Measuring Visual Social Engagement from Proxemics and Gaze

Nicola Webb[1], Manuel Giuliani[1] and Séverin Lemaignan[2]

Abstract— **When we approach a group, there is an exchange of a multitude of verbal or non-verbal social signals to indicate that we are looking to interact. We continue to share these signals throughout the interaction to portray our thoughts and motivations. We define an interaction by the signals we send; sending different signals evokes a different response. Giving social robots the knowledge of group social interaction, they will have the ability to more effectively participate in these interactions in the real world. In this paper, we present the results from an online data collection study looking at social group dynamics. We collected a dataset of social behaviours in a group using a socially interactive game played online by 88 participants. We also introduce a novel visual social engagement metric, which is derived from two social signals: proxemics (distance between interaction participants) and mutual gaze. We propose a mathematical formula of both mutual gaze as the product of the mutual distances to the optical axis, and the visual social engagement as mutual gaze divided by distance between participants. Additionally, we investigate the influence of personality traits on the resulting interaction patterns. Using the metric, we create unique interaction profiles which suggest that participants have an interaction 'style'. No clear correlation between personality and interaction patterns was found.**

I. Introduction: Measuring Social Cognition

We consider an interaction to have begun once there is a mutual understanding between both parties that there is the intention of starting an interaction [3]. As humans, this understanding comes naturally, requiring little conscious cognitive effort. Robots, however, are still unable to automatically identify the intention of their interaction partners. Giulio et al. [19] explain the importance of social cognition for successful human-robot symbiosis. This requires that both the human and robot have an understanding of their interaction partner's internal state. In this work, we introduce a metric to measure visual social engagement which uses two social signals (proximity and mutual gaze). We use these signals as they are easily obtained from visual observation. The aim of the metric is to ultimately improve the social awareness of a social robot and enable them to assess an interaction and proceed accordingly.

A. Approaching Behaviours and Personal Space

Llobera et al. [14] investigated whether the rules of proxemics defined in [9] were consistent within a virtual environment when the participant's avatar was stationary and the virtual characters were dynamic. They found that the closer the characters were to the participant, the higher the levels of physiological arousal experienced, which aligns with proxemics theory. This study, amongst others ([21], [7]), establish that proxemic rules are kept when navigating a virtual environment. Therefore, in our game environment, we utilise proxemic conventions to establish a social space in which communication between characters can occur.

Repiso et al. [18] created a framework to improve the way in which robots position themselves when walking in formation with people. When compared to a traditional teleoperated control, participants preferred the new framework as it adjusted its distance and positioning according to the person's movements. Extroverted participants complained that the robot kept too large of a distance between them.

B. Inferring Social Behaviours from Physical Behaviours

Bartlett et al. [1] looked into whether ones internal states could be identified from physical behavioural information. Participants were asked to identify the perceived internal states of data from the PInSoRo dataset and were shown either original clips or skeleton and facial landmarks that had been extracted. They found that the movement data alone was just as meaningful as full scene data, which suggests when utilising machine learning techniques for this problem space, raw low-dimensional data could be used to learn internal states. These findings suggest the possibility of using raw interaction data to create a social engagement metric.

C. Group Formations

When conversing, humans enter a shared inner space and arrange themselves into specific spatial orientations. The spatial arrangement of a formation can be affected by physical space, context, membership categories and socioeconomic status ([15], [6]). As highlighted in *The Hidden Dimension* [8], when a group of people are gathered in a social situation, distances between them can be classified into: *intimate*, *personal*, *social* and *public*. Then, based on the systematic behaviour of the group, a more complex formation would be formed, defined by [10] as facing-formations (F-Formations).

Setti et al. [20] developed an iterative approach to detect the number of F-formations within static images. They used images taken by a monocular camera and extracted each person's head and body pose from the image, to detect the centre of each formation. Another approach was taken by [5] for human activity detection with three layers of analysis: individual, pair and group. The first layer detected the trajectory of the movement of each group member. Then, the internal interactions between the group members are detected based on the pairs within the group. Lastly,

*This work was not supported by any organization

[1]Bristol Robotics Laboratory, University of the West of England, Bristol, United Kingdom nicola.webb@brl.ac.uk, manuel.giuliani@brl.ac.uk

[2]PAL Robotics, Barcelona, Spain severin.lemaignan@pal-robotics.com

the group layer detects the collective activity of the group by taking one of the group members as a reference and represents their association with the other group members.

