

# Dynamic Body-Garment Simulation to Characterise Wearable Activity Recognition Performance

Lars Ole HÄUSLER <sup>1,\*</sup>, Lena UHLENBERG <sup>1</sup>, Oliver AMFT <sup>1,2</sup>

<sup>1</sup> Intelligent Embedded Systems Lab, University of Freiburg, Freiburg, Germany;

<sup>2</sup> Hahn-Schickard, Freiburg, Germany

<https://doi.org/10.15221/25.40>

## Abstract

We propose an application of 4D modelling of human body and clothing to estimate the performance of wearable inertial sensors. While inertial sensors, e.g. accelerometers and gyroscopes, can be embedded into garments to capture activities of daily living (ADLs) of their wearer, clothes may move differently and have different orientations than the human body, potentially reducing human activity recognition (HAR) performance. In practice, it is challenging to estimate all error conditions that garments may introduce to wearable sensors, due to their varying body fit and sensor positioning. Thus, empirical evaluations provide limited insight into HAR performance in uncontrolled conditions.

Recent scientific advancements in 4D surface modelling may offer a novel simulation approach to estimate HAR performance for cloth-embedded wearable inertial sensors. Approaches so far have primarily used body surface models to simulate body-attached inertial sensors, and thus did not account for the additional movement dynamics of garments. For example, a smartphone captures different signal patterns for activities, including walking, sitting, or jumping, when placed in a loose-fitting trousers pocket rather than a tight-fitted belt pocket.

The goal of this work is to combine 4D garment and body surface models with inertial sensor models in a joint simulation approach that delivers HAR performance estimations ahead of any physical implementation of the wearable system. We employ textual ADL descriptions as specifications with a generative human motion model to obtain motion patterns. Subsequently, we use 3D Skinned Multi-Person Linear (SMPL) models parametrised for different body sizes to represent full volumetric body and garment motion. We place virtual inertial measurement units (vIMUs) at well-known positions of body and garment models to demonstrate how the effect of garments can be analysed. By simulating vIMUs in selected ADLs, we synthesise acceleration and angular velocity data, which is used to train a HAR model. To evaluate our approach, we generate synthetic inertial sensor data with and without garment simulations for various garment types and body sizes. We then examine the impact on HAR accuracy across specific ADLs in a public dataset, comparing performances between body and garment sensor mounts, as well as the effects of garment type, size, and the performed activity.

Our results show that inertial sensor synthesis is clearly affected by clothing, in particular for loose-fitting garments. We detail HAR performance differences between garment and body-mounts depending on ADLs, body-garment fit, and vIMU positioning. Our approach may offer an alternative to train robust HAR models with synthetic sensor data and deal with clothing-related artefacts.

**Keywords:** 4D body model, Machine learning, SMPL, Wearable sensors

## 1. Introduction

Wearable technologies and Human Activity Recognition (HAR) have become indispensable for continuous health and activity monitoring and provide insights into physiological and behavioural states in daily life [1], [2]. In particular, HAR based on IMUs has several advantages over, e.g., video-based HAR, including protecting privacy and avoiding occlusion [3]. In addition, modern IMU sensors are already integrated in common smartphones and smartwatches and thus are used as a data source for sensor-based HAR tasks [1]. Furthermore, IMU sensors can be integrated into garments [4], which has been an active area of research for years. Garment simulations are useful to investigate different sensor placements, material properties, and motion artefacts in advance without having to build prototypes. With garment simulations, design decisions can be made more quickly and potential problems can be identified at an early stage [5].

\*haeusler@informatik.uni-freiburg.de; <https://ies.cs.uni-freiburg.de/>

While there is a considerable amount of research and datasets for HAR [1], [3], [6], [7], most studies assume that IMU sensors are tightly attached to the body and maintain stable alignment. However, when inertial sensors are integrated into clothing, the fabric can shift, stretch and vibrate independently of the underlying body, which results in discrepancies between actual body movements and sensor measurements. Artefacts caused by clothing can impair HAR performance, but remain difficult to quantify systematically. In addition, synthetic sensor data augmentation for HAR has gained popularity, in particular in combination with deep learning approaches that require extensive training data [8], [9], [10], [11]. So far, sensor data synthesis and other augmentation strategies have neglected cloth-induced errors.

In this paper, we integrate garment simulations in a text-based IMU synthesis approach, which builds on 4D body models and motion synthesis. We investigate the impact of garments on synthesised IMU data, and evaluate HAR performance on public data with smartphone IMU recordings with varying attachment conditions, e.g., devices carried in a loose pocket.

In particular, the work provides the following contributions:

1. We combine 4D garment and body surface models with inertial sensor models in a joint simulation approach for text-based generated movements.
2. We systematically assess the influence of garment fit on synthesised IMU data and analyse how sensor location and activity type affect the IMU synthesis in tight-fitting and loose-fitting clothing compared to body-mounted IMUs.
3. We evaluate HAR performance differences between body-mount vs. tight-fitting and loose-fitting clothing on public HAR data against measured data.

## 2. Related Work

Deep learning architectures have enhanced recognition performance but depend on large training data. High-quality annotated training data is limited because the acquisition requires extensive expert labour. Considerable effort has been invested to utilise alternative data sources and augment model training. Common augmentation techniques include noise augmentation [5], mix-up augmentation [6], and generative adversarial networks [8], which aim to increase the robustness of HAR models. In addition, motion capture data and video recordings are employed for IMU sensor synthesis through 3D motion tracking [8], [10], [12]. Incorporating surface model personalisation also enabled the inclusion of body morphology into the sensor data synthesis [11].

