



Nymeria: A Massive Collection of Multimodal Egocentric Daily Motion in the Wild

Lingni Ma¹(✉), Yuting Ye¹, Fangzhou Hong², Vladimir Guзов³,
Yifeng Jiang⁴, Rowan Postyeni¹, Luis Pesqueira¹, Alexander Gamino¹,
Vijay Baiyya¹, Hyo Jin Kim¹, Kevin Bailey¹, David S. Fosas¹,
C. Karen Liu⁴, Ziwei Liu², Jakob Engel¹, Renzo De Nardi¹,
and Richard Newcombe¹

¹ Meta Reality Labs Research, Redmond, USA

lingni.ma@meta.com

² Nanyang Technological University, Singapore, Singapore

³ University of Tübingen and Max Planck Institute for Informatics,
Tübingen, Germany

⁴ Stanford University, Stanford, USA

<https://www.projectaria.com/datasets/nymeria>

Abstract. We introduce Nymeria - a large-scale, diverse, richly annotated human motion dataset collected in the wild with multiple multimodal egocentric devices. The dataset comes with a) full-body ground-truth motion; b) multiple multimodal egocentric data from Project Aria devices with videos, eye tracking, IMUs and etc.; and c) an third-person perspective by an additional “observer”. All devices are precisely synchronized and localized in one metric 3D world. We derive hierarchical protocol to add in-context language descriptions of human motion, from fine-grain motion narrations, to simplified atomic actions and high-level activity summarization. To the best of our knowledge, Nymeria dataset is the world’s largest human motion in the wild; first of its kind to provide synchronized and localized multi-device multimodal egocentric data; and the world’s largest motion-language dataset. It provides 300 hours of daily activities from 264 participants across 50 locations, total travelling distance over 399 Km. The language descriptions contain 310.5K sentences in 8.64M words from a vocabulary size of 6545. To demonstrate the potential of the dataset we evaluate several SOTA algorithms for egocentric body tracking, motion synthesis, and action recognition.

Keywords: human motion · egocentric · multimodal · dataset

F. Hong, V. Guзов and Y. Jiang—Work done during internships at Meta Reality Labs Research.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-72691-0_25.




Fig. 1. A glimpse of Nymeria dataset. The figure shows example indoor and outdoor activities captured on a campus, where the point clouds and trajectories are the SLAM output by tracking all egocentric devices *i.e.* the glasses and wristbands. Each sub-figure is a motion clip from a different participant, where the top left gives the latest egocentric view, the right is the 3D localized full-body motion synchronized with the headset and the bottom left provides an auxiliary third-person view.

1 Introduction

The advent of AI is leading to a surge of smart glasses [1–3, 5, 6, 9, 11, 30] and other wearables. These devices not only provide seamless access to LLM-based AI assistants, but also are multimodal data-capture vehicles that provide immediate and long-term personalized context, allowing AI assistants to evolve into the next generation of human-centric contextualized AI, and unlock a new era of *contextualized computing* combined with AR/VR technology.

In this paradigm, the wearer’s own body motion and action provides important context. The problem is challenging to solve given insufficient self-observations from wearable devices. In reality, a critical limiting factor to advance research is data. Currently, people either use real data with limited scale, diversity and modality [35, 39, 46, 56, 68, 78, 105, 112] or simulations that lack realism and completeness [12, 13, 45, 47, 59, 98]. There are three key technical challenges to create large motion datasets.

- *Obtaining long-term ground-truth motion in the wild.* There are two main motion capture (mocap) approaches. Vision-based solutions such as the ones relying on optical markers [56, 68, 78] or cameras [35, 55, 62, 100, 112] are adversely affected by line-of-sight visibility, and require a complex multi-camera-setup to cover limited range of motions inside a constrained volume. Inertial-based solutions [7, 10] suffer from dead-reckoning and thus are inferior in global positioning accuracy [69, 88].
- *Multi-device alignment.* Combining multiple capture devices or ground-truth systems requires accurate temporal and spatial alignment, which can be challenging with off-the-shelf hardware that cannot be modified or lack support

for universal synchronization protocols. The existing datasets work around this problem using visual cues [35, 55, 112] or audio [39]. Such approaches offer limited accuracy and reliability. In order to counter clock drift for long recordings, they can be intrusive and interrupt the natural activity. Consequently, existing datasets mostly record short motions (*cf.* Table 2).

- *Data processing and annotations.* These are critical for a dataset to develop its full potential. In addition to the body motion, device localization and scene representation, we believe natural language descriptions are crucial for future research directions. Existing work provides simple descriptions or action labels without scene context [26, 36, 82], and is of significantly smaller scale compared to the text corpus for training LLMs [16, 77, 99].

To fill the gap and accelerate the research, we introduce Nymeria - the world largest human motion dataset with 300 hours in-the-wild daily activities from 264 participants performing 20 scenarios from 50 indoor and outdoor locations. With average 15-min per recording, the data captures natural activities with spontaneous unscripted actions and authentic interactions. Nymeria is first-of-its-kind dataset recorded with multiple multimodal egocentric devices. Participants worn XSens mocap suit [7], Project Aria glasses [30] and Aria-alike wristbands to record egocentric motion, RGB, grayscale, eye tracking (ET) videos, inertial measurement units (IMUs), magnetometer, barometer and etc. Devices are synchronized with a non-intrusive hardware solution with sub-millisecond accuracy and localized into a single metric 3D leveraging Project Aria Machine Perception Service (MPS) [8]. We also developed novel algorithms to retarget XSens skeleton motion into a full parametric human model and correct the global drift with optimization. To connect human motion with natural languages, we developed a coarse-to-fine narration schema to describe in-context human motion at different granularity. With 310.5K sentences and 8.64M words from 6545 vocabulary size, Nymeria stands out as the world’s largest motion-language dataset (Table 1).

2 Related Works

Motion Datasets – Scale, Multimodal, in-the-Wild and Perspectives. Datasets are crucial ingredients in developing algorithms, particularly machine learning approaches. AMASS [68] is a pioneering effort in large motion dataset, which unifies multiple marker-based datasets into SMPL [64]. While AMASS provides diverse human motion, it lacks scene context. Recent works [46, 56, 78] extend the solution to include objects. Monocular [18, 33, 50, 62, 67, 80, 87, 89, 93, 107]

Table 1. Highlight statistics of Nymeria dataset. We capture 1200 sequences of 300-hour daily activity from 264 people performing 20 scenarios at 50 locations with 399 Km traveling distance, 260M body poses, 201.2M images, 11.7B IMUs, 10.8M gazes. The motion narrations contain 310.5K sentences in 8.64M words from 6545 vocabulary.