D. Using Games to Capture Behaviours

In recent years there has been some research into whether people's interactions with an online avatar are reflective of their real-life behaviours. Yee et al. used an online RPG to see if players adhere to social norms even in a virtual environment [22]. They determined that social norms were upheld within the virtual environment when groups of players were interacting. This demonstrates that even within a virtual environment, players replicate typical real-world behaviour making it possible to generalise virtual social interaction behaviour to the real world.

Additionally, Kozlov and Johansen illustrated the usefulness of using virtual environments within video games for psychological research on real-world behaviour [11]. They investigated whether the bystander phenomenon still presented itself within a virtual setup. Participants were less likely to help when there were more people in the surrounding area, demonstrating the bystander effect was still exhibited even within a game environment. In this work we do not claim that the outcomes can be directly applied to such interactions in the real world, however, we present a methodology for creating a representative metric for ones social behaviour from raw interaction data.

E. Social Presence

The field of social presence looks at ones ability to project themselves socially within a community. Oh et al. [17] conducted a review of the literature concerning social presence. They identified several studies covering visual representation in virtual or online interactions that determined that having a visual representation of an interaction partner affected the level of perceived social presence. In addition to a visible presence, the agent's behaviour can influence the levels of social presence. Lankes et al. [12] conducted a study comparing two video games of a similar nature, one with the inclusion of mutual gaze. Participants reported higher levels of social presence in the shared gaze condition, with players describing using gaze as an additional non-verbal communication method.

In our work, we do not measure the participant's perceived presence whilst playing. Instead, we record only the raw social signals during the entire game-play.

F. Measuring social engagement

Lemaignan et al. [13] looked to measure engagement and defined the concept of 'with-me-ness' in human-robot interaction, defined as how much a human is 'with' the robot whilst interacting. Their approach recorded children completing a collaborative task with a robot and estimated their focus of attention in real-time ('with-me-ness' metric). The resultant metric compared well against a baseline measure, providing a reasonable measure of social engagement.

Ben-Youssef et al. [2] looked to detect user engagement levels when interacting with a pepper robot. They recorded the participant's proximity to the robot, head and gaze direction and speech. Using this data, they created a recurrent neural network to be able to detect the moments when the user's social engagement was decreasing. These works manage to reasonably quantify the levels of engagement, however, the metrics do not take into account whether mutual gaze was present, or the perspective of the interaction partner. In this work, our metric incorporates proximity and mutual gaze from the perspectives of all parties.

G. Research questions and hypotheses

Building on this literature, the three specific research questions and hypotheses we study in this work are the following:

1) using readily available visual social signals (gaze and proxemics), how can we engineer a *synthetic measure of visual social engagement*, also suitable for group interactions?
2) building on this metric, can we generate *unique individual interaction profiles*, reflecting individual social behaviours and strategies?
3) finally, can we *correlate these profiles to personality traits*?

Our two key hypotheses are that (H1) using easy-to-observe social signals (gaze and proxemics), we can build a metric and derive individual interaction profiles that are unique interaction 'signature' for each participant; (H2) this signature should correlate to the participant's personality traits, in particular the level of extroversion.

II. Methodology

A. Game Development

The game has been developed using the open-source engine Godot [1]. The game has been developed in 3D as it provides a better way to capture social signals than in 2D as the additional dimension makes it easier to record signals, in particular, gaze direction. Within the game, 5 'groups' of non-player characters (NPCs) were created for the player to interact with. To determine whether an NPC was within speaking distance of the player, a 'personal' and 'social' space was established. Once the player entered the social space, the dialogue tree was triggered. From there the player could select from a series of dialogue options which were presented on the screen for the player to click on, to which the NPC would respond. We consider an interaction to have been initiated once this space has been entered by the player. If the player choose not to engage with the NPC and was still within the space in which an interaction was possible this was still considered ([9]). Note that, except for the player, the other characters were all controlled by the game itself.

