visuAAL

Privacy-Aware and Acceptable Video-Based Technologies and Services for Active and Assisted Living

Privacy Preservation in Video-based AAL Applications

Siddharth Ravi, University of Alicante

Stockholm, Sweden

20 Apr 2023

2nd Doctoral Seminar

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 861091".

Universitat d'Alacant Universidad de Alicante

Project Coordinator

RWTH AACHEN UNIVERSITY

Stockholm University

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

TU WIEN

Privacy by Context

Privacy by Context

Progress Review (2021-2022)

Progress Review (2021-2022)

- Created a comprehensive review of visual privacy preservation techniques [1].
    - Proposed a new taxonomy of visual privacy preservation techniques for AAL.

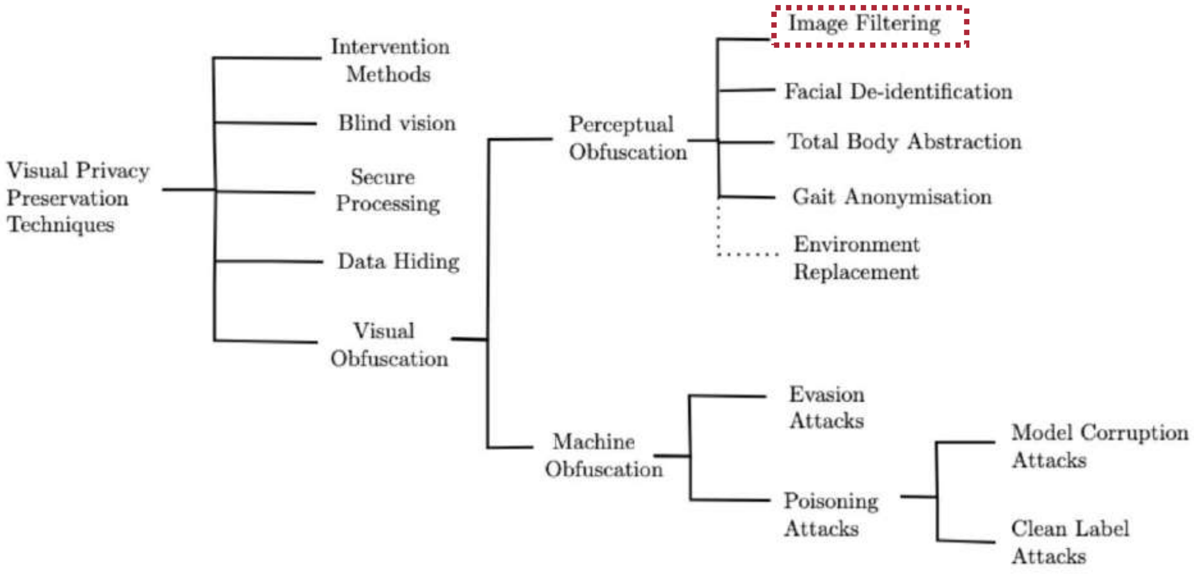

[1] Ravi, S., Climent-Pérez, P., & Florez-Revuelta, F. (2021). A review on visual privacy preservation techniques for active and assisted living. *arXiv preprint arXiv:2112.09422*.
[2] Mihailidis, A., & Colonna, L. (2020). A methodological approach to privacy by design within the context of lifelogging technologies. *Rutgers Computer & Tech. LJ*, *46*, 1.

Progress Review (2021-2022)

- Created a comprehensive review of visual privacy preservation techniques [1].
  - Proposed a new taxonomy of visual privacy preservation techniques for AAL.
  - Created a connection between the taxonomy for privacy by design proposed by Mihailidis and Colonna (2020) [2] and the proposed taxonomy.

