standing social interactions created by TU Delft's Socially Perceptive Computing Lab [6]. This data set was created because privacy-sensitive data was needed to analyze social networking where people were free to move into conversations or away from them. It used 49 Midges [5] during a social event. The Midges are labeled 1 to 50, however 38 is missing in the set. These Midges kept track on different kinds of data during the event, for example the RSSI data and the absolute rotation.

For this research the focus is on the orientation data of the Midge. The orientation data is stored in vectors for every Midge [7]. Every vector contains a timestamp, a value for the magnitude and the xyz coordinates. Where the magnitude is a real number and the xyz coordinates are imaginary. The xyz coordinates are based on the absolute orientation of the device, as can be seen in figure 1 from [8]. This is called a quaternion [9]. These vectors are first converted to Euler notation [10] and saved in a new data frame. Euler uses angles instead of coordinates, where the angles ϕ, ψ and θ represent the orientation of the Midge. Where ψ is a rotation about the Z-axis, θ is rotation about the Y-axis and ϕ is rotation about the X-axis For the conversion the method from [11] was used.

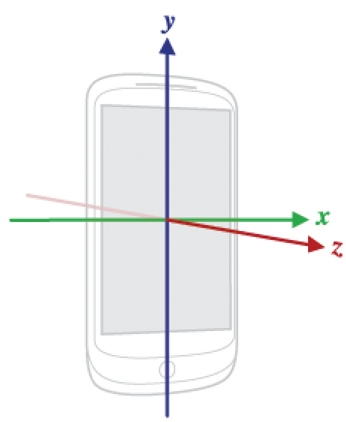

Figure 1: xyz coordinates visualized meaning [8]

For this research the ψ is the most interesting, as it indicates which direction the person is looking at. The ϕ indicates how far someone is leaning forward or backwards and the θ indicates how much they are leaning to the side. Because ϕ and θ do not tell much about who people are facing and as people were standing up wearing the Midge as a badge, ϕ and θ should be close to each other. Therefore, they will be not used for this research. However, for future research they are saved during the process of converting quaternion to Euler.

2.2 Filtering

Before the ConfLab can be used as input for a dominant set algorithm, some filtering and validating needs to be applied to the data. In the ConfLab data 49 Midges recordings are stored, however during the event 1 Midge was malfunctioning in its RSSI values. The malfunctioning Midge could potentially influence the results and therefore should be dealt with. In Dikker [1] it was managed to recreate the RSSI data for this Midge by using the data of the other Midges. Absolute orientations cannot be reconstructed the same way, however this was not necessarily as the orientation part of this Midge worked correctly.

2.3 Affinity Matrix

With the filtered data, an affinity matrix should be created as input for the dominant set algorithm (section 2.5). An affinity matrix is a matrix representing all pair-wise similarity scores [1]. Therefore it is an n by n matrix where n is the amount of Midges active at that timestamp. During the [1] research an affinity matrix was created to indicate which Midges where close to each other. This matrix is going to be updated by including the orientations of people (angle ψ) when determining the similarity scores.

2.4 Orientation for affinity

If two people are orientated in such a way that they are looking towards each other, there is a higher probability of social interaction. This, however, is not a hard requirement. During this research we will assume that people that are looking towards each other are in social interaction. A view frustrum (cone shape) can be used to indicate the visual limits of the eyes as shown in figure 2. For more explanation about this cone shape one can reference [12]. When someone is within the visual limits of the eyes, the person is visible for that person.

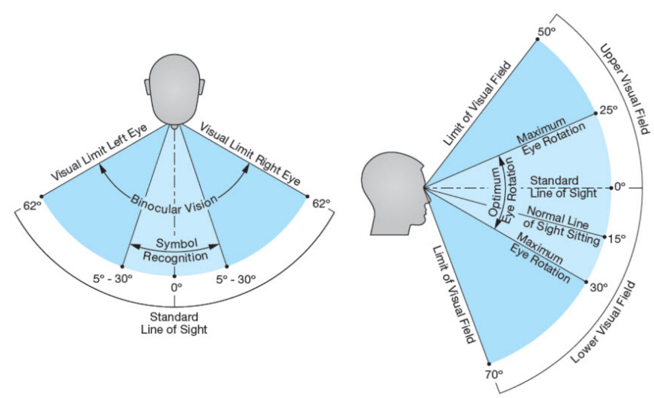

Figure 2: Field of vision and recommended head tilt and eye rotation angles [12]

