BodySLAM: Joint Camera Localisation, Mapping, and Human Motion Tracking

Dorian F. Henning¹⁽⁰⁾, Tristan Laidlow¹⁽⁰⁾, and Stefan Leutenegger^{1,2}⁽⁰⁾ *

¹ Imperial College London, UK {d.henning,t.laidlow15}@imperial.ac.uk ² Technische Universität München, Germany stefan.leutenegger@tum.de

Abstract. Estimating human motion from video is an active research area due to its many potential applications. Most state-of-the-art methods predict human shape and posture estimates for individual images and do not leverage the temporal information available in video. Many "in the wild" sequences of human motion are captured by a moving camera, which adds the complication of conflated camera and human motion to the estimation. We therefore present BodySLAM, a monocular SLAM system that jointly estimates the position, shape, and posture of human bodies, as well as the camera trajectory. We also introduce a novel human motion model to constrain sequential body postures and observe the scale of the scene. Through a series of experiments on video sequences of human motion captured by a moving monocular camera, we demonstrate that BodySLAM improves estimates of all human body parameters and camera poses when compared to estimating these separately.

Keywords: human pose estimation, human motion model, camera tracking, dynamic SLAM

1 Introduction

Estimating the orientation and translation, shape, and posture of human bodies from a sequence of images is an active research area in computer vision due to its many potential applications in domains such as safe human-robot interaction, biomechanical analysis, action recognition, and augmented and virtual reality (AR/VR). To fully realise these types of applications at scale, it is desirable that this estimation can be done "in the wild" from images captured by a moving monocular camera without relying on any special markers or motion-capture systems. This would even enable human motion analysis on image sequences that were not captured with these applications in mind, such as from videos found on YouTube or other online repositories. State-of-the-art estimation methods

^{*} Research presented in this paper has been supported by Dyson Technology Ltd. and the Technical University of Munich



Fig. 1: BodySLAM Problem Definition and Example Scene. The human body postures of the previous n frames are used to predict its motion; and a Human Mesh Regressor [28] is used for human centre pose, posture, and shape prediction. Then, a factor graph of camera trajectory, human states and natural landmarks are jointly optimised.

for human shape and posture involve training powerful deep neural networks (DNNs), but these methods require large amounts of annotated training data. As 3D manual annotations are extremely difficult to obtain, in-the-wild datasets are only labelled in 2D, with 3D annotations only being available in indoor settings where a motion capture system was used to produce ground truth labels. These restrictions have led to most methods only providing shape and posture estimates for individual frames and not leveraging the temporal information that is available from image sequences. Another critical limitation is that most of these training datasets only feature stationary cameras, and a true in-the-wild system will need to operate in environments where the camera is moving along an unknown trajectory. This further complicates the estimation procedure, as with a moving camera, the camera motion and human motion become conflated.

The standard method to incrementally estimate the position of a moving monocular camera in real time is with a visual simultaneous localisation and mapping (SLAM) system. While many visual SLAM methods have been shown to be both accurate and robust, image sequences containing people in motion pose a particular challenge, since most state-of-the-art visual SLAM methods assume a static environment. It is possible to use a separate network to segment and mask-out humans from the visual SLAM system, but in many sequences humans take up a large proportion of the frame leaving too few points for the SLAM system to track robustly. Importantly, even in static scenes with many observable points, classic monocular SLAM systems are not able to observe the scale of a scene, introducing ambiguity to the absolute poses of the cameras and any human motion observed by them.

With these limitations in mind, we propose a method, called BodySLAM, for jointly estimating the orientation and translation, shape, and posture of human bodies, as well as the trajectory of the moving monocular camera that observes them, in a single factor graph optimisation. We represent 3D human bodies using the shape and posture parameters of the SMPL mesh model [32], and the orientation and translation of the root of the corresponding kinematic tree underlying the SMPL model. In each frame, we obtain a measurement of these human body parameters using an off-the-shelf neural network. The body shape measurements are fused together in the factor graph optimisation as these should remain constant for a single individual over the image sequence. This constraint also helps to resolve the scale ambiguity from the monocular SLAM formulation.

Posture parameters are constrained between frames through the use of a novel motion model. The motion model uses body shape parameters and estimates of the human centre poses and postures from previous frames to form a prior on future body parameters. Unlike other models, we choose not to use a recurrent neural network architecture to allow for a simpler factor graph formulation and to enable the marginalisation of previous motions in future work.

Finally, as is standard in visual SLAM, we track and map a sparse set of feature-based landmarks in the background of the image sequences to help estimate the relative 6D camera transformations. The BodySLAM problem definition for an example scene is provided in Fig. 1.

We demonstrate through a series of quantitative experiments on image sequences featuring human motion and a moving camera that our approach of joint optimisation results in improvements to the estimates of all human body parameters and camera poses when compared to estimating these separately.³

In summary, the key contributions of this paper are:

- a monocular SLAM system, called BodySLAM, that jointly optimises human shape parameters, posture parameters, body centre poses, camera poses, and a set of 3D landmarks,
- a novel motion model that robustly predicts future human body parameters from previous ones, and
- an experimental evaluation of our joint estimation approach on our own dataset consisting of more than 8k frames with ground truth human body parameters and camera poses obtained by a Vicon motion tracking system, demonstrating that we can accurately recover human body parameters and camera trajectories at metric scale.

2 Related Work

2.1 3D Human Body Representations

3D representations for the human body generally fall into three broad categories: sparse sets of independent body joints, articulated body skeletons with rigidity constraints between joints, and (parametric) human mesh models.

Both of the first two categories involve the estimation of a set of 3D joint landmarks from 2D images. By using datasets such as Human3.6m [10] with ground truth 3D labels generated by motion capture systems, it is possible to directly regress from the input images to 3D joint positions using DNNs [40,39].