Furthermore, deeper neural network models can be trained with less measured data [11]. Therefore, researchers aim to replace not only measured sensor data but also human movements themselves with synthetic motion. Guo et al. [13] introduced HumanML3D, a large dataset that contains diverse motion sequences represented as joint rotations and positions, paired with textual descriptions of the corresponding activities. By leveraging latent embeddings of the textual input, models trained on the dataset can generate and combine a wide variety of human activities based on natural language processing. HumanML3D was used to train text-based human motion synthesis models, e.g. T2M-GPT [14]. Leng et al. [9], [15] used T2M-GPT and proposed IMUGPT, a framework to synthesise accelerometer data. ChatGPT was leveraged to generate activity text prompts, which were subsequently input into T2M-GPT. Generated joint representations were transformed into acceleration data following the motion trajectories and calibrated using measured data to perform human activity classification on public datasets. Ray et al. [16] used Joints2SMPL [17] to translate joint representations into SMPL models [18] and generated synthetic pressure maps to classify human activities.

In addition to research on data augmentation and motion synthesis, studies highlight the importance of smart garments and the integration of garments into existing simulations, e.g., [19]. Since IMUs are attached to the body surface rather than rigidly mounted, understanding the interactions between sensor placement, body motion, and garment behaviour is essential.

Harms et al. [20] estimated posture recognition performance in smart garments using geometric wrinkle modelling. The results showed that poorer body-garment fit led to decreased classification for shoulder rehabilitation exercise postures. Michael & Howard [21] investigated the differences between sensors mounted on clothing vs. directly on the body. The authors demonstrated that activity recognition can

benefit from cloth-mounted sensors compared to body-mounted placements. Harms et al. [5] presented a modelling and simulation approach to predict sensor performance in smart garments. The authors evaluated internal strain, body-garment distance, and angular deviation. Comparison with results from an experimental study showed that the modelling and simulation approach is suitable for planning sensor placement and garment design prior to implementing garment prototypes.

Patel et al. introduced TailorNet [22], a neural model that predicts 3D clothing deformations while retaining wrinkle details, and accounts for human shape, pose, and garment fit and size. The authors trained the model to predict fabric deformation on SMPL models [18] and showed that TailorNet can run 1000x faster than physics-based simulations while maintaining realistic results. Baronetto et al. [19] used TailorNet to dress 3D human body models with varying body shapes and three differently fitted smart T-shirts. The authors analysed the distance between the sensor and the body surface as well as the sensor displacement (sliding of the sensor) for a smart electrocardiogram T-shirt during the performance of dynamic ADLs. The study showed that the simulation method can be used to assess contact sensor performance for different body shapes, ADLs, and garment designs.

Loose Inertial Poser [23] used TailorNet to simulate a T-shirt on body models from the AMASS dataset and synthesised IMU data. The synthesised data was used to train an autoencoder that estimated human poses for IMUs attached to a jacket. Experimental results demonstrate that the approach maintains posture estimation even in loose-fitting scenarios. However, the study only used tall, thin SMPL preset body models.

### 3. Methods

#### 3.1 Text-based human motion synthesis

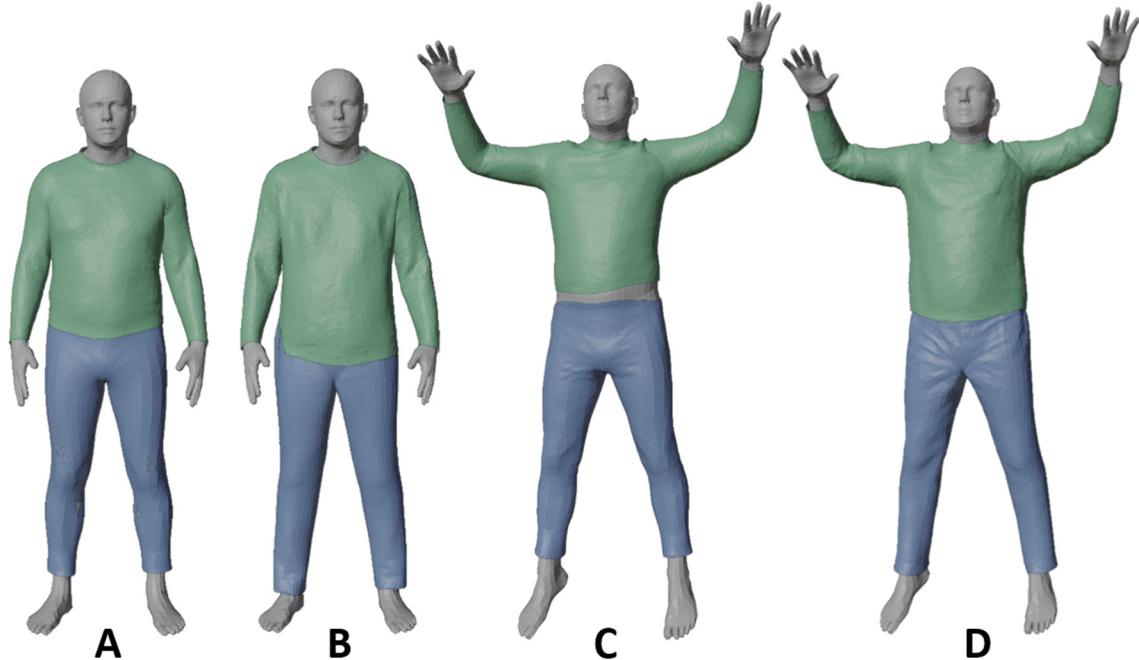
To generate textual descriptions of simulated ADLs, we provided ChatGPT with a general structure, length constraints and a few examples of the desired descriptions. We obtained ten descriptions for each ADL with variations that depend on, e.g., the simulated person's trajectories and speed. As an example, two generated prompts for the running activity are: "a person is slowly running", "a person is running through a corner". Similar to Leng et al. [9], [15] and Ray et al. [16], we used T2M-GPT [14] to generate human motion representations from the obtained descriptions. Motion representations consist of a sequence of 22 3D joint positions per second, with length variations between three and ten seconds.

