

Legs Over Arms: On the Predictive Value of Lower-Body Pose for Human Trajectory Prediction from Egocentric Robot Perception

Nhat Le, Daeun Song, and Xuesu Xiao



Fig. 1: Human Trajectory Prediction Using Human Skeletal Keypoints from 360° Robot Egocentric Camera View.

Abstract—Predicting human trajectory is crucial for social robot navigation in crowded environments. While most existing approaches treat human as point mass, we present a study on multi-agent trajectory prediction that leverages different human skeletal features for improved forecast accuracy. In particular, we systematically evaluate the predictive utility of 2D and 3D skeletal keypoints and derived biomechanical cues as additional inputs. Through a comprehensive study on the JRDB dataset and another new dataset for social navigation with 360° panoramic videos, we find that focusing on lower-body 3D keypoints yields a 13% reduction in Average Displacement Error and augmenting 3D keypoint inputs with corresponding biomechanical cues provides a further 1-4% improvement. Notably, the performance gain persists when using 2D keypoint inputs extracted from equirectangular panoramic images, indicating that monocular surround vision can capture informative cues for motion forecasting. Our finding that robots can forecast human movement efficiently by watching their legs provides actionable insights for designing sensing capabilities for social robot navigation.

I. INTRODUCTION

Accurate multi-agent trajectory prediction is a key enabler for social robot navigation. As robots are extensively deployed in human-centric environments such as hospitals, campuses, and public venues, it increases the demand for robots' ability to anticipate human trajectories. This capability is essential for navigating safely and efficiently while respecting social norms and comfort of nearby people [1]–[4].

All authors are with the Department of Computer Science, George Mason University. {nle47, dsong26, xiao}@gmu.edu

While much of the literature still represents humans as points on 2D maps or bird's-eye view (BEV) tracks, multi-agent trajectory prediction can benefit from human visual features. Data-driven methods [5]–[8] perform well on BEV datasets but abstract away visual cues such as body pose, gaze, and gait. This abstraction is limiting for human-robot interaction with short histories and occlusions due to corners and doorways [9]. Recent work begins to close this gap by injecting human visual cues into forecasting models.

Similar to humans who use head and eye movement to observe the surrounding environments, robots in human-crowded environments can also benefit from surround-view perception to ensure safety and other social navigation objectives. While a single front-facing camera can miss people approaching from sides and behind with a limited field of view, a 360° RGB camera offers full coverage for perceiving nearby humans and capturing incoming trajectories from all directions. Recent development has made 360° RGB camera a low-cost sensor modality for mobile robots. Although its equirectangular images introduce distortion and horizontal wrap-around, these effects are well understood and tractable with standard engineering.

Motivated by the potential of rich first-person-view visual cues and 360° RGB camera (Fig. 1), this paper presents a study on which visual information, particularly skeletal regions and derived cues, matter the most for multi-agent trajectory prediction, and how useful 2D keypoints from equirectangular images are in practice. Based on Human Scene Transformer (HST) [9], we compare the added values of skeletal keypoints and derived biomechanical gait cues

from 2D and 3D pose for the trajectory prediction problem on two egocentric robot perception datasets. From our study, we observe the following:

- Lower-body 3D skeletal keypoints hold the most predictive values for our trajectory prediction task.
- 2D skeletal keypoints from equirectangular images improve trajectory prediction despite distortion.

We translate these findings into practical recommendations for feature selection and camera placement and outline future directions.

II. RELATED WORK

We review related work in human trajectory prediction in social navigation, human pose detection and intent prediction, and existing egocentric perception datasets.

A. Human Trajectory Prediction for Social Navigation

Predicting human motion is crucial to social robot navigation. Mavrogiannis et al. [10], [11] showed that predicting human trajectories can reduce discomfort around robots, although current models need significant improvements to increase robot navigation metrics. Deep learning approaches [5]–[8] model multi-agent interactions effectively, but they represent humans as 2D points and are benchmarked on bird’s-eye datasets like ETH/UCY [12], [13] and SDD [14]. These approaches ignore rich human visual features from robot egocentric perspective which are potentially beneficial for predicting human motion.