³ We encourage the reader to view our supplementary video available at: https://youtu.be/0-SL3VeWEvU

Another option is to use the abundance of 2D ground truth labels to learn to detect 2D joint keypoints [13,9], and then "lift" these to 3D using methods such as dictionary of skeletons [2,43,48] or DNNs [14,51,34]. While the direct approach tends to overfit to indoor lab environments, the lifting strategy discards a lot of potentially useful information from the captured images.

More recently, parametric 3D human body models [32,3,36] have been successfully used as output targets for human body shape and posture estimation. Human body models encapsulate statistics of body shape and other priors to reduce the ambiguity of the 3D estimation problem, and can provide high-level details such as facial expressions and hand or foot articulation [38]. It is for these advantages that we choose to use the SMPL parametric human mesh model [32] as the 3D representation for human bodies in BodySLAM.

These parametric mesh models can be fitted end-to-end using either an optimisation-based top-down approach, a regression-based bottom-up approach, or a combination of both [8,23,28]. SMPLify [8] proposes an optimisation-based fitting routine where the 3D mesh model is fitted to predicted 2D keypoints. While top-down approaches such as SMPLify require a 2-step process of keypoint prediction and mesh optimisation, regression-based bottom-up approaches start directly from pixel-level information. [23] develops a direct parametric mesh regressor trained on 2D and 3D data, predicting mesh and camera parameters from an image input. [28] uses an optimisation routine in the training loop of a deep convolutional neural network to improve the supervisory signal and quality of the predictions. In our work, we use the method of [28] to initialise the human shape and posture parameters, but then refine these estimates through our joint factor graph optimisation procedure.

2.2 Human Motion Models

While there are many approaches to estimate human motion from videos looking only at joint locations [11,18,41], our focus is on methods that use parametric human body models like SMPL.

A common approach is to enforce a smoothness prior on the posture parameters rather than define an explicit motion model. For example, [4] does this to include temporal information in their extension of the previously introduced SMPLify routine. By adding smoothness constraints on posture parameters and enforcing consistent body shape, the task of human shape and posture estimation from video is essentially formulated as a bundle adjustment problem. [19] uses multiple views of the same scene and silhouette constraints to improve single-frame fitting. While smoothness constraints help improve estimates of the human body parameters, they are agnostic to the direction of motion and do not help to estimate scale. As scale recovery is an important part of our work, we use a model to predict future body position and posture parameters from previous estimates, in which smoothness is implicitly included.

There are other examples that use DNNs to learn a human motion model [24,26,31,44]. In [24], a network learns human kinematics by prediction of past and future image frames. Recently, [26] introduced both a temporal encoder

with a gated recurrent unit (GRU) to estimate temporally consistent motion and shapes, and a training procedure with a motion discriminator that forces the network to generate feasible human motion. While such a recurrent network is capable of generating very accurate human motion predictions, including it in a factor graph optimisation would be difficult as each hidden state would need to become a state variable. Instead, we opt to use a multi-layer perceptron that takes the body shape parameters and the past n body posture estimates as input to predict the current posture parameters. This formulation allows us to greatly simplify the factor graph and will allow for much easier marginalisation of past states as part of future work.

2.3 Visual SLAM with Dynamic Objects

Most visual SLAM systems assume that the scene they are mapping and tracking against is static (e.g. ORB-SLAM [35]). Many of these systems are still capable of handling some small dynamic regions of the scene as any keypoints in these areas will be treated as outliers when using robust methods; however, they will fail if the dynamic objects occupy large segments of the frame as is often the case in image sequences capturing human motion. Some systems are designed to handle dynamic environments through the use of elimination strategies, where moving objects are detected and removed from the tracking and mapping processes (e.g. [47,7,20,5,12,21]). While this strategy offers an improvement over the static scene assumption in standard SLAM systems, it still requires that enough static background is visible for robust tracking and throws away any information contained in the dynamic parts of the environment. Alternative methods explicitly model motion with techniques such as rigid object-level maps [50,45,46,6].

More recently, there have been some approaches that estimate the camera motion and dynamic scenes simultaneously, without any a priori knowledge of the environment (e.g. [15,22,16]). While these systems are limited to estimating rigid-body motion, AirDOS [42], the work most closely related to our own, extends this line of work to include articulated objects such as humans. Air-DOS uses a combination of rigidity and motion constraints to create articulated models for humans and vehicles within a factor graph framework. By jointly optimising for the object motion, object 3D structure and the camera trajectory, improved camera tracking over an elimination strategy baseline is shown. While AirDOS presents impressive results, we are able to make a number of key improvements over their approach. Firstly, whereas AirDOS depends on stereo input, BodySLAM works with only a monocular camera, making it suitable to more "in the wild" applications. Unlike stereo SLAM systems, standard monocular SLAM systems are unable to observe the scale of the scene; however, we demonstrate that by simultaneously optimising for the human posture, shape and centre poses, BodySLAM is able to recover the metric scale. Secondly, since the focus of AirDOS is on improved camera tracking and not human motion estimation, it uses a simple skeleton-based model of human rigidity. It furthermore does not use a a learned temporal model, but uses motion constraints for the rigid body segments. BodySLAM, on the other hand, uses the SMPL parametric

mesh model, allowing for full human body reconstructions, and uses a temporal human motion model to predict translation, orientation, and posture change. Finally, while AirDOS provides quantitative camera tracking results, only a few qualitative results for the object tracking and structure estimation are provided. In our experiments on real-world datasets, we quantitatively demonstrate that BodySLAM improves not only camera tracking, but the estimation of the human body centre poses, shape and posture parameters.