#### 3.2 Surface model and garment generation

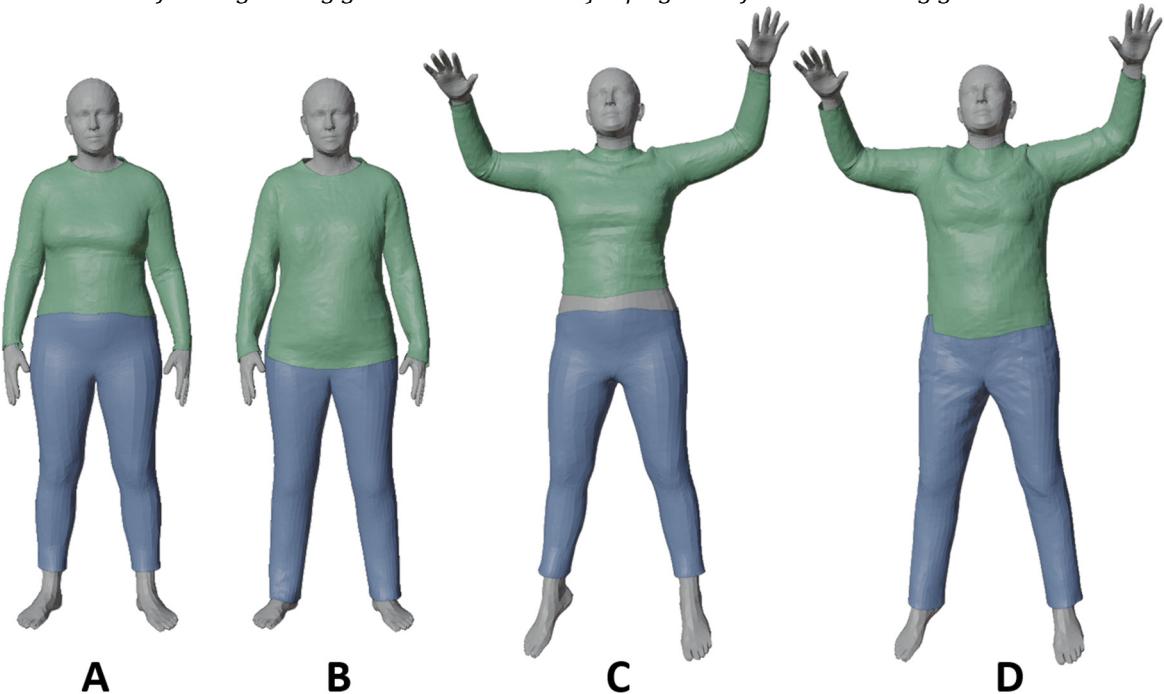
To accurately model 4D human surfaces and garments, we applied generated motions to different SMPL models and dressed them accordingly. We used joints2smpl [17] to obtain the pose parameters  $\theta$  corresponding to the generated motions by T2M-GPT. To simulate different participants performing the prompted activities, we used a custom set of shape parameters  $\beta$  to acquire eight different SMPL models (four female and four male). Body proportions were chosen to present the shortest and tallest models for each gender, one model for each gender with average measurements<sup>1</sup>, and one model for each with average height but an adipose Body Mass Index of 35. Pose parameters  $\theta$  and shape parameters  $\beta$  were used to create triangulated meshes that represent our human surface models, which perform respective activities.

To simulate the garments, we used TailorNet [22] due to scalability and accurate clothing deformations [19], [23]. For each of the eight SMPL models, we applied the male and female versions of a long-sleeve shirt and long pants. We selected the respective TailorNet PCA parameters and generated two sets of garments for each model: one set with a tight fit and one with a loose fit. Fig. 1 and 2 depict examples of a male (see Fig. 1) and female (see Fig. 2) SMPL model with average measurements in two different poses, each with a tight-fitting and loose-fitting set of garments. Due to a lack of relative size scaling with the SMPL models, we excluded certain body models from our garment sets. For the tight-fitting set, we removed the short male and female models, as the clothes still appeared loose on them due to their body height. For the loose-fitting set, we excluded the adipose models, as the largest size parameters still resulted in a tight fit.

<sup>1</sup><https://www.destatis.de/DE/Themen/Gesellschaft-Umwelt/Gesundheit/Gesundheitszustand-Relevantes-Verhalten/Tabellen/liste-koerpermasse.html>



*Fig. 1. Male SMPL model with average measurements fitted with a long-sleeve shirt and long pants.*  
 A: Static pose with tight-fitting garments. B: Static pose with loose-fitting garments. C: Frame of jumping activity with tight-fitting garments. D: Frame of jumping activity with loose-fitting garments.



*Fig. 2. Female SMPL model with average measurements fitted with a long-sleeve shirt and long pants.*  
 A: Static pose with tight-fitting garments. B: Static pose with loose-fitting garments. C: Frame of jumping activity with tight-fitting garments. D: Frame of jumping activity with loose-fitting garments.

### 3.3 IMU synthesis

For IMU data synthesis from the human body of garment models, we selected nine vertices for each sensor location. A selection of the chosen virtual placements is depicted in Fig. 3. We selected the nine vIMUs for each sensor location to include variation in sensor placement and orientation. To obtain acceleration and angular velocity signals from vIMUs, we followed a similar approach as introduced by Uhlenberg et al. [11], [24]. Each vertex of the respective mesh can be described as a position vector  $\vec{k}(t) = (k_x(t), k_y(t), k_z(t))$  for any time  $t$  in a coordinate frame  $G$  in  $R^{3 \times 3}$ . The linear acceleration of

each vertex was calculated as the sum of the dynamic sensor acceleration  $\vec{a}_d(t)$  and the gravitational acceleration  $\vec{a}_g(t)$  exerted on the sensor device:  $\vec{a}(t) = \vec{a}_d(t) + \vec{a}_g(t)$ . The dynamic acceleration  $\vec{a}_d(t)$  was calculated as the second time derivative of the position vector  $\vec{k}(t)$ , and the gravitational acceleration  $\vec{a}_g(t)$  was obtained by multiplying the gravity contribution along each axis with the gravity constant:

$$\vec{a}_d(t) = \frac{d^2\vec{k}(t)}{dt^2}, \quad \vec{a}_g(t) = \vec{k}_j(t) \cdot g$$