Salzmann et al. [9] introduced the Human Scene Transformer (HST), showing that full-body 3D keypoints and head orientation improve prediction from robot egocentric perception. Gao et al. [15] proposed a portable pose encoder that boosts trajectory forecasting, but evaluated mainly on synthetic data or single-agent scenarios and focused on egomotion of single agents. Taking full-body pose as input, however, no prior work has examined how different skeletal regions and derived biomechanical cues affect multi-agent human trajectory prediction in egocentric robot navigation.

Based on these insights, our study is the first to systematically compare different skeletal regions and biomechanical cues for multi-agent trajectory prediction from robot egocentric perspective, aiming to identify which human body signals matter most for socially aware prediction.

B. Human Pose Detection and Intent Prediction

Human skeletal pose detection has become practical through 2D keypoint detectors [16]–[18]. These models localize skeletal keypoints from RGB image frames, which can then be lifted to 3D pose sequences by fitting parametric human body models [19], [20] or learning human motion representations [21]. This process provides temporally consistent 3D skeletal keypoints usable for downstream tasks.

In autonomous driving, human pose is utilized to predict path direction and street-crossing intention. Minguez et al. [22] showed that focusing on leg and shoulder joints improves path direction prediction. Zhang et al. [23] and Li et al. [24] demonstrated the usage of derived biomechanical

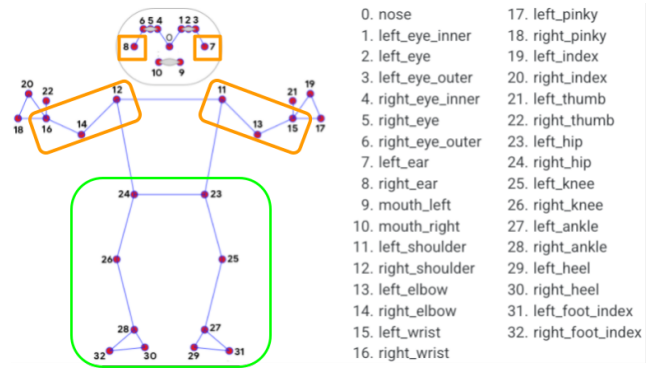


Fig. 2: Diagram for 33-keypoint 3D skeletal pose with K_L^{3D} and K_U^{3D} enclosed in green and orange boxes, respectively. Adapted from MediaPipePose [30].

cues for intention prediction. While informative, these studies focus on single-agent pedestrian direction and crossing intent in driving scenarios and do not address trajectory prediction.

C. Egocentric Perception Datasets

Several datasets such as JRDB [25], SCAND [26], MuSoHu [27], GND [28], and CODa [29] provide perception data from egocentric-view sensors. However, SCAND, CODa, most of MuSoHu, and part of GND do not provide full surround RGB views, which are crucial for detecting skeletal keypoints of all nearby humans. SCAND, MuSoHu, and GND additionally lack trajectory annotation and features relatively sparse interaction with humans. Although JRDB also lacks social navigation behaviors as the robot is often stationary, it remains a useful benchmark for our study with trajectory and keypoint annotations from full surround views. To better capture the dynamics of mobile service robots operating in human environments, we introduce in Section III-B a new dataset that emphasizes realistic navigation interactions with desired perception data.

III. METHODOLOGY

We formulate the problem of multi-agent trajectory prediction from egocentric perception and set the stage for our study on the values of different visual features in solving this prediction problem. We also present details of our datasets and implementations to facilitate our study.