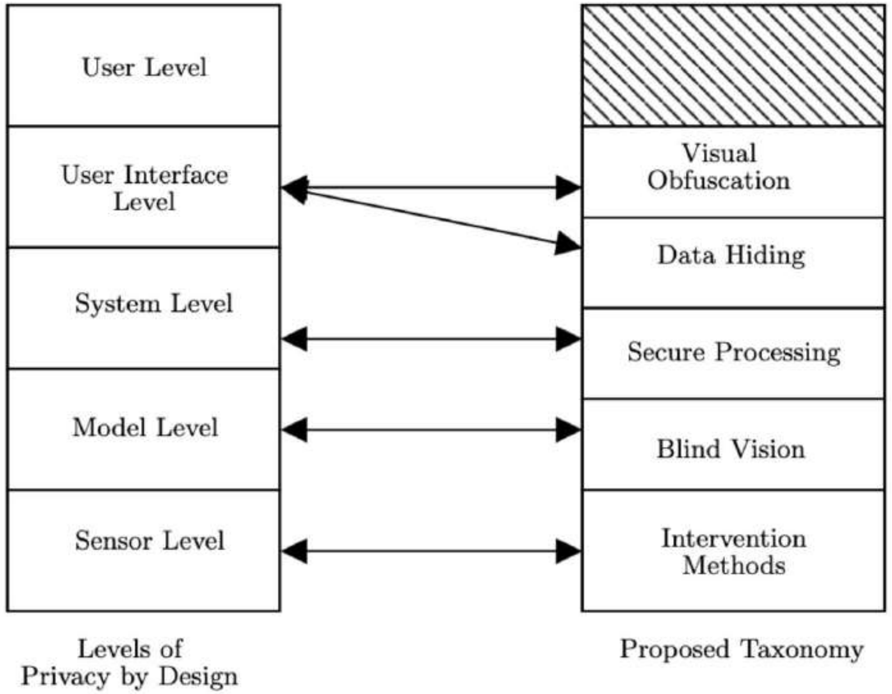

[1] Ravi, S., Climent-Pérez, P., & Florez-Revuelta, F. (2021). A review on visual privacy preservation techniques for active and assisted living. *arXiv preprint arXiv:2112.09422*.
[2] Mihaildis, A., & Colonna, L. (2020). A methodological approach to privacy by design within the context of lifelogging technologies. *Rutgers Computer & Tech. LJ*, *46*, 1.

Progress Review (2021-2022)

- Collaborated on studies of fairness of commonly-used visual privacy preservation methods with Sophie Noiret et al. [1, 2]

[1] Noiret, S., Ravi, S., Kampel, M., & Florez-Revuelta, F. (2022, June). On The Nature of Misidentification With Privacy Preserving Algorithms. In *Proceedings of the 15th International Conference on PErvasive Technologies Related to Assistive Environments* (pp. 422-424).
[2] Noiret, S., Ravi, S., Kampel, M., & Florez-Revuelta, F. (2023). Fairly Private: Investigating The Fairness of Visual Privacy Preservation Algorithms. *arXiv preprint arXiv:2301.05012*.

Privacy Preservation Reimagined: Top-view Omnidirectional Cameras in AAL

Privacy Preservation Reimagined: Top-view Omnidirectional Cameras in AAL

- Visual privacy preservation algorithms currently focus on lateral-view RGB(-D) camera images.
- Omnidirectional cameras with fisheye lenses provide a better alternative -
  - Less occlusions due to wide field of view, generally cheaper, and less obtrusive.
- A different set of challenges - human behaviour understanding (HBU) algorithms don't work due to the heavy distortions from the fisheye lens.
- No datasets exist specifically for Omnidirectional HBU challenges.
- Existing omnidirectional datasets are also mostly either synthetic, or staged.
- Created ODIN - a large-scale OmniDirectional INdoor Dataset capturing Activities of Daily Living from multiple synchronised viewpoints.

ODIN - An OmniDirectional INdoor dataset

ODIN - An OmniDirectional INdoor dataset

- Recorded activities of daily living in real indoor environments with varying levels of occlusion.
- 4 locations, 5 environment types, 15 participants, wide range of activities and poses.
- All modalities are synchronised, static cameras are all calibrated.
- First dataset aimed at HBU with top-view omnidirectional images.
- Made to be used for tasks as varied as activity recognition, person tracking and monitoring, scene understanding, biometric monitoring, novel view synthesis, generative modelling, 3D scene reconstruction, and image registration.
- First version of ODIN aimed at omnidirectional 3D human pose estimation.

Modality/characteristic	Amount
Omnidirectional RGB images	332K
Lateral-view RGB images	1.464M
Lateral-view infrared images	1.464M
Lateral-view depth images	1.453M
Environment meshes	3
Egocentric videos	52
Physiological readings	39
Accelerometer measurements	39
Participants	15
Locations	4
Types of environments	5

ODIN - An OmniDirectional INdoor dataset

- Recorded activities of daily living in real indoor environments with varying levels of occlusion.
- 4 locations, 5 environment types, 15 participants, wide range of activities and poses.
- All modalities are synchronised, static cameras are all calibrated.
- First dataset aimed at HBU with top-view omnidirectional images.
- Made to be used for tasks as varied as activity recognition, person tracking and monitoring, scene understanding, biometric monitoring, novel view synthesis, generative modelling, 3D scene reconstruction, and image registration.
- First version of ODIN aimed at omnidirectional 3D human pose estimation.