with the unit vector  $\vec{k}_j(t)$  describing the gravity contribution along each axis of the sensor and  $g$  the gravity constant  $g = 9.81 \text{ m/s}^2$ . To acquire  $\vec{k}_j(t)$ , the gravity-carrying (“up”) axis  $j$  was determined and multiplied by the inverse of the rotation matrix  $Q_S^G$  to align the gravity vector with the sensor frame  $S$ , according to  $\vec{k}_j(t) = (Q_S^G)^{-1} \cdot \vec{j}$ . To calculate the angular velocity  $\vec{\omega}(t)$  of the simulated sensor, the quaternion orientation estimates  $q_S^G$  of the rotation matrix  $Q_S^G$  were derived over time according to:

$$\vec{\omega}(t) = \frac{q_S^G \vec{k}(t)}{dt}.$$

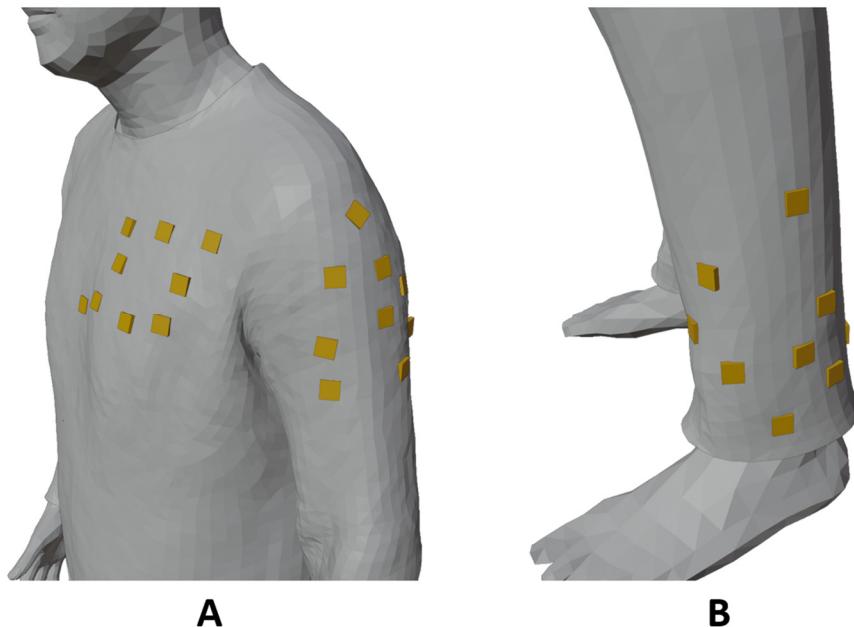


Fig. 3. Visualisation of vIMUs placed on the garments. A: Upper body with all nine vIMU locations of the chest and left upper arm visible. B: Lower body with all nine vIMUs of the left shank visible.

### 3.4 HAR evaluation

As a downstream application of the introduced joint simulation approach, we use the generated synthetic IMU data to train a machine learning model and evaluate it on a real-world HAR dataset. We chose the RealWorld dataset [7] for its detailed documentation of used sensors and their positioning. The vIMU locations (see Section 3.3) were chosen to match descriptions provided in the dataset. The head location was excluded due to a lack of available garments. Consequently, we use six locations for our analysis: left forearm, left upper arm, left shank, left thigh, chest, and waist. In total, we generated three synthetic datasets based on vIMU positioning: One with the vIMUs directly placed on the body of the SMPL models, one placed on the tight-fitting garment set, and one placed on the loose-fitting garments set. The generated activities included walking, running, sitting, lying, stairs down and jumping. We excluded the activities stairs up and standing from our evaluation because T2M-GPT frequently confused them with stairs down and sitting. Therefore, we removed the activities to focus our evaluation on the impact of garments. To improve the model's generalisation to axis rotations and noise, we employed data augmentation with all six possible axis permutations (x, y, z) and added Gaussian noise with a zero mean to the synthetic data, following [25]. We employed the widely used DeepConvLSTM [26] model architecture for our classification. The measured data of the RealWorld dataset was downsampled to 20 Hz to match the synthetic data and used in its entirety as the test set in our evaluation. All data were divided into 2-second segments, corresponding to 40 samples each, using a sliding window with 50% overlap. For classification, we use balanced accuracy of the model predictions as a reliable performance measure on imbalanced datasets.

## 4. Results

### 4.1 Garment-induced IMU error

Fig. 4 illustrates the influence of garment simulation and fit on the synthesis of IMU data for a jumping activity. The linear acceleration signals from the body-mounted and garment-mounted sensors followed similar patterns, with only minor deviations at the signal maxima. The loose-fitting garments produced the largest deviations from the body-mounted synthesised signals. In contrast, the tight-fitting garments exhibited similar signal patterns, with intermediate values between those of the body-mounted and loose-fitting conditions. Angular velocity of loose-fitting garments showed smaller z-axis maxima compared to the body-mounted placement. Moreover, smaller local maxima were observed for both tight-fitting and loose-fitting garments.