A. Problem formulation

For our study on the predictive value of skeleton-based keypoint features, we formulate the problem of multi-agent trajectory prediction from robot egocentric perception as follows: Let \mathcal{A}_t denote the set of agents observed at time t . For each agent $i \in \mathcal{A}_t$, let $x_t^{(i)} \in \mathbb{R}^2$ be its 2D position in the robot-egocentric frame. Given a history window of length H , define the past trajectories

$$X_{t-H+1:t} = \{x_{t-H+1:t}^{(i)} \mid i \in \mathcal{A}_t\}. \quad (1)$$

The prediction target is each agent’s future positions over a horizon of F steps,

$$Y_{t+1:t+F} = \{x_{t+1:t+F}^{(i)} \mid i \in \mathcal{A}_t\}. \quad (2)$$

We learn a predictor f_θ that models a (potentially multi-modal) distribution over futures conditioned on past motion and optional egocentric visual features:

$$p_\theta\left(Y_{t+1:t+F} \mid X_{t-H+1:t}, G_{t-H+1:t}(s)\right), \quad (3)$$

where $G_{t-H+1:t}(s)$ denotes additional information detected from onboard perception (such as skeletal keypoints) and subsequently derived for observed agents based on feature configuration s . In particular, for observed agent $i \in \mathcal{A}_t$,

$$G_{t-H+1:t}^{(i)}(s) = K_{t-H+1:t}^{(i)}(s) \cup C_{t-H+1:t}^{(i)}(s),$$

where $K_{t-H+1:t}^{(i)}(s)$ denotes the set of keypoints available for configuration s and $C_{t-H+1:t}^{(i)}(s)$ denotes corresponding derived cues.

In our study with two datasets (see Section III-B), G instantiates different feature configurations s to enable controlled comparisons under a fixed training protocol. For 3D skeletal pose, we define K^{3D} , K_L^{3D} , K_U^{3D} to be the set of all 33 keypoints, the subset of 10 lower-body keypoints, and the subset of 10 upper-body keypoints, respectively (see Fig. 2). Subsequently, C_L^{3D} derived from K_L^{3D} includes leg articulation angles and step length, C_U^{3D} includes arm articulation angles and head orientation derived from K_U^{3D} , and $C^{3D} = C_L^{3D} \cup C_U^{3D}$. For the 2D skeletal pose with standard COCO 17-keypoint format, we define K^{2D} and K_L^{2D} to be the set of all 17 keypoints and the subset of 6 lower-body keypoints from the hips down, respectively.

On JRDB, we conduct two sets of experiments. In the first experiment, we evaluate the impact of different 3D skeletal feature configurations to quantify the relative predictive value of skeletal regions and whether derived indicators provide complementary gains. In the second experiment, we directly compare 2D and 3D skeletal inputs, training on the publicly available 2D annotations from JRDB-Pose alongside 3D pose estimates, to assess how much predictive performance is lost when only 2D features are available.

For our new dataset, we focus on the predictive utility of 2D skeletal inputs extracted from equirectangular panoramic images. We specifically evaluate whether 2D keypoints from distorted panoramic imagery provide useful predictive cues, and whether restricting inputs to lower-body keypoints yields stronger improvements than using the full 2D skeleton. This tests our hypothesis that leg motion provides the most informative signals for short-term trajectory forecasting, even under uncorrected panoramic distortion.

B. Datasets

We base part of our evaluation on JRDB [25], a large-scale dataset collected from a mobile social robot with LiDAR and 360° stereo camera in pedestrian zones and indoor environments. JRDB provides multimodal egocentric perception and ground-truth annotations for human detection

and tracking. A later version JRDB-Pose supplies 2D skeletal keypoints labels for all humans around the robot in COCO 17-keypoint format. Building on this foundation, HST used parametric human body fitting [19], [31] to estimate 3D skeletal keypoints for humans around the robots with a 33-keypoint format and made them publicly available. Although these methods may not be suitable for onboard real-time detection and tracking capabilities, they provide high-quality ground truth for skeletal representations that can convey intent and gait. We adopt JRDB and its pose extensions as a baseline to compare the gains from 2D and 3D skeletal features as well as derived mechanical cues and the predictive power of lower-, upper-, and full-body features for trajectory prediction.

