Does Full-Body Motion Reveal More? Comparing Human Inference of Personal Attributes from Upper-Body and Full-Body Motion

Jiayi Liu* UNC Chapel Hill
Arran Zeyu Wang† UNC Chapel Hill
Mark Roman Miller‡ Illinois Institute of Technology
Danielle Szafir§ UNC Chapel Hill
William Payne¶ UNC Chapel Hill

Abstract

Motion tracking in extended reality (XR) enables immersive experiences but also raises privacy concerns. Prior work has shown that motion data enables model-based personal attribute inference, with richer body tracking improving accuracy. However, less is known about whether tracking more body regions can also improve human perceptual inference. This study compares upper-body and full-body motion tracking in a between-subjects experiment ($N = 120$), where participants inferred gender, age, height, weight, and ethnicity from skeletal motion clips. While participants reliably inferred certain attributes in both conditions, there was no statistically significant difference in accuracy between conditions.

Index Terms: Social XR, privacy, HCI.

1 Introduction

Extended reality (XR) systems commonly use body motion tracking to support immersive interaction. Prior work has shown that motion from the head and hands, as well as from additional body regions, can be used for identification and inference of personal attributes such as gender and age [7, 8], indicating that motion data can pose privacy risks. Prior work usually focuses on algorithmic inference. In social XR, however, motion cues are visible to other users, who may become potential adversaries by inferring personal attributes from movement [5], potentially violating contextual integrity [9]. Such (mis)perceptions can shape social interaction and may cause harassment or other safety risks [10, 12]. Less is known about whether the scope of motion tracking affects perceptual risks. Consumer XR systems mainly track the head and hands, while support for full-body tracking has become more common. Full-body tracking exposes additional cues from the legs and hips and could potentially improve observers' ability to infer personal attributes.

This paper presents a user study examining differences in human attribute inference between upper-body and full-body skeletal motion. Motion sequences involving full-body coordination were selected from the Nymeria dataset [6], from which upper-body and full-body skeletal representations were constructed. Participants were assigned to either an upper-body or full-body condition and asked to infer five personal attributes of the motion subjects (gender, age, height, weight, and ethnicity) and report their confidence. We hypothesized that these five attributes could be inferred from skeletal motion, and that accuracy and confidence would increase when viewing full-body motion than upper-body motion alone. Our results show that gender and height were inferred reliably in both viewing conditions, while age and weight were inferred above random chance only in the full-body condition. However, comparisons between conditions did not provide statistically reliable evidence that full-body motion improves inference accuracy or confidence relative to upper-body motion for any of the five attributes. Through an exploratory analysis, we found that participants relied on stereotypes when inferring, associated gender with body-size inferences, and were influenced by their own demographics. This work contributes a controlled comparison of upper-body and full-body motion for attribute inference, and empirical evidence that increased tracking scope does not necessarily increase human inference performance. All supplemental materials are available at OSF.

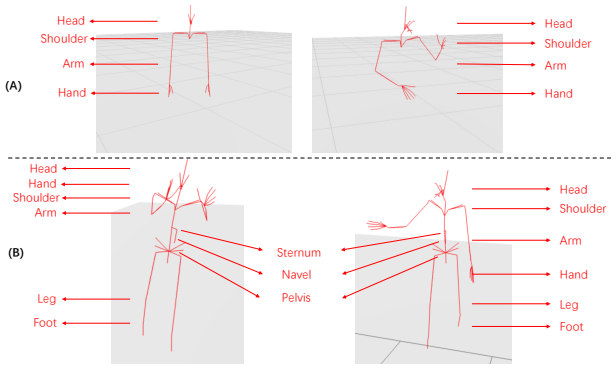


Figure 1: Example of upper-body (A) and full-body (B) rendering.

