Improving Real-Time Omnidirectional 3D Multi-Person Human Pose Estimation with People Matching and Unsupervised 2D-3D Lifting

Pawel Knap, Peter Hardy, Alberto Tamajo, Hwasup Lim, Hansung Kim
University of Southampton, UK, pmk1g20@soton.ac.uk

Introduction

Motivation

- Address a major limitation in existing models, namely the focus on single-person pose estimation
- Lack of omnidirectional Human Pose Estimation systems
- Cheap and Robust system utilizing mmWave radars and 360 camera
- Applications in in healthcare, entertainment, surveillance, sports, education, and beyond

Contributions

- Robust detection system, combining OpenPose [1] and off-the-shelf 2D-3D lifting algorithm [2]
- Performs consistently regardless of the number of individuals and it could theoretically can handle any number of them
- The only limitations being, the speed of off-the-shelf 2D detectors and the range of the radar sensor

Workflow

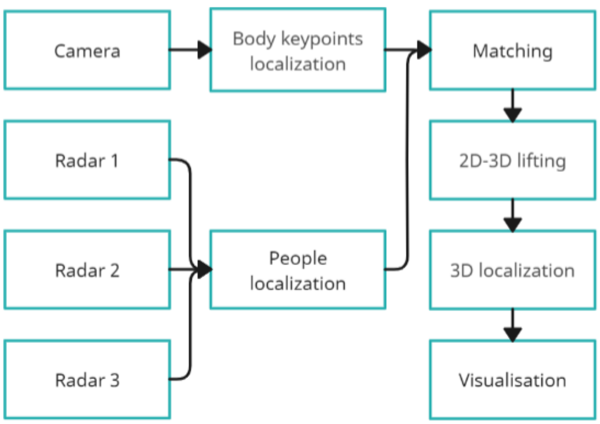

System setup

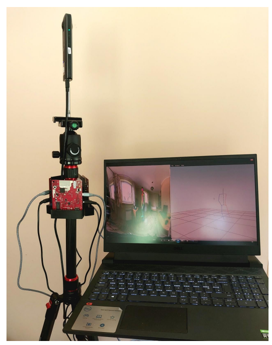

Input image

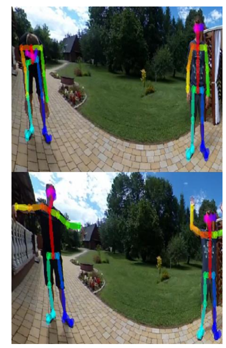

Reconstructed poses

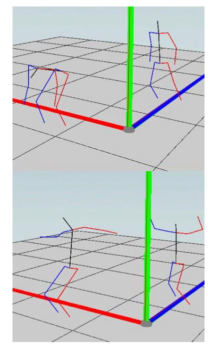

Experiments

Qualitative pose reconstruction on the Human3.6M dataset

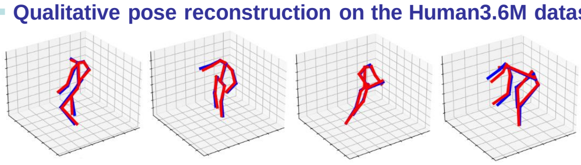

Radar detection errors

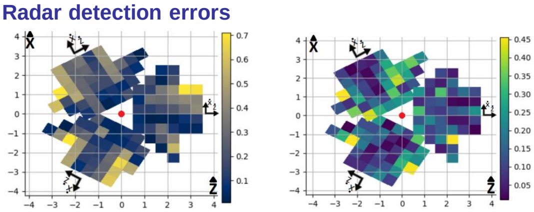

Pose reconstruction in the wild

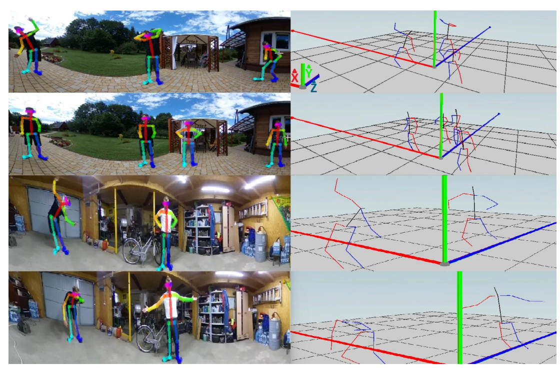

Results

Evaluation again preliminary work

- Radar calibration – Table I
- Matching of camera and radar detected individuals – Table II
- 2D-3D lifting algorithm – Table III

TABLE II

	Radar 1 ↓	Radar 2 ↓	Radar 3 ↓
Preliminary Work [3]	23.89% ± 6.57%	33.57% ± 50.55	66.89% ± 263.89
Ours	2.52% ± 2.51	9.44% ± 13.27	1.94% ± 1.52

TABLE I

Radar	Direction	Preliminary [3]	Ours
1	x	20.65	16.45
	z	11.41	11.45
2	x	26.19	24.86
	z	15.39	10.77
3	x	16.88	15.94
	z	13.83	13.46

TABLE III

Method	Occlusion	PA-MPJPE	N-MPJPE
LInKs [2]	None	33.8	61.6
Ours (Recreation)	None	37.2	61.7
Ours (Recreation)	Left Arm	52.1	78.1
Ours (Recreation)	Left Leg	46.0	73.2
Ours (Recreation)	Right Arm	49.8	75.7
Ours (Recreation)	Right Leg	44.5	71.6
Ours (Recreation)	Left Arm & Leg	62.0	86.0
Ours (Recreation)	Right Arm & Leg	60.2	83.7
Ours (Recreation)	Both Legs	69.3	99.8
Ours (Recreation)	Torso	88.4	122.0

* This work was partially supported by the EPSRC Programme Grant Immersive Audio-Visual 3D Scene Reproduction (EP/V03538X/1) and partially by the Korea Institute of Science and Technology (KIST) Institutional Program (Project No. 2E32303).

Conclusion

Summary

- Low average matching error of 4.63%
- Localisation errors reduced with radar and camera calibrations
- Lifting algorithm [2] results of a PA-MPJPE of 37.2 and N-MPJPE of 61.7 on the GT 2D poses in the Human3.6M dataset

Future work

- Optimizing algorithm speed
- Expanding the system's operation range
- Improving occlusion handling

[1] Z. Cao, G. Hidalgo Martinez, T. Simon, S.-E. Wei, and Y.A. Sheikh. "Openpose: Realtime multi-person 2d pose estimation using part affinity fields" IEEE Transactions on Pattern Analysis and Machine Intelligence, pages 1–1, 2019.

[2] Peter Hardy and Hansung Kim. "LInKs - Lifting Independent Keypoints - Partial Pose Lifting for Occlusion Handling with Improved Accuracy In 2D-3D Human Pose Estimation", 2023.

[3] Aarti Amin, Alberto Tamajo, "Real-time 3d multi-person pose estimation using an omnidirectional camera and mmwave radars," in Proc. ICEET, October 2023