

Transformer Inertial Poser: Real-time Human Motion Reconstruction from Sparse IMUs with Simultaneous Terrain Generation

Yifeng Jiang
Stanford University
United States of America
yifengj@stanford.edu

Jungdam Won
Meta Reality Labs Research
United States of America
jungdam@fb.com

Yuting Ye
Meta Reality Labs Research
United States of America
yuting.ye@fb.com

Alexander W. Winkler
Meta Reality Labs Research
United States of America
winklera@fb.com

Deepak Gopinath
Meta Reality Labs Research
United States of America
dgopinath@fb.com

C. Karen Liu
Stanford University
United States of America
karenliu@cs.stanford.edu

ABSTRACT

Real-time human motion reconstruction from a sparse set of (e.g. six) wearable IMUs provides a non-intrusive and economic approach to motion capture. Without the ability to acquire position information directly from IMUs, recent works took data-driven approaches that utilize large human motion datasets to tackle this under-determined problem. Still, challenges remain such as temporal consistency, drifting of global and joint motions, and diverse coverage of motion types on various terrains. We propose a novel method to simultaneously estimate full-body motion and generate plausible visited terrain from only six IMU sensors in real-time. Our method incorporates 1. a conditional Transformer decoder model giving consistent predictions by explicitly reasoning prediction history, 2. a simple yet general learning target named "stationary body points" (SBPs) which can be stably predicted by the Transformer model and utilized by analytical routines to correct joint and global drifting, and 3. an algorithm to generate regularized terrain height maps from noisy SBP predictions which can in turn correct noisy global motion estimation. We evaluate our framework extensively on synthesized and real IMU data, and with real-time live demos, and show superior performance over strong baseline methods.

CCS CONCEPTS

• Computing methodologies → Motion capture.

KEYWORDS

Wearable Devices, Inertial Measurement Units, Human Motion

ACM Reference Format:

Yifeng Jiang, Yuting Ye, Deepak Gopinath, Jungdam Won, Alexander W. Winkler, and C. Karen Liu. 2022. Transformer Inertial Poser: Real-time Human Motion Reconstruction from Sparse IMUs with Simultaneous Terrain

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SA '22 Conference Papers, December 6–9, 2022, Daegu, Republic of Korea

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9470-3/22/12...\$15.00

<https://doi.org/10.1145/3550469.3555428>

Generation. In *SIGGRAPH Asia 2022 Conference Papers (SA '22 Conference Papers)*, December 6–9, 2022, Daegu, Republic of Korea. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3550469.3555428>

1 INTRODUCTION

Real-time reconstruction of 3D human motion is crucial for applications in various domains, such as biomechanics and sports analysis, motion-based video games, and virtual presence in VR/AR systems. While marker-based optical motion capture systems [Vicon n d] remain an ideal option for research labs and professional studios due to the superior accuracy, more and more applications demand a portable, less costly, and minimally-intrusive mocap system that reconstructs human movements in real-time and can be used anywhere by everyone.




Figure 1: We develop an attention-based deep learning method to reconstruct full-body motion from six IMU sensors in real-time. Our method simultaneously generates plausible terrain maps that can explain the reconstructed motions, for a variety of motion types.

Among many sensing modalities, such as RGB cameras [Cao et al. 2019; Güler et al. 2018; Kanazawa et al. 2019], depth cameras [Taylor et al. 2012; Wei et al. 2012], or wearable electromagnetic sensors [Kaufmann et al. 2021], inertial measurement sensors (IMUs) [Huang et al. 2018; von Marcard et al. 2017] provide many unique advantages. IMU-based mocap is self-contained, egocentric, and untethered, applicable to both indoor and outdoor activities in all lighting conditions. For example, an IMU system can capture long-distance motion of a hiker or a mountain climber, but third-person cameras would be limited by their field of view. In addition, IMU-based mocap is unsusceptible to occlusion and is less sensitive to privacy issues. Recent works [Huang et al. 2018; Yi et al. 2021] have advanced the field from commercial systems using 17 IMUs [Xsens n d] to machine-learning-based systems using only six IMUs, demonstrating the great potential of the IMU-based approach as a practical and reliable solution to "everyday motion capture".

With all the recent progress, there are still a few critical issues, such as temporal inconsistency of prediction, and global translational drifts, that prevent IMU-based mocap systems to achieve their full potential. In response to these challenges, we propose a novel method to enhance the reconstruction of full-body human motion in real-time from only six IMU sensors. First, inspired by recent success of Transformer models for free-form natural language generation [Radford et al. 2018], we cast the IMU reconstruction as a constrained motion generation problem and solve it by learning a conditioned Transformer decoder model. Comparing with RNN-based models used by previous methods, our model is able to achieve more accurate joint angle estimation, especially for those motions with almost identical IMU signals such as sitting and standing, thanks to Transformer's improved ability to reason about its own past predictions by taken them explicitly as input.

Second, we introduce a general technique to address the gradual drift of global translation and joint motion in an unified way. We train our Transformer decoder model to additionally predict the points on the character that have near-zero velocity, which we call "Stationary Body Points" (SBPs). Once learned, SBPs can be used by analytical routines to counteract drifts in motion, which is particularly helpful when there is a distribution mismatch between test and training sets due to noisy or corrupted IMU sensors, or unseen motions. Concurrent to our work, [Yi et al. 2022] also tackles the drifting problems using a combination of physics-based optimization and improved initialization techniques for RNNs. Our work is interestingly complementary to [Yi et al. 2022] on this front.

The IMU-based approach opens the door to motion capture in large environments with a variety of terrains, posing a new research challenge of terrain reconstruction in real-time. As our aforementioned motion estimators are terrain-agnostic, by leveraging the predicted SBPs, our method can be used to generate plausible height maps of the terrain traversed by the human. Moreover, instead of treating terrain reconstruction as a byproduct of our method, we simultaneously estimate the human motion and the corresponding terrains such that the generated terrains can also be used to regularize the reconstructed global motion in real-time.