Seq	Qty	Pts	Sc	Loc	Pose	Img	IMU	Gaze	Traj	Sent	Word	Voc
1200	300h	264	20	50	260M	201M	11.7B	10.8M	399 Km	310.5K	8.64M	6545

and multi-view cameras [35, 48, 49, 55, 112] are common mocap alternatives. For monocular camera, Motion-X [62] stands out as a comprehensive large collection of whole-body motion with facial expressions and hand gestures. For multi-view setting, EgoExo4D [35] stands out as a large dataset of skilled activities. Vision-based algorithms require good line-of-sight. Consequently, they are better suited to record motion with clear body observations bounded by a volume. To capture data in the wild, mocap suit is a popular candidate [22, 39, 52, 53, 56, 58, 69, 100, 105]. To address dead-reckoning for inertial-based tracking, previous works fuse IMUs with vision [69, 100, 109], optimize motion with 3D scenes [39, 58] and limit the range of motion locally [105]. Simulation as a more scalable solution, offers valuable supplement to real data. Existing works leverage gaming engine for character animation [19], render marker-based mocap with virtual characters and scenes [12, 15, 46], simulate motion with VR [13] or by generative algorithms [59] etc. Simulations often struggle to present noise characteristics of real data, result in domain gaps. While many solutions are developed for third-person views, egocentric motion datasets remain sparse, leaving a gap to the recent advance of egocentric perception [6, 30, 34]. Existing works focus on hands with object-interaction [14, 24, 25, 57, 90], lack ground-truth [34] or parametric body motion [35], or is limited in scale, diversity, and modality [12, 39, 55, 58, 78, 102, 103, 112]. Nymeria is designed to fill the gaps with significant delta to existing datasets (*cf.* Table 2).

Table 2. Human motion datasets by releasing date. Columns 2 to 5 show activity in hour (q/h), pose frames in millions (p/M), mean sequence duration in minute (μ /m) and number of participants (pp). We then compare language narrations w.r.t. number of descriptions (tt/K) and vocabulary size (voc). The remaining columns mark following features: parametric model for motion representation (pm), egocentric head-mounted device (hd), third-person perspectives (3p), wristbands (wd), global positioning (gp), outdoor scene (od), gaze (gz), 3D scene representations (sr), multi-people scenarios (mp) and human-human interactions (hh). Note EgoExo4D [35] reports 1422 h by summing per camera recording time, where the total activity is 180 h, with 88.8 h annotated with MSCOCO keypoints. The underline numbers are reported for the full dataset.

[illegible]

Motion with Natural Language. Adding language descriptions to motion leads to unique perspective in motion understanding, especially given the powerful capability of large language models (LLMs). Early datasets [17, 32, 38, 40, 42, 92] offer sparse annotations as action categories or semantic labels [83]. KIT [82] is the first attempt in using complete sentences to describe locomotion. HumanML3D [36] enriches HumanAct12 [38] and AMASS [68] with multiple descriptions per motion clip. Motion-X [62] leverages narration algorithm to obtain large-scale fine-grained descriptions at sequence and frame levels. Compared to available text corpus, motion-language data is rather sparse with brief text on brief motion [22]. Ego4D [34] and EgoExo4D [35] offer large-scale descriptions of atomic actions, however, ground-truth motion is missing or without parametric representations. In our work, we provide in-context coarse-to-fine narration by annotators, with the amount of data a magnitude larger than prior works. Since Closely related to tracking is motion synthesis.

Egocentric Body Tracking, Synthesis and Action Recognition. Tracking one’s own body motion with wearable devices is challenging. Early approaches leverage dense body-worn IMUs [41, 69, 70]. Recent methods reduce IMUs to improve practicality [47, 106, 110, 111], assume IMU-enabled AR/VR devices [44, 45], combine IMU with cameras [39, 102, 103, 109, 117], mobile phones [72] and wristbands [58]. While sparse sensors are more practical, our work centers on constructing an offline dataset to serve “ground truth”. Consequently, we use dense IMUs to ensure high accuracy and correct global drift via optimization. Since full-body motion is ill-posed given insufficient observation from egocentric headset, motion synthesis is often used to produce plausible motion. To condition the generation, research explore sparse motion measures [20, 104, 105], headset motion [60] and eye gaze [116] and text [37, 81, 95]. The success of diffusion models lead to active develop in text-driven motion synthesis [21, 23, 51, 91, 96, 115]. Similarly LLMs inspire novel motion understanding algorithms [31, 43, 65, 114] that tightly couple motion with natural language. Our work constructs a dataset with rich hierarchical narrations to inspire further breakthroughs in the field.

3 Building Nymeria Dataset

3.1 Data Collection Setup

Hardware. Each participant wears a mocap suit, a pair of glasses, two wristbands, and a synchronization device (*cf.* Fig. 2. XSens MVN Link [7] is adopted for mocap, which is a tight-fit body suit wired with 17 inertial trackers and a magnetometer. MVN Link supports on-device recording, making it ideal to collect in-the-wild data. We use Project Aria glasses [30] as a lightweight headset to record multimodal data. The sensor suite includes 1 RGB camera, 2 grayscale peripheral cameras, 2 ET cameras, 2 IMUs, 1 barometer, 1 magnetometer, 7 microphones, 1 thermometer, GNSS, WiFi and BT. We repackaged the electronics and sensors of Project Aria into a new wristband device called *miniAria*, in order to closely resemble current AR/VR headsets and provide data to better




Fig. 2. Capture setup. (a) A full-dressed participant. (b) The set of hardware including Project Aria glasses, two miniAria wristbands and synchronization device. (c) The sensor suite of Project Aria and (d) the miniAria wristband.

constrain body tracking algorithms. The wristband matches Project Aria’s sensing ability, with exclusion of microphones, barometer, magnetometer and ET cameras. The supplementary provides detailed sensor configuration and recording profiles.

Synchronization. Project Aria can record an externally provided time signal to aid synchronization. We further enable MVN Link to accept the same signal. A synchronization device is developed to supply the timestamps for all devices, which can optionally receive time from a wireless server located in radio range (~ 100 m). This facilitates synchronizing multiple devices with sub-millisecond accuracy. The alignment between XSens and Aria is within 1 motion frame *i.e.* 4.2 ms. To capture multiple people simultaneously, we replicate the described setup per participant and leverage a common time server.