ODIN: An OmniDirectional INdoor dataset capturing Activities of Daily Living from multiple synchronized modalities

Siddharth Ravi[1[0000-0002-2301-569X]], Pau Climent-Perez[1], Théo Morales[2], Carlo Huesca-Spairani[1], Kooshan Hashemifard[1], Francisco Florez-Revuelta[1]

[1]University of Alicante [2]Trinity College Dublin

{siddharth.ravi, pau.climent, k.hashemifard, francisco.florez}@ua.es, moralest@tcd.ie, chs17@alu.ua.es

Abstract

We introduce ODIN (the OmniDirectional INdoor dataset), the first large-scale multi-modal dataset aimed at spurring research using top-view omnidirectional cameras in challenges related to human behaviour understanding. Recorded in real-life indoor environments with varying levels of occlusion, the dataset contains images of participants performing various activities of daily living. Along with omnidirectional images, additional synchronized modalities of data are provided. These include (1) RGB, infrared, and depth images from multiple RGB-D cameras, (2) egocentric videos, (3) physiological signals and accelerometer readings from a smart bracelet, and (4) 3D scans of the recording environments. To the best of our knowledge, ODIN is also the first dataset to provide camera-frame 3D human pose estimates for omnidirectional images, which are obtained using our novel pipeline. The project is open sourced and available at https://odin-dataset.github.io.

1. Introduction

Challenges relating to the analysis of Activities of Daily Living (ADL) have become essential topics of research in computer vision and active and assisted living [3, 7, 20]. Examples of these challenges include human pose estimation and activity recognition. For the rest of the paper, these will be referred to as human behaviour understanding (HBU) challenges. Most of the research in these fields is done using lateral-view RGB(-D) images as inputs. However, record-

thy solution to these problems. These cameras are generally unobtrusive, have a larger field of view, and can provide largely unoccluded views of the environments being monitored. However, HBU challenges such as pose estimation become all the more challenging due to the viewpoint and due to the heavy distortions introduced by the lens when compared to wide-angle lenses.

The aim of this work is to introduce a new large-scale omnidirectional dataset which contains numerous synchronized modalities. This includes images and videos from cameras of different types recording participants carrying out various activities of daily living, along with their physiological data. ODIN will support research in areas as varied as human pose estimation, activity recognition, person tracking and monitoring, scene understanding, privacy preservation, biometric monitoring, novel view synthesis, generative modelling, 3D scene reconstruction, and image registration. Through our first release, we aim to promote research on 3D human pose estimation using omnidirectional cameras. Research in this area is scarce, arguably due to the difficulty of the problem and the dearth of datasets. For the omnidirectional camera images, the dataset provides associated camera-frame 3D pose estimates. We propose a novel unsupervised pipeline for obtaining these pose estimates in real-life indoor settings while preserving the state of the environment, and also without the use of expensive equipment.

The main contributions of this work are as follows:

- This work introduces a large-scale dataset of omnidirectional images capturing a diverse range of activi-