The cone shape is defined mathematically in equation 2. If the reference angle is close to 360° the cone would exceed this. This causes the upper bound to be smaller than the lower bound. If we then would take the cone from the lower bound to the upper bound it would be to big. The problem is solved with the introduction of a condition in equation 1.

$$(\psi - \alpha) \bmod 360 > (\psi + \alpha) \bmod 360 \qquad (1)$$

$$Cone(y) = \begin{cases} 1 & (\psi - \alpha) \bmod 360 < y < (\psi + \alpha) \bmod 360 \\ & \text{if not condition } 1, \\ 1 & y > (\psi - \alpha) \bmod 360 \text{ if condition } 1, \\ 1 & y < (\psi + \alpha) \bmod 360 \text{ if condition } 1, \\ 0 & else \end{cases} \quad (2)$$

One problem is that absolute orientation does not reflect the orientation between two Midges. The absolute value is based on a point that every Midge uses as the zero point of the orientation. Then the angle between the current orientation and this zero point is the absolute orientation. Two absolute orientations only tell their orientation compared to the zero point and not what the angle between them is. This problem has been ran into before. In Montanari et all [13] the researchers discuss how to get relative orientations from the absolute orientations. However, Montanari et all [13] also mentions that this might lead to undesired behavior. They mention that different positions can lead to the same absolute orientation and distance. This can be problematic for finding the relative orientation. This latter will be further discussed in this paper.

To determine the affinity value (used for the affinity matrix) between Midge M_1 and Midge M_2, equation 2 is used to create their cones. The cone of M_1 is then flipped. Then if M_1 and M_2 do not have overlapping cones, they are not in that specific view area of each other in any configuration which equation 3 is used for. In this case the affinity between these people is too low and should therefore not be treated as one. This should remove most of the false positives from the RSSI based values.

$$Overlap(x, y) = \begin{cases} 1 & Cone(x.\phi) \cap Cone(y.\phi) \\ 0 & else \end{cases} \quad (3)$$

In figure 3 an example is shown of two configurations when the absolute angles are given. Here the cone shape is an angle of $-45°$ and $+45°$ from the orientation direction. In the left figure it is impossible for a and b to see each other, since at most one of the two can see the other no matter how we rotate a around b. For the right figure this is not the case, there is a configuration possible where a can see b and the other way around. This still does not need to be the case, however we do not have enough information to know if they are facing each other or not. Here we would need the relative angles to find out if this is the case or not.

The method can be done for any angle up to $90°$. Any value higher gives that both cones are larger then $180°$. Which means that there is always a possibility to see each other.

With these insights it is now possible to go over the RSSI value based affinity matrix and if they have a possibility to be in each others cone. If there is no possibility to face each other, they should not be classified as having a high affinity.

2.5 Dominant set algorithm

The dominant set algorithm is a clustering algorithm. For this research the dominant set algorithm designed in [4] was used. It takes an affinity matrix as input and then looks who

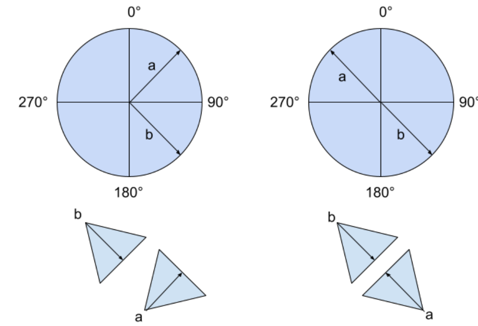

Figure 3: Using absolute angles it is possible to discard people who are for sure not facing at each other.

have a close affinity to each other. The practical benefit of this algorithm is that it doesn't require a specific number of groups to find. It can analyse the data and creates a group if the affinity between two data points is high enough. This is useful because the used data does not indicate the number of groups, in this way the algorithm can decide this based on the data provided.

2.6 Evaluation method

Earlier papers mostly use the F1 score (e.g. [1]) or something similar for their metric. To compare the results gathered in this experiment with the ones from Dikker [1], the decision has been made to use the F1 score as well, which can be found in equation 4. The F1 score also fits this experiment as it only takes into account the found groups and the missed groups.