We evaluate our system, which we call "Transformer Inertial Poser" (TIP), extensively on synthetic and real datasets, and show-case clear improvement over recent state-of-the-art methods on flat terrains. As no real IMU datasets with non-flat terrains exist

for training or testing, we qualitatively evaluate our reconstructed motions along with the simultaneously generated terrains both in simulation and with our own live demos. We additionally show real-time live demos covering a wide variety of motion types, many of which unseen in previous works.

2 RELATED WORK

Human motion reconstruction from various sensor inputs has been studied for a long time particularly in Computer Graphics and Computer Vision communities. We mainly review prior works that use IMU sensors as part of or the sole input modality. We also review motion generation models based on the Transformer [Vaswani et al. 2017] as it constitutes the core of our reconstruction model.

A typical IMU sensor includes an accelerometer measuring 3-axis linear acceleration, a gyroscope measuring 3-axis angular velocity, and a magnetometer identifying the vector pointing Earth's magnetic north. From these raw signals, sensor fusion algorithms based on Kalman filter or its extended version are used to provide more robust measures of the orientation [Bachmann et al. 2001; Del Rosario et al. 2018; Foxlin 1996; Vitali et al. 2021]. IMU sensors have been used along with vision-based sensors such as RGB or RGB-D cameras for motion estimation. Some work regard IMU signals as extra constraints to regularize motions predicted from vision, where those constraints are formulated either in offline optimizations [Helten et al. 2013; Pons-Moll et al. 2011, 2010; von Marcard et al. 2018, 2016; Zheng et al. 2018], online (often per-frame) optimizations [Charles Malleson 2020; Malleson et al. 2017; Zhang et al. 2020], or learning deep neural networks [Gilbert et al. 2019; Trumble et al. 2017]. There have also been works combining IMUs with other modalities such as optical markers [Andrews et al. 2016] or ultrasonic [Liu et al. 2011; Vlastic et al. 2007]. Although these systems can produce plausible motions, they also suffer from inherent limitations of vision-based sensors such as narrow capture region, occlusion, or sensitivity to the light condition.

With IMU sensors getting more compact and inexpensive, they have received increasing attentions from both industry and research communities for a standalone body tracking solution. Popular commercial products such as [Xsens n d] and [Rokoko n d] can generate high-quality human motions ready to be used in real-time game engines. However, requiring a sophisticated full-body setup with at least 17 IMUs hinders their accessibility to everyday users. Researchers have therefore proposed body tracking systems with a small number of IMUs sparsely placed on the body, usually utilizing statistical body models and/or high quality optical mocap data as prior to mitigate input signals being under-specified. Marcard et al. [2017] developed an offline system (SIP) with only six IMUs, which optimizes poses and the parameters of the SMPL body model [Loper et al. 2015] to fit the sparse sensor input. Huang et al. [2018] learned a deep neural-net model (DIP) from a large amount of motion capture data to directly map the IMU signals to poses. Their model is based on bidirectional recurrent neural networks (BRNN), so the system can run in an online manner while considering both the past and future sensor inputs with a negligible latency, outperforming previous non-learning online methods. An ensemble of BRNNs was further adopted by Nagaraj et al. [2020] to improve upon the results. However, the two real-time solutions

mostly focus on reconstructing the local joint motion without global translation. Yi et al. [2021] proposed a new neural model (TransPose) where the progressive upscaling of joint position estimation showed more accurate pose estimation. The model can additionally generate accurate global root motions by combining a supporting-foot heuristics and a small learned deep network, similar to [Rempe et al. 2021]. Recently, an extension of this system (PIP) [2022] has been introduced, which is concurrent to our paper, where the predicted motions are further optimized to reduce violations of physics laws [Shimada et al. 2020]. We explore this problem domain with a set of drastically different techniques, and with much more relaxed assumptions on the environment geometry, while producing comparable or better reconstruction results.

The blooming AR/VR industry draws attention to full-body motion synthesis, using IMUs on headsets and controllers. With no sensors available on the lower body, utilizing deep reinforcement learning, Luo et al. [2021] is able to synthesize physically valid locomotion from only the 6D egocentric headset pose. Cha et al. [2021] complements pose estimation from headset cameras with IMUs when the hands are out-of-view. Choutas et al. [2021] and Dittadi et al. [2021] experimented with deep generative models conditioned on headset and controllers poses to synthesize full-body poses. More similar to our work is LoBSTR [Yang et al. 2021], where they include an IMU on the waist in addition to IMUs on the headset and controllers. Using a recurrent network, they can synthesize both sitting and running motions from only 4 sensors. We opt to use 6 IMUs with lower body information for accurate motion reconstruction rather than synthesis, but these sparser setups are fruitful future directions.

Since the inception of the Transformer, attention-based models have become the state-of-the-art on many problems involving sequence data, such as language translation [Brown et al. 2020] and audio generation [Dhariwal et al. 2020]. It is natural to also apply Transformer models to synthesizing human motion. Aksan et al. [2021] developed a generative model using dual attention mechanism to capture spatial and temporal correlations, which predicts future full-body locomotion given a short history. Petrovich et al. [2021] used a Transformer and a variational autoencoder conditioned on action labels, such as *walking* or *jumping*, to generate full body motions. Valle-Pérez et al. [2021] instead combined a Transformer with normalizing flows to synthesize dancing motion from music, building on a similar prior work [Li et al. 2021]. For motion reconstruction, Kim et al. [2021] experimented with a Transformer encoder-decoder model with sparse synthetic input features, and found it more effective than recurrent networks. In our case, we face the additional challenge of handling noise from real IMU sensors.

3 TRANSFORMER INERTIAL POSER (TIP)

We introduce a real-time human motion reconstruction technique from six IMU sensors placed on the user’s legs, wrists, head, and waist. Our approach combines a learning-based model and an analytical routine to estimate full-body joint angles \mathbf{q} and the root velocity \mathbf{v} from a real-time stream of IMU orientation \mathbf{R} and acceleration \mathbf{A} signals. Summarized in Figure 2, our method depends on a learned *Transformer decoder* to estimate the motion and an

analytical *drift stabilizer* to refine the estimation. In addition to \mathbf{q} and \mathbf{v} , the Transformer decoder also predicts stationary body points (SBPs) \mathbf{c} which are used by the drift stabilizer during run time to improve the accuracy of the reconstructed motion. The drift stabilizer utilizes the predicted SBPs to mitigate drifts in motion over time and a non-learning *terrain update* module estimates the corresponding terrain while further improving motion reconstruction. A plausible terrain height map is generated and updated in real time, as an additional product of our algorithm.