B. Game Play

Before a participant entered the game, they were informed that they will be playing the role of 'detective' in a murder mystery-themed game. Their challenge was to speak to

[1] https://godotengine.org/

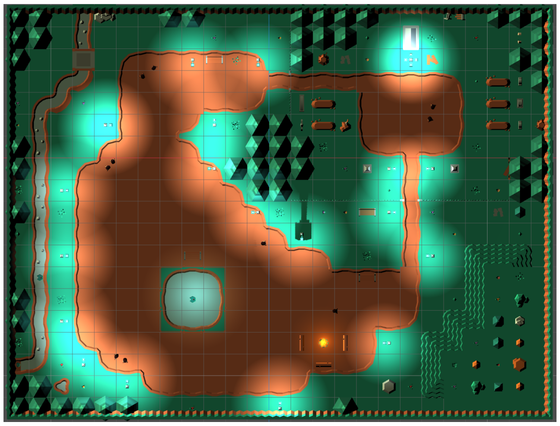

Fig. 1. Aerial view of game map

Fig. 2. Screen Capture from game

the 'townspeople' to gather information and determine who was guilty. When they entered the game, they were free to navigate the map, shown in Figure 1, in any way they wish, but were reminded of the 15-minute time limit, which was displayed on the screen at all times. The player could navigate the map using the arrow keys and controlled their gaze direction using their mouse, which used the move-look mechanism.

When a player approached a group of NPCs, they were given the option to click on an NPC and initiate a conversation, shown in Figure 2. Dialogue options were predetermined for simplicity, but players were given a choice of 4 options. Alongside each dialogue option was a choice of three emotions (happy, neutral and angry) to portray their chosen message on-screen. This in turn influence how the NPCs would respond to the player. Ultimately, the emotion dialogue selections were not used further in this work. When the timer runs out, the player was then asked to decide which NPC they suspect the most.

C. Data Collection

Participants were crowd-sourced using the Prolific Academic platform. The prerequisites for participants to sign onto the study required that they were fluent English speakers, 18 or over and were not exclusively using the macOS system (due to problems running the game on this platform). Basic demographic information was captured by Prolific which included participants' age, sex and country of residence and all other information was discarded. During the personality test, two attention check questions were added within the questions

During game-play, in-game analytics were recorded once the participants were informed and gave consent. Data that was collected included: frame rate, location, rotation degrees, dialogue and emotion selections and players' guess of who the 'murderer' was. Those whose average frame rate fell below 15fps had their data removed as it would have made game-play difficult. In total, 10 participants' data were removed due to this issue but were still compensated through Prolific, leaving 88 participants used in the analysis. The final participant pool consisted of 17 nationalities with a 65/23 male/female split.

III. Data Processing

A. Visual social engagement metric and interaction profiles

To quantify how 'social' a participant was, their in-game analytics were used to measure how 'engaged' they were during game-play. Firstly, by recording the location of all characters throughout the game, we can calculate how the proximity between NPCs and the player changes. A smaller distance measure could imply the likelihood of an incoming interaction or formation forming. Conversely, a greater distance reduces the possibility of the two parties initiating an interaction. Secondly, we collect each character's gaze direction throughout the game. With this, an estimated field of view (FOV) for the player was calculated using the typical measurements of the human binocular range of roughly 120° ([4]). Using this FOV, we can estimate which characters would be visible to the player and vice versa. If an NPC appears in the player FOV, in particular, close to the centre of the line of sight, it is likely that an interaction will or has been initiated. If an NPC was well out of the FOV, there was little chance of an interaction. The lower the distance metric and the closer the gaze direction then the higher the likelihood of an interaction taking place, and therefore the higher the visual social engagement score is.

A similar metric used by [16] determined whether non-verbal behaviour can influence proximity in virtual space during dyadic interactions. To examine the behaviours of both interaction partner's they proposed a new method of proxemic analysis proxemic imaging. This analysis utilises both the interpersonal distance of the partners and the gaze direction to determine participants' proxemic responses to different non-verbal behaviour. They begin the proxemic imaging once the participant enters the personal space of a virtual agent and discard all other data. In our work, we too initiate the computation of the metric once the social area was entered, however, it was calculated for all characters in the virtual space. In this way, we see the engagement level of all characters at different stages of interaction.