3 Methodology

3.1 Preliminaries and Notation

We use the following notation throughout this work: a reference frame is denoted as $\underline{\mathcal{F}}_A$, and vectors expressed in it are written as $_A\mathbf{r}$ or translations $_A\mathbf{r}_{OP}$ with O and P as start and endpoints, respectively. The homogeneous transformation from reference frame $\underline{\mathcal{F}}_B$ to $\underline{\mathcal{F}}_A$ is denoted as \mathbf{T}_{AB} . Its rotation matrix is denoted as \mathbf{C}_{AB} , whose minimal representation we parameterise as an axisangle rotation vector in the respective frame, $_A\alpha_{AB}$. Commonly used reference frames are the static world frame $\underline{\mathcal{F}}_W$, the camera centric frame $\underline{\mathcal{F}}_C$, and the human body centric frame $\underline{\mathcal{F}}_H$. Camera frames are indexed by time steps k, and body joints by j. Measurements of quantity \mathbf{z} are denoted with a tilde, $\tilde{\mathbf{z}}$.

3.2 System Overview



Fig. 2: The BodySLAM framework.

An overview of the BodySLAM framework is provided in Fig. 2. BodyS-LAM can be split into two main components: a front-end that performs the pre-processing of the input image stream, and the back-end that performs the graph optimisation (bundle adjustment).

In the front-end, BRISK keypoints are extracted from each image for use in camera tracking [29]. The monocular bundle adjustment problem is then initialised with static landmarks and camera poses obtained by applying RANSAC on the keypoints. Each image is also processed by an instance segmentation network [49] to detect human bounding boxes, and by OpenPose [9] to extract human keypoints. From these, human mesh parameters are estimated using a Human Mesh Regressor (HMR) network [28]. Finally, human motion and posture estimates are obtained from our deep motion model (see Section 3.5).

Many state-of-the-art methods perform bundle adjustment only on keyframes, a selected subset of the total frames; however, to avoid complications to the training and inference of the motion model by having a variable frame rate, we assume a constant frame rate of 30 FPS and estimate the camera pose, human centre pose, and human shape and posture parameters at every frame.

3.3 3D Representation of the Human Body

The SMPL model [32] is a parametric human mesh model that supplies a function, $\mathcal{M}(\beta, \theta)$, of the shape parameters, β , and posture parameters, θ , and returns a mesh, ${}_{H}\mathbf{M} \in \mathbb{R}^{N \times 3}$, as a set of N = 6890 vertices in the humancentric frame \mathcal{F}_{H} . The body shape is parameterised as a 10-dimensional vector, $\beta \in \mathbb{R}^{10}$, where the components are determined by a principal component analysis to capture as much of the variability in human body shape as possible. The body posture parameters, θ , represent each of the 23 joint orientations in the respective parent frame in axis-angle representation with dimensionality $\dim(\theta) = 23 \times 3 = 69$. With this kinematic tree representation, it is easy to encode the same posture for different human shapes, as the joint positions can be defined as a linear combination of a set of mesh vertices. For J joints, a linear regressor $\mathbf{W} \in \mathbb{R}^{J \times N}$ is pre-trained to regress the J 3D body joints as ${}_{H}\mathbf{X} = \mathbf{W}_{H}\mathbf{M}$ with ${}_{H}\mathbf{X} \in \mathbb{R}^{J \times 3}$.

The 6D human centre pose is defined as $\mathbf{x}^{\text{human}} = [_W \mathbf{r}_{WH}, _W \boldsymbol{\alpha}_{WH}] \in \mathbb{R}^6$, where $_W \mathbf{r}_{WH}$ is the body centre position with respect to the static world frame, and $_W \boldsymbol{\alpha}_{WH}$ is the body centre orientation represented as a rotation vector.

3.4 Measurements of Human Body Parameters

SMPL Parameter Prediction The state-of-the-art method from Kolotouros *et al.* [28] is used to generate measurements of the SMPL parameters $\{\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\theta}}\}$ in each frame. This method uses an HMR network to regress the SMPL parameters from a normalised and cropped image of a human. The cropped human image is obtained by detecting human bounding boxes using Detectron2 [49].

Human Centre Pose Prediction We also use the network in [28] to generate measurements of the 6D transformation between the camera and human centre frames. In addition to the SMPL parameters, the HMR network also predicts a weak perspective camera model. As detailed in [17], after rectifying and removing focal distortion from the images of a calibrated camera, we can use the weak perspective camera parameters and the rotation between the camera centre and the centre of the detected bounding box to derive the transformation \tilde{T}_{CH} .

3.5 Human Motion Model

Formulation For the motion model, we use a deep neural network. Its design follows the architecture from Martinez *et al.* [34], but instead of a recurrent unit used for the temporal encoding, we use a sequential backbone only consisting of a multi-layer perceptron. The motion model uses the past *n* body posture estimates $\boldsymbol{\Theta} = [\boldsymbol{\theta}_{k-n}, \ldots, \boldsymbol{\theta}_{k-1}]$ and the (constant) body shape $\boldsymbol{\beta}$ as inputs, to predict the current body posture, $\boldsymbol{\theta}_k^{\text{pred}}$, and relative human centre translation, $H_{k-1}\mathbf{r}_{H_kH_{k-1}}^{\text{pred}}$. This prediction is used as a motion measurement in the optimisation problem discussed in Section 3.6. The model is trained using the loss:

$$\mathcal{L} = c_p + \lambda c_\theta, \tag{1}$$

where $c_p = \|_{H_{k-1}} \mathbf{r}_{H_{k-1}H_k}^{\text{true}} - _{H_{k-1}} \mathbf{r}_{H_{k-1}H_k}^{\text{pred}} \|^2$ is the error on the predicted relative position, and $c_{\theta} = \|\boldsymbol{\theta}_k^{\text{true}} - \boldsymbol{\theta}_k^{\text{pred}}\|^2$ is the error on the human body posture. The global orientation change of the human centre frame is predicted as part of the 24 body joint angles. The relative weighting between the two error terms λ is set to 1/24, weighting the position error as much as each rotation in the kinematic tree. Since the input and output values of the network were normalised with the respective means and standard deviations of the dataset, the loss values quickly converge. No robust cost function was used for the optimisation loss.