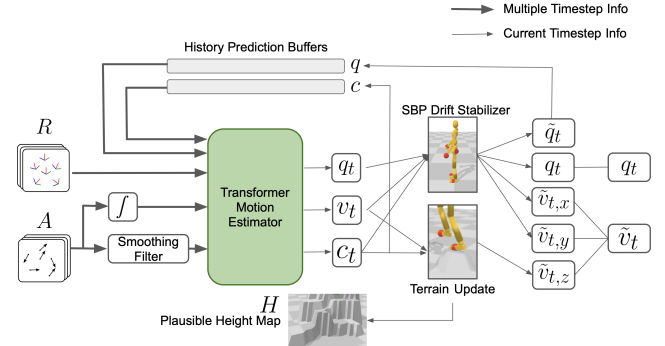


Figure 2: Overview of our pose estimation algorithm. Thicker arrows are "pipes" flowing multi-timestep information at each model step, while thinner arrows flow current-step information.




Figure 3: Our conditional Transformer decoder architecture during training, showing $M = 39$. Following GPT training conventions, during training, input and output contain the same (but time-shifted by 1) ground-truth motion (\mathbf{q} & \mathbf{c}). During testing, output from $t - 39$ to $t - 1$ are discarded and the last ones $\mathbf{v}_t, \mathbf{c}_t, \mathbf{q}_t$ are used as current prediction, and input (\mathbf{q} & \mathbf{c} , i.e. the history buffers in Figure 2) is autoregressively updated by most recent outputs $\mathbf{v}_t, \mathbf{c}_t$ at each step.

3.1 Transformer Motion Estimator

To simplify the notations, consider the problem of reconstructing only the joint angles \mathbf{q} . The standard learning-based approach can be summarized as $\max P(\mathbf{q}_t | \mathbf{R}_{t-N:t}, \mathbf{A}_{t-N:t})$, where a neural network is trained to output the most probable current joint angle given the IMU readings. Note that since 6 noisy IMUs do not provide enough signals to fully determine the whole-body motion, we follow previous work [Huang et al. 2018] to include a time window from $t - N$ to t , mitigating under-specification.

Such models could still have trouble distinguishing motions with similar IMU readings over the entire window, such as sitting vs. standing, and the transition in between (if slow). To principally address the issue, we draw an analogy between motion generation and language modeling: $\max P(\mathbf{q}_{t-M:t})$, where the goal is to free-form generate all motion (language) sequences that are natural. Applying recursively, a neural network can be used to generate plausible \mathbf{q}_t conditioned on all the previous generations from itself:

$$\max P(\mathbf{q}_{t-M:t}) = \max P(\mathbf{q}_t | \mathbf{q}_{t-M:t-1}) P(\mathbf{q}_{t-1} | \mathbf{q}_{t-M:t-2}) \cdots P(\mathbf{q}_{t-M}).$$

A Transformer decoder model [Radford et al. 2018; Vaswani et al. 2017] excels at this task as it explicitly takes previous predictions as input and is capable of reasoning about such contexts fast enough in parallel with temporal attention mechanisms. The Transformer decoder model is able to take variable-sized history as input but in practice, a maximum window length M is usually set.

Unlike free-form language generation [Radford et al. 2018], we have additional constraints from the IMU sensors. As such, we model our constrained motion generation problem with a *conditional* Transformer decoder, i.e. $\max P(\mathbf{q}_{t-M:t} | \mathbf{R}_{t-M:t}, \mathbf{A}_{t-M:t})$, where we feed the model the sequences of IMU readings \mathbf{R} and \mathbf{A} in parallel to its past predictions, with the same sequence length at each step, up to a maximum of M .

Model. The input to our model includes the IMU acceleration readings after smoothing filtering or integration $\tilde{\mathbf{A}} \in \mathbb{R}^{36 \times (M+1)}$ (Appendix A), and the IMU orientations $\mathbf{R} \in \mathbb{R}^{54 \times (M+1)}$ represented as flattened rotation matrix (length 9) from 6 sensors (Figure 3). The output includes 18 joint angles $\mathbf{q}_t \in \mathbb{R}^{108}$ defined in the SMPL [Loper et al. 2015] human model (excluding joints such as toes following previous works), the root linear velocity, $\mathbf{v}_t \in \mathbb{R}^3$, and the stationary body points (SBPs) $\mathbf{c}_t \in \mathbb{R}^{20}$. Each joint in \mathbf{q}_t is represented redundantly as first two columns of its local rotation matrix for unique and numerically stable ground-truth labels [Zhou et al. 2019]. The root orientation is given directly by one of the IMUs placed on the waist. We empirically found that adding an recurrent layer to "summarize" output embedding of the Transformer accelerates convergence during training.

Following the standard practice of Transformers (GPT [Radford et al. 2018]) training, given the shifted ground-truth sequences from $t - M - 1$ to $t - 1$, the model learns to predict *in parallel* the whole sequence from $t - M$ to t for efficiency (Figure 3). To prevent the model from learning simply to shift the input by one timestep, a causal mask [Vaswani et al. 2017] is added to hide future attention information, mimicking the test-time setting. During test time, since we only care about the most recent prediction, output from $t - M$ to $t - 1$ will be discarded at each prediction step.

There is still a large asymmetry between training and testing times. During training, the model sees ground-truth motion as the history input, but during testing, the history is noisily accumulated from its own predictions. Previous works, such as [Radford et al. 2018], found in practice that such "teacher-forcing" training will not cause overfitting to the clean history. In our case, different from language where neighboring words are distinct, neighboring poses are usually similar to each other, providing much duplicate information more susceptible to overfitting. We add a 80% dropout [Srivastava et al. 2014] to the history of \mathbf{c} and \mathbf{q} , effectively dropping

4 frames out of 5. We also found that excluding \mathbf{v} from history is important to prevent test-time autoregressive divergence, possibly because \mathbf{v} is usually nearly constant in a time window, and the model could easily exploit and overfit to history \mathbf{v} without truly reasoning about IMUs or history joint angles.