ODIN - environments and modalities

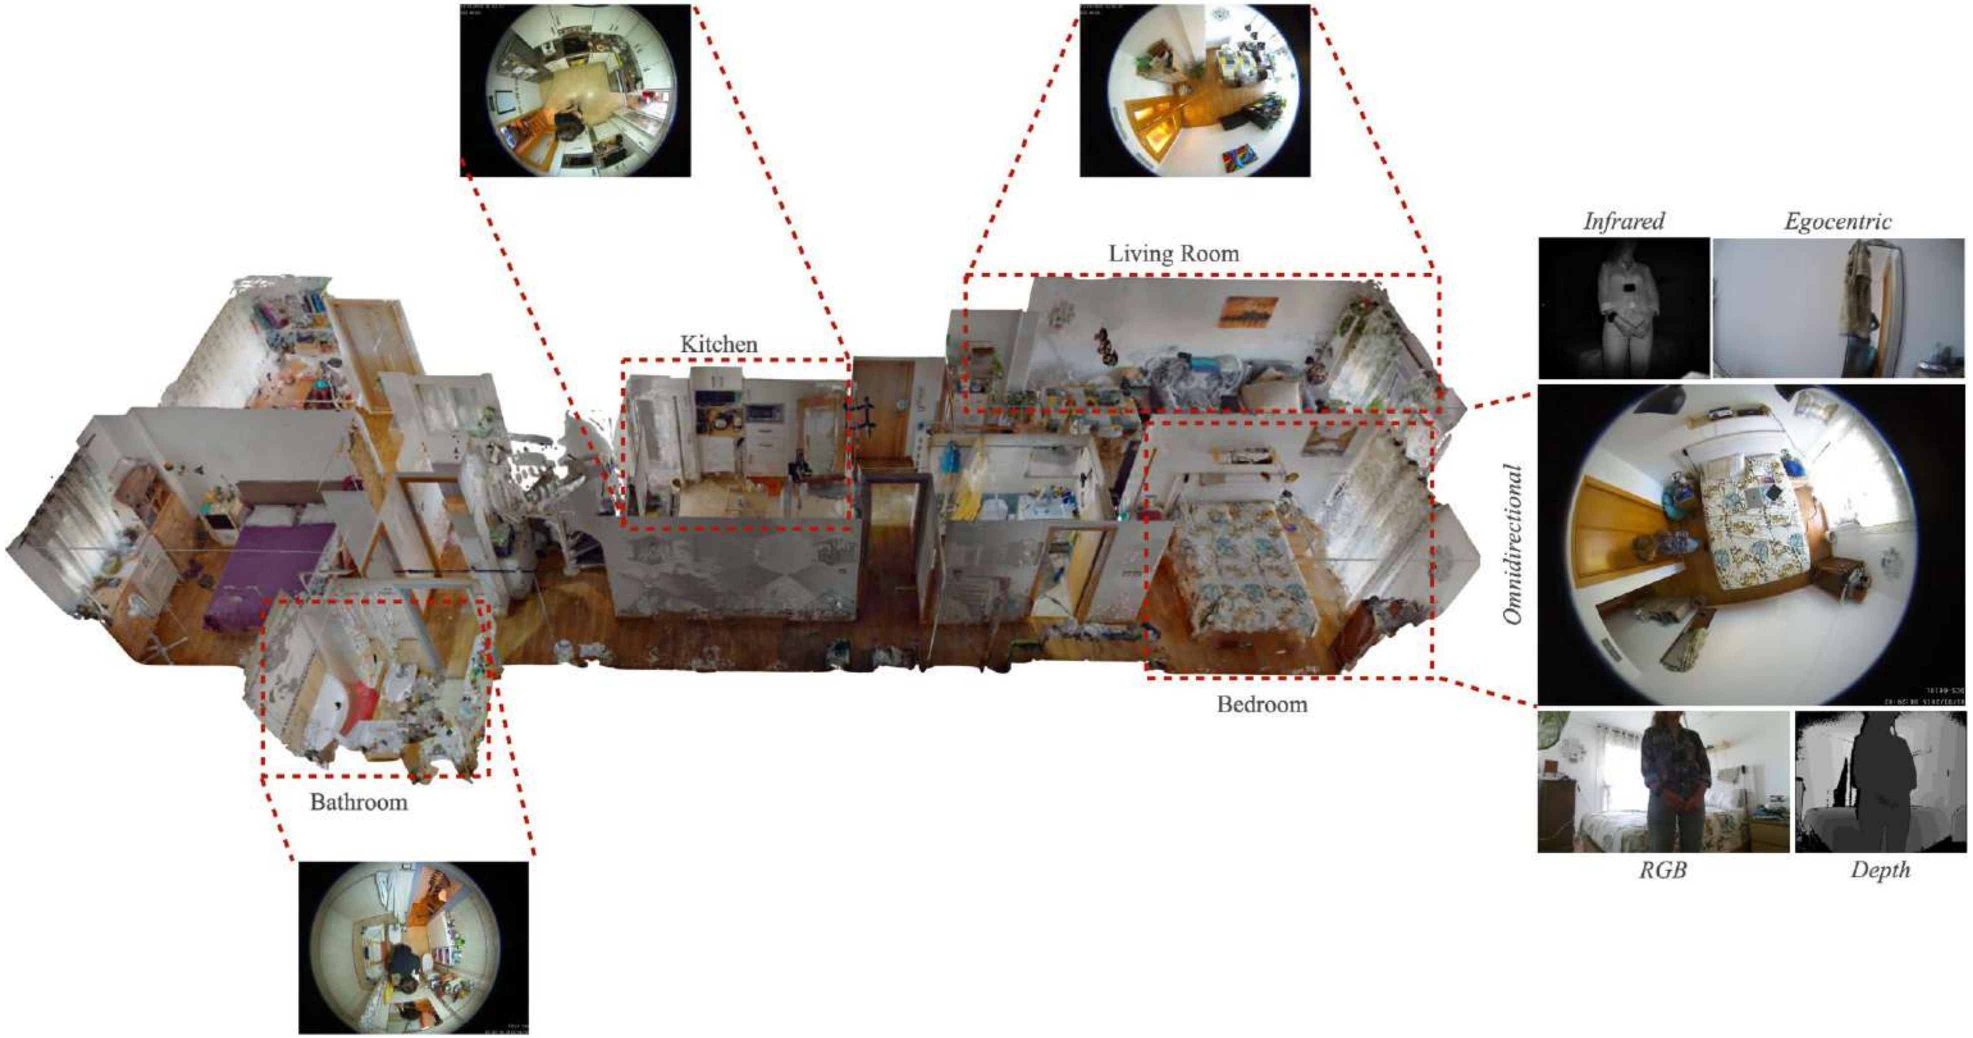

To be open sourced at https://odin-dataset.github.io

Environment layout

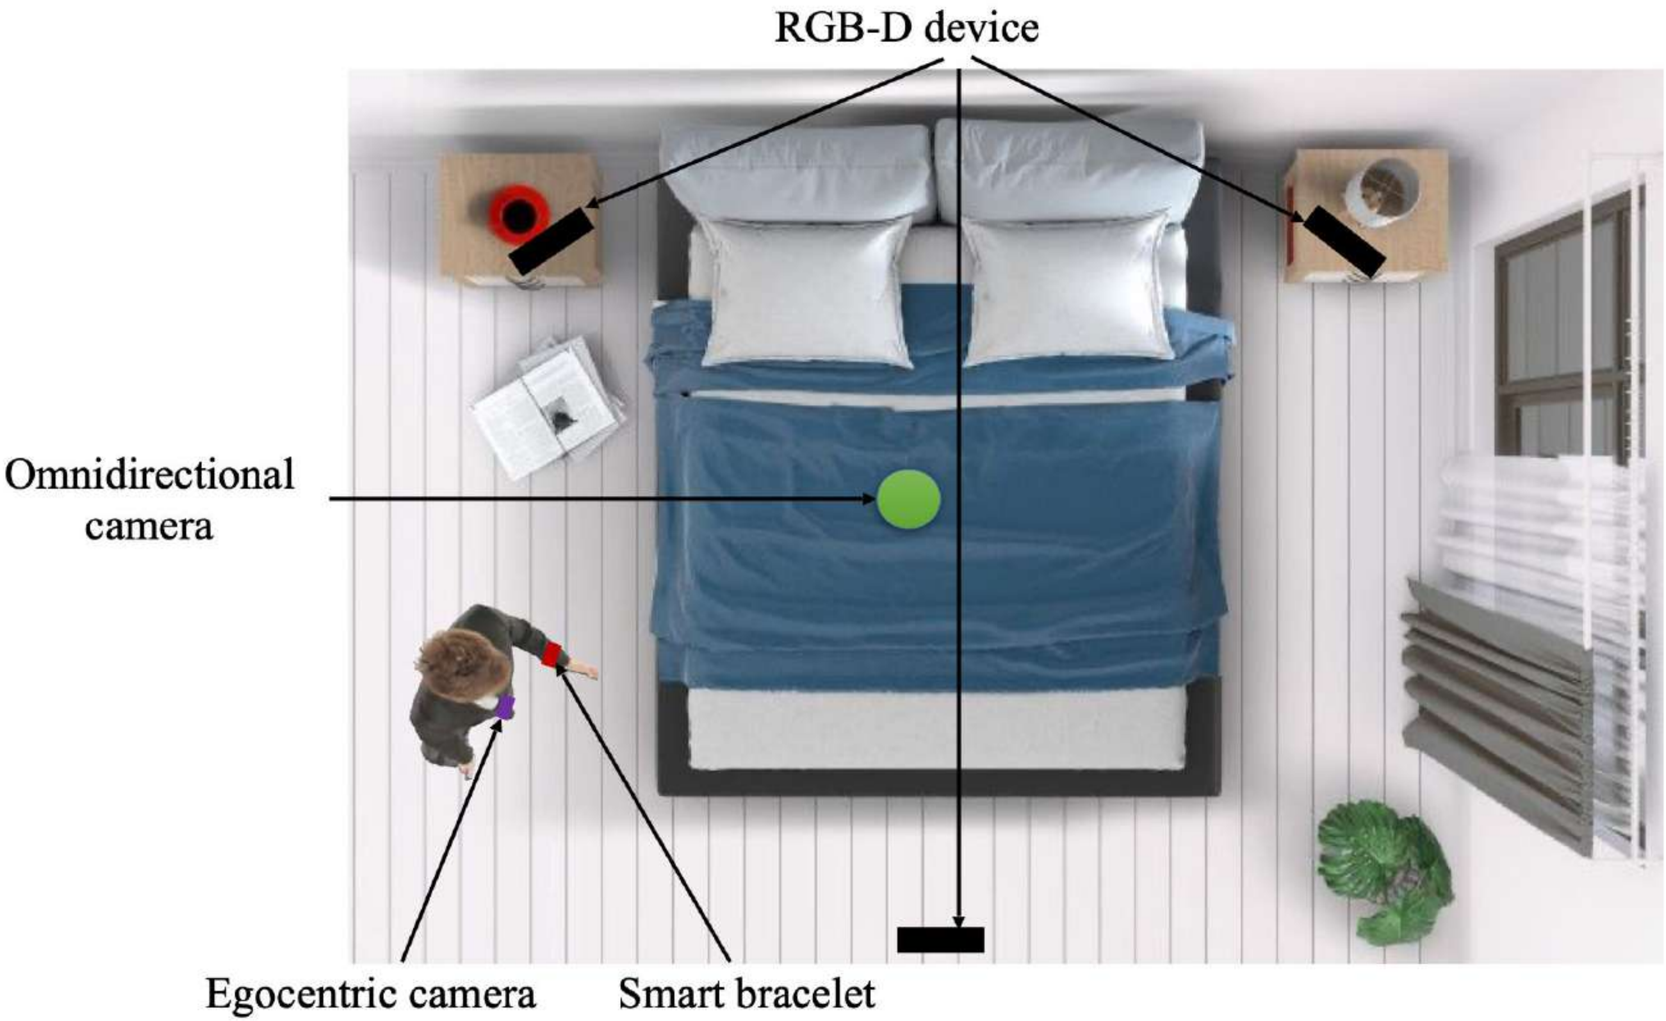

Equipment used

Kinect V2

Xiaomi Mi Action Camera 4K

D-Link DCS-6010L

Empatica E4

Matterport Pro2

A synchronized multi-modal setting

Comparison to related work

Dataset	Omni	Ego	RGB	3D scans	Stereo	IMU	Synced-cam	Phys. signals	Pose	Activity labels	Audio
ODIN	✓	✓	✓	✓	✓	(Partial)	✓	✓	✓	(×)	×
PIROPO Database	✓	×	×	×	×	×	✓	×	×	✓	×
WEPDTOF	✓	×	×	×	×	×	×	×	×	×	×
Fisheye dataset	✓	×	×	×	×	×	×	×	×	×	×
MPII Human Pose	×	×	✓	×	×	×	×	×	✓	×	×
Human3.6M	×	×	✓	×	×	×	✓	×	✓	×	×
Toyota Smarthome	×	×	✓	×	✓	×	✓	×	✓	✓	×
NTU RGB+D Dataset	×	×	✓	×	✓	×	✓	×	✓	✓	×
ADL Dataset	×	×	✓	×	×	×	×	×	×	✓	×
EPIC KITCHENS	×	✓	×	×	×	×	×	×	×	✓	×
Ego4D	×	✓	✓	✓	✓	✓	✓	×	×	✓	✓