To complement JRDB, we additionally use a new social navigation dataset that will be released later. In this dataset, a teleoperated robot navigates through dense human crowds in indoor campus canteens and hallways in a socially compliant manner to reach goals implicitly determined by the operators. The platform is an AgileX Scout Mini equipped with an Insta360 X4 camera, producing equirectangular panoramic images for full surround-view coverage, along with other onboard sensors. We employ HRNet [17] with MMPose [32] to detect and track humans and their 2D skeletal keypoints in image space, then convert the trajectories to robot Cartesian space using the pipeline by Bacchin et al. [33], which leverages the properties of equirectangular images and the flat-ground assumption. This yields continuous 2D tracks of humans together with their 2D skeletal poses in COCO 17-keypoint format. Despite panoramic distortions near upper and lower image boundaries, the Insta360 X4 camera provides a full vertical field of view and reliably captures full-body poses at close range (see Fig. 4). By contrast, the stitched 360° images in JRDB have a limited vertical field of view and often crop out the lower body and upper body of nearby pedestrians (see Fig. 4). In total, our dataset comprises 3 hours of socially-compliant and goal-oriented navigation, about three times the duration of JRDB. Unlike JRDB, where the robot frequently remains stationary or operates in relatively sparse outdoor settings, our dataset better reflects dynamic social navigation scenarios where the robot interacts with and navigates in dense human flows, making trajectory prediction a more challenging problem while imitating the real scenarios that mobile service robots encounter.

C. Implementation Details

We use the HST as the backbone for all experiments. It offers a flexible architecture for the multi-agent trajectory prediction problem as its modular design and attention mechanism make it straightforward to augment the model with additional modalities like skeletal keypoints and derived biomechanical cues. This extensibility allows us to perform controlled comparisons under a consistent modeling framework to ensure that observed differences in performance arise from the features themselves rather than architectural discrepancies.




Fig. 3: Our robot AgileX Scout Mini setup with onboard sensors: Insta360 X4 360° camera, Zed 2 RGB camera, Velodyne VLP-16 LiDAR, and WitMotion IMU.

We re-implement the HST in PyTorch, closely following the design and reference code released with the original work. Each additional feature stream, such as skeletal keypoints, their subsets, and derived biomechanical cues, is encoded with a lightweight MLP before being fused with human trajectories for full self-attention across agent and timestep dimensions. Biomechanical cues include head orientation, step length, and limb articulation angles for both arms and legs, which serve as compact indicators of motion intent [34], [35]. All models are trained with the AdamW optimizer, a batch size of 64, and configured to predict six consistent future modes per sequence.

For both datasets, we subsample the sequences to 3 Hz and train models to use information from six past timesteps (2s) to predict human trajectories in the next 12 timesteps (4s), ensuring consistent temporal resolution. We extract robot positions from odometry data and transform all trajectories into the robot’s egocentric frame at the first timestep of each sequence.

D. Evaluation

During evaluation, we restrict metric computation to pedestrians for whom at least partial keypoint observations and position histories are available, ensuring feature parity across comparisons. We report our results using four trajectory prediction metrics, which capture geometric accuracy, quality of multimodal predictions, and probabilistic soundness of the predictive distribution:

- Minimum Average Displacement Error ($MinADE$): the lowest average Euclidean distance between predicted and ground-truth trajectories across all sampled futures;
- Minimum Final Displacement Error ($MinFDE$): the lowest Euclidean distance between the predicted and ground-truth positions at the final timestep;

TABLE I: Evaluation on 3D skeletal keypoints and cues on JRDB

$G(s)$	$MinADE$	$MinFDE$	$MLADE$	NLL_{pos}
\emptyset (baseline)	0.39	0.64	0.71	0.86
K^{3D}	0.37	0.62	0.68	0.88
$K^{3D} \cup C^{3D}$	0.36	0.60	0.62	0.80
K_U^{3D}	0.38	0.58	0.83	0.82
$K_U^{3D} \cup C_U^{3D}$	0.36	0.57	0.61	0.75
K_L^{3D}	0.34	0.57	0.61	0.65
$K_L^{3D} \cup C_L^{3D}$	0.34	0.57	0.69	0.67

- Most Likely Average Displacement Error ($MLADE$): the average Euclidean distance between the most likely mode and ground-truth trajectories; and
- Negative Log-Likelihood of Positions (NLL_{pos}): the negative log-likelihood of the ground-truth trajectory under the model’s predicted probabilistic distribution.