1) Visual social engagement: We define the *Visual Social Engagement metric* $\mathcal{S}_{AB}$ between two persons A and B as:

$$\mathcal{S}_{AB} = min(1, \frac{\mathcal{M}_{AB}}{d_{AB}}))$$

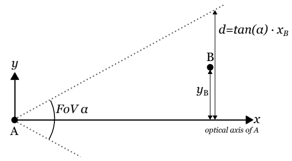

Fig. 3. Calculation of $Gaze_{AB}$, i.e. the normalised distance of B to A's optical axis.

where $\mathcal{M}_{AB}$ is the estimation of mutual gazing and d_{AB} the Euclidian distance between A and B.

$\mathcal{M}_{AB}$ is calculated as $\mathcal{M}_{AB} = \mathcal{M}_{BA} = Gaze_{AB} \cdot Gaze_{BA}$, with:

$$Gaze_{AB} = \{ m\,ax(0, 1 - \frac{y_B}{tan(\alpha) \cdot x_B}), if x_B > 0 0, otherwise$$

with x_B, y_B the coordinates of B in the reference frame of A, and α the field of view of A (see Figure 3).

As a result of this calculation:

- $\mathcal{S}_{AB}$ is symmetrical ($\mathcal{S}_{AB} = \mathcal{S}_{BA}$);
- $0 \leq \mathcal{S}_{AB} \leq 1$;
- $\mathcal{S}_{AB} = 0$ iff either A is outside of B's field of view, or B outside of A's field of view;

2) Interaction profiles: An interaction window is defined as the 15 frames leading up to the 'start' of an interaction and the 15 frames following. We consider an interaction to have been initiated once the player enters the 'social space' of an NPC. The size of the interaction window was decided as the point at which there was little change in the metric at the boundaries of the window. For each participant, the metric was averaged over that window of interaction to create an 'interaction profile'. The profiles give an idea of how someone interacts. The expectation was that we would see differences in the profiles between participants, as people may have different interaction styles.

B. Mutual gazing estimation

To determine whether characters were within the field of view (FOV) of another, we first computed the position of the NPCs in the frame of reference of the player and vice versa. With this, those within the FOV could be identified. Additionally, we calculated how far from the 'optical axis' a given character was. The further from the axis the closer to the edge of the FOV a character was. Those closer to the centre of the FOV, the higher their estimation was weighted. In the circumstance that a character was facing another from behind, and therefore had them within their FOV, the weighting would be low.

Figure 4 shows two frames of a plot of all characters' gaze orientation and FOV. In the left frame, we see the player approaching a single NPC within their FOV but was facing away. In the subsequent frame (right), the player was within a close enough proximity that the NPC turns to face them, as per their pre-programmed behaviour.

C. Visual social engagement estimation

The engagement estimation, as described in Section 2.3.1, was calculated as the product of the mutual optical axis

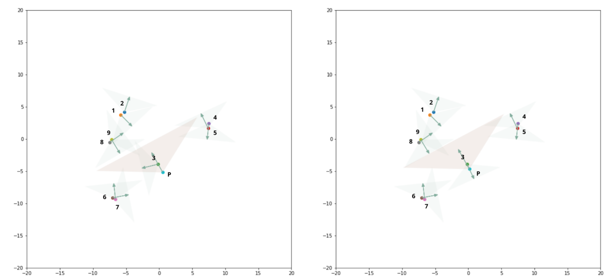

Fig. 4. Example plots of player (light blue) and NPC gaze orientation and FOV

distance of two persons, divided by the distance between these two persons. The distance between the player and NPCs was calculated using the Euclidean distance. This was computed for every participant and each NPC at each frame. Figure 5 shows the visual social engagement score for a single participant for one NPC (NPC 3, in green) for the duration of the game, alongside a snapshot diagram of the positioning and orientation of all characters. The peaks show the points in which the player was potentially interacting with the NPC. When the engagement score reaches 1, the player was currently in an interaction, as shown by the related diagram. When the peaks start to fall, the player was either moving away from the group or beginning to direct their gaze elsewhere. The second snapshot diagram shows a circumstance where the player was near NPC 3 and was facing them, but the NPC was facing away. Here the visual social engagement score was not at its peak as there was no mutual gaze. Finally, when both the player and NPC 3 were at a distance from one another and were not sharing a mutual gaze, the engagement score was 0.