Architecture For the human posture and shape encoder, we use a 3-layer linear encoder, with an input layer of size $(n \times 3 \times 24) + 10$ for the *n* previous human body posture and shape estimates, and two fully connected layers of size 1024. We use 2 separate decoders for the human position change and human body posture, each with two fully connected layers, and a final layer of size 3 and 72, respectively. The human posture decoder has an additional residual connection to the most recent posture $\theta_k = \theta_{k-1} + \theta^{\text{pred}}$.

Training Procedure The human motion model is trained on a subset of the AMASS dataset [33], only considering sequences with walking, running, or jogging motions. Since we predict the pose change between two consecutive camera frames, we processed the data of the AMASS dataset to match the most common camera frame rate of 30 FPS. The network uses a fixed number of input frames for its prediction. To avoid overlapping sequences, we processed the dataset into chunks containing the required number of input frames and a target frame. We experimented with different sequence lengths of $n = \{2, 4, 8, 16\}$, but did not notice a drastic change in performance with longer sequences, and so trained the model with n = 2. This left us with more than 340,000 sequences, from which we reserved 20% for validation. We trained the network until convergence (100 epochs), with a batch size of 4096 on a single Nvidia GeForce GTX1080Ti GPU. We used the Adam optimiser [25] with a learning rate of 1×10^{-3} .

3.6 Factor Graph Formulation

Overview During bundle adjustment, we jointly optimise camera poses, $\mathbf{x}^{cam} = [_W \mathbf{r}'_{WC_k}, _W \boldsymbol{\alpha}_{WC_k}]$, human centre poses, $\mathbf{x}^{human} = [_W \mathbf{r}_{WH_k}, _W \boldsymbol{\alpha}_{WH_k}]$, body shape, $\boldsymbol{\beta}$, postures, $\boldsymbol{\theta}_k$, and L static landmarks, $_W \boldsymbol{l}_l$ with $l \in 1, \ldots, L$. For faster convergence, we over-parameterise the camera position as $_W \mathbf{r}_{WC_k} = s_W \mathbf{r}'_{WC_k}$, where the scale s is estimated along with the (scaled) position vector \mathbf{r}'_{WC_k} . Thus we estimate $\mathbf{x} = [s, \boldsymbol{\beta}, \mathbf{x}_1^{cam}, \mathbf{x}_1^{human}, \boldsymbol{\theta}_1, \ldots, \mathbf{x}_K^{cam}, \mathbf{x}_K^{human}, \boldsymbol{\theta}_K, _W \boldsymbol{l}_1, \ldots, _W \boldsymbol{l}_L]$.

Fig. 3 shows the respective factor graphs for three different SLAM formulations. The variables to be estimated are drawn as circles and the measurements are drawn as boxes. The measurements include 3D landmarks, OpenPose keypoint measurements and relative motion factors. The leftmost factor graph describes the traditional monocular SLAM problem without temporal constraints between the camera poses. The middle factor graph shows the naïve formulation, where human reprojection error terms are included in blue, but there are still no temporal constraints. The rightmost factor graph includes measurements from our learned motion model, creating temporal constraints between the estimated camera and human centre poses, rendering scale observable. Through our experimental results, we will demonstrate that using the rightmost factor graph leads to the most accurate estimates of both human body parameters and camera poses. Therefore, the total cost function optimised by BodySLAM is given by:

$$\begin{split} c(\mathbf{x}) &= \lambda^{\mathrm{lm}} \sum_{k=1}^{K} \sum_{l=1}^{L} \mathbf{e}_{k,l}^{\mathrm{lm}^{\mathrm{T}}} \mathbf{e}_{k,l}^{\mathrm{lm}} + \sum_{k=1}^{K} \sum_{j=1}^{J} \lambda_{k,j}^{\mathrm{joints}} \mathbf{e}_{k,j}^{\mathrm{joints}^{\mathrm{T}}} \mathbf{e}_{k,j}^{\mathrm{joints}} + \lambda^{\mathrm{mm}} \sum_{k=2}^{K} \mathbf{e}_{k}^{\mathrm{mm}^{\mathrm{T}}} \mathbf{e}_{k}^{\mathrm{mm}} \\ &+ \lambda^{\mathrm{posture}} \sum_{k=1}^{K} \mathbf{e}_{k}^{\mathrm{posture}^{\mathrm{T}}} \mathbf{e}_{k}^{\mathrm{posture}} + \lambda^{\mathrm{shape}} \sum_{k=1}^{K} \mathbf{e}_{k}^{\mathrm{shape}^{\mathrm{T}}} \mathbf{e}_{k}^{\mathrm{shape}}, \end{split}$$

where $\mathbf{e}_{k,l}^{\text{lm}}$, $\mathbf{e}_{k,j}^{\text{joints}}$, $\mathbf{e}_{k}^{\text{mm}}$, $\mathbf{e}_{k}^{\text{posture}}$, and $\mathbf{e}^{\text{shape}}$ are the error terms for the structural landmarks, human joint landmarks, human motion, body posture, and body shape, respectively. Each of these error terms is discussed in detail below. The structural landmarks are weighted by $\lambda^{\text{lm}} = 64/b^2$, where *b* is the size of the BRISK keypoint (b = 2.25 px in our implementation). The human joint landmarks are weighted by $\lambda_{k,j}^{\text{joints}} = \sigma_{k,j}^{-2}$, where $\sigma_{k,j}$ is confidence predicted by OpenPose. Finally, in our implementation, we use $\lambda^{\text{mm}} = 100$, $\lambda^{\text{posture}} = 1$, and $\lambda^{\text{shape}} = 1$, the values of which were determined experimentally.