Recording Protocols. Data collection is managed by 2–3 onsite operators. In addition to participants, a trained observer wearing Project Aria is present to record participants from third-person perspective. All people interact naturally as per activity requires, contributing to rich dynamics as opposed to staged motion. To complete each recording, participants first perform a brief mocap calibration, then gaze calibration, and finally 15–20 min activity. We collect 4–8 recordings per person, where a bulk of data is captured at family houses.

Scenarios. We define 20 scenarios (*cf.* examples in Fig. 1 and Fig. 3). For indoor activities, scenarios include cooking, working, entertaining, searching objects etc. For outdoor activities, scenarios include hiking, biking, dining, sports etc. To encourage natural interactions and authentic motions, participants are instructed with high-level guidelines *e.g.* “grab food in the cafeteria and eat on the patio”. Operators also prompt in-context actions to increase dynamics.

Privacy Considerations. We follow Project Aria research guideline for responsible innovation. Prior to data collection, consents were obtained from participants and home owners for recording and data usage. The SOTA de-identification algorithm EgoBlur [85] is used to blur faces and license plates for all videos.




Fig. 3. Diverse scenarios by diverse people. We show different participants performing different indoor/outdoor activities at different locations. In each subfigure, we show an egocentric view on the top left, a third-person view on the bottom left, and motion rendering on the right.

3.2 Data Processing

To process data, we first use XSens software to obtain the skeleton motion, and Project Aria MPS [8] for device localization, scene representation and gaze estimation. Then motion is retarget to a parametric human model [4] and registered into the coordinates of Aria devices via optimization.

Full-Body Mocap and Retargetting. We record motion at 240 Hz, following the recommended procedures: 1) carefully measuring body dimensions of participants; 2) performing calibration prior to every recording; and 3) processing XSens with the highest quality with single- or multi-floor specification.

XSens represents the body motion as the global transformation and 3D local joint angles of a template skeleton. The skeleton consists of 23 segments, where each segment matches the measured body dimensions of the subject. In addition, a set of $K = 79$ anatomical landmarks are defined on the skeleton model [74]. Their global positions, $\{\mathbf{p}_i\}_K$, can be computed by evaluating the forward kinematics at each frame. We utilize these landmarks to retarget the body motion onto an anatomically-inspired human model for improved realism and visual validation. Our human model is parameterized by $\{\theta, \phi\}$, where the pose parameters




Fig. 4. Global aligned trajectories and point clouds by locations. We show examples of split-level residential house with gardens, where each contain ≈ 5 h of recording. The left shows the top-down views of accumulated trajectories where red, green and blue indicate the head, the left and right wrist. On the right we sample closed-up views where the clusters emerge from human 3D motion distribution.

θ define the global transformation and local joint angles, and the shape parameters ϕ represent a global body scale and individual bone length. Given a motion of N frames, we solve the following inverse kinematics optimization problem:

$$\arg \min_{\phi, \theta_0 \dots \theta_{T-1}, \mathbf{v}^0 \dots \mathbf{v}^{K-1}} \sum_{t=0}^{N-1} \sum_{i=0}^{K-1} \|T^i(\phi, \theta_t) \mathbf{v}^i - \mathbf{p}_t^i\|^2, \quad (1)$$

where $\mathbf{v}^i$ is the local offset of the i th landmark defined on our model, and T^i is the global transformation of its parent joint. We initialize $\mathbf{v}^i$ by manually placing them on the model. The supplementary provides more details about our human model and motion retargeting.

6DoF Localization and Mapping with Global Alignment. Data recorded at the same location are globally aligned into a single metric 3D world via Project Aria MPS [8], which employs state-of-the-art visual inertial odometry (VIO), SLAM and mapping algorithms [29, 73, 75]. In a nutshell, first SLAM is run for each individual recording independently. Subsequently, the resulting maps are loop-closed and jointly optimized via visual-inertial bundle adjustment. The output are highly accurate 1KHz trajectories (*cf.* Fig. 4 and supplementary) – allowing for example to visualize head- and wrist-motion-clusters respectively.




Fig. 5. End-to-end quality assessment. We uniformly sample a 20 min recording over 1.5 Km moving distance and project skeleton in observer’s camera. The rendering and image aligns well, due to precise tracking and synchronization.

Given the device trajectories, we align the body poses into this same reference coordinates by correlating the dead-reckoned trajectory from XSens. Since the latter accumulates significant drift, there exists no static global transformation to align them. Instead, we assume a constant transformation T_{HD} between the user’s head segment from XSens H and the Aria device D (Aria is firmly held in place with straps and participants are asked to avoid adjusting the glasses during the recording). We then cut the trajectory into a large number of short 4.2 ms segments, and solve the following optimization problem

$$\arg \min_{T_{HD}} \sum_t \left\| \log \left(\left(T_{OH}^t \right)^{-1} T_{OH}^{t+1} \right) \cdot \left(T_{HD} T_{WD}^t \right)^{-1} T_{WD}^{t+1} T_{HD}^{-1} \right)^{-1} \right\|^2, \quad (2)$$

where O is the drifting odometry frame of XSens and W is the world coordinates of the MPS output. The is a HandEye calibration problem with closed-form solution [94]. Note the formulation effectively aligns a large number of local motion clips by comparing the local velocity. In practice, Aria is not completely rigid during recording, resulting in a main source of inaccuracy. Precision can be improved with a rolling window optimization. Figure 5 provides a qualitative end-to-end assessment of our multi-device location, XSens motion registration, and time synchronization.

3.3 In-Context Motion-Language Description

To build the connection between body motion, natural language and activity recognition, we ask human annotators to write textual descriptions of in-context human motion by viewing playback videos of the dataset. The annotators segment the video into clips to write descriptions by answering predefined questions. To give annotators a holistic understanding of the motion, the playback video contains synchronized views of the egocentric video, third-person video, and human motion rendered with 3D scene.

We define three annotation tasks to describe motion coarse to fine, and scale up human efforts in a meaningful way. The finest level is *motion narration* for detailed body posture, *e.g.* motion direction, velocity, interactions and attention. Next, we annotate for *atomic action*. Compared to motion narration, annotators are encouraged to use verbs whenever possible, *e.g.* using “dancing” instead of “swing both arms while rotating the body to the right with legs slightly apart”.