3.2 Stabilizing Drift with SBPs

Combating drifts is one of the biggest challenges for IMU-based motion reconstruction. Unlike motion estimation methods using external cameras, IMUs have no direct sense of relative position, and learning or optimization based algorithms can all to some extent be seen as relying on double integration of the noisy acceleration readings to estimate position. If there is any biased error due to calibration or environment interference, the drifts of root translation or joint angles over time will result in artifacts, such as frequent foot skating or locking, or erroneous motion transitions.

We present a simple and general technique, combining learned and analytical components, which does not require body-part-specific heuristics, nor rely on assumptions of the environment terrains. Besides motion predicting, our model will also predict Stationary Body Points (SBPs) \mathbf{c}_t on the character, the representative locations on the human body with near-zero velocity. For example, the heel-to-toe rolling contact during walking results in moving SBPs across the foot, and rolling on the floor motion results in SBPs moving across the lower back and pelvis (Figure 4). Once learned, SBPs could be used for analytical correction of the predicted motion. Though the learning of SBPs could also be subject to inaccuracies, the analytical usage of the SBPs constrains the amount of errors they can produce.




Figure 4: During heel-to-toe contact in locomotion, contact patch is stationary but the foot link velocity is not zero. Similarly, a person rolling on the ground can be seen as rotating around changing stationary points.

Discover Ground-truth SBPs for Learning. SBPs can be discovered on any body part, but in practice we choose the hands, feet, and pelvis to be the only candidate regions for SBPs in our implementation. We parameterize each SBP location as a vector offset $\mathbf{r}_t \in \mathbb{R}^3$ from the center of mass of the body part it belongs to. We use one additional bit for each SBP $b_t \in \{0, 1\}$ to indicate the existence of SBP at the current moment. Therefore, SBP is represented as $\mathbf{c} \in \mathbb{R}^{20}$, where $\mathbf{c}_t = [b_t^{(i)}, \mathbf{r}_t^{(i)}]_{i=1, \dots, 5}$.

We devise an efficient sampling-based method to search for the point on the rigid body with minimal velocity,

$$\operatorname{argmin}_r \|\omega \times \mathbf{R}_B \mathbf{r} + \mathbf{v}\|,$$

where ω and \mathbf{v} are angular and linear velocities of the body part, and \mathbf{R}_B is the orientation of the body, which can all be easily obtained

from ground-truth $\mathbf{q}_t, \mathbf{v}_t$ in training data. If the minimal velocity is below a small threshold, a SBP for that body part is found.

We evaluate the velocity of every points in a 3D grid that encloses the body part of interest and find the point with the minimal velocity. For example, the 3D grid for a foot is a 6cm-thick box covering the bottom of the foot and that for the pelvis covers a larger region on the back. Evaluation for all candidate points can be done in parallel efficiently using matrix operations.

In practice, directly searching for the point with minimal velocity can lead to jittery SBPs over time, due to noise in motion capture data [Le Calennec and Boulic 2006] and approximating human body parts with rigid bodies. We thus add a temporal regularizer to the search criteria:

$$l(\mathbf{r}) = \|\omega \times \mathbf{R}_B \mathbf{r} + \mathbf{v}\| + 0.3 \|\mathbf{r} - \mathbf{r}_{t-1}\|,$$

where \mathbf{r}_{t-1} is the solution in the previous frame, if existing. If minimal $l(\mathbf{r}^*)$ is less than a manually chosen threshold (0.25), we label $\mathbf{c}_t(k) = [1, \mathbf{r}^*]$. Otherwise, $\mathbf{c}_t(k) = [0, (0, 0, 0)]$.

We did not find our results sensitive to the choices of regularization weight (0.3) or threshold, and use the same formula for all five SBPs, across the entire training dataset with various terrains.

Run-time Root Correction. The predicted SBPs provide constraints to correct the root velocity. We set the corrected root velocity to be the average of $\tilde{\mathbf{v}}_t^{(i)}$, where each $\tilde{\mathbf{v}}_t^{(i)}$ is the root velocity that makes i -th (active) SBP exactly stationary in world space at current step. In practice, we only use this technique to correct root translation in the horizontal plane, and leave root correction in the height direction to the terrain estimation module described in Sec. 3.3.

Run-time Joint Correction. We can also modify the joint angle estimation \mathbf{q}_t using SBPs. When a pair of SBPs are on for consecutive frames, we can solve for the joint angle correction that maintains the distance vector between the SBPs. We adopt two-bone inverse kinematics (IK) [Holden et al. 2020], which gives numerically stable solutions and changes only a minimal number of joints, $\tilde{\mathbf{q}}_t$. Directly replacing the joint angles \mathbf{q}_t estimated by the Transformer decoder with $\tilde{\mathbf{q}}_t$ can result in discontinuity in motion when SBPs switch on and off. The problem is more visible as SBP onsets are imperfect model predictions. We adopt a new "soft-IK" technique utilizing the Transformer decoder's dependency on its own history of prediction. Intuitively, next model predictions will be improved if the history buffer is filled with more accurate motion. As such, our idea is to accept the current prediction \mathbf{q}_t from the Transformer, but feed the corrected joint angles, $\tilde{\mathbf{q}}_t$, back to the Transformer's history buffer. This essentially creates a soft-IK constraint which does not enforce SBP pairs immediately in the current time step, but do so gradually in subsequent frames.

3.3 Plausible Terrain Generation

Inspired by SLAM (Simultaneous Localization And Mapping) algorithms [Durrant-Whyte and Bailey 2006; Guзов et al. 2021], our method generates a plausible terrain consistent with the reconstructed motion. In our case, simultaneously predicting the terrain and the motion is mutually beneficial in achieving improved results for both. The algorithm takes as input current pelvis and feet SBP locations, as well as Transformer-estimated vertical root velocity

$\mathbf{v}_{t,z}$, and outputs a height map $\mathbf{H} \in \mathbb{R}^{L \times L}$ and the proposed vertical root correction $p \in \mathbb{R}$ ($\tilde{\mathbf{v}}_{t,z} = \mathbf{v}_{t,z} + p$). We do not consider hand SBPs here because in most cases it is hard to know if hand rests on something or simply stays stationary in the air.