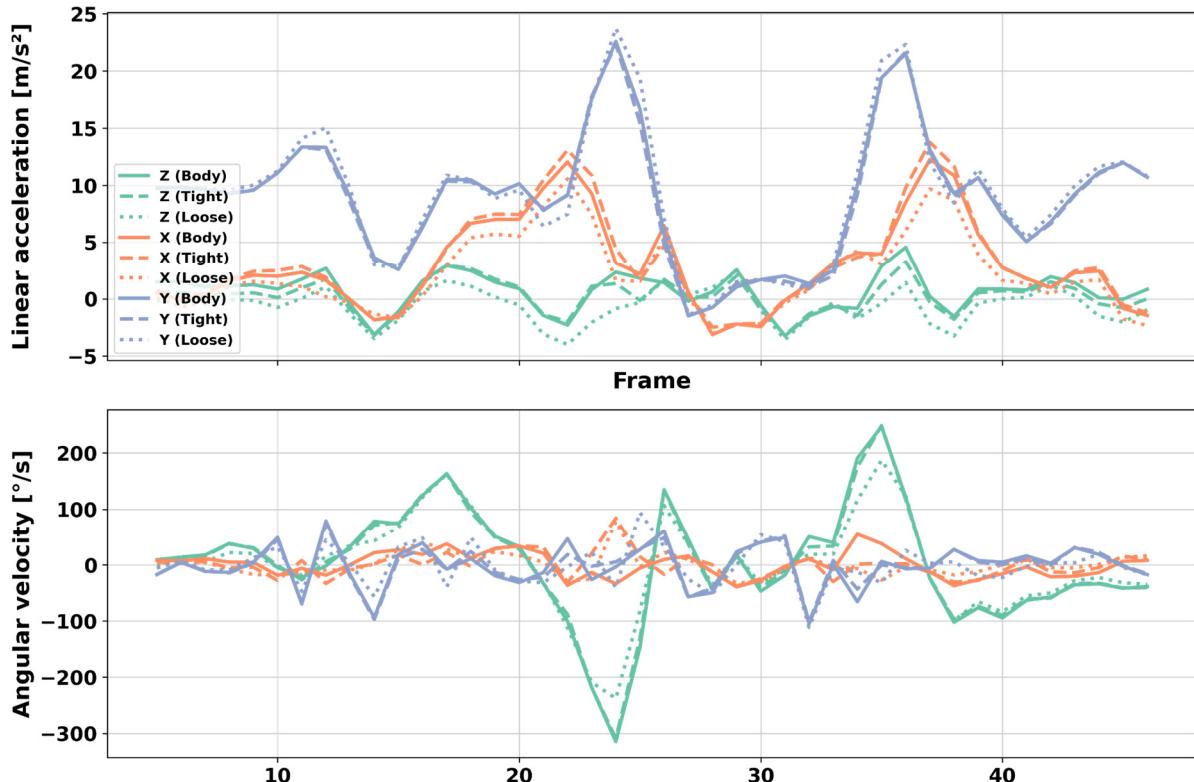


Fig. 4. Comparison of synthesised data for linear acceleration and angular velocity of a vIMU placed at body, tight-, or loose-fitting garments. The timeseries excerpt shows jumping of a female model with average height and weight for a vIMU at the left shank.

As the synthesised angular velocity data exhibits higher deviation and variance relative to the body-mounted signals, we focus on the gyroscope modality for the subsequent analysis.

Fig. 5 depicts the Mean Absolute Error (MEA) of the angular velocity between the tight-fitting garments and the body-mounted vIMUs. The average MAE between both sensor placements resulted in  $12.9^{\circ}/s$ , but was highly dependent on the performed activity and vIMU location. The lowest error was observed in the sitting activity, with an average of  $4.4^{\circ}/s$  and the highest error in the running and jumping activities, with an average MAE of  $29.4^{\circ}/s$  each. Concerning the vIMU location, the error ranges from  $8.3^{\circ}/s$  at the left upper arm to  $17.9^{\circ}/s$  at the left forearm.

Fig. 6 depicts MAE when comparing loose-fitting garments with their corresponding body-mounted locations. For loose-fitting garments, the overall MAE increased to  $19.0^{\circ}/s$ . Similar to tight-fitting garments, the smallest error was observed for sitting (MAE:  $8.2^{\circ}/s$ ), and the highest error for running and jumping (MAE:  $20.3^{\circ}/s$  and  $20.2^{\circ}/s$ , respectively). The error varied depending on vIMU location, ranging from  $11.9^{\circ}/s$  at the left upper arm to  $25.7^{\circ}/s$  at the left forearm.

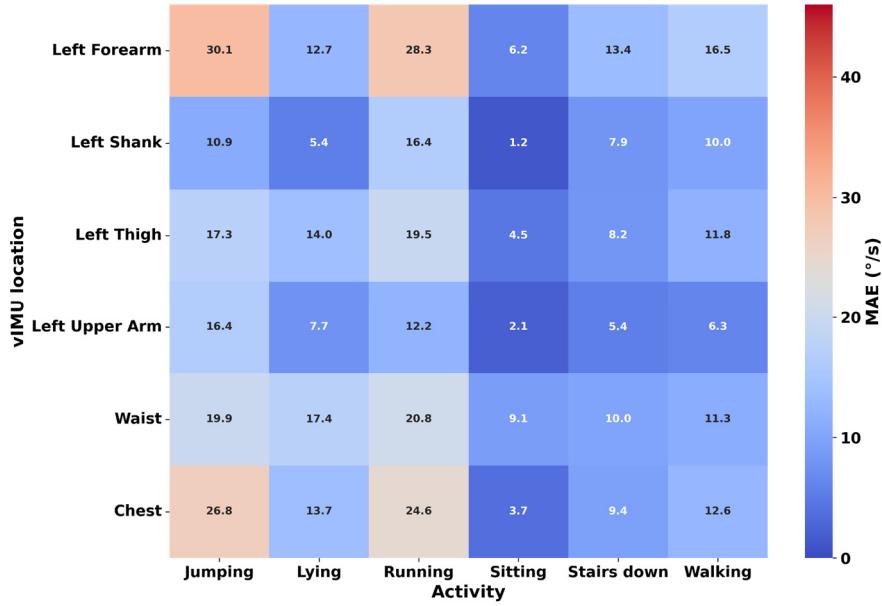


Fig. 5. Heatmap of the MAE of the synthesised angular velocity data for tight-fitting garments in comparison to body-mounted vIMUs for the same motion, depending on the vIMU location and performed activity.