Reprojection Error for Structural Landmarks For the structural landmarks, ${}_{W}\boldsymbol{l}_{l}$, we use the standard reprojection error term (as described in [30]):

$$\mathbf{e}_{k,l}^{\mathrm{lm}} = \tilde{\mathbf{z}}_{k,l} - \mathbf{u}(T_{WC_{k}}^{-1} \mathbf{w} \boldsymbol{l}_{l}), \tag{2}$$

where $\mathbf{u}(\cdot)$ denotes the camera projection model and $\tilde{\mathbf{z}}_{k,l}$ is the 2D keypoint measurement of the static landmark ${}_W \mathbf{l}_l$ in frame k. The measurements are BRISK keypoints filtered for outliers by 3D-2D RANSAC. To handle any remaining outliers, we use the robust Cauchy loss function on the error term.



Fig. 3: Baseline and proposed Factor Graphs of the visual SLAM problem. Without human bodies, the monocular SLAM Factor Graph has no relation between successive camera poses (left). A naive formulation of SLAM problem with human 2D keypoint measurements as dynamic keypoints fails to create a relation between successive frames (middle). Our proposed BodySLAM formulation adds a motion model that predicts human centre pose and posture, and creates a temporal constraint for camera and human center poses (right). Additionally, we overparameterise the camera pose with an additional scale factor s, to achieve faster convergence of our Bundle Adjustment.

Reprojection Error for Human Joint Landmarks As done in [8,28], we use OpenPose keypoints and their corresponding confidence values as measurements in the joint reprojection error term, employing the same standard reprojection error term as for the structural keypoints:

$$\mathbf{e}_{k,j}^{\text{joints}} = \tilde{\mathbf{z}}_{k,j} - \mathbf{u}(\mathbf{T}_{WC_k}^{-1} \mathbf{T}_{WH_k} H \mathbf{l}_{j,k}(\boldsymbol{\theta}_k, \boldsymbol{\beta})),$$
(3)

where $\tilde{\mathbf{z}}_{k,j}$ denotes the OpenPose keypoint measurement, and ${}_{B}\boldsymbol{l}_{j,k}(\boldsymbol{\theta}_{k},\boldsymbol{\beta})$ denotes the SMPL joint j in frame k expressed as homogeneous points in the human centered frame \mathcal{F}_{H} as a function of the posture $\boldsymbol{\theta}_{k}$ and shape $\boldsymbol{\beta}$. To increase robustness in particular with (self-)occluded joints and noisy keypoint detections, we use the Geman-McLure robust cost function in line with [8].

Note that the SMPL parameters $\{\theta, \beta\}$ can be directly initialised from the predictions as explained in Section 3.4. In turn, the human poses can be initialised as $T_{WH_k} = T_{WC_k} \tilde{T}_{C_k H_k}$.

Human Motion Error Our motion model is used to predict the expected human translation in the next camera frame $_{H_{k-1}}\tilde{\mathbf{r}}_{H_{k-1}H_k} := _{H_{k-1}}\mathbf{r}_{H_{k-1}H_k}^{\text{pred}}$. This prediction is used in the human motion error term:

$$\mathbf{e}_{k}^{\mathrm{mm}} = {}_{H_{k-1}}\tilde{\mathbf{r}}_{H_{k-1}H_{k}} - \mathbf{C}_{WH_{k-1}}^{T} \left({}_{W}\mathbf{r}_{WH_{k}} - {}_{W}\mathbf{r}_{WH_{k-1}} \right), \tag{4}$$

where $\mathbf{C}_{WH_{k-1}}$ denotes orientation of the human pose. One can notice the similarity of the human motion error term to the position error term in Eq. (1) used during the training of the human motion model.

Human Body Posture Error Our motion model predicts the expected human body posture in the current frame, $\tilde{\theta}_k$. The posture error is formulated as:

$$\mathbf{\hat{\theta}}_{k}^{\text{posture}} = \tilde{\boldsymbol{\theta}}_{k} - \boldsymbol{\theta}_{k}.$$
 (5)

Human Shape Error In each frame, we use an HMR network to predict the SMPL parameters $\{\tilde{\beta}_k, \tilde{\theta}_k\}$. As β should be constant over the image sequence, we constrain this to be a constant using the body shape error term:

$$\mathbf{e}_{k}^{\mathrm{shape}} = \tilde{\boldsymbol{\beta}}_{k} - \boldsymbol{\beta}. \tag{6}$$

Optimisation Procedure and Implementation Details To initialise the full BodySLAM optimisation, we perform a straightforward implementation of bundle adjustment using the BRISK keypoints, RANSAC-based pose initialisation and optimisation using the Google Ceres optimiser [1] to obtain estimates for the camera poses and landmark locations up to an arbitrary scale. After this initialisation, we use a first-order nonlinear least squares optimisation in Py-Torch [37] for bundle adjustment. However, to increase convergence speed and performance, one could easily implement a second-order optimisation.

Since the first-order optimisation is prone to converge at local minima, we use a two-step optimisation schedule often used in human mesh optimisation [4,8,28]. First, we optimise only the relative transformations between the cameras and the human body centres, the transformations between the cameras and the world frame, and the static landmarks, while keeping all other variables constant. The purpose of this optimisation step is to roughly estimate the scale of the given scene, and the camera and human centre trajectories. We optimise until convergence, finding approximately 200 iterations to be sufficient. In a second step, we optimise all variables for another 100 iterations after which we always reached convergence in our experiments.