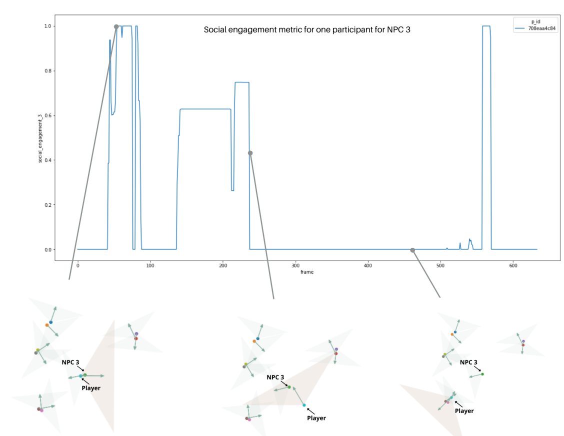

Fig. 5. Example visual social engagement metric for one participant for NPC 3

D. Personality profiles

The responses to the big 5 personality questionnaire, were accumulated to give each participant a score in each of the

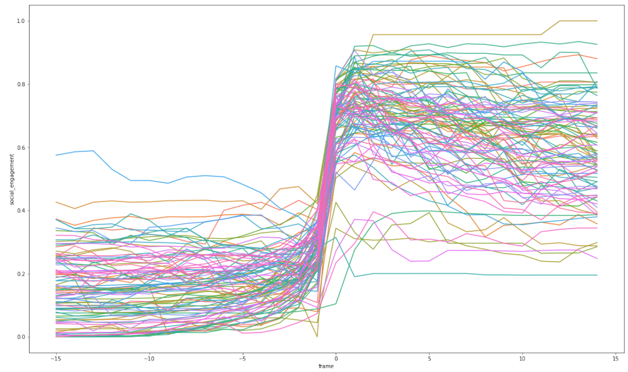

Fig. 6. Graph showing the interaction profiles for each participant

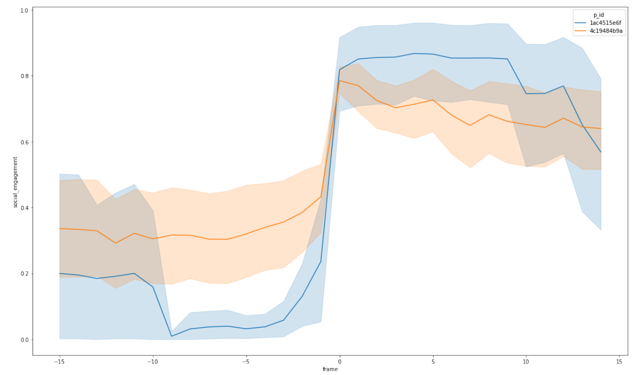

Fig. 7. Two participants mean visual social engagement score in all interaction windows. Shaded areas represent the 95% confidence interval.

following traits: emotional stability, extraversion, agreeableness, conscientiousness and intellect.

IV. Results

A. Interaction profiles

Figure 6 shows the interaction profiles for each participant. It shows a wide distribution of profiles, with differing levels of engagement before and after the point of interaction. We can see that most participants follow similar trends before and after the interaction is initiated, but with varying degrees of engagement. The 'step' at frame 0 shows the point at which players enter an NPC's social space. Figure 7 shows the interaction profiles of two participants that were randomly selected. Both participants follow a general trend of low engagement levels before interacting, a sharp increase at the point of interacting, and then levelling off. On average, the orange participant was more socially engaged leading up to an interaction compared to the blue but less engaged once the interaction had started, showing two unique interaction styles.

B. Distribution of the interaction profiles

Figure 6 shows the interaction profiles for each participant. Each profile was plotted across all frames of the interaction window, with each frame representing a second. The wide distribution in profiles suggests that participants do have different approaches to interactive engagement, and therefore different interaction styles. However, the pre-programmed behaviour of the game characters needs to be considered. As explained in Section 3.4, once the player enters the social space, the NPC turn to face the player, causing a spike in the level of interaction. Despite this, it was still expected that there would be a similar rise in the level of engagement once the player moves towards a potential interaction partner and initiates an interaction.