2 Methods

2.1 Stimulus Design

This study investigated whether full-body motion improves human perceptual inference of personal attributes compared to upper-body motion alone. First, to avoid confounding factors, we sought a stimulus set representing a diverse yet balanced distribution of personal attributes with minimized inter-attribute correlations. To this end, we selected a subset from Nymeria [6] optimized to meet these criteria. Second, to quantitatively evaluate the effect of tracking scope on perceptual inference, we employed a between-subjects design in which participants viewed either full-body or upper-body motion videos drawn from the same set of motion subjects.

Nymeria is a large full-body motion tracking dataset comprising 1,200 sequences from 264 subjects, with associated personal attributes. From this dataset, we excluded sequences containing multiple subjects and discretized continuous attributes into categorical bins to support balanced stimulus sampling: height into 10 cm intervals and weight into 10 kg intervals, chosen to reflect the coarseness of human perceptual granularity. Age was retained in its original bins, and ethnicity labels were mapped to standardized U.S. Census categories [11], as shown in Appendix A.1.

Additionally, the original dataset exhibited natural inter-attribute correlations that reflect real-world distributions but risk introducing perceptual confounds, where participants might infer one attribute from cues belonging to another (e.g., inferring gender from perceived height). To reduce such statistical dependencies and attribute imbalance, we applied a decoupling procedure that jointly minimized inter-attribute correlation and bin imbalance across at-

*e-mail: liujiayi@unc.edu
†e-mail: zeyuwang@cs.unc.edu
‡e-mail: mmiller30@iit.edu
§e-mail: danielle.szafir@cs.unc.edu
¶e-mail: William.Payne@unc.edu

tribute categories, formalized as:

$$\text{Score} = \sum_{\substack{(a_i,a_j)\in\mathscr{A}\times\mathscr{A}\\ i<j}} |\rho(a_i,a_j)| + \lambda \cdot \text{Imbalance}, \tag{1}$$

where $\rho(a_i,a_j)$ denotes Pearson's correlation coefficient between attributes a_i and a_j, and $\lambda = 1$ weights the relative contribution of the imbalance penalty. The attribute set is defined as $\mathscr{A} = \{\text{Gender}, \text{Height}, \text{Weight}, \text{Age}, \text{Ethnicity}\}$. The imbalance term is the average standard deviation of subject counts across bins:

$$\text{Imbalance} = \frac{1}{|\mathscr{A}|} \sum_{a\in\mathscr{A}} \sigma\big(\{n_{a,b}\}_{b\in B_a}\big), \tag{2}$$

where B_a is the set of bins for attribute a, and $n_{a,b}$ is the number of selected motion subjects falling into bin b.

We ran a stochastic hill-climbing algorithm for 10,000 iterations. At each iteration, a randomly selected subject within the sample was swapped with an excluded subject; the swap was accepted only if it decreased the score and preserved full bin coverage. This yielded a final subset of 18 motion subjects with reduced inter-attribute correlations and improved balance across attribute bins. Attribute distributions and correlation matrices before and after stimulus construction are provided in Appendix A.2.

To examine the effect of tracking scope on human perceptual inference, we selected walking as the target stimulus motion. Walking involves coordinated full-body dynamics that are informative across both upper- and lower-body segments. Prior work found walking to yield the highest model-based identification accuracy [1], making it a valid starting point for examining human perceptual inference. For each of the 18 selected individuals, we extracted a 15-second clip from continuous walking segments. We manually selected clips to ensure walking was the dominant action while preserving natural auxiliary movements (e.g., arm swings, head turns).

We rendered motion sequences as skeletal animation videos, following prior work employing video-based motion stimuli for perceptual inference tasks [2, 3]. We loaded Nymeria's full-body Xsens motion capture data (provided as `.glb`) via GLTFLoader in Three.js, and represented them as skeletal animations rather than avatar-based renderings to prevent body-shape cues from influencing attribute inference, especially weight inference. To ensure that inferred height judgments were based on motion rather than visual body size, we uniformly scaled each model to the same height before rendering. We centered and grounded each model on a fixed horizontal plane, and rendered it with a perspective camera (FOV = 75°) that followed the subject at a fixed viewing distance of 3.2 m, with smoothed interpolation to maintain stable framing across sequences. We recorded videos at 1110 × 1372 resolution and exported them at 30 frames per second using the MediaRecorder API. Within each session, we randomized the presentation order of videos to minimize potential order effects.