This algorithm is based on two mild assumptions, albeit crucial to regularizing noisy SBPs from real IMU data. First, if two SBPs are horizontally nearby ($< 1\text{m}$) and have sufficiently similar heights ($< 0.1\text{m}$), we assume they have the same height on the terrain. As such, we assume there is no gradual slope in the scene, since under sensor drifting it is extremely difficult to distinguish mild slopes from flat ground. By this assumption, we cluster nearby SBPs with similar heights into buckets and store the mean height for each bucket. When a new SBP is detected, if an existing nearby cluster is similar in height, this SBP will join that cluster and update its mean height. $p = -k * d$ will be proposed to drag the root so that the SBP is closer to the updated mean height, where k is a constant correction coefficient and d is the difference between SBP height and cluster mean height.

Second, the terrain is assumed to be a Voronoi diagram with different vertical levels. That is, we assign each unvisited grid to the same height as its closest SBP-visited neighbor. To build Voronoi online, upon each new SBP visitation, we need to check if each unvisited grid is now closer to the new SBP than all existing SBPs. Instead of storing all distances, we only need to store the closest distance up to now, which we call (inverse) confidence map $\mathbf{C} \in \mathbb{R}^{L \times L}$. If the new SBP is closer, the height map \mathbf{H} at the unvisited grid will be updated to the new SBP's height (precisely, the corresponding SBP cluster's height, from the first assumption).

Further implementation details can be found in Appendix C.

4 EVALUATIONS

We organize our experiments in this section to demonstrate that:

- Using the same held-out datasets of real or synthesized sparse IMU signals as benchmarks, our method improves over recent works quantitatively by a significant margin.
- Without training on any real non-flat-terrain IMU data, and without any ground-truth terrain supervision, our method can reconstruct motions on different terrains both in simulation and on real sensor data.
- Both real-time motion correction and terrain reconstruction using SBPs contribute to our ability to stably mitigating root drift during a variety of human activities.
- Modifying joint history with SBP correction mitigates drift between joints and qualitatively improves motion reconstruction in challenging motions.
- Our Transformer decoder model taking history as input facilitates stable and consistent learning of the SBPs.

4.1 Quantitative Evaluation on Flat Terrain

We evaluate our results on existing IMU datasets that cover diverse types of motions and are paired with ground-truth full-body motions, for quantitative comparisons.

Datasets. Following previous works, evaluation datasets include DIP and TotalCapture of real IMU data. We match the exact training data and evaluation settings of baseline methods. This mainly

serves to test the sim-to-real transfer of the models, which are all predominantly (~95%) trained on synthesized AMASS data (Appendix A). Additionally, we hold out a synthetic IMU dataset, DanceDB, to evaluate cross-motion-type generalization.

- **DIPEval** (*real heldout*): Data from two held-out subjects in the ten-subject DIP dataset.
- **TotalCapture** (*real heldout*): We held out real IMU measurements from the TotalCapture dataset [Trumble et al. 2017] for evaluation, but still use its ground truth and synthesized IMU readings as part of the AMASS training set, following the same practice as previous works.
- **DanceDB** (*synthetic heldout*): A large dataset of contemporary dances, therefore containing unique motion types to any other training dataset. Note that DanceDB is part of AMASS but we intentionally hold it out from our training data. Previous works likely did not include them in training either since they predate DanceDB’s release in AMASS.

Metrics. We use the following metrics common for evaluating motion reconstruction quality. We randomly sample 600 consecutive frames (10s) from each motion in the evaluation datasets, to prevent very long motions biasing the statistics, and to avoid evaluation of root translation error in the beginning of motions where they mostly start from a stationary standing pose.

- **Mean Joint Angle Error** (*in degrees*): Joint angle (represented in axis-angles) difference between reconstruction and ground-truth, averaged over all joints.
- **Mean Root-Relative Joint Position Error** (*in centimeters*): Global joint Cartesian position difference (Euclidean norm) between the reconstruction and ground-truth by aligning at the root, averaged over all joints.
- **Root Error 2s/5s/10s** (*in meters*): Root translation error measured in Euclidean norm during a continuous period of 2s/5s/10s. Note that existing works have no errors from the vertical axis as they are designed for flat-ground motions.
- **Mean Joint Position Jitter** (*in m/s^3*): Joint position jitter computed using the same formula as in TransPose, averaged over all joints.
- **Root Jitter** (*in m/s^3*): Root position jitter computed using the same formula as above.

Results. We present quantitative metrics on the evaluation datasets between our model, TransPose, and DIP in Table 1. We used the best performing models published by the authors in this comparison. Overall, our system achieves better accuracy in almost all evaluations. Our system does not have an offline mode, and all reported metrics reported are from online inference (i.e. pretending existing data files as streaming in).

We see most significant improvements from previous works on root translation, on both TotalCapture and DanceDB datasets. Sec. 4.3 shows the effects of our key design choices on root drifting mitigation. Thanks to our preprocessing filtering (Appendix A), we can treat the DIP training split as ordinary ~5% subset of the whole training data with the rest ~95% being purely synthetic IMU. By using a simple one-stage training procedure, instead of the two-stage training-then-finetuning as in previous works, we reduce the risk of overfitting to activities in the DIP dataset during finetuning.

Comparing to results from concurrent PIP work [Yi et al. 2022], our method achieves a similar level of improvement over TransPose. On flat terrain activities, our work complements the PIP model by exploring drastically different techniques. We are not able to compare with this new work quantitatively as the code and model have not been released yet.