Fig. 6. Hierarchical narration example. (left) motion narration of 5s clip – “C straightens the body to receive a badminton racket from his peer. C extends the right arms towards the bat, with the left arm slightly bent. C takes a step forward with the left foot. C focuses on the racket.” (middle) Atomic action of 5s clip – “C lifts the right leg to kick a soccer ball in front of a building.” (right) Activity summarization of 30s clip – “C bikes on the road with a peer”.




Fig. 7. Distribution of language descriptions. The word cloud visualizes annotations of each task in separation and with all data combined. The tuple (N, X, Y, Z) means N hours of data is described by X sentences, Y words, and Z vocabulary size.

The two tasks are done for <5s clips. The last task *activity summarization* is to give one-sentence summary over 30s activity. Figure 6 provide examples of each annotation task. The figure shows the benefit of providing annotators three synchronized views. While the egocentric view captures closed-up hand-object interactions, the third-person and motion rendering help annotators grasp a holistic understanding of the actions.

3.4 Statistics

Data Statistics. We collected 300 daily activities from 264 participants, which amounts to 1200 sequences with average 15min duration. The accumulated trajectory from all participants is 399 Km for headset and 1053 Km for both wristbands. Figure 8 shows the participant demographics w.r.t. the self-reported ethnicity, age, height and weight. The statistics is split by gender, where 48.5% participants self-identified as female, and 51.4% as male. The dataset captures 47 houses, where 31 are multi-floor. In total, there are 201 rooms and 45 gardens. We also capture three locations from an open campus, including 1 cafeteria with an outdoor patio, 1 office building, and a public parking connected to multiple




Fig. 8. Demographics by female and male. Meaning of symbols: Caucasian (Cau), Hispanic (His), African American (AA), East/South/Southeast Asian (EA/SA/SEA).

biking/hiking trails. Approximately 15% recordings contain outdoor activities. Among all scenarios, the highest occurrences are cooking, searching objects and improvised actions. We provide detailed breakdown in supplementary.

Annotation Statistics. For language descriptions, we annotated 38.6 h motion narration, 207 h atomic action and 196 h activity summarization. The average video segment is 5 s for narration and atomic action and 30 s for summarization. In total, the dataset provides 310.5K sentences 8.64M words from a vocabulary of 6545 distinctive words. Note the average word per sentence is 27.8, which is significant longer than existing motion-language narrations. Figure 7 visualizes the language distribution.

4 Benchmark Tasks and Baselines

4.1 Research Opportunities

With an enormous amount of contextualized human motion, Nymeria dataset provides unprecedented research opportunities. Following, we highlight a few research domains. This is by no means complete, but to inspire novel directions and ideas by scratching the surface of full potential.

Motion Tasks. Nymeria is created to assist human motion understanding, with an emphasis on, but not limited to, egocentric perception. The data supports various topics, *e.g.* full-body tracking, motion synthesis, motion forecasting, path planning, action recognition, human behavior analysis and etc. With a multimodal multi-device dataset, we encourage exploring unique algorithms with novel settings, *e.g.* gaze-conditioned motion prediction, action recognition from headset and wristbands, interaction generation from language etc.

Multimodal Spatial Reasoning and Video Understanding. In addition to body motion, Nymeria also provides extensive egocentric videos with precise localization and synchronization. Nymeria therefore is a great asset for algorithms

Table 3. AvatarPoser [45] train/test with real vs. synthetic poses from Nymeria. The first row reports AvatarPoser train/test on AMASS [68] for reference.

	MPJPE (cm)			Hand PE (cm)	MPJVE (cm/s)
	Mean	Lower	Upper		
AMASS	4.20	8.06	1.88	2.34	28.23
Real	7.97	16.74	3.13	6.25	16.71
Synthetic	7.31	15.97	2.51	3.47	16.63

Table 4. Motion synthesis with EgoEgo [60] and BoDiffusion [20] on Nymeria. BoDiffusion(A) reports results train/test on AMASS [68].

	MPJPE (cm)			FID
	Mean	Lower	Upper	
BoDiffusion(A)	3.63	7.07	1.53	–
BoDiffusion	7.98	15.27	5.28	2.32
EgoEgo	13.22	19.03	10.00	5.14

requiring camera poses as priori, *e.g.* scene reconstruction [54, 71]. In this regard, the rich dynamics raise novel real-world challenges. Nymeria also facilitates video understanding with numerous in-context narration. Image retrieval and relocalization also benefit from Nymeria, since we capture multiple videos per location and align them in global coordinates.

Simulation. Simulation is a promising technique to gather massive data, however, synthetic data is typically seed in real world [15, 46]. By design, Nymeria is naturally useful to drive in-context character animation. Since human motion is a function of the environment, we believe Nymeria also aid simulating 3D scene [108]. Sensor simulation can benefit from our data as well, especially for IMUs, magnetometer, and barometer in combination with motion priors.

4.2 Motion Tasks Baselines

We use three case studies to showcase motion algorithms on Nymeria. The goal is to validate the data and provide future research with common baselines. Due to page limit, details of model training is present in supplementary.

Motion Tracking and Synthesis from Sparse Inputs. In this case study, we take the 1-point/3-point body tracking problem in modern AR/VR, where people wear a headset and optionally controllers in VR, and a pair of glasses and optionally wristbands in AR. The task is to recover wearer’s full-body motion using only headset (1-point) [60, 66] or additional wrist devices (3-point) [20, 28, 44, 45, 105]. The problem is considered mixture of tracking and synthesis, due to insufficient lower-body observation. Existing works simulate sensor inputs from motion datasets, which lack the noise characteristic of real data. Nymeria is the first large dataset with multimodal real data to support model training and evaluation. In particular, we provide raw IMU and device poses, as opposed to filtered [88] acceleration and velocity released in previous datasets [39, 100, 105]. In the first experiment, we adapt the 3-point regression method AvatarPoser [45] to train on Nymeria with two variants – one model trained with “real” device poses from SLAM and the other trained with “synthetic” input mocked-up from body motion same as in the original

work. The evaluation uses the same metric as AvatarPoser, including mean per joint position error (MPJPE), hand position error and mean per joint velocity error (MPJVE). Results are shown in Table 3. As expected, “real” data lead to worse performance due to additional error-sources. However the gap is small, which indicates the quality of device tracking. Compare to the original AvatarPoser trained on AMASS [68], Nymeria yields reasonable but higher MPJPE, mainly from worse lower-body tracking. We hypothesize our data contain harder motion, *e.g.* hiking uneven terrain, taking stairs, playing sports and etc. The second experiment evaluates two diffusion models for motion synthesis, *i.e.* 3-point Bodiffusion [20] and 1-point EgoEgo [60]. In addition to MPJPE, we report the Frechet Inception Distance (FID) which measures the distribution distance between generated motion and real motion, following the same procedure in [38] (*cf.* Table 4). Results trained from our data are comparable with the original works trained on subset of AMASS, where 3-point diffusion yields better performance as expected.