The experiment consisted of two viewing conditions. In the full-body condition, we rendered the complete skeletal structure. In the upper-body condition, the pelvis, lower spine, legs, and feet were hidden (see Figure 1). This separation was informed by pilot testing, which revealed that retaining the pelvis and lower spine preserved discernible lower-body motion signals. Excluding these segments ensured a clean separation between conditions. A full list of included and excluded bones is provided in Appendix B.

2.2 Participants and Procedure

We recruited 120 participants (60 per condition) from the US via Prolific, applying a minimum 95% approval rating and an age range of 18–65 as eligibility criteria (M=36.72, SD=11.17). The participants represented diverse ethnic backgrounds; demographic characteristics are reported in Appendix C.1. Three participants each

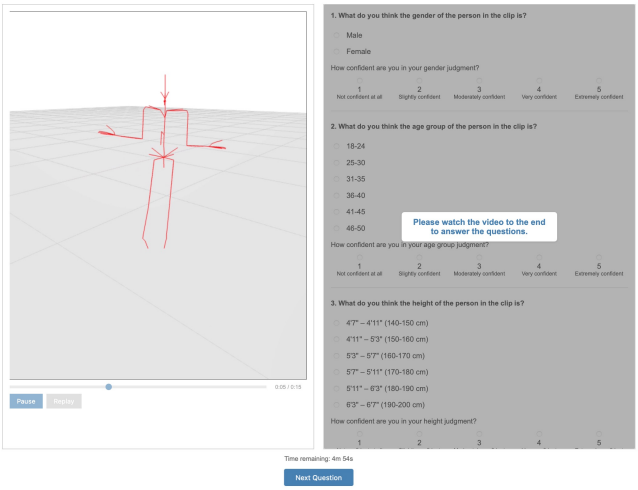


Figure 2: Interface. Left: 15-second skeletal movement clip with playback controls. Right: Questionnaire enabled after video is watched.

provided one incomplete response (17 of 18 trials answered): one in the full-body condition and two in the upper-body condition . All remaining responses from these participants were retained, and analyses were conducted on all available data. Participants were compensated $2.67 for their participation. The study took an average of 21.3 minutes to complete (SD=10.4 min).

Prior to the experiment, participants provided informed consent and completed a demographic questionnaire covering age, gender, race/ethnicity, and familiarity with motion capture, virtual reality, augmented reality, and animation systems (see Appendix C.2 for full details). Participants then completed a tutorial trial identical in structure to the formal trials but drawn from outside the experimental stimulus set, (see Fig 2) and response format. Participants could not begin the formal study until they completed the tutorial.

The formal study consisted of 18 trials, each corresponding to one motion sequence. Participants were required to view the full 15-second clip before response options became available. For each trial, participants inferred five attributes (gender, age group, height, weight, and race/ethnicity) and rated their confidence on a 5-point Likert scale ranging from 1 (*Not confident at all*) to 5 (*Extremely confident*) (see Appendix C.3 for full response format). Participants could replay and scrub through the video during the response phase. No feedback was provided during the experiment.

3 Results

Comparison to Random Chance: Before evaluating whether full-body motion conferred an advantage over upper-body motion, we first assessed whether skeletal motion supported above-chance inference of each attribute, aiming to confirm that both full-body and upper-body motions contain human-perceivable cues. Following the approach of prior work [3], we evaluated participant performance against a random-guess baseline, generated for each of the five attributes using the demographic distribution of the full Nymeria dataset. We then used Chi-square tests to compare participants' inference accuracy in each condition against their corresponding baselines. For each attribute and condition, we report observed accuracy, BH-FDR-corrected p-values across the ten attribute × condition tests, and Cohen's w as the effect size measure.