4.2 Quality of Simultaneously Generated Terrains

Since our reconstructed terrains are only one *plausible* result among infinite possibilities, we perform the following two *qualitative* experiments to evaluate their consistency with the human motion:

- **Simulation**: we evaluate our terrain generation algorithm by running our system on terrain navigation motions included in the AMASS synthesized IMU training data. (Video 1m10s, note AMASS does not provide terrain ground-truths.)
- **Live Demos**: though our model has never seen any real IMU signals on terrains, we push its limit and showcase simultaneous motion reconstruction and terrain generation for walking on stairs (Video 1m56s), and climbing/jumping off chairs (Video 2m23s).

The supplementary video demonstrates plausible terrains both in live demos and on the AMASS dataset. Figure 5 presents one of the generated terrains from simulation. Although we do not have access to the ground truth terrain, our Voronoi-graph based algorithm generates plausible staircases that are consistent with the human motion. Body parts may still penetrate the terrain briefly

Table 1: Comparison of model performance on evaluation datasets. Bold numbers indicate the best performing entries.

Our TIP Model			
	DIPEval	TotalCapture	DanceDB
joint angle errors (degree)	12.09586	8.91642	15.57031
joint position errors (cm)	5.82242	5.14566	8.50089
root errors in 2s (meter)		0.08031	0.20295
root errors in 5s (meter)		0.1351	0.29681
root errors in 10s (meter)		0.19446	0.35759
joint position jitter (m/s^3)	0.8823	0.75075	1.43867
root jitter (m/s^3)	0.66211	0.61108	0.98474

TransPose Model			
	DIPEval	TotalCapture	DanceDB
joint angle errors (degree)	13.10898	11.37459	17.42089
joint position errors (cm)	6.28771	5.60905	8.46368
root errors in 2s (meter)		0.16928	0.20812
root errors in 5s (meter)		0.32677	0.31241
root errors in 10s (meter)		0.35323	0.46058
joint position jitter (m/s^3)	0.63113	0.66816	1.49544
root jitter (m/s^3)	0.53972	0.63572	1.31713

DIP Model			
	DIPEval	TotalCapture	DanceDB
joint angle errors (degree)	14.14957	16.54206	23.27462
joint position errors (cm)	7.79099	9.36553	13.83818
joint position jitter (m/s^3)	1.85047	1.74527	2.15585

Table 2: Comparison of ablation models on TotalCapture evaluation dataset.

	No history, w/ SBP	No SBPs	Predict SBPs wo/ usage	wo/ terrain corrections	TIP
root errors in 2s	0.15833	0.16987	0.13927	0.07817	0.08031
root errors in 5s	0.28575	0.32543	0.32972	0.15462	0.13509
root errors in 10s	0.41114	0.5051	0.51706	0.24196	0.19445

while it is being generated (Figure 5, second to the left), because the algorithm ignores each newly active SBP for a short period of time (Appendix C). Final terrain should see little penetration when playing back the motion.




Figure 5: Generated terrain. Many equally plausible solutions exist for same motion, the solution from our algorithm being one of them.

4.3 Ablation Studies on Root Drifting

Multiple key design choices improved our root translation accuracy, which is also crucial to the success of terrain generation. To gain a better understanding, we created the following ablation models:

- **No history, with SBP:** a Transformer *encoder* can be used as a sequence summarizing model, taking time windows of IMU signals and output only the current full-body pose q_t , v_t and SBPs c_t , with no consideration of its past predictions. We still used the SBPs to correct the root in this case.
- **No SBPs:** same network as Ours without learning SBPs.
- **Predict SBPs without Usage:** no test-time usage of the predicted SBPs. This is different from "No SBPs" since learning an unused auxiliary task may help the main task learning.
- **Without Terrain Corrections:** Do not use terrain algorithm for vertical drift correction.

Table. 2 summarizes the models' performances on the real-IMU TotalCapture dataset. Both "no SBPs" and "predicting SBPs without usage" rely on the model's raw prediction of v_t , and perform the worst in this test, showing the importance of both learning the SBPs and enforcing them at run time. "No history, with SBP" also fails, interestingly because the SBPs predicted by a Transformer encoder becomes intermittent in this case and therefore less helpful during run time. This shows the importance of explicitly including output history as input. Intuitively, the onset history of SBPs are important information for next SBP onset prediction.

Finally, without terrain-based root drifting correction, in some testing motions, vertical drift may be large for reasons such as biased sensor error from calibrations. Such errors in vertical direction also negatively affects generated terrains if not corrected in time.

4.4 Qualitative Comparisons of Run-time Joint IK Correction

Introduced in Sec. 3.2, we use IK to correct joint prediction histories with a pair of active SBPs. We showcase its usage in one of the most common joint drifting scenarios of long sitting [Yi et al. 2022, 2021], where the character may easily transition to standing over an extended period of time due to similar orientations and accelerations on legs and waist IMUs. With our SBP IK correction, we can stably reconstruct a two-minute sitting sequence, both with our own sensors (Video 0m42s) or using public real-IMU sequences (Video 3m10s) in DIP dataset, showing the value of correcting history buffer with IK to be consistent with predicted SBPs.

4.5 Live Demo

We test our system live with 6 Xsens IMU sensors. Our video visualizes live performance side-by-side with real-time reconstructions, with a slight latency caused by our pre-processing filter. Besides the aforementioned tasks, we cover a variety of motion tasks in our demos, both common ones such as locomotion and whole-body manipulation, and more challenging ones such as jumping from a high place, "swimming" on a stool, dancing, hand on floor movements, or swirl kicks. We tested our system on one male and one female subjects, and observed degraded performance on the female subject which we did not expect (Appendix F).

5 CONCLUSION

This paper presents a new data-driven method for human motion reconstruction from six wearable IMUs, with simultaneous plausible terrain generation. By combining a conditional Transformer decoder model for consistent prediction, a hybrid drift stabilizer utilizing learned stationary information across human body, and an algorithm to simultaneously generate regularized terrain and correct noisy global motion estimation, new downstream applications can be made possible with this self-contained and economic setup of motion capture. For future work, our method could be largely improved by collecting real IMU datasets with motions on various terrains. Personalized finetuning or calibration may also improve reconstruction for each individual user.