Representing Human Motion Manifold.

Learning the embedding space to model human motion manifold has many benefits, *e.g.* dimension reduction, learning motion priors, motion denoising by projection, motion synthesis by sampling and interpolation etc. Previous attempts mainly rely on AMASS [68] to learn the representation [63, 80, 86, 97], where the data heavily focused on isolated locomotion or professional motion. By capturing rich daily activities of common people interacting with real world, our data distribution can better represent the motion manifold concerning everyday human activities. To take a

stab in this direction, we to train Vector-Quantized Variational Autoencoder (VQ-VAE) [27, 76] for Nymeria motion data, following the previous work [43]. VQ-VAE can be used as a “motion tokenizer” to generate motion with auto-regression [65], and to de-noise motion by projecting the input onto the manifold [86]. An ablation is perform to study the impact of product quantization, codebook size and latent dimension. For evaluation, we adopt the same metrics in [43] and include the VQ-VAE trained on AMASS for comparison. As shown in Table 5, Nymeria motion data can be well tokenized to achieve similar performance as VQ-VAE trained with AMASS data. By leveraging product quantization, increasing the codebook size and decreasing the latent dimension, the reconstruction quality is further improved. The resulting motion tokenizer can therefore be used with LLMs akin to language tokenizer to foster motion understanding [43, 114] as detailed in the next case study.

Table 5. Ablation of motion VQ-VAE trained on Nymeria (metric unit: mm). We compare product quantization (PQ), codebook (CB) size and latent dimensions (Dim). The first row show results of AMASS as reported in [43]. PA stands for Procrustes-aligned and ACC for joint position acceleration.

PQ	CB	Dim	MPJPE	PA-MPJPE	ACC
*	512	—	55.80	40.10	7.50
1	2048	512	51.60	37.55	1.09
2	2048	512	39.63	29.77	0.71
2	4096	512	39.20	29.66	0.82
2	16384	64	34.49	26.83	0.67

Motion and Language. While parametric human motion is useful for machine algorithms, natural language description is a better interface with human. An valuable feature of Nymeria dataset is the high-quality hierarchical narrations. Compared with existing data [36, 62], our narrations are in longer natural sentences with context descriptions of objects and environments. It can be used to learn models for text-driven motion generation and motion-to-text descriptions. More importantly, the contextual descriptions are not only paired with the human motion, but also with videos, point clouds, and other sensory data and annotations. The corroboration of both 2D and 3D environment information with language offers exciting opportunities in grounding language and motion research in the physical world. Leveraging VQ-VAE experiment, we train MotionGPT [43] and TM2T [37] for the motion-to-text task with a subset of 30h motion narrations from Nymeria (*cf.* Table 6) to be directly comparable with previous results, using the same metrics of BERT [113], BLEU [79], CIDEr [101] and ROUGE-L [61]. TM2T performs worse than MotionGPT since it lacks strong language prior with the T5 [84] backbone. Compare to the original works trained with HumanML3D [36] and KIT [82], our results are worse as expected. Given similar hours of data, our narrations are much more complex and diverse. By using the full narration data, we expect the performance to be significantly better.

Table 6. Evaluation of motion-to-text. Models are trained with small subset of Nymeria.

	Bert	Bleu@1	Bleu@4	CIDEr	RougeL
TM2T	11.08	40.11	8.99	20.85	30.70
MotionGPT	14.09	42.22	10.31	37.27	32.33

5 Conclusions and Discussions

We propose Nymeria dataset to accelerate research in egocentric motion understanding. The dataset is the world’s largest collection of human motion in the wild with 300 hours daily activity, 260M body poses of 264 participants across 50 locations. We provide accurate 6DoF tracking, 3D scene points and gaze, with all modalities synchronized and aligned into one metric 3D world. Collectively, the dataset captured 399 Km of travel by the participants for a total of 201.2M egocentric images, 11.7B IMU samples and 10.8M gaze point. The Nymeria dataset also stands out as the world largest motion-language dataset with 310.5K sentences in 8.64M words with 38.6 h of fine-grained motion narration, 207 h atomic actions and 196 h activity summarization.

Limitations. The mocap suit and wristbands lead to unnatural appearance in videos, and restrict certain range of motion. XSens quality is known to be affected by motion calibration and body measurements. Our dataset only covers a portion of daily activities, leaving out common public scenarios.

Social Impact. Understanding egocentric full-body motion is crucial towards contextual AI, however it heavily involves personal data. We make best effort to respect privacy via consent, de-identification, minimum data retention and permissive research license.

Contribution Statements. Lingni Ma led the project, developed the pipeline to construct the dataset, coordinated data collection/processing, and led the baseline evaluations. Yuting Ye developed the solution for human motion retargeting, advised data collection and evaluations. Fangzhou Hong, Vladimir Guзов and Yifeng Jiang, validated the dataset and implemented baselines. Rowan Postyeni supported daily operations, processed XSens motion data and performed quality assessment for XSens motion and narration. Luis Pesqueria served the program manager for data collection and narration. Alexander Gamino was responsible for multi-device tracking. Vijay Baiyya was responsible to Project Aria MPS scaled processing. Hyo Jin Kim supported narration annotations. Kevin Bailey and David Soriano Fosas lead hardware development of miniAria wristband. C. Karen Liu and Ziwei Liu served as technical advisor to data collection, annotation and baseline evaluations. Jakob Engel, Renzo De Nardi and Richard Newcombe were the senior technical and scientific advisors.