Our results suggest that people could reliably infer gender and height in both conditions, but age and weight in the full-body condition only. However, these advantages over random guessing were small, and we did not find evidence that ethnicity could be reliably perceived in either condition. Gender and height were inferred significantly better than random in both conditions. Gender accuracy was .56 in the upper-body condition and .59 in the full-body

condition, compared to a baseline accuracy of .50 (upper-body: $\chi^2(1, N = 1078) = 16.8$, $p < .001$, $w = 0.12$; full-body: $\chi^2(1, N = 1079) = 38.6$, $p < .001$, $w = 0.19$). Height accuracy was .28 in the upper-body condition and .27 in the full-body condition, compared to a baseline accuracy of .22 (upper-body: $\chi^2(1, N = 1078) = 22.6$, $p < .001$, $w = 0.14$; full-body: $\chi^2(1, N = 1079) = 15.7$, $p < .001$, $w = 0.12$). Age and weight, were inferred significantly better than random in the full-body condition only: age accuracy was .20 compared to a baseline of .17 ($\chi^2(1, N = 1079) = 8.41$, $p = .004$, $w = 0.09$), and weight accuracy was .20 compared to a baseline of .18 ($\chi^2(1, N = 1079) = 3.60$, $p = .048$, $w = 0.06$).

Full-body vs. Upper-body Comparison: We next examined whether full-body motion improved inference accuracy or confidence compared to upper-body motion. For each attribute, we fit a binomial generalized linear mixed model (GLMM) predicting the binary inference accuracy and a linear mixed-effects model (LMM) predicting the continuous confidence rating. Each model included condition as the fixed effect and crossed random intercepts for participant and video stimulus ($(1 \mid \text{participant}) + (1 \mid \text{video})$). Three accuracy models (age, weight, ethnicity) produced boundary-singular fits (participant-level variance ≈ 0); refitting without the participant intercept left all significance conclusions unchanged. We tested the condition effect in each model and Bonferroni-adjusted the five resulting p-values across the five attributes ($k = 5$).

We failed to find evidence that full-body motion improved participants' inference accuracy or confidence relative to upper-body motion for any of the five attributes. For accuracy, the models did not provide statistically significant evidence of a condition effect for gender ($\beta = 0.16$, $SE = 0.10$, $\chi^2(1) = 2.41$, $p = .604$), height ($\beta = -0.05$, $SE = 0.10$, $\chi^2(1) = 0.26$, $p = 1.00$), age ($\beta = 0.26$, $SE = 0.11$, $\chi^2(1) = 5.15$, $p = .116$), weight ($\beta = 0.14$, $SE = 0.11$, $\chi^2(1) = 1.48$, $p = 1.00$), or ethnicity ($\beta = 0.12$, $SE = 0.10$, $\chi^2(1) = 1.42$, $p = 1.00$). For confidence, the models likewise did not provide statistically reliable evidence of a condition effect for gender ($\beta = 0.02$, $SE = 0.16$, $t(118) = 0.11$, $p = 1.00$), height ($\beta = -0.01$, $SE = 0.15$, $t(118) = -0.07$, $p = 1.00$), age ($\beta = -0.07$, $SE = 0.16$, $t(118) = -0.45$, $p = 1.00$), weight ($\beta = -0.06$, $SE = 0.16$, $t(118) = -0.41$, $p = 1.00$), or ethnicity ($\beta = -0.33$, $SE = 0.18$, $t(118) = -1.81$, $p = .367$).

Exploratory Findings: Beyond the primary analyses reported above, we conducted exploratory analyses to understand response patterns and stimulus-level variability.