IV. RESULTS AND DISCUSSION

We present our experiment results and discuss the findings to make recommendations for future feature selection and camera placement for human trajectory prediction.

TABLE II: Evaluation of 3D vs. 2D skeletal feature configurations on JRDB

$G(s)$	$MinADE$	$MinFDE$	$MLADE$	NLL_{pos}
\emptyset (baseline)	0.42	0.62	0.72	1.08
K^{3D}	0.39	0.56	0.65	0.88
K_L^{3D}	0.37	0.54	0.64	0.77
K^{2D}	0.41	0.6	0.74	1.08
K_L^{2D}	0.41	0.59	0.72	0.97

TABLE III: Evaluation of 2D keypoints detected from equirectangular images on our dataset

$G(s)$	$MinADE$	$MinFDE$	$MLADE$	NLL_{pos}
\emptyset (baseline)	1.1	1.44	1.86	3.1
K^{2D}	1.02	1.34	1.6	2.8
K_L^{2D}	1.03	1.35	1.73	2.86

A. Predictive Value of Lower-body Keypoints

Table I presents the experimental results for different configurations of 3D skeletal keypoints on the JRDB dataset. The result shows that restricting additional inputs $G(s)$ to lower-body 3D keypoints K_L^{3D} yields the largest combined gains over the baseline across all metrics. In particular, K_L^{3D} reduces $MinADE$ by 13%, $minFDE$ by 11%, $MLADE$ by 14%, and NLL_{pos} by 24%, indicating that leg motion features carry the strongest predictive signal for our trajectory prediction task. Full-body K^{3D} and upper-body K_U^{3D} configurations also outperform the baseline, but both trail K_L^{3D} . We further observe that augmenting keypoints




Fig. 4: Sample panoramic images from JRDB (left) and our dataset (right). In JRDB, nearby humans are cropped due to the camera’s limited vertical field of view, while in our dataset, full 2D keypoints can be detected even at distances below 1m.

with derived biomechanical cues improves the corresponding configurations by 1-4%. This suggests that while the cues are beneficial, the primary information is already well captured by 3D keypoints encoding under this training regime.

B. 3D versus 2D Keypoints on JRDB

We compare the 3D and 2D skeletal inputs on a paired subset of JRDB in which both features are available to ensure parity. As shown in Table II, the 3D input provides better predictive value for the trajectory prediction task. Both 2D keypoints K^{2D} and their lower subset K_L^{2D} yield only marginal changes on $MinADE$, whereas K^{3D} and K_L^{3D} yields 7% and 12% respectively. This is consistent with what we observe in Section IV-A. Two factors likely explain the minimal 2D gains on JRDB: (i) training and evaluation are done on a smaller dataset, and (ii) JRDB’s 360° camera often crops lower and upper body parts of nearby pedestrians, degrading the quality of image-based 2D pose more than that of 3D pose, which is obtained via sequence-level human parametric fitting.

C. 2D Keypoints from Equirectangular Images

Table III reports results on our social navigation dataset collected with an Insta360 X4. Full-body 2D keypoints K^{2D} , extracted from equirectangular images without distortion correction, reduce $MinADE$ by 7% relative to the baseline, with similar trends across other metrics. We attribute part of these gains to full vertical field of view of equirectangular images, which captures complete pose even at close range, mitigating the issues observed on JRDB. However, the difference between K^{2D} and K_L^{2D} is not significant, which is consistent with what we see in Section IV-B and motivating future analysis on when using lower-body 2D keypoints is preferable.