Acknowledgements. We gratefully acknowledge the following colleagues for their valuable discussions and technical support. Genesis Mendoza, Jacob Alibadi, Ivan Soeria-Atmadja, Elena Shchetinina, and Atishi Bali worked on data collection. Yusuf Mansour supported gaze estimation on Project Aria. Ahmed Elabbasy, Guru Somasundaram, Omkar Pakhi, and Nikhil Raina supported EgoBlur as the solution to anonymize video and explored bounding box annotation. Evgeniy Oleinik, Maien Hamed, and Mark Schwesinger supported onboarding Nymeria dataset into Project Aria dataset explorer and data release. Melissa Hebra helped with coordinating narration annotation. Edward Miller served as research program manager. Pierre Moulon provided valuable guidance to open source code repository. Tassos Mourikis, Maurizio Monge, David Caruso, Duncan Frost, and Harry Lanaras provided technical support for SLAM. Daniel DeTone, Dan Barnes, Raul Mur Artal, Thomas Whelan, and Austin Kukay provided valuable discussions on annotating semantic bounding box. Julian Nubert adopted the dataset for early dogfooding. Pedro Cancel Rivera, Gustavo Solaira, Yang Lou, and Yuyang Zou provided support from Project Aria program. Svetoslav Kolev provided frequent feedback. Arjang Talattof supported MPS. Gerard Pons-Moll served as senior advisor. Carl Ren and Mingfei Yan served as senior managers.

References

1. Apple Vision Pro. <https://www.apple.com/apple-vision-pro/>
2. HTC VIVE. [vive.com](https://www.vive.com)
3. Magic Leap 2. <https://www.magicleap.com/magic-leap-2>
4. Meta momentum library. <https://github.com/facebookincubator/momentum/>
5. Meta Quest. <https://www.meta.com/quest/>
6. Microsoft HoloLens. <https://learn.microsoft.com/en-us/hololens/>
7. Movella XSens MVN Link motion capture. <https://www.movella.com/products/motion-capture/xsens-mvn-link>
8. Project Aria Machine Perception Services. https://facebookresearch.github.io/projectaria_tools/docs/ARK/mps
9. Ray-Ban Meta smart glasses. <https://www.meta.com/smart-glasses/>
10. Rokoko. <https://www.rokoko.com/>

11. Vuzix smart glasses. <https://www.vuzix.com/pages/smart-glasses>
12. Akada, H., et al.: UnrealEgo: a new dataset for robust egocentric 3D human motion capture. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022. LNCS, vol. 13666, pp. 1–17. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20068-7_1
13. Araujo, J.P., et al.: CIRCLE: capture in rich contextual environments. In: CVPR (2023)
14. Banerjee, P., et al.: Introducing HOT3D: an egocentric dataset for 3D hand and object tracking (2024)
15. Black, M.J., Patel, P., Tesch, J., Yang, J.: BEDLAM: a synthetic dataset of bodies exhibiting detailed lifelike animated motion. In: Proceedings IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
16. Brown, T., et al.: Language models are few-shot learners. In: Advances in Neural Information Processing Systems, vol. 33, pp. 1877–1901 (2020)
17. Cai, Z., et al.: HuMMan: multi-modal 4D human dataset for versatile sensing and modeling. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022. LNCS, vol. 13667, pp. 557–577. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20071-7_33
18. Cai, Z., et al.: SMPLer-X: scaling up expressive human pose and shape estimation (2023)
19. Cai, Z., et al.: Playing for 3D human recovery. arXiv preprint [arXiv:2110.07588](https://arxiv.org/abs/2110.07588) (2021)
20. Castillo, A., et al.: BoDiffusion: diffusing sparse observations for full-body human motion synthesis. In: ICCV (2023)
21. Chen, X., et al.: Executing your commands via motion diffusion in latent space. In: CVPR (2023)
22. Cong, P., et al.: LaserHuman: language-guided scene-aware human motion generation in free environment (2024)
23. Dabral, R., Mughal, M.H., Golyanik, V., Theobalt, C.: MoFusion: a framework for denoising-diffusion-based motion synthesis. In: Computer Vision and Pattern Recognition (CVPR) (2023)
24. Damen, D., et al.: Scaling egocentric vision: the epic-kitchens dataset. In: ECCV (2018)
25. Damen, D., et al.: The epic-kitchens dataset: collection, challenges and baselines. IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI) **43**(11), 4125–4141 (2021)
26. Delmas, G., Weinzaepfel, P., Lucas, T., Moreno-Noguer, F., Rogez, G.: PoseScript: 3D human poses from natural language. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022. LNCS, vol. 13666, pp. 346–362. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20068-7_20
27. Dhariwal, P., Jun, H., Payne, C., Kim, J.W., Radford, A., Sutskever, I.: Jukebox: a generative model for music. arXiv preprint [arXiv:2005.00341](https://arxiv.org/abs/2005.00341) (2020)
28. Du, Y., Kips, R., Pumarola, A., Starke, S., Thabet, A., Sanakoyeu, A.: Avatars grow legs: generating smooth human motion from sparse tracking inputs with diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 481–490 (2023)
29. Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry (2016)
30. Engel, J., et al.: Project aria: a new tool for egocentric multi-modal AI research (2023)
31. Feng, Y., Lin, J., Dwivedi, S.K., Sun, Y., Patel, P., Black, M.J.: ChatPose: chatting about 3D human pose. In: CVPR (2024)

32. Ghorbani, S., et al.: MoVi: a large multi-purpose human motion and video dataset. *PLoS ONE* **16**(6), e0253157 (2021)
33. Goel, S., Pavlakos, G., Rajasegaran, J., Kanazawa, A., Malik, J.: Humans in 4D: reconstructing and tracking humans with transformers. In: *International Conference on Computer Vision (ICCV)* (2023)
34. Grauman, K., et al.: Ego4D: around the world in 3,000 hours of egocentric video. In: *CVPR*, pp. 18995–19012 (2022)
35. Grauman, K., et al.: Ego-Exo4D: understanding skilled human activity from first- and third-person perspectives. In: *CVPR* (2024)
36. Guo, C., et al.: Generating diverse and natural 3D human motions from text. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5152–5161 (2022)
37. Guo, C., Zuo, X., Wang, S., Cheng, L.: TM2T: stochastic and tokenized modeling for the reciprocal generation of 3D human motions and texts. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022. LNCS*, vol. 13695, pp. 580–597. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19833-5_34
38. Guo, C., et al.: Action2Motion: conditioned generation of 3D human motions. In: *Proceedings of the 28th ACM International Conference on Multimedia*, pp. 2021–2029 (2020)
39. Guzun, V., Mir, A., Sattler, T., Pons-Moll, G.: Human pose estimation system (HPS): 3D human pose estimation and self-localization in large scenes from body-mounted sensors. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021)
40. Harvey, F.G., Yurick, M., Nowrouzezahrai, D., Pal, C.: Robust motion in-betweening. *ACM Trans. Graph. (TOG)* **39**(4), 60–1 (2020)
41. Huang, Y., Kaufmann, M., Aksan, E., Black, M.J., Hilliges, O., Pons-Moll, G.: Deep inertial poser: learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Trans. Graph. (TOG)* **37**(6), 1–15 (2018)
42. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6M: large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(7), 1325–1339 (2013)
43. Jiang, B., Chen, X., Liu, W., Yu, J., Yu, G., Chen, T.: MotionGPT: human motion as a foreign language. In: *Advances in Neural Information Processing Systems* (2024)
44. Jiang, J., Strel, P., Meier, M., Fender, A., Holz, C.: EgoPoser: robust real-time ego-body pose estimation in large scenes. *arXiv preprint [arXiv:2308.06493](https://arxiv.org/abs/2308.06493)* (2023)
45. Jiang, J., et al.: AvatarPoser: articulated full-body pose tracking from sparse motion sensing. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022. LNCS*, vol. 13665, pp. 443–460. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20065-6_26
46. Jiang, N., et al.: Scaling up dynamic human-scene interaction modeling. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024)
47. Jiang, Y., Ye, Y., Gopinath, D., Won, J., Winkler, A.W., Liu, C.K.: Transformer inertial poser: real-time human motion reconstruction from sparse IMUs with simultaneous terrain generation. In: *SIGGRAPH Asia 2022 Conference Papers*, pp. 1–9 (2022)
48. Joo, H., et al.: Panoptic studio: a massively multiview system for social motion capture. In: *The IEEE International Conference on Computer Vision (ICCV)* (2015)