C. Interaction between personality and visual social engagement

From the personality questionnaire, each trait shows a fairly normal distribution with a central tendency. They also all share similar levels of variability, perhaps with the exception of intellect. The Pearson correlation coefficient was used to measure the relationship between features extracted from the interaction profiles defined in Section 2.3 and the personality traits collected through the big 5 questionnaire. The features extracted from the interaction profiles included: the mean engagement score, the standard deviation of visual social engagement scores, the mean engagement score before the start of interaction, the mean engagement score after the start of interaction and the delta of engagement scores. No strong correlations were found. We also investigated a cluster-based correlation between visual social engagement and personality profiles. We used the silhouette coefficient to determine the validity and the optimal number of clusters for our k-means algorithm and avoid a false-positive result. The optimal k value was 4, but no obvious clusters were generated.

V. Discussion

The visual social engagement metric, defined in Section 3.1, is an aggregation of a few social signals available within the game. The metric currently incorporates only mutual gaze and proxemics, however it could be easily extended to other signals, such as dialogue, in future. Using the visual social engagement metric, we create interaction profiles based on windows of interaction. The profile captures the average of how a user initiates an interaction. Even on impoverished data like our online game where social signals are limited, we observe a broad range of social interaction profiles, showing that players demonstrate differing interaction styles. Importantly, we can talk of individual profiles as the standard deviation of each profile is low, meaning their levels of engagement were fairly consistent during each interaction, showing that we have 'unique' interaction styles.

By extending the metric to include more signals, this may increase the accuracy of the interaction profiles. From the interaction profiles, we can extract features to enable comparison, in particular the delta (pre-interaction, post-interaction) and the standard deviation. Using these features, we looked for correlations between visual social engagement profiles and personality traits. Unfortunately, no strong correlations were found between any of the interaction profile features and personality traits. Additionally, there were no obvious clusters after conducting k-means clustering on these features. As we do see diversity in the interaction profiles, we can assume that other personality traits or other individual characteristics were responsible for the variation.

A. *Limitations*

The main limitation of the study is the lack of evidence regarding the real-world significance of our results. Indeed, the in-game NPCs were programmed to have a social behaviour, automatically turning towards the player as soon as the player enters their social space. It is unlikely that we would observe that with human-human interaction; emphasising the need to test with real-world data. However, the NPC behaviours were the same for all the players. As such, the diversity of interaction profiles that we observe here still indicates that these profiles are unique markers of the interaction style of each participant. Additionally, by using the online game platform, we were limited to the behaviours that we could record, in particular, any verbal communication methods. With this limitation, we may miss out on engagement behaviours that may ultimately affect their interaction style.

VI. Conclusion

This work introduces a novel *visual social engagement* metric, which is derived from two social signals: proxemics (distance between interaction participants) and mutual gaze. We propose a mathematical formula of both mutual gaze as the product of the mutual distances to the optical axis, and the visual social engagement as mutual gaze divided by distance between participants. By measuring the engagement over a window of time, centred around interaction onsets, and averaging these values over every interaction instance, we create for each participant *interaction profiles*. These profiles would suggest that participants have an interaction 'style': how quickly they engage their partners, how close to each other they stand and whether or not they look straight at each other.

Using a specially developed 3D online game, we recorded the behaviours of 90 participants. Crowd-sourced participants played for about 10min each a crime-solving game where they have to interact with nine other non-playing game characters (NPC) in order to solve the game. We build for each of the participants their interaction profile. We show that, even within the constraints of a virtual game with identical, pre-programmed behaviours for the NPCs, the participants' interaction profiles cover a broad spectrum of behaviours.

References

[1] Madeleine Bartlett, Charlotte Edmunds, Tony Belpaeme, Serge Thill, and Severin Lemaignan. What can you see? identifying cues on internal states from the movements of natural social interactions. *Frontiers in Robotics and AI*, 6:49, 2019.

[2] Atef Ben-Youssef, Giovanna Varni, Slim Essid, and Chloé Clavel. On-the-fly detection of user engagement decrease in spontaneous human–robot interaction using recurrent and deep neural networks. *International Journal of Social Robotics*, 11(5):815–828, 2019.

[3] Niklas Bergstrom, Takayuki Kanda, Takahiro Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. Modeling of natural human-robot encounters. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2623–2629, 2008.