D. Discussion

1) *Sensor placement for surround-view:* Fig. 3 shows our robot hardware setup. The 360° camera is mounted at 0.85m from the ground, approximately half of average adult height. This placement reduces extreme equirectangular distortion near upper and lower image boundaries on nearby humans and likely contributes to the gains we see with 2D keypoints. Practitioners can take it into account together with this paper’s key results to consider proper sensor placements that balance perception capabilities and human’s comfort.

2) *Sensor choice:* While a single off-the-shelf 360° camera is simple and low-cost, a multi-fisheye rig can integrate more easily with different robot form factors and retain maximum vertical field of view. For social navigation, this can improve leg visibility and tracking continuity, especially near corners and obstacles.

3) *Human 3D Keypoint Detection On Equirectangular Images:* Recent work [36] outlines a practical path: learn an accurate fisheye-to-equirectangular projection model, fine-tune 2D keypoint detectors on synthetic panoramic data, and then perform 2D-to-3D lifting. This pipeline applies to both single 360° and multi-fisheye configurations and can improve 3D pose quality and achieve real-time processing capability for downstream tasks.

4) *Future Work:* While we standardize our evaluation on HST to isolate feature effects as embedded additional inputs are treated equally with simple encoders, our results could still be backbone-sensitive and confirming the results across distinct suitable backbones remains future work. Furthermore, our study does not explicitly consider other social cues such as gestures and group behaviors that could affect pose and alter feature importance. To explore further improvements, one can consider using more complex encoding scheme for skeletal representations [15], [37], [38]

V. CONCLUSION

This study investigates the predictive values of different sets of human skeletal keypoints for multi-agent trajectory prediction from robot egocentric perception. From the study, we observe two key findings. First, lower-body 3D skeletal keypoints provide the most predictive value. Incorporating the set of 3D lower-body keypoints yields the highest gain of 13% ADE reduction over a baseline that uses only 2D position histories for trajectory prediction task on JRDB. Using biomechanical cues such as limb articulation angles and head orientation, derived from corresponding sets of 3D skeletal keypoints, yields an additional improvement of 1-4%. Second, 2D skeletal keypoints from equirectangular images from the 360° camera improve prediction, even with distortion. On our collected social navigation dataset, 2D body keypoints extracted from equirectangular panoramic images improve performance by 7% over the same baseline. These findings indicate that skeletal keypoints, especially in lower-body region, and corresponding biomechanical cues capture

important information about human motion for trajectory prediction. We expect this paper to serve as a reference for designing robust human trajectory prediction system for mobile robots to achieve more socially compliant navigation.

ACKNOWLEDGMENT

This work has taken place in the RobotiXX Laboratory at George Mason University. RobotiXX research is supported by National Science Foundation (NSF, 2350352), Army Research Office (ARO, W911NF2320004, W911NF2520011), Google DeepMind (GDM), Clearpath Robotics, FrodoBots Lab, Raytheon Technologies (RTX), Tangenta, Mason Innovation Exchange (MIX), and Walmart.