4 Experimental Results

To evaluate BodySLAM, we collected a dataset of 10 video sequences of human motion captured by a moving monocular camera. A motion capture system was used to collect ground truth camera trajectories and human body parameters. This dataset is described in more detail in Section 4.1.

For evaluation, we consider the standard metric of average trajectory error for the camera trajectory (C-ATE), the human centre trajectory (H-ATE), and the human joint trajectories (J-ATE). Furthermore, we align the estimated camera trajectories in Sim(3), and interpret the aligned scale $s_{\text{Sim}(3)}$ as scale error. With a perfect scale estimation, this estimated scale would be 1.0.

As one of the main contributions of BodySLAM is its ability to estimate metric scale from a monocular image sequence, we first evaluate how well it can recover the scale of a scene when it is given the scale-perturbed ground truth camera trajectory. The results of this preliminary study are provided in Section 4.2. In Section 4.3, we evaluate the ability of the full BodySLAM system to estimate the camera poses, human body centre poses, and human joint positions for all sequences in our dataset. Finally, in Section 4.4, we perform an ablation study on our motion model, showing that our model is able to provide meaningful estimates of human motion and improves estimates of the trajectories.

4.1 Custom Human Pose Estimation Dataset

Current real-world evaluation datasets are lacking at least one of the following properties: a moving camera, known camera intrinsics, ground truth camera motion, a moving human subject, or ground truth human body parameters. An extensive but incomplete list of these can be found in Table 1 in the Appendix. Due to these limitations, we perform the evaluation on our own custom dataset that captures the ground truth camera poses of a moving camera with known intrinsics, along with human centre poses and posture parameters for moving humans. Ground truth values were obtained using a Vicon motion capture system. The dataset consists of 10 separate sequences, each between 600 to 1000 frames, totalling 8174 images. The dataset is split into easy, medium, and hard sequences, categorised by the speed and complexity of the camera motion.

4.2 Study of Scale Recovery Through Motion Factors

To evaluate the ability of BodySLAM to recover metric scale, we perform an ablation study where structural landmark measurements are ignored and where we initialise the camera trajectory with the Vicon ground truth poses after scaling the camera positions with an arbitrary perturbation factor between 0.1 and 5.0. We then freeze the Sim(3) transformations of the camera poses in the graph, only allowing the scale of the camera trajectory, s, to be optimised along with the human body parameters. As shown in Fig. 4, after initial perturbation by 1/s', we are able to reduce the H-ATE to approximately the value it had been before perturbation. Furthermore, we show that the scale is correctly estimated with a maximum error of 26.5% for the most extreme perturbation factor of 5.0 (as found in Table 3 in the Appendix). For smaller initial perturbations, the estimated scale deviates by only roughly 5% from the true scale.

4.3 BodySLAM: Optimisation of Camera Poses and Human States

In Table 1, we present the main results for our full BodySLAM system. As a baseline for comparison, we apply classic monocular bundle adjustment with naïve human motion tracking (i.e. only considering the unary OpenPose measurements and no motion model). This baseline, represented by the the middle factor graph in Fig. 3, also does not use the scale overparameterisation of BodySLAM.

Our results show that both the C-ATE and the H-ATE, after SE(3) alignment, are significantly lower than the baseline method. The average improvement through the introduction of the motion model is 48.3% for the root trajectory, and 35.6% for all joints combined.

Furthermore, we show that the motion model factor is able to accurately recover the scale, with scale estimation errors in the range of 0.93 to 1.03, compared to errors between 0.35 to 0.74 for the baseline, as shown in Table 1. However, one must note that the values of the scale error are completely arbitrary, since they depend only on the random initialisation of the camera poses. The C-ATE with SE(3) alignment for BodySLAM shows a huge improvement compared to **Table 1:** Comparison of absolute trajectory errors: camera (C-ATE), human (H-ATE), and joints (J-ATE) plus camera trajectory scale on test sequences for the baseline without motion model factor ("no MM") and the full BodySLAM (ours). Position errors are computed after SE(3) alignment to account for the difference in the Vicon global frame and visual world frame \mathcal{F}_W . We also aligned the scale-unaware baseline using Sim(3), which reveals that our method still outperforms it regarding human centre position and joint locations, and performs on-par regarding camera pose accuracy. Note, however, that **in the wild Sim(3) alignment is not possible**, since it requires access to the ground truth camera trajectory – illustrating the usefulness of our method.

	C-ATE [mm]				H-ATE [mm]				J-ATE [mm]				Seels Emen	
	SE(3)		Sim(3)		SE(3)		Sim(3)		SE(3)		Sim(3)		Scale Error	
Seq.	no MM	(ours)	no MM	(ours)	no MM	(ours)	no MM	(ours)	no MM	(ours)	no MM	(ours)	no MM	(ours)
E 1	108.1	44.3	42.4	43.5	114.0	56.3	106.5	56.0	154.8	118.8	149.8	115.9	0.671	0.940
E_2	193.0	50.8	50.7	49.1	178.6	55.0	182.8	53.0	210.9	124.2	214.6	120.5	0.581	1.030
E_{3}	122.3	40.9	40.8	38.9	220.8	160.9	217.1	160.4	247.7	196.1	244.3	192.8	0.738	0.981
M 1	314.7	92.6	93.3	91.0	326.9	201.4	314.6	200.0	368.8	236.2	355.3	236.2	0.502	0.978
M 2	241.6	64.0	57.9	61.1	368.1	295.4	353.5	297.1	408.6	323.3	392.0	325.4	0.507	0.936
M 3	392.1	128.3	128	128.1	311.2	121.7	305.9	122.2	335.5	170.2	326.4	171.9	0.354	0.985
M 4	203.1	92.4	91.3	90.9	153.7	70.4	142.8	71.7	184.2	123.0	175.0	125.3	0.635	0.951
D 1	310.4	108.1	107	103.5	303.7	136.7	293.9	135.9	330.6	180.6	321.1	181.1	0.445	0.951
D_2	316.0	100.0	98.9	95.2	297.6	91.2	283.5	90.6	320.6	142.6	306.8	143.8	0.430	0.933
D 3	246.9	115.3	113	114.0	235.4	149.9	228.8	149.5	264.5	183.9	258.0	185.0	0.494	0.942