The F1 scores take the True Positives (TP), False Positives (FP), False Negatives (FN) in as an input and returns a score between 0 and 1 (equation 4). Where 1 is the prefect classification and 0 is no correct classifications. TP are correct classifications of the algorithm while FP and FN are incorrect classifications. A TP is a group found by the algorithm that is exactly the same as a group in the ground-truth. It is also possible to use the T=2/3 rule, which states that if 2/3 of the group is correctly identified it will count as a TP. It was chosen to not include this rule, since Dikker [1] did not do so. To make the comparison of the results as fair as possible, it was chosen to not include the T=2/3 rule. A FP is a group found by the algorithm, but not containing the same people as in the ground-truth. And finally the FN is a group in the ground-truth, which was not found by the algorithm. True Negatives would indicate the people that are not considered as a group correctly, according to the ground-truth. True negatives do not exist, as the group size is allowed to be of size 1. With this method the group size of 1 is counted as a True Positive if a person is not in a social interaction at all.

$$F_1 = \frac{TP}{TP + \frac{1}{2}(FP + FN)} \quad (4)$$

3 Results

With the given method from section 2 the values for the cone can differ. The full results can be found in table 2. For the results in the table the following settings were used which can be found in table 1.

The results show also the intermediate results of calculating the precision (equation 5) and the recall (equation 6).

$$Precision = \frac{TP}{TP + FP} \tag{5}$$

$$Recall = \frac{TP}{TP + FN} \tag{6}$$

Constant	Value
Threshold	-55
Timeout	30
Symmetrisation	average
Reconstruction	True

Table 1: The constants used for every test. Based on the results gathered from Dikker [1].

ϕ	F1 Score	Precision	Recall
baseline	0.613614831	0.690084067	0.552402313
10°	0.177108703	0.257907989	0.134859074
15°	0.225018224	0.295565945	0.181658677
20°	0.256230892	0.319662213	0.213805055
25°	0.288835441	0.350259441	0.245740611
30°	0.345672503	0.408656313	0.299510673
35°	0.410922931	0.477805406	0.360465516
40°	0.475396523	0.542073174	0.423326121
45°	0.510465832	0.57503357	0.458934316
50°	0.541863393	0.610615441	0.487026771
60°	0.601400083	0.676681468	0.541192015
63°	0.611099105	0.684090816	0.552181924
65°	0.617870481	0.69094506	0.558774323
70°	0.62044782	0.696065039	0.559650078
75°	**0.622020091**	**0.696782798**	**0.561746387**
80°	0.616634675	0.692464973	0.55577319
85°	0.617112723	0.693874938	0.555642886
90°	0.613513099	0.693725773	0.549927259
91°	0.618257125	0.694434675	0.557140377

Table 2: The results of the different angles and the difference compared to the baseline.

The results show that when the angle of the cone is small, the precision and F1 score will be very low. On the other hand when the angle is very big, it comes closer to the baseline values. When the angles is $91°$, which will cause that the angle is big enough that there always will be overlap, the F1 Score and precision are very close to the baseline. As the rotation data will have no influence on the true or false positives. The F1 score is maximal at an angle of $75°$ and the precision is maximal at the same angle.

4 Discussion and Future Work

The method used to gather the results is not a perfect solution. It can only discard values if it is impossible by the method to find an angle where two Midges are facing each other. This does not mean that the two Midges are looking at each other or are in a social interaction group together. Furthermore it can still happen that two Midges who are not facing each other still have a high affinity to be in a social interaction.
An improvement would be using orientation data that is relative to each Midge. With absolute orientations this cannot be found without some more information about the location of the Midge. Therefore, it might be a good idea to measure this data. To gather this data, more research is needed in the field of orientations.
Another option would be tracking location data. Knowing where a Midge is on a grid together with the absolute orientation, gives the possibility to calculate the relative angle. The relative angle can then be used to calculate if two midges are facing each other. This could be achieved by getting positions from video data for example. Using the positions from the video data could be an interesting follow-up research.
Other research questions would be more in dept of when people are facing each other. The correlation between facing each other and having social interaction. When is someone facing another person? Is there a relative angle that social interaction is impossible? Further research can be done in the field of psychology to answer these questions. With those answers a more reliable definition can be made for the ideal cone shape.