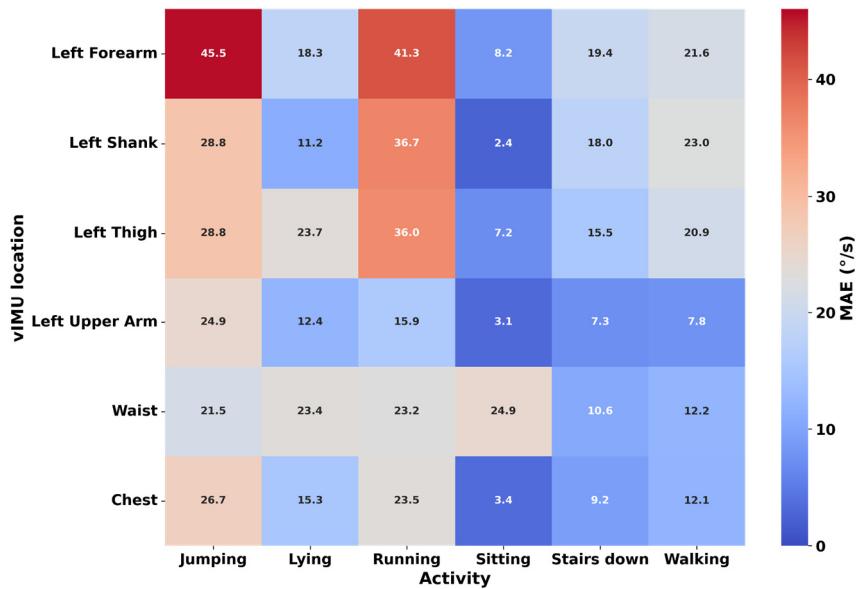


Fig. 6. Heatmap of the MAE of the synthesised angular velocity data for loose-fitting garments in comparison to body-mounted vIMUs for the same motion, depending on the vIMU location and performed activity.

When comparing the error differences between tight-fitting and loose-fitting garments, we observed that the largest discrepancies occurred at location-activity combinations where the error of loose-fitting garments was the highest. The highest difference was observed at the left shank vIMU, with a maximum of 20.3°/s during the running activity. A notable exception was the chest location, where the error remained roughly the same for the tight- and loose-fitting garments.

## 4.2 Downstream HAR Application

Fig. 7 shows the HAR performance for ten runs when synthetic datasets from vIMU placed on body, tight-, or loose-fitting garments were used as training data for a realistic HAR classification task (see Sec. 3.4). The trained DeepConvLSTM models were subsequently evaluated on measured data from the RealWorld [7] dataset. The DeepConvLSTM model trained on synthetic data with the body-fitted vIMUs yielded an average balanced accuracy of 69.3%. The DeepConvLSTM model trained on the vIMUs placed on tight-fitting garments improved balanced accuracy to 76.3 %, which is a 10.1%

increase compared to the body-mounted vIMUs. Lastly, the model trained on loose-fitting garments achieved a similar average balanced accuracy of 75.6 %, which is a 9.1% increase compared to the body-mounted vIMUs. The standard deviation of the performance results from loose-fitting garments increased compared to the body-mounted results, but was reduced in comparison to the ones obtained from tight-fitting garments.

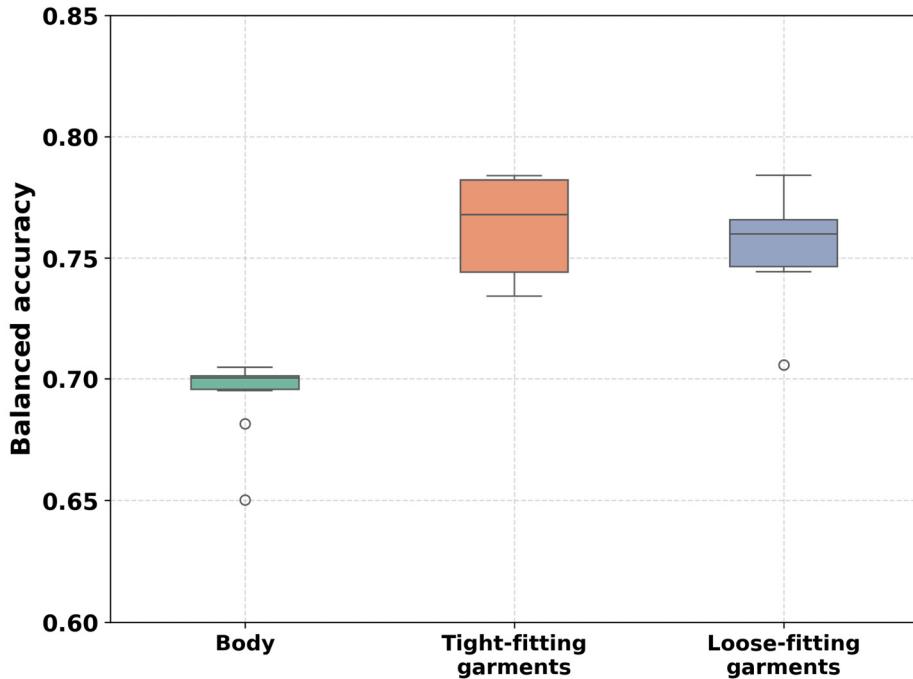


Fig. 7. HAR performance results when training the DeepConvLSTM model on synthetic data of body, tight-, or loose-fitting garments and tested on measured data from the RealWorld [7] dataset on the following activities: Jumping, lying, running, sitting, walking, stairs down.

## 5. Discussion

We showed that garment simulations, in particular for loose-fitting clothes, clearly change IMU data compared to body-mounted sensors. We can observe that, especially for angular velocity, a looser garment fit translates into higher deviations from body-mounted sensors when high velocity movements occur (see Fig. 4). Our observation is consistent with the distribution of MAEs visualised in Fig. 5 and Fig. 6. High-velocity activities with a great range of movements, such as jumping and running, exhibit the highest MAEs. The lowest errors can be found for the sitting activity, where only minimal movement occurs. We attribute the differences to the complexity of garment dynamics. High-velocity movements result in additional effects, e.g., inertial lag and aerodynamic drag, which cause the garments to exhibit secondary motions independent from the body. In contrast, during static or low-velocity activities, the garment movements remained relatively close to the body, which results in a lower MAE compared to high-acceleration activities. In addition to secondary motions, another effect of garments is the occurrence of wrinkles and folds (see Fig. 1 and Fig. 2). Depending on the garment fit and body pose, certain areas of the garments become compressed or stretched, which results in the formation and smoothing out of wrinkles or even folds. When a vIMU was located on a wrinkle, additional signal variation in the angular velocity could be observed, which further increased the overall MAE. An example of the impact of a small wrinkle on the synthesised angular velocity can be observed in Fig. 4, between the 20th and 30th frames, on the X and Y axes.