REFERENCES

- [1] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core Challenges of Social Robot Navigation: A Survey," *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 3, pp. 1–39, Sep. 2023. [Online]. Available: <https://dl.acm.org/doi/10.1145/3583741>
- [2] A. Francis, C. Pérez-d'Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra *et al.*, "Principles and guidelines for evaluating social robot navigation algorithms," *ACM Transactions on Human-Robot Interaction*, vol. 14, no. 2, pp. 1–65, 2025.
- [3] R. Mirsky, X. Xiao, J. Hart, and P. Stone, "Conflict avoidance in social navigation—a survey," *ACM Transactions on Human-Robot Interaction*, vol. 13, no. 1, pp. 1–36, 2024.
- [4] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Motion planning and control for mobile robot navigation using machine learning: a survey," *Autonomous Robots*, vol. 46, no. 5, pp. 569–597, 2022.
- [5] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2255–2264.
- [6] B. Xu, X. Wang, S. Li, J. Li, and C. Liu, "Social-cvae: Pedestrian trajectory prediction using conditional variational auto-encoder," in *International Conference on Neural Information Processing*. Springer, 2023, pp. 476–489.
- [7] J. Yue, D. Manocha, and H. Wang, "Human trajectory prediction via neural social physics," in *European conference on computer vision*. Springer, 2022, pp. 376–394.
- [8] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *European Conference on Computer Vision*. Springer, 2020, pp. 683–700.
- [9] T. Salzmann, H.-T. L. Chiang, M. Ryll, D. Sadigh, C. Parada, and A. Bewley, "Robots That Can See: Leveraging Human Pose for Trajectory Prediction," *IEEE Robotics and Automation Letters*, vol. 8, no. 11, pp. 7090–7097, Nov. 2023. [Online]. Available: <https://ieeexplore.ieee.org/document/10239318/>
- [10] S. Poddar, C. Mavrogiannis, and S. S. Srinivasa, "From Crowd Motion Prediction to Robot Navigation in Crowds," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2023, pp. 6765–6772, iSSN: 2153-0866. [Online]. Available: <https://ieeexplore.ieee.org/document/10341464/?arnumber=10341464>
- [11] A. Stratton, P. T. Singamaneni, P. Goyal, R. Alami, and C. Mavrogiannis, "How human motion prediction quality shapes social robot navigation performance in constrained spaces," *arXiv preprint arXiv:2601.09856*, 2026.
- [12] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *2009 IEEE 12th International Conference on Computer Vision*, 2009, pp. 261–268.
- [13] A. Lerner, Y. Chrysanthou, and D. Lischinski, "Crowds by example," in *Computer graphics forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 655–664.
- [14] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *European conference on computer vision*. Springer, 2016, pp. 549–565.
- [15] Y. Gao, S. Saadatnejad, and A. Alahi, "Social-pose: Enhancing trajectory prediction with human body pose," *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [16] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, "Openpose: Realtime multi-person 2d pose estimation using part affinity fields," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 1, pp. 172–186, 2019.
- [17] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep High-Resolution Representation Learning for Human Pose Estimation," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Long Beach, CA, USA: IEEE, Jun. 2019, pp. 5686–5696. [Online]. Available: <https://ieeexplore.ieee.org/document/8953615/>
- [18] Y. Xu, J. Zhang, Q. Zhang, and D. Tao, "Vitpose: Simple vision transformer baselines for human pose estimation," *Advances in neural information processing systems*, vol. 35, pp. 38 571–38 584, 2022.
- [19] I. Grishchenko, V. Bazarevsky, A. Zafir, E. G. Bazavan, M. Zafir, R. Yee, K. Raveendran, M. Zhdanovich, M. Grundmann, and C. Sminchisescu, "BlazePose GHUM Holistic: Real-time 3D Human Landmarks and Pose Estimation," Jun. 2022, arXiv:2206.11678. [Online]. Available: <http://arxiv.org/abs/2206.11678>
- [20] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "Smpl: A skinned multi-person linear model," in *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 2023, pp. 851–866.
- [21] W. Zhu, X. Ma, Z. Liu, L. Liu, W. Wu, and Y. Wang, "Motionbert: A unified perspective on learning human motion representations," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 15 085–15 099.
- [22] R. Quintero Mínguez, I. Parra Alonso, D. Fernandez-Llorca, and M. A. Sotelo, "Pedestrian Path, Pose, and Intention Prediction Through Gaussian Process Dynamical Models and Pedestrian Activity Recognition," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 5, pp. 1803–1814, May 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8370119/>
- [23] S. Zhang, M. Abdel-Aty, Y. Wu, and O. Zheng, "Pedestrian Crossing Intention Prediction at Red-Light Using Pose Estimation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 3, pp. 2331–2339, Mar. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9423518/>
- [24] H. Li, Y. Jin, and G. Ren, "Interpretable Prediction of Pedestrian Crossing Intention: Fusion of Human Skeletal Information in Natural Driving Scenarios," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–18, 2024, conference Name: IEEE Transactions on Intelligent Transportation Systems. [Online]. Available: <https://ieeexplore.ieee.org/document/10623908/?arnumber=10623908>
- [25] R. Martín-Martín, M. Patel, H. Rezatofighi, A. Shenoi, J. Gwak, E. Frankel, A. Sadeghian, and S. Savarese, "JRDB: A Dataset and Benchmark of Egocentric Robot Visual Perception of Humans in Built Environments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 6, pp. 6748–6765, Jun. 2023, arXiv:1910.11792 [cs]. [Online]. Available: <http://arxiv.org/abs/1910.11792>
- [26] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone, "Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 807–11 814, 2022.
- [27] D. M. Nguyen, M. Nazeri, A. Payandeh, A. Datar, and X. Xiao, "Toward human-like social robot navigation: A large-scale, multi-modal, social human navigation dataset," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 7442–7447.
- [28] J. Liang, D. Das, D. Song, M. N. H. Shuvo, M. Durrani, K. Taranath, I. Penskiy, D. Manocha, and X. Xiao, "Gnd: Global navigation dataset with multi-modal perception and multi-category traversability in outdoor campus environments," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 2383–2390.
- [29] A. Zhang, C. Eranki, C. Zhang, J.-H. Park, R. Hong, P. Kalyani, L. Kalyanaraman, A. Gamare, A. Bagad, M. Esteva *et al.*, "Toward robust robot 3-d perception in urban environments: The ut campus object dataset," *IEEE Transactions on Robotics*, vol. 40, pp. 3322–3340, 2024.
- [30] MediaPipe Team, "Mediapipe pose," <https://github.com/google-ai-edge/mediapipe/blob/master/docs/solutions/pose.md>, Google AI Edge, 2025, documentation; Apache License 2.0. Accessed: 2025-09-15.
- [31] H. Xu, E. G. Bazavan, A. Zafir, W. T. Freeman, R. Sukthankar,