5 Responsible Research

The Midge is a small recorder with a lot of (personal) data. This data should be well protected. Currently the ConfLab data set is the only data set containing this data. Only a limited set of people can access the data and signed a End User License Agreement (EULA), which states the use of the data and how the data should be handled. A copy of the full EULA can be found in appendix A. There are also smaller sets of data for testing and such, however these are not based on real personal data.
Finding F-formations can be used in many ways, however there are cases where it could be used in a way that might not be intended. To gather the data of F-formations a lot of private data is needed from all individuals. For example who they have a social interaction with. This data can tell a lot about the individuals. The data can then also be used for targeted advertisement. Which might not be the wished effect by people who participate in these activities.
However, it can also be used to learn from. People can learn how to have social interactions. This can help people who struggle in their social life to find a better connection/social interactions with others. Knowing where there is social interaction, these people who find it hard to interact or recognise interaction can join these groups, based on the information provided by the program. An example would be someone with Autism Spectrum Disorder (ASD). People with ASD try to avoid social interaction as they find it hard to find and maintain social interaction.

6 Conclusions

The goal of this report is to research the performance of including orientation data for finding F-formations in the social interaction data from the Conflab data set and compare the

results with the proximity-based data as the baseline. To find the performance from orientation data as addition to proximity data, a method was developed that used absolute orientations and the proximity data from the Midges. With the absolute value and the proximity alone, the relative orientation cannot be calculated. However, given the absolute orientation it can be calculated if two midges have a possibility for an orientation in which they can face each other. If this angle does not exist, then it is impossible for those Midges to face each other. When this is impossible, the probability of having affinity between these Midges is low.
For the experiment a view frustum (cone shape) was used to identify when two Midges are facing each other. If the absolute values of the cone never can overlap, the affinity is too low. Then the value from the proximity-based approach is overwritten to the minimal value. This removes false positives found by the proximity-based approach.
From the results it can be concluded that a small cone will remove too many true positives, rather than false positives. This will lower the F1 score significantly as the true positives are now false negatives. As the angle of the cone becomes larger the F1 score as well as the precision improves. At an angle of 75° the F1 score is maximal and slightly better then the F1 score of the baseline. This means that here the most false positives are removed and the most amount of true positives are kept. Going higher then 75° leads to a more closer score to the baseline, as almost no value will be overwritten.
In conclusion the orientation data can improve the F1 score compared to the proximity-based approach. However, there is not a significant high difference.
Different improvements can be made to the current work. Using location data gathered from other available data, the relative orientation can be calculated/estimated. The relative orientation can then directly reason about the orientation between two Midges without the need of the distance. Combining the RSSI and relative orientation into one method to find a new threshold for social interaction.

References

[1] Bram Dikker. "Finding f-formations using dominant sets in the Bluetooth proximity data of the Conflab data set". In: (2022).

[2] Igor Nesteruk. "SIR-simulation of Corona pandemic dynamics in Europe". In: *medRxiv* (2020). DOI: 10.1101/2020.04.22.20075135. eprint: https://www.medrxiv.org/content/early/2020/04/25/2020.04.22.20075135.full.pdf. URL: https://www.medrxiv.org/content/early/2020/04/25/2020.04.22.20075135.

[3] Prof. Camerla Troncoso, et all. "Decentralized Privacy-Preserving Proximity Tracing". In: 25-05-2020 (2020), pp. 8–12.

[4] Hayley Hung and B. Krose. "Detecting F-formations as dominant sets". In: Nov. 2011, pp. 231–238. DOI: 10.1145/2070481.2070525.

[5] *GitHub - TUDelft-SPC-Lab/spcl midge hardware: Hardware design files for the SPC lab MIDGE wearable devices for in-the-wild social data collection.* https://github.com/TUDelft-SPCLab/spcl_midge_hardware. 2019.

[6] Chirag Raman et al. "ConfLab: A Rich Multimodal Multisensor Dataset of Free-Standing Social Interactions In-the-Wild". In: (May 2022).

[7] TDK InvenSense. "World's Lowest Power 9-Axis MEMS MotionTracking™ Device". In: (2017), pp. 83–84.