[4] Randolph Blake and Robert Fox. The psychophysical inquiry into binocular summation. *Perception & Psychophysics*, 14(1):161–185, 1973.

[5] Zhongwei Cheng, Lei Qin, Qingming Huang, Shuicheng Yan, and Qi Tian. Recognizing human group action by layered model with multiple cues. *Neurocomputing*, 136:124–135, Jul 2014.

[6] Yasuharu Den. F-formation and social context: how spatial orientation of participants' bodies is organized in the vast field. In *Proceedings of LREC 2018 Workshop: Language and Body in Real Life (LB-IRL2018) and Multimodal Corpora (MMC2018) Joint Workshop*, pages 35–39, 2018.

[7] Doron Friedman, Anthony Steed, and Mel Slater. Spatial social behavior in second life. In *International Workshop on Intelligent Virtual Agents*, pages 252–263. Springer, 2007.

[8] Edward T Hall. *The hidden dimension*. Peter Smith Pub, 1992.

[9] Edward T Hall, Ray L Birdwhistell, Bernhard Bock, Paul Bohannan, A Richard Diebold Jr, Marshall Durbin, Munro S Edmonson, JL Fischer, Dell Hymes, Solon T Kimball, et al. Proxemics [and comments and replies]. *Current anthropology*, 9(2/3):83–108, 1968.

[10] Adam Kendon. *Conducting interaction: patterns of behavior in focused encounters*. Cambridge University Press, 2009.

[11] Michail D Kozlov and Mark K Johansen. Real behavior in virtual environments: Psychology experiments in a simple virtual-reality paradigm using video games. *Cyberpsychology, behavior, and social networking*, 13(6):711–714, 2010.

[12] Michael Lankes, Bernhard Maurer, and Barbara Stiglbauer. An eye for an eye: Gaze input in competitive online games and its effects on social presence. In *Proceedings of the 13th International Conference on Advances in Computer Entertainment Technology*, ACE '16, New York, NY, USA, 2016. Association for Computing Machinery.

[13] Séverin Lemaignan, Fernando Garcia, Alexis Jacq, and Pierre Dillenbourg. From real-time attention assessment to "with-me-ness" in human-robot interaction. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 157–164, 2016.

[14] Joan Llobera, Bernhard Spanlang, Giulio Ruffini, and Mel Slater. Proxemics with multiple dynamic characters in an immersive virtual environment. 8(1), November 2010.

[15] Paul Marshall, Yvonne Rogers, and Nadia Pantidi. Using f-formations to analyse spatial patterns of interaction in physical environments. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*, CSCW '11, page 445–454, New York, NY, USA, 2011. Association for Computing Machinery.

[16] Cade McCall and Tania Singer. Facing off with unfair others: Introducing proxemic imaging as an implicit measure of approach and avoidance during social interaction. *PLOS ONE*, 10(2):1–14, 02 2015.

[17] Catherine S. Oh, Jeremy N. Bailenson, and Gregory F. Welch. A systematic review of social presence: Definition, antecedents, and implications. *Frontiers in Robotics and AI*, 5:114, 2018.

[18] Ely Repiso, Anaís Garrell, and Alberto Sanfeliu. Adaptive side-by-side social robot navigation to approach and interact with people. *International Journal of Social Robotics*, 12(4):909–930, 2020.

[19] Giulio Sandini, Vishwanathan Mohan, Alessandra Sciutti, and Pietro Morasso. Social cognition for human-robot symbiosis—challenges and building blocks. *Frontiers in Neurorobotics*, 12:34, 2018.

[20] Francesco Setti, Chris Russell, Chiara Bassetti, and Marco Cristani. F-formation detection: Individuating free-standing conversational groups in images. *PLOS ONE*, 10(5):e0123783, May 2015.

[21] Laurie Wilcox, Robert Allison, Samuel Elfassy, and Cynthia Grelik. Personal space in virtual reality. *TAP*, 3:412–428, 10 2006.

[22] Nick Yee, Jeremy N Bailenson, Mark Urbanek, Francis Chang, and Dan Merget. The unbearable likeness of being digital: The persistence of nonverbal social norms in online virtual environments. *CyberPsychology & Behavior*, 10(1):115–121, 2007.