- and C. Sminchisescu, "GHUM & GHUML: Generative 3D Human Shape and Articulated Pose Models," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Seattle, WA, USA: IEEE, Jun. 2020, pp. 6183–6192. [Online]. Available: <https://ieeexplore.ieee.org/document/9157563/>
- [32] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu *et al.*, "Mmdetection: Open mmlab detection toolbox and benchmark," *arXiv preprint arXiv:1906.07155*, 2019.
- [33] A. Bacchin, F. Berno, E. Menegatti, and A. Pretto, "People tracking in panoramic video for guiding robots," in *International Conference on Intelligent Autonomous Systems*. Springer, 2022, pp. 407–424.
- [34] C. Prakash, R. Kumar, and N. Mittal, "Recent developments in human gait research: parameters, approaches, applications, machine learning techniques, datasets and challenges," *Artificial Intelligence Review*, vol. 49, no. 1, pp. 1–40, Jan. 2018. [Online]. Available: <http://link.springer.com/10.1007/s10462-016-9514-6>
- [35] J. Hart, R. Mirsky, X. Xiao, S. Tejada, B. Mahajan, J. Goo, K. Baldauf, S. Owen, and P. Stone, "Using human-inspired signals to disambiguate navigational intentions," in *International Conference on Social Robotics*. Springer, 2020, pp. 320–331.
- [36] K. Aso, D.-H. Hwang, and H. Koike, "Portable 3d human pose estimation for human-human interaction using a chest-mounted fisheye camera," in *Proceedings of the Augmented Humans International Conference 2021*, ser. AHs '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 116–120. [Online]. Available: <https://doi.org/10.1145/3458709.3458986>
- [37] A. Mohamed, K. Qian, M. Elhoseiny, and C. Claudel, "Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 14 424–14 432.
- [38] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.