[8] *SensorEvent — Android Developers.* Mar. 2022. URL: https://developer.android.com/reference/android/hardware/SensorEvent.html#sensor.type_rotation_vector-:. accessed: 17.05.2022.

[9] Sir William Rowan Hamilton LL.D. V.P.R.I.A. F.R.A.S. "XLVI. On quaternions; or on a new system of imaginaries in algebra". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 31.208 (1847). DOI: 10.1080/14786444708562643. eprint: https://doi.org/10.1080/14786444708562643. URL: https://doi.org/10.1080/14786444708562643.

[10] *Euler angles — Encyclopedia of Mathematics*. Feb. 7, 2011. URL: http://encyclopediaofmath.org/index.php?title=Euler_angles&oldid=15870. accessed: 25.05.2022.

[11] *GitHub - Simple Quaternions (squaternion).* https://github.com/MomsFriendlyRobotCompany/squaternion/tree/582ff5e7edc650e5505d78a530d25e4e010bfc63. 2018.

[12] *Environmental Considerations and Human Factors for Videowall Design*. URL: https://www.extron.com/article/environconhumanfact. accessed: 26.05.2022.

[13] et all Alessandro Montanari. "Measuring Interaction Proxemics withWearable Light Tags". In: *Proceedings of the ACM on Interactive, Mobile,Wearable and Ubiquitous Technologies* 2.1 (2018).

A End User License Agreement

In the following pages the EULA is given for the ConfLab Dataset. This EULA has to be signed before anyone can access the data. As the data contains personal information from the participants.

ConfLab Dataset
End User License Agreement (01/2020)

1 Scope

This End User License Agreement (hereinafter: “License”) is granted by the Socially Perceptive Computing Lab, Pattern Recognition and Bioinformatics group, Delft University of Technology, Van Mourik Broekmanweg 6, 2628XE, Delft, The Netherlands (hereinafter: “Licensor”) to the End User. This License is an Open Source License.

2 Material under License

The material under License (hereinafter: “Dataset”) is a Delft University of Technology dataset described in Schedule 1 composed of videos, wearable acceleration, proximity, low-frequency audio and surveys recorded during a mingling event. The dataset contains the processed data as well as any derivative work, products or services based on all or part of the data. All data contained within the Dataset have been collected and processed in accordance with the laws applicable in The Netherlands.

3 Copyright

The Dataset is the sole property of the Licensor and is protected by copyright. The Licensor reserves all rights to use in any way and distribute the Dataset. The Dataset shall remain the exclusive property of the Licensor. The End User acquires no ownership, rights or title of any kind in all or any parts with regard to the Dataset.

4 License

Subject to prior identification of the End User and signature of this License, the Dataset is freely available to the benefit of the End User. The Licensor grants to the End User the right to use the Dataset, **for academic non-governmental research with non-commercial purposes only.** The End User shall be responsible for any infringement of the present License by one of their subsidiaries and/or student’s. The End User may only disclose, give access, and/or transfer the rights related to the Dataset to subsidiaries and/or students, under the following conditions:

1. a copy of the present License has already been transferred to them;
2. the subsidiaries/students have fully read and understood all terms and conditions of the present License;
3. the access to the Dataset is granted under the close supervision of the End User;
4. the access to the Dataset is granted under the sole responsibility of the End User.

The Licensor grants to the End User the right to reproduce temporarily, to adapt, arrange and modify by any means the Dataset. The Licensor grants to the End User the right to rework and build upon the Dataset, or any component thereof, as necessary or desirable for research or technology development activity and create derivative products or services for the End User's own internal research and development. The End User is permitted to make a copy of the Dataset for archiving only. This License gives no right of any kind to the End User over the Dataset. The License is deemed non-exclusive and non-transferable to third parties.

5 Access

The End User may only use the Dataset after this License has been signed and returned to the Licensor. The End User must return the signed and dated License by email, in PDF format to the Licensor at the following address: `h.hung@tudelft.nl`.

The End User will receive access to pseudonomised data and will not attempt to establish the identity of, or attempt to contact any of the experiment subjects. The End User will not attempt to make direct contact with ConLab PIs or staff at sites concerning the specific results of individual subjects.