The results of the error analysis generally agree with expectations and related literature that greater, faster movements cause greater errors between body- and garment-mounted sensors than static poses [19], [27]. We also provide quantitative insights into the errors, depending on the IMU and activities, which would be infeasible to derive from empirical evaluation on a large scale.

As a downstream application of 4D modelling of the human body and garments, we used the synthetic IMU data generated from vIMU placements on the body, tight- and loose-fitting garments each to train

a machine learning model. The trained models were then used to classify activities from a public HAR dataset with measured data. Our results showed that incorporating garment simulations into the IMU data synthesis substantially increased the HAR performance (see Fig. 7). The highest balanced accuracy was achieved when training with vIMUs positioned on loose-fitting garments, with an improvement of 10.1% compared to body-mounted sensors. We attribute the increased performance to the characteristics of the RealWorld dataset [7] that was used as the test set. In the Realworld dataset, IMU data was collected using regular smartphones. The devices were positioned on participants' bodies either with a sports armband or by being placed inside a trouser pocket, shirt pocket, or bra. [7]. In particular, trousers and shirt pocket placements differ from a rigidly attached IMU on the body. Therefore, we believe that our garment simulation generates more realistic IMU data, which better matches the real-world measurements from the dataset. However, the model performance of the loose-fitting garments was 0.9% lower than that of the tight-fitting garments. We believe the difference can be attributed to a combination of two factors: (1) Most sensors in the RealWorld dataset are still placed using sports armband, which are functionally more like a tight-fitting sensor, (2) while garment simulation generally improves realism, excessively loose clothing can introduce a high level of noise from secondary motions and fabric artefacts, which can negatively impact the classification task.

Overall, we showed that our combined simulation provides an alternative method to train robust HAR models that are resilient to garment-related artefacts. The simulation further offers the estimation of performance changes on different garments by simulating specific body types, activities, and sensor placements before manufacturing physical prototypes.

Future work may investigate the differences and impact of garments on HAR performance in more detail and determine the optimal garment fit for improved HAR results based on measured data. Furthermore, the introduced combined 4D modelling approach can be used to generate different synthetic datasets and analyse the impact of garment fit on both the training and testing sets. For instance, future work could investigate whether a model trained exclusively on synthetic body-mounted sensor data can generalise well to data from loose-fitting garments, and vice versa.

Further improvements of the simulation approach could likely be achieved by incorporating more physics-based garment simulations. The physical modelling of forces, e.g., gravity and inertia, in the simulation process would enhance the realism and consistency of garment motions, and thereby generate more accurate synthetic data. Likewise, the properties of the virtual sensor could be included in the garment simulation with a defined size and mass, which impacts, e.g., wrinkle formation at the sensor locations. Another improvement could be achieved by enhancing the human motion generation and extending its training data to increase the collection of correctly synthesizable activities.

## 6. Conclusion

In this work, we introduced a joint simulation approach that combines 4D garment and body surface models with inertial sensor models. From text-based movement descriptions, we simulate diverse body–garment sets across different activities. We systematically analysed the influence of garment fits on the synthesised IMU data. In particular, we assessed how the location of a sensor and the performed activity affect IMU data synthesis in tight-fitting and loose-fitting clothing compared to body-mounted IMUs. In addition, we evaluated the performance differences of a HAR model between body-mounted vs. tight- and loose-fitting garments on public HAR data to predict the effect of clothing on model training ahead of any physical implementation. Our results indicate that garment simulations, particularly in the case of loose-fitting clothing, change synthetic IMU data relative to body-mounted sensors. IMU signal deviations increase as the garment fit becomes looser and activities become faster and involve a greater range of movement, compared to tight-fitting garments and slower activities. Furthermore, in our application example, the inclusion of garment simulations into the IMU data synthesis substantially increased HAR performance. Our combined simulation approach offers an alternative method to train HAR models that remain robust against garment-induced artefacts and estimate classification performances without any measured training data.