Response patterns and tendencies. Across both conditions, participants exhibited perceptual skews, favoring younger, taller, and heavier profiles alongside bias toward male and white attributes. For age, height, and weight, we measured signed bin distance—the difference between the selected and true bin, where positive values indicate overestimation. Results show that participants tended to infer younger ages (upper-body -0.60, full-body -0.59), taller heights (upper-body $+0.26$, full-body $+0.35$) and heavier weights (upper-body $+0.10$, full-body $+0.18$). For categorical attributes, response distributions deviated from the true stimulus distribution in both viewing conditions. Although the stimulus set was gender-balanced (50% male, 50% female), responses skewed toward male in both the upper-body ($+12.7$pp) and full-body ($+13.0$pp) conditions. Also, ethnicity responses showed a larger skew in both conditions, with white substantially over-selected (upper-body: $+19.9$pp; full-body: $+20.7$pp) and Asian substantially under-selected (upper-body: -17.8pp; full-body: -17.0pp). Confusion matrices are provided in Appendix D.

We also examined whether participants' gender inference were associated with body-size inference, finding that participants' choices of height and weight were heavily dependent on inferred gender. Among the motion subjects, male and female motion subjects differed modestly in height (169.9cm vs. 165.7cm) and weight (69.1kg vs. 69.0kg). However, participant responses revealed a much larger gap: subjects inferred as male were perceived to be 10.8 cm taller (175.9 cm vs. 165.2 cm) and 11.2 kg heavier (75.0 kg vs. 63.8 kg) than those inferred as female. LMMs showed that inferred gender was more strongly associated with inferred weight in the full-body condition than in the upper-body condition (12.15kg vs. 8.99kg; $\beta = -3.16$, $SE = 0.84$, $p < .001$). Inferred height showed a similar but marginally significant pattern (10.82cm vs. 9.43cm; $\beta = 1.39$, $SE = 0.81$, $p = .085$).

Finally, we examined whether participants' own demographic group shaped their responses in two ways: whether participants showed a tendency to infer their own demographic group, and whether accuracy or confidence was higher when the video's true demographic group matched the participant's own. We tested response tendencies using GLMMs comparing observed match rates against chance, and tested accuracy and confidence matching effects using the same GLMM/LMM framework as above, with own-group match and condition (and their interaction) as fixed effects, separately for gender, age, and ethnicity. We report results that reached statistical significance. For gender, confidence was significantly higher when the video's true gender matched the participant's gender in the upper-body condition ($\beta = 0.11$, $SE = 0.05$, $p = .028$). For age, accuracy was significantly higher when the video's true age matched the participant's age in the full-body condition ($\beta = 0.61$, $SE = 0.20$, $p = .002$). For ethnicity, participants were significantly more likely to infer their own ethnicity than chance in the upper-body condition (40.0% vs. 32.3%, $p = .014$), though not in the full-body condition. Accuracy was significantly higher when the video's true ethnicity matched the participant's own ethnicity in the upper-body condition ($\beta = 0.78$, $SE = 0.16$, $p < .001$), and this match effect interacted significantly with condition ($\beta = -0.61$, $SE = 0.22$, $p = .005$), such that the own-ethnicity accuracy advantage was smaller in the full-body condition.

Stimulus outliers. Accuracy varied substantially across the 18 video stimuli within each attribute, ranging from 18.3% to 84.9% for gender, 2.5% to 33.6% for age, 4.2% to 43.7% for height, 2.5% to 34.2% for weight, and 7.5% to 50.8% for ethnicity. Two videos ran counter to the overall response patterns described previously. Although responses skewed toward male overall (62.9%), one video with a true female motion subject elicited 71.7% female responses, and was also the video with the largest full-body accuracy advantage for gender (60.0% upper-body vs. 83.0% full-body). Although responses generally skewed toward younger age bins, one video with a true young motion subject (18–24) elicited an unusually high rate of older responses (41–45 or 46–50: 21.7% vs. 15.1% overall), and was among the harder videos for age (20.0% accuracy). Videos are available for viewing in supplementary materials.