ACKNOWLEDGMENTS

To members of the Stanford Movement Lab, the Stanford Human Performance Lab, Meta Reality Labs Research, and to Josh Cooley, Yinghao Huang, Manuel Kaufmann, Ari Tamari, Xia Wu, Di Xia, Xinyu Yi, Eris Zhang, for helpful discussions and technical assistance. To anonymous reviewers whose feedback substantially helped refine this work. Yifeng Jiang is partially supported by the Wu Tsai Human Performance Alliance at Stanford University.

REFERENCES

- Emre Aksan, Manuel Kaufmann, Peng Cao, and Otmar Hilliges. 2021. A Spatio-temporal Transformer for 3D Human Motion Prediction. *International Conference on 3D Vision (3DV)* (2021).
- Sheldon Andrews, Ivan Huerta, Taku Komura, Leonid Sigal, and Kenny Mitchell. 2016. Real-Time Physics-Based Motion Capture with Sparse Sensors. In *Proceedings of the 13th European Conference on Visual Media Production (CVMP 2016)* (CVMP 2016). Article 5.
- Eric R. Bachmann, Robert B. McGhee, Xiaoping Yun, and Michael J. Zyda. 2001. Inertial and Magnetic Posture Tracking for Inserting Humans into Networked Virtual Environments. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '01)*. 9–16.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. (2020). arXiv:2005.14165 [cs.CL]
- Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. 2019. OpenPose: Real-time Multi-Person 2D Pose Estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019).
- Young-Woon Cha, Husam Shaik, Qian Zhang, Fan Feng, Andrei State, Adrian Ilie, and Henry Fuchs. 2021. Mobile. Egocentric Human Body Motion Reconstruction Using Only Eyeglasses-mounted Cameras and a Few Body-worn Inertial Sensors. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*.
- Adrian Hilton Charles Malleon, John Collomosse. 2020. Real-Time Multi-person Motion Capture from Multi-view Video and IMUs. *International Journal of Computer Vision* 128 (06 2020).
- Vasileios Choutas, Federica Bogo, Jingjing Shen, and Julien Valentin. 2021. Learning to Fit Morphable Models. *CoRR* abs/2111.14824 (2021). arXiv:2111.14824 <https://arxiv.org/abs/2111.14824>
- Michael B. Del Rosario, Heba Khamis, Phillip Ngo, Nigel H. Lovell, and Stephen J. Redmond. 2018. Computationally Efficient Adaptive Error-State Kalman Filter for Attitude Estimation. *IEEE Sensors Journal* 18, 22 (2018), 9332–9342.
- Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. 2020. Jukebox: A Generative Model for Music. arXiv:2005.00341 [eess.AS]
- Andrea Dittadi, Sebastian Dziadzio, Darren Cosker, Ben Lundell, Tom Cashman, and Jamie Shotton. 2021. Full-Body Motion From a Single Head-Mounted Device: Generating SMPL Poses From Partial Observations. In *International Conference on Computer Vision 2021*.
- H. Durrant-Whyte and T. Bailey. 2006. Simultaneous localization and mapping: part I. *IEEE Robotics Automation Magazine* 13, 2 (2006), 99–110. <https://doi.org/10.1109/MRA.2006.1638022>
- E. Foxlin. 1996. Inertial head-tracker sensor fusion by a complementary separate-bias Kalman filter. In *Proceedings of the IEEE 1996 Virtual Reality Annual International Symposium*. 185–194.
- Andrew Gilbert, Matthew Trumble, Charles Malleon, Adrian Hilton, and John Collomosse. 2019. Fusing Visual and Inertial Sensors with Semantics for 3D Human Pose Estimation. *International Journal of Computer Vision* 127 (04 2019), 1–17.
- Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. 2018. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7297–7306.
- Vladimir Guvoz, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. 2021. Human POSEitioning System (HPS): 3D Human Pose Estimation and Self-localization in Large Scenes from Body-Mounted Sensors. In *CVPR*.
- Thomas Helten, Meinard Müller, Hans-Peter Seidel, and Christian Theobalt. 2013. Real-Time Body Tracking with One Depth Camera and Inertial Sensors. In *2013 IEEE International Conference on Computer Vision*. 1105–1112.
- Daniel Holden, Oussama Kanoun, Maksym Perepichka, and Tiberiu Popa. 2020. Learned Motion Matching. *ACM Trans. Graph.* 39, 4, Article 53 (jul 2020), 13 pages. <https://doi.org/10.1145/3386569.3392440>
- Yinghao Huang, Manuel Kaufmann, Emre Aksan, Michael J. Black, Otmar Hilliges, and Gerard Pons-Moll. 2018. Deep Inertial Poser: Learning to Reconstruct Human Pose from Sparse Inertial Measurements in Real Time. *ACM TOG* 37, 6 (12 2018).
- Angjoo Kanazawa, Jason Y. Zhang, Panna Felsen, and Jitendra Malik. 2019. Learning 3D Human Dynamics from Video. In *Computer Vision and Pattern Recognition (CVPR)*.
- Manuel Kaufmann, Yi Zhao, Chengcheng Tang, Lingling Tao, Christopher Twigg, Jie Song, Robert Wang, and Otmar Hilliges. 2021. EM-POSE: 3D Human Pose Estimation from Sparse Electromagnetic Trackers. In *International Conference on Computer Vision (ICCV)*.
- Seong Uk Kim, Hanyoung Jang, Hyeonseung Im, and Jongmin Kim. 2021. Human motion reconstruction using deep transformer networks. *Pattern Recognition Letters* 150 (2021), 162–169.
- Benoît Le Calennec and Ronan Boulic. 2006. Robust kinematic constraint detection for motion data. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on Computer animation*. 281–290.
- Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. 2021. AI Choreographer: Music Conditioned 3D Dance Generation with AIST++.
- Huajun Liu, Xiaolin Wei, Jinxiang Chai, Inwoo Ha, and Taehyun Rhee. 2011. Realtime Human Motion Control with a Small Number of Inertial Sensors. In *Symposium on Interactive 3D Graphics and Games (I3D '11)*. 133–140.
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. 2015. SMPL: A Skinned Multi-Person Linear Model. *ACM TOG* 34, 6 (Oct. 2015), 248:1–248:16.
- Zhengyi Luo, Ryo Hachiuma, Ye Yuan, and Kris Kitani. 2021. Dynamics-Regulated Kinematic Policy for Egocentric Pose Estimation. In *NeurIPS*.
- Charles Malleon, Marco Volino, Andrew Gilbert, Matthew Trumble, John Collomosse, and Adrian Hilton. 2017. Real-time Full-Body Motion Capture from Video and IMUs. In *Int. Conf. 3D Vis.*
- Deepak Nagaraj, Erik Schake, Patrick Leiner, and Dirk Werth. 2020. An RNN-Ensemble Approach for Real Time Human Pose Estimation from Sparse IMUs. In *Proceedings of the 3rd International Conference on Applications of Intelligent Systems* (Las Palmas de Gran Canaria, Spain) (APPIS 2020). Article 32, 6 pages.
- Mathis Petrovich, Michael J. Black, and Gül Varol. 2021. Action-Conditioned 3D Human Motion Synthesis with Transformer VAE. In *International Conference on Computer Vision (ICCV)*. 10985–10995.
- Gerard Pons-Moll, Andreas Baak, Juergen Gall, Laura Leal-Taixé, Meinard Müller, Hans-Peter Seidel, and Bodo Rosenhahn. 2011. Outdoor human motion capture using inverse kinematics and von mises-fisher sampling. In *2011 International Conference on Computer Vision*. 1243–1250.
- Gerard Pons-Moll, Andreas Baak, Thomas Helten, Meinard Müller, Hans-Peter Seidel, and Bodo Rosenhahn. 2010. Multisensor-fusion for 3D full-body human motion capture. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 663–670.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018).
- Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J. Guibas. 2021. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 11488–11499.
- Rokoko. n. d. Rokoko <https://www.rokoko.com/>. Last visited: 08/26/2022.
- Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. 2020. PhysCap: Physically Plausible Monocular 3D Motion Capture in Real Time. *ACM TOG* 39, 6 (12 2020).
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *The Journal of Machine Learning Research* 15, 1 (2014), 1929–1958. <http://dl.acm.org/citation.cfm?id=2627435.2670313>
- Jonathan Taylor, Jamie Shotton, Toby Sharp, and Andrew Fitzgibbon. 2012. The Vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. 103–110.
- Matt Trumble, Andrew Gilbert, Charles Malleon, Adrian Hilton, and John Collomosse. 2017. Total Capture: 3D Human Pose Estimation Fusing Video and Inertial Sensors. In *BMVC*.
- Guillermo Valle-Pérez, Gustav Eje Henter, Jonas Beskow, Andre Holzapfel, Pierre-Yves Oudeyer, and Simon Alexanderson. 2021. Transflower: Probabilistic Autoregressive Dance Generation with Multimodal Attention. *ACM Trans. Graph.* 40, 6, Article 195 (dec 2021), 14 pages. <https://doi.org/10.1145/3478513.3480570>
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, Vol. 30.
- Vicon. n. d. Vicon Motion Systems <https://www.vicon.com/>. Last visited: 08/26/2022.
- Rachel V. Vitali, Ryan S. McGinnis, and Noel C. Perkins. 2021. Robust Error-State Kalman Filter for Estimating IMU Orientation. *IEEE Sensors Journal* 21, 3 (2021), 3561–3569.
- Daniel Vlasic, Rolf Adelsberger, Giovanni Vannucci, John Barnwell, Markus Gross, Wojciech Matusik, and Jovan Popović. 2007. Practical Motion Capture in Everyday Surroundings. *ACM Trans. Graph.* 26, 3 (2007).
- Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. 2018. Recovering Accurate 3D Human Pose in The Wild Using IMUs and a Moving Camera. In *European Conference on Computer Vision (ECCV)*.
- Timo von Marcard, Gerard Pons-Moll, and Bodo Rosenhahn. 2016. Human Pose Estimation from Video and IMUs. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)* (jan 2016).
- Timo von Marcard, Bodo Rosenhahn, Michael Black, and Gerard Pons-Moll. 2017. Sparse Inertial Poser: Automatic 3D Human Pose Estimation from Sparse IMUs. *Computer Graphics Forum* 36(2), *Proceedings of the 38th Annual Conference of the European Association for Computer Graphics (Eurographics)* (2017), 349–360.
- Xiaolin Wei, Peizhao Zhang, and Jinxiang Chai. 2012. Accurate Realtime Full-Body Motion Capture Using a Single Depth Camera. *ACM Trans. Graph.* 31, 6, Article 188 (nov 2012).
- Xsens. n. d. Xsens <https://www.xsens.com/>. Last visited: 08/26/2022.