49. Joo, H., et al.: Panoptic studio: a massively multiview system for social interaction capture. *IEEE Trans. Pattern Anal. Mach. Intell.* (2017)
50. Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3D human dynamics from video. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5614–5623 (2019)
51. Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: Guided motion diffusion for controllable human motion synthesis. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2151–2162 (2023)
52. Kaufmann, M., et al.: EMDb: the electromagnetic database of global 3D human pose and shape in the wild. In: *International Conference on Computer Vision (ICCV)* (2023)
53. Kaufmann, M., et al.: EM-pose: 3D human pose estimation from sparse electromagnetic trackers. In: *The IEEE International Conference on Computer Vision (ICCV)* (2021)
54. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4) (2023). <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
55. Khirodkar, R., Bansal, A., Ma, L., Newcombe, R., Vo, M., Kitani, K.: EgoHumans: an egocentric 3D multi-human benchmark. In: *ICCV* (2023)
56. Kim, J., Kim, J., Na, J., Joo, H.: ParaHome: parameterizing everyday home activities towards 3D generative modeling of human-object interactions (2024)
57. Kwon, T., Tekin, B., Stühmer, J., Bogo, F., Pollefeys, M.: H2o: two hands manipulating objects for first person interaction recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 10138–10148 (2021)
58. Lee, J., Joo, H.: Mocap everyone everywhere: lightweight motion capture with smartwatches and a head-mounted camera. *arXiv preprint [arXiv:2401.00847](https://arxiv.org/abs/2401.00847)* (2024)
59. Li, G., Zhao, et al.: EgoGen: an egocentric synthetic data generator (2024)
60. Li, J., Liu, K., Wu, J.: Ego-body pose estimation via ego-head pose estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17142–17151 (2023)
61. Lin, C.Y.: ROUGE: a package for automatic evaluation of summaries. In: *Text Summarization Branches Out*, pp. 74–81 (2004)
62. Lin, J., et al.: Motion-X: a large-scale 3D expressive whole-body human motion dataset. In: *Advances in Neural Information Processing Systems* (2023)
63. Ling, H.Y., Zinno, F., Cheng, G., van de Panne, M.: Character controllers using motion VAEs. *ACM Trans. Graph.* **39**(4), 1–40 (2020)
64. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)* **34**(6), 248:1–248:16 (2015)
65. Lucas, T., Baradel, F., Weinzaepfel, P., Rogez, G.: PoseGPT: quantization-based 3D human motion generation and forecasting. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022. LNCS*, vol. 13666, pp. 417–435. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20068-7_24
66. Luo, Z., Hachiuma, R., Yuan, Y., Kitani, K.: Dynamics-regulated kinematic policy for egocentric pose estimation. In: *Neural Information Processing Systems* (2021)
67. Luvizon, D., Habermann, M., Golyanik, V., Kortylewski, A., Theobalt, C.: Scene-aware 3D multi-human motion capture from a single camera. *Comput. Graph. Forum* **42**(2), 371–383 (2023)

68. Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: AMASS: archive of motion capture as surface shapes. In: *International Conference on Computer Vision*, pp. 5442–5451 (2019)
69. von Marcard, T., Henschel, R., Black, M.J., Rosenhahn, B., Pons-Moll, G.: Recovering accurate 3D human pose in the wild using IMUs and a moving camera. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2018)
70. Marcard, T., Rosenhahn, B., Black, M., Pons-Moll, G.: Sparse inertial poser: automatic 3D human pose estimation from sparse IMUs. *Comput. Graph. Forum* **36**(2) (2017). *Proceedings of the 38th Annual Conference of the European Association for Computer Graphics (Eurographics)* (2017)
71. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: representing scenes as neural radiance fields for view synthesis. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) *ECCV 2020. LNCS*, vol. 12346, pp. 405–421. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58452-8_24
72. Mollyn, V., Arakawa, R., Goel, M., Harrison, C., Ahuja, K.: IMUPoser: full-body pose estimation using IMUs in phones, watches, and earbuds. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI 2023*. Association for Computing Machinery, New York (2023)
73. Mourikis, A.I., Roumeliotis, S.I.: A multi-state constraint Kalman filter for vision-aided inertial navigation. In: *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pp. 3565–3572. IEEE (2007)
74. Movella: MVN user manual. https://www.movella.com/hubfs/MVN_User_Manual.pdf
75. Mur-Artal, Raúl, M.J.M.M., Tardós, J.D.: ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE Trans. Robot.* **31**(5), 1147–1163 (2015)
76. van den Oord, A., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems* (2017)
77. OpenAI: Achiam, J., et al.: GPT-4 technical report (2023)
78. Pan, X., et al.: Aria digital twin: a new benchmark dataset for egocentric 3D machine perception (2023)
79. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318 (2002)
80. Pavlakos, G., et al.: Expressive body capture: 3D hands, face, and body from a single image. In: *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10975–10985 (2019)
81. Petrovich, M., Black, M.J., Varol, G.: TEMOS: generating diverse human motions from textual descriptions. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022. LNCS*, vol. 13682, pp. 480–497. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20047-2_28
82. Plappert, M., Mandery, C., Asfour, T.: The KIT motion-language dataset. *Big Data* **4**(4), 236–252 (2016)
83. Punnakkal, A.R., Chandrasekaran, A., Athanasiou, N., Quiros-Ramirez, A., Black, M.J.: BABEL: bodies, action and behavior with English labels. In: *Proceedings IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 722–731 (2021)
84. Raffel, C., et al.: Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* **21**(1), 5485–5551 (2020)