4 Discussion

Taken together, our results suggest that motion contains limited yet perceivable cues for inferring certain personal attributes, but that increasing tracking scope from upper-body to full-body motion does not produce a reliable perceptual advantage for either inference accuracy or confidence. This finding aligns with prior biological-motion work showing that observers can infer gender from motion cues [4] and extends this line of work by examining a broader set of personal attributes. However, our results differ from prior work on VR avatar motion representations: whereas full-body motion-captured representations improved emotion perception in prior work [2], we did not observe a comparable full-body advantage for inferring personal attributes.

A closer look at individual attributes reveals distinct patterns. People inferred gender and height above the random-guess baseline in both conditions, while full-body movement provided no measurable benefit to either accuracy or confidence, suggesting that upper-body motion alone may carry sufficient cues for gender and height inference. These results are particularly notable because our

design forced participants to rely solely on motion dynamics by isolating skeletal movement and scaling all models at a uniform height. These findings suggest that changing a user's physical appearance may be insufficient to prevent the leakage of gender and height attributes. Consequently, effective privacy preservation cannot rely on anonymizing visual appearance alone. Future work should therefore prioritize motion anonymization while actively educating users about these overlooked privacy risks.

For age and weight, people performed above the random-guess baseline only in the full-body condition, suggesting that lower-body motion or full-body coordination may provide additional usable cues for these attributes. However, the full-body and upper-body comparisons did not show reliable full-body improvements, suggesting that these additional cues may have been weak or stimulus-dependent. This lack of a reliable perceptual advantage for full-body tracking suggests that for observers, the exposure of age and weight might stabilize at the upper-body stage, indicating that omitting the lower body does not significantly mitigate privacy risks.

Consistent with prior work [3], we did not find reliable evidence that ethnicity was inferable in either condition. Full-body motion also did not improve ethnicity inference, possibly because ethnicity-related motion cues are less directly tied to kinematics. However, this null result does not guarantee safety across other behavioral channels (e.g., voice), and should be interpreted with caution due to several limitations in our study design. First, our stimuli were restricted to walking, leaving it unclear how these inferences generalize to other motion. Second, our participant pool and video length were limited. We expect longer clips would better simulate real-world scenarios, where continuous exposure gives observers more time to infer attributes. Furthermore, participants viewed stimuli on desktop displays rather than through a VR headset, leaving how spatial presence and different animation approaches impact the perception of motion cues to future work.

Response Tendencies and Individual Variability: One explanation for the observed perceptual skews toward younger, taller, and heavier profiles, as well as a bias toward male and white attributes, is that when motion cues are weak or ambiguous, observers retreat to socially salient default categories rather than grounding their inference in individual motion differences. The gender-body-size association further illustrates this phenomenon, where inferences were driven by stereotypical heuristics. This gap, expanded in the full-body condition, suggests richer tracking data can counterintuitively yield more subjective, biased human inferences, where full-body motion amplified these stereotypes, leading participants to exaggerate physical differences based on their gender inferences.

Own-group analyses revealed three distinct mechanisms interacting with available information . For ethnicity, own-group matching produced an accuracy advantage that was strongest when visual information was limited (upper-body) and diminished with full-body information, suggesting that own-group recognition functions as a compensatory heuristic that observers rely on when diagnostic cues are scarce. For age, however, own-group matching produced an accuracy advantage only in the full-body condition, suggesting a genuine perceptual sensitivity that required richer motion cues. Gender showed neither pattern: own-group matching increased confidence without improving accuracy in either condition, indicating overconfidence rather than a perceptual advantage.

Finally, accuracy varied widely across individual stimuli, with some videos eliciting near-ceiling accuracy and others near-floor for the same attribute. This indicates that privacy risk for each user may deviate considerably, with some individuals possessing highly identifiable motion signatures, making them more vulnerable.

5 Conclusion

This study examined whether full-body skeletal motion increases human inference of personal attributes compared to upper-body motion alone. We found that people inferred gender and height from motion above a random-guess baseline, but that full-body motion did not significantly improve accuracy or confidence over upper-body motion. These findings suggest that upper-body motion alone can expose perceptible personal cues, and that adding lower-body motion does not necessarily increase human inference.