Recording methodology

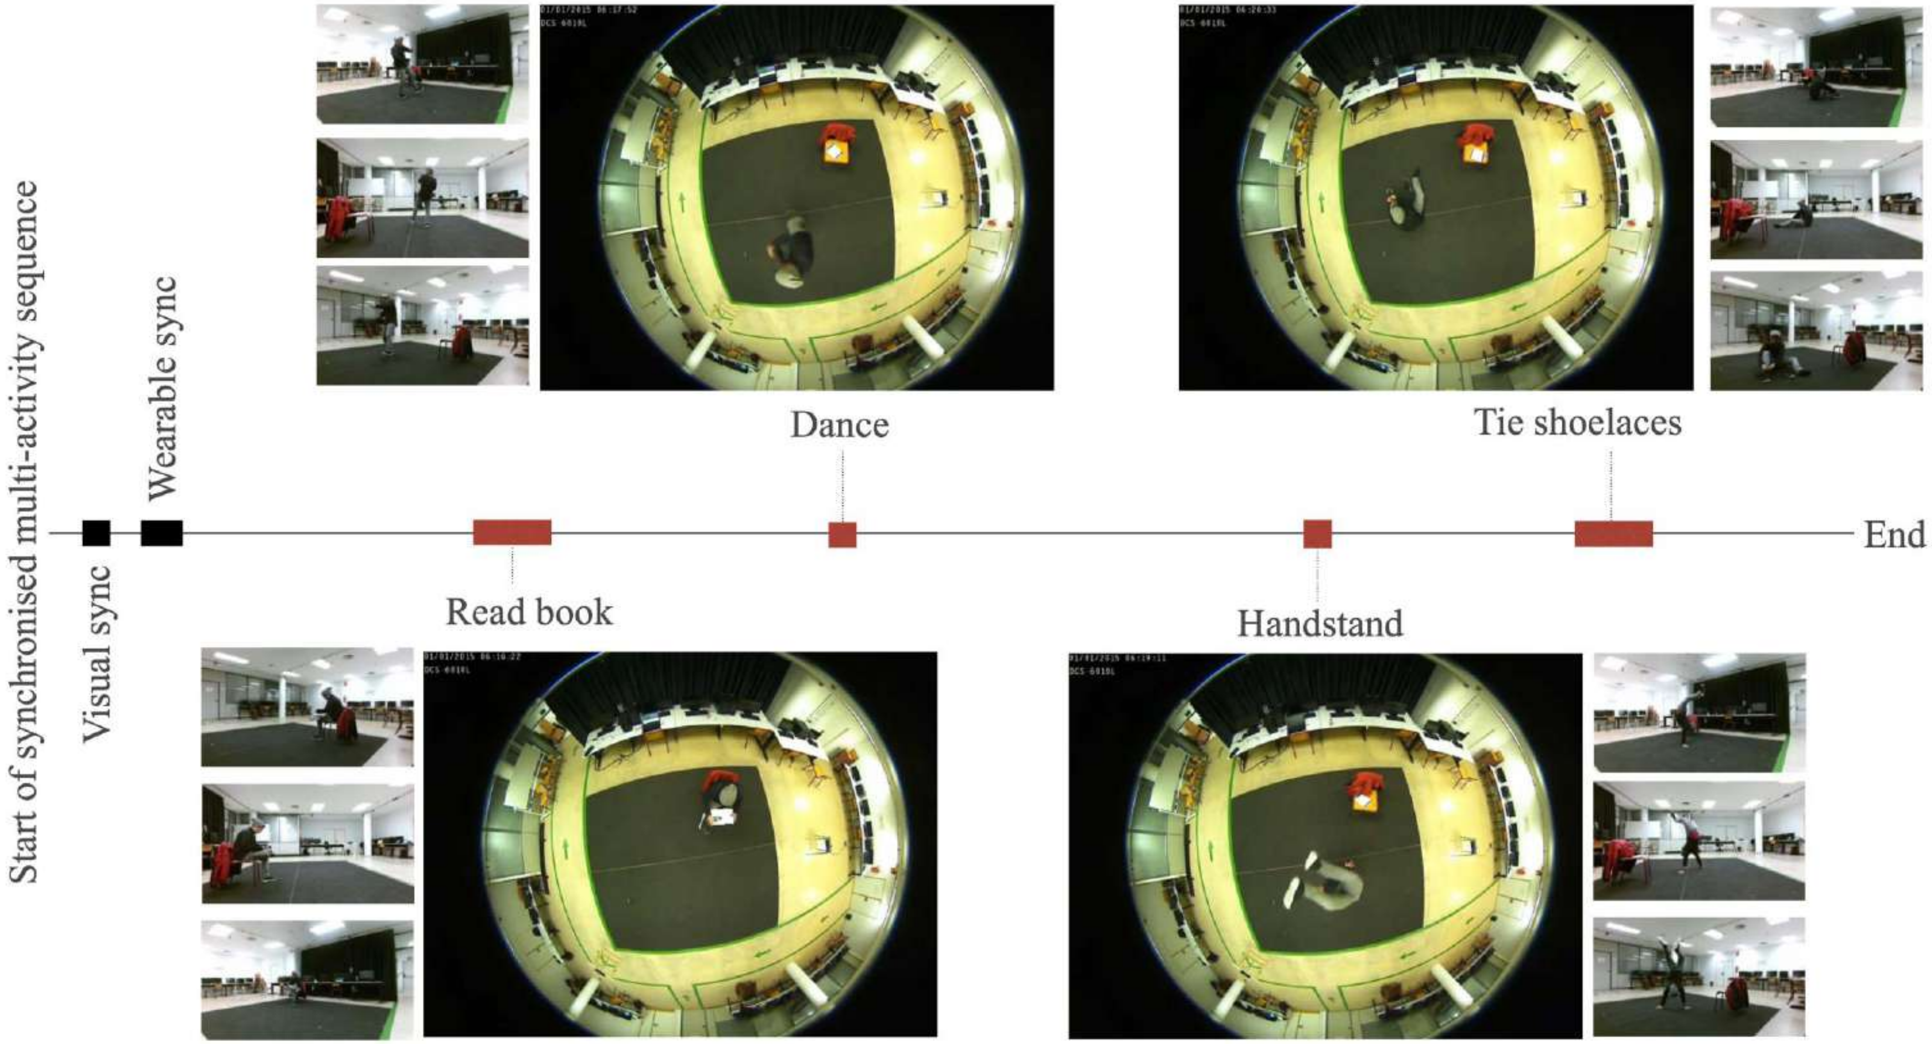

Pose annotation

- Each omnidirectional image is accompanied by a camera-frame full-body 3D human pose estimate.
- The pose estimate is obtained by the perspective projection of pose estimates obtained using lateral-view cameras.
  - We calibrate each static cameras against the omnicam to obtain extrinsics.
- A pose annotation pipeline is created for the purpose.