## References

- [1] F. Gu, M.-H. Chung, M. Chignell, S. Valaee, B. Zhou, and X. Liu, 'A Survey on Deep Learning for Human Activity Recognition', *ACM Comput. Surv.*, vol. 54, no. 8, p. 177:1-177:34, Oct. 2021, <https://doi.org/10.1145/3472290>.
- [2] J. Seiter, O. Amft, M. Rossi, and G. Tröster, 'Discovery of activity composites using topic models: An analysis of unsupervised methods', *Pervasive Mob. Comput.*, vol. 15, pp. 215--227, Dec. 2014, <https://doi.org/10.1016/j.pmcj.2014.05.007>.
- [3] O. D. Lara and M. A. Labrador, 'A Survey on Human Activity Recognition using Wearable Sensors', *IEEE Commun. Surv. Tutor.*, vol. 15, no. 3, pp. 1192–1209, 2013, <https://doi.org/10.1109/SURV.2012.110112.00192>.
- [4] H. Harms, O. Amft, D. Roggen, and G. Tröster, 'Rapid prototyping of smart garments for activity-aware applications', *J. Ambient Intell. Smart Environ.*, vol. 1, no. 2, Art. no. 2, Apr. 2009, <https://doi.org/10.3233/AIS-2009-0015>.
- [5] H. Harms, O. Amft, and G. Tröster, 'Does loose fitting matter? Predicting sensor performance in smart garments', in *Bodynets 2012: Proceedings of the International Conference on Body Area Networks*, ACM, 2012, pp. 1–4. <https://doi.org/10.1145/2802083.2802093>.
- [6] A. Reiss and D. Stricker, 'Introducing a New Benchmarked Dataset for Activity Monitoring', in *2012 16th International Symposium on Wearable Computers*, Jun. 2012, pp. 108–109. <https://doi.org/10.1109/ISWC.2012.13>.
- [7] T. Sztyler and H. Stuckenschmidt, 'On-body localization of wearable devices: An investigation of position-aware activity recognition', in *2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, Sydney, Australia: IEEE, Mar. 2016, pp. 1–9. <https://doi.org/10.1109/PERCOM.2016.7456521>.
- [8] H. Kwon *et al.*, 'IMUTube: Automatic Extraction of Virtual on-body Accelerometry from Video for Human Activity Recognition', Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., May 2020, <https://doi.org/10.1145/3411841>
- [9] Z. Leng, H. Kwon, and T. Ploetz, 'Generating Virtual On-body Accelerometer Data from Virtual Textual Descriptions for Human Activity Recognition', in *Proceedings of the 2023 ACM International Symposium on Wearable Computers*, in ISWC '23. New York, NY, USA: Association for Computing Machinery, Oct. 2023, pp. 39–43. <https://doi.org/10.1145/3594738.3611361>.
- [10] L. Pei *et al.*, 'MARS: Mixed Virtual and Real Wearable Sensors for Human Activity Recognition With Multidomain Deep Learning Model', *IEEE Internet Things J.*, vol. 8, no. 11, pp. 9383–9396, Jun. 2021, <https://doi.org/10.1109/JIOT.2021.3055859>.
- [11] L. Uhlenberg, L. Ole Haeusler, and O. Amft, 'SynHAR: Augmenting human activity recognition with synthetic inertial sensor data generated from human surface models', *IEEE Access*, vol. 12, pp. 194839–194858, 2024, <https://doi.org/10.1109/ACCESS.2024.3513477>.
- [12] H. Kwon, B. Wang, G. D. Abowd, and T. Plötz, 'Approaching the Real-World: Supporting Activity Recognition Training with Virtual IMU Data', *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 3, Art. no. 3, Sep. 2021, <https://doi.org/10.1145/3478096>.
- [13] C. Guo *et al.*, 'Generating Diverse and Natural 3D Human Motions from Text', in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2022, pp. 5142–5151. <https://doi.org/10.1109/CVPR52688.2022.00509>.
- [14] J. Zhang *et al.*, 'T2M-GPT: Generating Human Motion from Textual Descriptions with Discrete Representations' in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2023, pp. 14730–14740, <https://doi.org/10.1109/CVPR52729.2023.01415>.
- [15] Z. Leng *et al.*, 'IMUGPT 2.0: Language-Based Cross Modality Transfer for Sensor-Based Human Activity Recognition', *Proc ACM Interact Mob Wearable Ubiquitous Technol*, vol. 8, no. 3, p. 112:1-112:32, Sep. 2024, <https://doi.org/10.1145/3678545>.
- [16] L. S. S. Ray, B. Zhou, S. Suh, L. Krupp, V. F. Rey, and P. Lukowicz, 'Text me the data: Generating Ground Pressure Sequence from Textual Descriptions for HAR', presented at the 2024 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops), IEEE Computer Society, Mar. 2024, pp. 461–464. <https://doi.org/10.1109/PerComWorkshops59983.2024.10503379>.

- [17] X. Zuo *et al.*, 'SparseFusion: Dynamic Human Avatar Modeling From Sparse RGBD Images', *IEEE Trans. Multimed.*, vol. 23, pp. 1617–1629, 2021, <https://doi.org/10.1109/TMM.2020.3001506>.
- [18] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, 'SMPL: A Skinned Multi-Person Linear Model', *ACM Trans. Graph.*, vol. 34, no. 6, Art. no. 6, Nov. 2015, <https://doi.org/10.1145/2816795.2818013>.
- [19] A. Baronetto, L. Uhlenberg, D. Wassermann, and O. Amft, 'Simulation of Garment-Embedded Contact Sensor Performance under Motion Dynamics', in *ISWC '21: Proceedings of the ACM International Symposium on Wearable Computers*, in ISWC '21. Virtual Conference: ACM, Sep. 2021, pp. 73–77. <https://doi.org/10.1145/3460421.3480423>.
- [20] H. Harms, O. Amft, and G. Tröster, 'Estimating rehabilitation exercise recognition performance in sensing garments', *IEEE Trans. Inf. Technol. Biomed.*, vol. 14, no. 6, Art. no. 6, 2010, <https://doi.org/10.1109/TITB.2010.2076822>.
- [21] B. Michael and M. Howard, 'Activity recognition with wearable sensors on loose clothing', *PLOS ONE*, vol. 12, no. 10, p. e0184642, Oct. 2017, <https://doi.org/10.1371/journal.pone.0184642>.
- [22] C. Patel, Z. Liao, and G. Pons-Moll, 'TailorNet: Predicting Clothing in 3D as a Function of Human Pose, Shape and Garment Style', presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 7365–7375. <https://doi.org/10.1109/CVPR42600.2020.00739>.
- [23] C. Zuo *et al.*, 'Loose Inertial Poser: Motion Capture with IMU-attached Loose-Wear Jacket', in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, USA: IEEE, Jun. 2024, pp. 2209–2219. <https://doi.org/10.1109/CVPR52733.2024.00215>.
- [24] L. Uhlenberg and O. Amft, 'Comparison of Surface Models and Skeletal Models for Inertial Sensor Data Synthesis', in *BSN '22: Proceedings of the 18th IEEE-EMBS International Conference on Wearable and Implantable Body Sensor Networks*, Ioannina, Greece: IEEE, Sep. 2022. <https://doi.org/10.1109/BSN56160.2022.9928504>.
- [25] X. Zhang *et al.*, 'UniMTS: Unified Pre-training for Motion Time Series', In *Proceedings of the 38th International Conference on Neural Information Processing Systems (NIPS '24)*, Vol. 37. <https://doi.org/10.1145/3411841>.
- [26] F. J. Ordóñez and D. Roggen, 'Deep Convolutional and LSTM Recurrent Neural Networks for Multimodal Wearable Activity Recognition', *Sensors*, vol. 16, no. 1, Art. no. 1, Jan. 2016, <https://doi.org/10.3390/s16010115>.
- [27] H. Harms, O. Amft, and G. Tröster, 'Modeling and simulation of sensor orientation errors in garments', in *Bodynets 2009: Proceedings of the 4th International Conference on Body Area Networks*, ACM, 2009, pp. 1–8. <https://doi.org/10.4108/ICST.BODYNETS2009.5977>.