Acknowledgments

This work was supported by the Carnegie Grant and the Edward G. Holley Student Research Fund at UNC SILS.

The work was exempted by the Institutional Review Board of the University of North Carolina at Chapel Hill (IRB NO. 26-0032).

References

[1] Y. Abdrabou, L. Mecke, R. Rivu, S. Prange, Q. D. Nguyen, V. Voigt, F. Alt, and K. Pfeuffer. How unique do we move? understanding the human body and context factors for user identification. In *Proceedings of Mensch Und Computer 2023*, MuC '23, p. 127–137. Association for Computing Machinery, New York, NY, USA, 2023. doi: 10.1145/3603555.3603574 2

[2] K. Brandstätter, B. J. Congdon, and A. Steed. Do you read me? (e)motion legibility of virtual reality character representations. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 299–308, 2024. doi: 10.1109/ISMAR62088.2024.00044 2, 3

[3] Y. Jang, J. A. Guridi, D. Molitor, R. Kwon, and S. Nagarajan. Are you anonymous? inferring personal information from nonverbal behavior data tracked in immersive virtual reality. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing*, CSCW '23 Companion, p. 172–176. Association for Computing Machinery, New York, NY, USA, 2023. doi: 10.1145/3584931.3607004 2, 4

[4] L. T. Kozlowski and J. E. Cutting. Recognizing the sex of a walker from a dynamic point-light display. *Perception & Psychophysics*, 21:575–580, Nov. 1977. doi: 10.3758/BF03198740 3

[5] J. Liu, H. Chen, J. Kandel, M. R. Miller, and D. Szafir. Your motion, your identity: Perceiving personal traits from full-body motion data. In *2025 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 873–874, 2025. doi: 10.1109/ISMAR-Adjunct68609.2025.00233 1

[6] L. Ma, Y. Ye, F. Hong, V. Guzov, Y. Jiang, R. Postyeni, L. Pesqueira, A. Gamino, V. Baiyya, H. J. Kim, K. Bailey, D. S. Fosas, C. K. Liu, Z. Liu, J. Engel, R. D. Nardi, and R. Newcombe. Nymeria: A massive collection of multimodal egocentric daily motion in the wild, 2024. 1

[7] M. R. Miller, F. Herrera, H. Jun, J. A. Landay, and J. N. Bailenson. Personal identifiability of user tracking data during observation of 360-degree VR video. *Scientific Reports*, 10(1):17404, Oct. 2020. doi: 10.1038/s41598-020-74486-y 1

[8] V. Nair, C. Rack, W. Guo, R. Wang, S. Li, B. Huang, A. Cull, J. F. O'Brien, M. Latoschik, L. Rosenberg, and D. Song. Inferring private personal attributes of virtual reality users from ecologically valid head and hand motion data. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pp. 477–484, 2024. doi: 10.1109/VRW62533.2024.00094 1

[9] H. Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1):119–158, 2004. 1

[10] K. Schulenberg, G. Freeman, L. Li, and B. P. Knijnenburg. Does who you are or appear to be matter?: Understanding identity-based harassment in social vr through the lens of (mis)perceived identity revelation. *Proc. ACM Hum.-Comput. Interact.*, 8(CSCW2), Nov. 2024. doi: 10.1145/3686923 1

[11] U.S. Census Bureau. Updates to race/ethnicity standards for our nation. https://www.census.gov/about/our-research/race-ethnicity/standards-updates.html, Dec. 2024. Accessed: 2026-05-13. 1

[12] Q. Zheng, S. Xu, L. Wang, Y. Tang, R. C. Salvi, G. Freeman, and Y. Huang. Understanding safety risks and safety design in social vr environments. *Proc. ACM Hum.-Comput. Interact.*, 7(CSCW1), Apr. 2023. doi: 10.1145/3579630 1