the baseline results. This improvement is attributed to the correctly estimated scale, as the results after Sim(3) alignment clearly show. After Sim(3) alignment, the C-ATEs differ only by a few millimeters, which can be caused by small convergence differences during the first order optimisation.

We performed an additional experiment, ablating the necessity of the structural landmarks in the joint optimisation. For this, we performed an optimisation of the camera poses and human states, with and without structural landmarks to constrain the camera poses. The initial camera poses were randomly perturbed ($\pm 100 \text{ mm}, \pm 0.01 \text{ rad}$) and optimised until convergence. The optimisation reduced the mean camera position error from 228.96 mm after perturbation to 155.5 mm with landmarks, and 158.6 mm without landmarks (1.94% error increase), highlighting BodySLAM's versatility. The complete results are found in Table 2 in the Appendix.

4.4 Motion Model Prediction Performance

To evaluate the contribution of our motion model, we analyse the performance of the model on the test split of our preprocessed AMASS dataset.

The mean translation prediction error is 3.4mm for a mean per-frame position change of 18.4mm. The distribution of the error vector components is shown in Figure 5. Please note that in the SMPL body model, the human body coordinate frame $\underline{\mathcal{F}}_{H}$ is aligned to the sagittal axis (forward) in the z-direction, and the longitudinal axis (upward) in y-direction. Despite a relatively large standard deviation, our prediction errors are centered around zero, and, as previous results

have shown, can significantly improve the estimation of the scale and the human centre and joint trajectories when included in the graph optimisation.





Fig. 4: Average Human Trajectory Error, SE(3) aligned. Without Motion Model, the scale is not recovered.

Fig. 5: Human motion prediction error histogram in x, y, and z direction.

5 Conclusions and Future Work

We have presented BodySLAM, a novel monocular SLAM system that jointly optimises a set of 6D camera poses, 3D landmarks, human body centre poses, and human body shape and posture parameters in a factor graph formulation. To the best of our knowledge, this is the first system to jointly estimate for the configuration of a dense human body mesh model and the trajectory of a moving monocular camera. We have also introduced a human motion model to constrain sequential body postures and to help observe scale. To validate our contribution, we have collected a series of video sequences of human motion captured by a moving monocular camera, with ground truth camera trajectories and human body parameters from a motion capture system. Through a number of experiments, we have demonstrated that our joint formulation consistently results in improvements to both camera tracking and human body estimation.

In future work, we plan to extend the BodySLAM method into an incremental and real-time SLAM system. We would also like to explore more complex architectures for the human motion model and evaluate whether or not any gains in accuracy offset the increased complexity of the factor graph formulation that these architectures would require. Finally, one limitation of our method is that it assumes the intrinsic parameters of the camera are given, which may not always be the case. Some very recent work [27] has demonstrated the ability to estimate the parameters of a perspective camera from a single image. Adding the camera parameters to our factor graph formulation and initialising with this method may allow us to apply BodySLAM to true "in the wild" datasets.

References

- 1. Agarwal, S., Mierle, K., Others: Ceres Solver. http://ceres-solver.org
- Akhter, I., Black, M.J.: Pose-conditioned joint angle limits for 3D human pose reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
- Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: SCAPE: Shape Completion and Animation of People. ACM Transactions on Graphics (2005)
- Arnab, A., Doersch, C., Zisserman, A.: Exploiting temporal context for 3D human pose estimation in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Barnes, D., Maddern, W., Pascoe, G., Posner, I.: Driven to Distraction: Self-Supervised Distractor Learning for Robust Monocular Visual Odometry in Urban Environments. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2018)
- Bârsan, I.A., Liu, P., Pollefeys, M., Geiger, A.: Robust Dense Mapping for Large-Scale Dynamic Environments. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2018)
- Bescos, B., Fácil, J.M., Civera, J., Neira, J.: DynaSLAM: Tracking, Mapping and Inpainting in Dynamic Scenes. Tech. rep. (2018)
- Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it SMPL: Automatic Estimation of 3D Human Pose and Shape from a Single Image. In: Proceedings of the European Conference on Computer Vision (ECCV) (2016)
- Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2D pose estimation using part affinity fields. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Catalin Ionescu Fuxin Li, C.S.: Latent Structured Models for Human Pose Estimation. In: Proceedings of the International Conference on Computer Vision (ICCV) (2011)
- Dabral, R., Mundhada, A., Kusupati, U., Afaque, S., Sharma, A., Jain, A.: Learning 3D Human Pose from Structure and Motion. In: Proceedings of the European Conference on Computer Vision (ECCV) (2017)
- Dai, W., Zhang, Y., Li, P., Fang, Z., Scherer, S.: RGB-D SLAM in Dynamic Environments using Point Correlations. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) (2020)
- Fang, H.S., Xie, S., Tai, Y.W., Lu, C.: RMPE: Regional Multi-person Pose Estimation. In: Proceedings of the International Conference on Computer Vision (ICCV) (2017)
- Habibie, I., Xu, W., Mehta, D., Pons-Moll, G., Theobalt, C.: In the Wild Human Pose Estimation Using Explicit 2D Features and Intermediate 3D Representations. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Henein, M., Kennedy, G., Mahony, R., Ila, V.: Exploiting Rigid Body Motion for SLAM in Dynamic Environments. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2018)
- Henein, M., Zhang, J., Mahony, R., Ila, V.: Dynamic SLAM: The Need for Speed. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2020)