6 Copyleft and Open Science

The Licensor grants to the End User the License defined under section 2 under the sole condition that any derivative work, product, services, scientific developments of any kind or any improvement or modification of any kind based on all or part of the Dataset **shall be made freely available to other End Users**. A copy of any modification of the Dataset shall be submitted to the Licensor at the following address: `h.hung@tudelft.nl`. The End User shall ensure that, in the event of a modification of the Dataset, the substance of the Material is not altered and that, the facility still operates, and performs whatever part of its purpose remains meaningful. Any violation of this clause will give rise to immediate legal prosecution.

7 Distribution

The End User shall not, without prior authorization of the Licensor, transfer in any way, permanently or temporarily, distribute or broadcast all or part of the Dataset to third parties. The End User shall send all requests for the distribution of the Dataset to the Licensor at the following address: `h.hung@tudelft.nl`.

8 Research

Research includes all type of scientific research, irrespective of the object under scrutiny, aimed at achieving a progress in science.

9 Commercial use

Any commercial use of the Dataset is strictly prohibited.
Commercial use of the Dataset includes, but is not limited to:

- Proving the efficiency of commercial systems;
- Testing commercial systems;
- Using screenshots of subjects from the Dataset in advertisements;
- Selling data or making any commercial use of the Dataset;

- Broadcasting data from the Dataset.

Any violation of this clause will give rise to immediate legal prosecution by the Licensor. Any damages and/or unfair enrichment or the like of the End User due to the breach of the License shall be immediately restituted to the Licensor together with the derivative works, products and services based on all or part of the Dataset.

10 Publications

The End User shall reference the Dataset, or results obtained with it, in publications. Publications include, but are not limited to:

- Research papers,
- Reports,
- Articles,
- Presentations for conferences or educational purposes.

All uses of images and videos from the Dataset for demonstration purposes (articles, presentations, posters) must ensure the privacy of the participants in the Dataset, such as covering their faces.

All publications that report on research that use the Dataset will acknowledge this as follows: **(Portions of) the research in this paper used the ConfLab Dataset made available by the Delft University of Technology, Delft, The Netherlands.**

11 Illegal or criminal use of the dataset

Any illegal or criminal use of the Dataset by the End User is strictly prohibited.

12 Legal Disclaimer

The Dataset is granted without any warranty. Licensor shall not be held responsible for any damage (physical, financial or otherwise) caused by the use of the Dataset. Licensor shall not be held responsible of any illegal or criminal use of the Dataset by the End User.

13 Jurisdiction

This License is subject to and interpreted in accordance with Dutch Law. Any claim arising on the basis of this License shall exclusively be submitted to the Courts of Delft, The Netherlands.

14 Amendments

The Licensor is allowed to amend this License at any time without prior consent of the End User. The End User shall be informed about changes and given the choice to opt out of this License within 10 days to the Licensor at the following address: `h.hung@tudelft.nl`. Without any notice within 10 days and provided the amendment is not substantial, the amendment to the License will be fully applicable to the End User.

15 Warranties

The End User warrants that they are authorized signatory, adult and not legally forbidden to enter into this License. The End User warrants that they have read and understood all elements contained herein and that the signature apposed hereunder is the result of a fully aware decision. Explicitly, by signing, the End User agrees to the following conditions:

- The Dataset includes personal data with privacy protection.
- The End User is responsible for the correct use of the Dataset.
- The Dataset may not be further distributed.
- The Dataset may only be used for research purposes.
- The Dataset may not be used with the intention of identifying persons.
- The End User will be excluded in the case of abuse.
- The End User should take sufficient security measures for protecting the personal data.

By signing this License, the End User engages to strictly respect the conditions set forth herein and to respect all the laws applicable in The Netherlands in relation to data and personality protection with regard to the data contained within the Daaset collected and processed by the Licensor:

End-User (Authorized Signatory[1]):

Full Name (Block Letters):	
Title / Function :	
Email :	
Organization Address :	
Intended use of the data:	
Date :	
Signature :	
People with access under this EULA (person signing above is responsible for their actions and the correct use of the data)	Name: Position: Email: Name: Position: Email: Name: Position: Email:

[1]The person signing must have a permanent position in the institution. Access to MSC and PhD students is possible if their supervisor signs this EULA and includes them as people with access.