Pose Estimates

Pose annotation pipeline

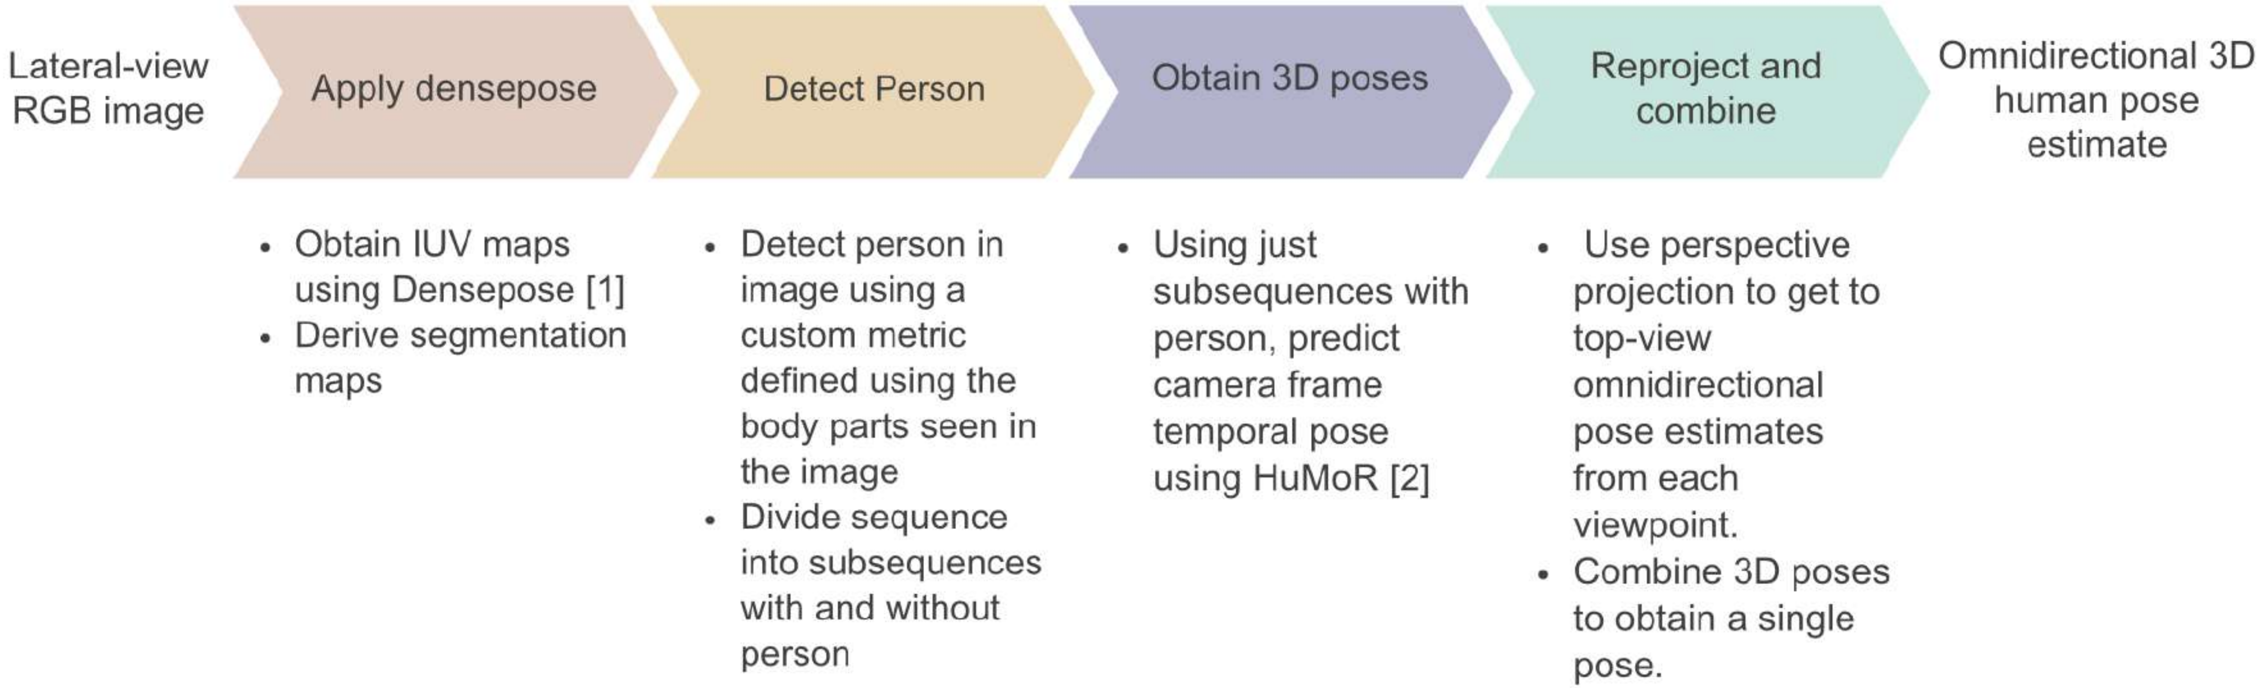

[1] Güler, R. A., Neverova, N., & Kokkinos, I. (2018). Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7297-7306).
[2] Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., & Guibas, L. J. (2021). Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 11488-11499).

Incorrect Pose

Future Work and Collaborations

Future research plans

- 3D pose estimation benchmarks for ODIN. (end of May 2023)
- Creating improved omnidirectional pose estimation models. (Sep 2023)
- Activity recognition benchmark for ODIN with Kooshan Hashemifard to advance privacy by context. (Oct 2023)
- Secondment in Sweden starting very soon -
  - Investigate privacy preservation from a law perspective.

Thank you!

Siddharth Ravi

University of Alicante

siddharth.ravi@ua.es