- 16 D. F. Henning et al.
- 17. Henning, D., Guler, A., Leutenegger, S., Zafeiriou, S.: HPE3D: Human Pose Estimation in 3D. https://github.com/dorianhenning/hpe3d (2020)
- Hossain, M.R.I., Little, J.J.: Exploiting temporal information for 3D pose estimation. In: Proceedings of the European Conference on Computer Vision (ECCV) (2017)
- Huang, Y., Bogo, F., Lassner, C., Kanazawa, A., Gehler, P.V., Romero, J., Akhter, I., Black, M.J.: Towards Accurate Marker-less Human Shape and Pose Estimation over Time. In: Proceedings of the International Conference on 3D Vision (3DV) (2017)
- Jaimez, M., Kerl, C., Gonzalez-Jimenez, J., Cremers, D.: Fast Odometry and Scene Flow from RGB-D Cameras based on Geometric Clustering. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2017)
- Ji, T., Wang, C., Xie, L.: Towards Real-Time Semantic RGB-D SLAM in Dynamic Environments. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2021)
- Judd, K.M., Gammell, J.D., Newman, P.: Multimotion Visual Odometry (MVO): Simultaneous Estimation of Camera and Third-Party Motions. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (2018)
- Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end Recovery of Human Shape and Pose. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3D Human Dynamics from Video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings (2015)
- Kocabas, M., Athanasiou, N., Black, M.J.: VIBE: Video Inference for Human Body Pose and Shape Estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Kocabas, M., Huang, C.H.P., Tesch, J., Müller, L., Hilliges, O., Black, M.J.: SPEC: Seeing People in the Wild with an Estimated Camera. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021)
- Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to Reconstruct 3D Human Pose and Shape via Model-fitting in the Loop. In: Proceedings of the International Conference on Computer Vision (ICCV) (2019)
- Leutenegger, S., Chli, M., Siegwart, R.Y.: BRISK: Binary Robust invariant scalable keypoints. In: Proceedings of the International Conference on Computer Vision (ICCV) (2011)
- Leutenegger, S., Lynen, S., Bosse, M., Siegwart, R., Furgale, P.: Keyframe-based visual-inertial odometry using nonlinear optimization. International Journal of Robotics Research (IJRR) (2015)
- Ling, H.Y., Zinno, F., Cheng, G., van de Panne, M.: Character controllers using motion vaes. ACM Transactions on Graphics (2020)
- 32. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM Transactions on Graphics (2015)
- Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: AMASS: Archive of Motion Capture as Surface Shapes. In: Proceedings of the International Conference on Computer Vision (ICCV) (2019)

- Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3d human pose estimation. In: Proceedings of the International Conference on Computer Vision (ICCV) (2017)
- Mur-Artal, R., Tardos, J.D.: ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras. IEEE Transactions on Robotics (2017)
- Osman, A.A.A., Bolkart, T., Black, M.J.: STAR: Sparse Trained Articulated Human Body Regressor. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
- Paszke, A., Chanan, G., Lin, Z., Gross, S., Yang, E., Antiga, L., Devito, Z.: Automatic differentiation in PyTorch. In: Neural Information Processing Systems (NIPS) (2017)
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive Body Capture: 3D Hands, Face, and Body from a Single Image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Pavlakos, G., Zhou, X., Daniilidis, K.: Ordinal Depth Supervision for 3D Human Pose Estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Pavlakos, G., Zhou, X., Derpanis, K.G., Daniilidis, K.: Coarse-to-fine volumetric prediction for single-image 3D human pose. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
- Pavllo, D., Feichtenhofer, C., Grangier, D., Auli, M.: 3D human pose estimation in video with temporal convolutions and semi-supervised training. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- 42. Qiu, Y., Wang, C., Wang, W., Henein, M., Scherer, S.: AirDOS: Dynamic SLAM Benefits from Articulated Objects. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2022)
- Ramakrishna, V., Kanade, T., Sheikh, Y.: Reconstructing 3D Human Pose from 2D Image Landmarks. In: Proceedings of the European Conference on Computer Vision (ECCV) (2012)
- Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., Guibas, L.J.: HuMoR: 3D Human Motion Model for Robust Pose Estimation. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021)
- Rünz, M., Agapito, L.: Co-Fusion: Real-Time Segmentation, Tracking and Fusion of Multiple Objects. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2017)
- Rünz, M., Buffier, M., Agapito, L.: MaskFusion: Real-Time Recognition, Tracking and Reconstruction of Multiple Moving Objects. In: Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR) (2018)
- Scona, R., Jaimez, M., Petillot, Y.R., Fallon, M., Cremers, D.: StaticFusion: Background Reconstruction for Dense RGB-D SLAM in Dynamic Environments. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2018)
- Valmadre, J., Lucey, S.: Deterministic 3D human pose estimation using rigid structure. In: Proceedings of the European Conference on Computer Vision (ECCV) (2010)
- Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. https://github.com/facebookresearch/detectron2 (2019)

- 18 D. F. Henning et al.
- Xu, B., Li, W., Tzoumanikas, D., Bloesch, M., Davison, A., Leutenegger, S.: MID-Fusion: Octree-based Object-Level Multi-Instance Dynamic SLAM. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2019)
- Zhao, R., Wang, Y., Martinez, A.: A Simple, Fast and Highly-Accurate Algorithm to Recover 3D Shape from 2D Landmarks on a Single Image. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) (2016)