85. Raina, N., et al.: EgoBlur: responsible innovation in aria. ArXiv abs/2308.13093 (2023)
86. Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., Guibas, L.J.: Humor: 3D human motion model for robust pose estimation. In: International Conference on Computer Vision (ICCV) (2021)
87. Alp Gueler, R., Natalia Neverova, I.K.: Densepose: dense human pose estimation in the wild. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
88. Roetenberg, D., Luinge, H., Slycke, P.: Xsens MVN: full 6DOF human motion tracking using miniature inertial sensors. Xsens Motion Technol. BV Technical report **3** (2009)
89. Rong, Y., Shiratori, T., Joo, H.: FrankMocap: a monocular 3D whole-body pose estimation system via regression and integration. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1749–1759 (2021)
90. Sener, F., et al.: Assembly101: a large-scale multi-view video dataset for understanding procedural activities. In: CVPR (2022)
91. Shafir, Y., Tevet, G., Kapon, R., Bermano, A.H.: Human motion diffusion as a generative prior. In: ICLR (2023)
92. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: NTU RGB+D: a large scale dataset for 3D human activity analysis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1010–1019 (2016)
93. Shimada, S., Golyanik, V., Xu, W., Theobalt, C.: PhysCap: physically plausible monocular 3d motion capture in real time. ACM Trans. Graph. (ToG) **39**(6), 1–16 (2020)
94. Sorkine-Hornung, O., Rabinovich, M.: Least-squares rigid motion using SVD. Computing **1**(1), 1–5 (2017)
95. Tevet, G., Gordon, B., Hertz, A., Bermano, A.H., Cohen-Or, D.: MotionCLIP: exposing human motion generation to CLIP Space. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022, Part XXII. LNCS, vol. 13682, pp. 358–374. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20047-2_21
96. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-or, D., Bermano, A.H.: Human motion diffusion model. In: ICLR (2023)
97. Tiwari, G., Antic, D., Lenssen, J.E., Sarafianos, N., Tung, T., Pons-Moll, G.: PoseNDF: modeling human pose manifolds with neural distance fields. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022. LNCS, vol. 13665, pp. 572–589. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20065-6_33
98. Tome, D., et al.: SelfPose: 3D egocentric pose estimation from a headset mounted camera. IEEE Trans. Pattern Anal. Mach. Intell. **45**(6), 6794–6806 (2020)
99. Touvron, H., et al.: LLaMA: open and efficient foundation language models (2023)
100. Trumble, M., Gilbert, A., Malleson, C., Hilton, A., Collomosse, J.: Total capture: 3D human pose estimation fusing video and inertial sensors. In: Proceedings of 28th British Machine Vision Conference, pp. 1–13 (2017)
101. Vedantam, R., Lawrence Zitnick, C., Parikh, D.: CiDER: consensus-based image description evaluation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4566–4575 (2015)
102. Wang, J., Liu, L., Xu, W., Sarkar, K., Luvizon, D., Theobalt, C.: Estimating egocentric 3D human pose in the wild with external weak supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 13157–13166 (2022)

103. Wang, J., Luvizon, D., Xu, W., Liu, L., Sarkar, K., Theobalt, C.: Scene-aware ego-centric 3D human pose estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13031–13040 (2023)
104. Wouwe, T., Lee, S., Falisse, A., Delp, S., Liu, C.: Diffusion inertial poser: human motion reconstruction from arbitrary sparse IMU configurations. In: *CVPR* (2024)
105. Yang, D., Kang, J., Ma, L., Greer, J., Ye, Y., Lee, S.H.: DivaTrack: diverse bodies and motions from acceleration-enhanced three-point trackers. In: *EuroGraphics* (2024)
106. Yang, D., Kim, D., Lee, S.H.: LoBSTR: real-time lower-body pose prediction from sparse upper-body tracking signals. In: *Computer Graphics Forum*, vol. 40, pp. 265–275. Wiley Online Library (2021)
107. Ye, V., Pavlakos, G., Malik, J., Kanazawa, A.: Decoupling human and camera motion from videos in the wild. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2023)
108. Yi, H., Huang, C.H.P., Tripathi, S., Hering, L., Thies, J., Black, M.J.: MIME: human-aware 3D scene generation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12965–12976, June 2023
109. Yi, X., et al.: EgoLocate: real-time motion capture, localization, and mapping with sparse body-mounted sensors. *ACM Trans. Graph. (TOG)* **42**(4), 1–17 (2023)
110. Yi, X., et al.: Physical inertial poser (PIP): physics-aware real-time human motion tracking from sparse inertial sensors. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13167–13178 (2022)
111. Yi, X., Zhou, Y., Xu, F.: Transpose: real-time 3D human translation and pose estimation with six inertial sensors. *ACM Trans. Graph. (TOG)* **40**(4), 1–13 (2021)
112. Zhang, S., et al.: EgoBody: human body shape and motion of interacting people from head-mounted devices. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022. LNCS*, vol. 13666, pp. 180–200. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20068-7_11
113. Zhang, T., Kishore, V., Wu, F., Weinberger, K.Q., Artzi, Y.: BERTScore: evaluating text generation with bert. *arXiv preprint arXiv:1904.09675* (2019)
114. Zhang, Y., et al.: MotionGPT: finetuned LLMs are general-purpose motion generators. *arXiv preprint arXiv:2306.10900* (2023)
115. Zhang, Z., Liu, R., Aberman, K., Hanocka, R.: TEDi: temporally-entangled diffusion for long-term motion synthesis (2023)
116. Zheng, Y., et al.: GIMO: gaze-informed human motion prediction in context. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022. LNCS*, vol. 13673, pp. 676–694. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19778-9_39
117. Zheng, Z., et al.: HybridFusion: real-time performance capture using a single depth sensor and sparse IMUs. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *ECCV 2018. LNCS*, vol. 11213, pp. 389–406. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01240-3_24