- Dongseok Yang, Doyeon Kim, and Sung-Hee Lee. 2021. LoBSTr: Real-time Lower-body Pose Prediction from Sparse Upper-body Tracking Signals. *Computer Graphics Forum* (2021). <https://doi.org/10.1111/cgf.142631>
- Xinyu Yi, Yuxiao Zhou, Marc Habermann, Soshi Shimada, Vladislav Golyanik, Christian Theobalt, and Feng Xu. 2022. Physical Inertial Poser (PIP): Physics-aware Real-time Human Motion Tracking from Sparse Inertial Sensors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xinyu Yi, Yuxiao Zhou, and Feng Xu. 2021. TransPose: Real-time 3D Human Translation and Pose Estimation with Six Inertial Sensors. *ACM TOG* 40, 4 (8 2021).
- Zhe Zhang, Chunyu Wang, Wenhui Qin, and Wenjun Zeng. 2020. Fusing Wearable IMUs with Multi-View Images for Human Pose Estimation: A Geometric Approach. In *CVPR*.
- Zerong Zheng, Tao Yu, Hao Li, Kaiwen Guo, Qionghai Dai, Lu Fang, and Yebin Liu. 2018. HybridFusion: Real-Time Performance Capture Using a Single Depth Sensor and Sparse IMUs. In *European Conference on Computer Vision (ECCV)*, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). 389–406.
- Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. 2019. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5745–5753.