

SOMA: Solving Optical Marker-Based MoCap Automatically

Nima Ghorbani Michael J. Black

Max Planck Institute for Intelligent Systems, Tübingen, Germany

{ngorbani, black}@tuebingen.mpg.de

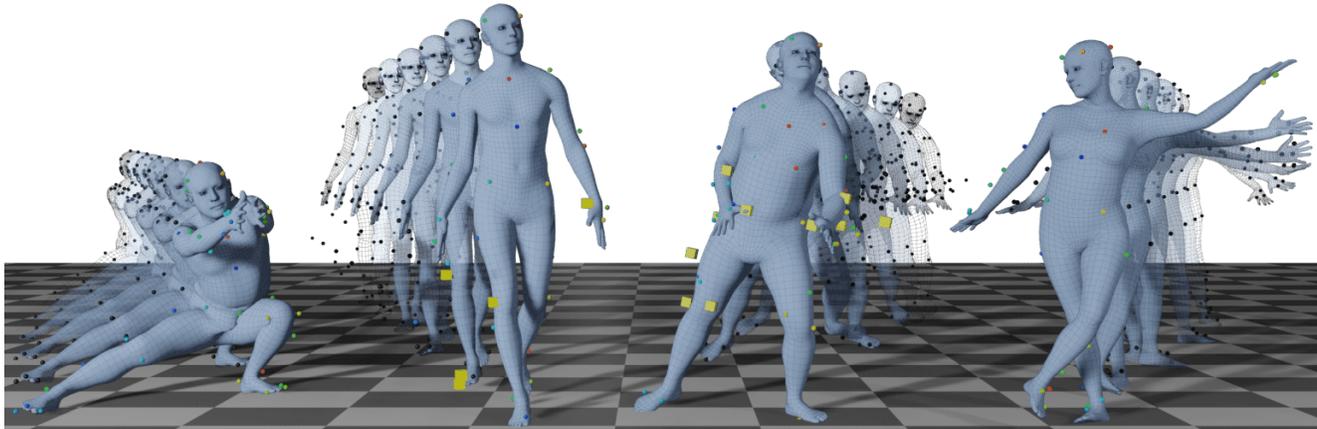


Figure 1: SOMA transforms raw motion capture (mocap) point clouds (black dots) to labeled markers (colored dots). The cubes in yellow are detected “ghost” points; e.g. spurious reflections, non-body markers, or unidentified points. With labeled mocap we fit SMPL-X bodies (blue mesh) using MoSh [26, 28].

Abstract

Marker-based optical motion capture (mocap) is the “gold standard” method for acquiring accurate 3D human motion in computer vision, medicine, and graphics. The raw output of these systems are noisy and incomplete 3D points or short tracklets of points. To be useful, one must associate these points with corresponding markers on the captured subject; i.e. “labelling”. Given these labels, one can then “solve” for the 3D skeleton or body surface mesh. Commercial auto-labeling tools require a specific calibration procedure at capture time, which is not possible for archival data. Here we train a novel neural network called SOMA, which takes raw mocap point clouds with varying numbers of points, labels them at scale without any calibration data, independent of the capture technology, and requiring only minimal human intervention. Our key insight is that, while labeling point clouds is highly ambiguous, the 3D body provides strong constraints on the solution that can be exploited by a learning-based method. To enable learning, we generate massive training sets of simulated noisy and ground truth mocap markers animated by 3D bodies from AMASS. SOMA exploits an architecture with stacked self-attention elements to learn the spatial structure of the 3D body and an optimal transport layer to constrain the assignment (labeling) problem while rejecting outliers. We extensively evaluate SOMA both quantitatively and quali-

tatively. SOMA is more accurate and robust than existing state of the art research methods and can be applied where commercial systems cannot. We automatically label over 8 hours of archival mocap data across 4 different datasets captured using various technologies and output SMPL-X body models. The model and data is released for research purposes at <https://soma.is.tue.mpg.de/>.

1. Introduction

Marker-based optical motion capture (mocap) systems record 2D infrared images of light reflected or emitted by a set of markers placed at key locations on the surface of a subject’s body. Subsequently, the mocap systems recover the precise position of the markers as a sequence of sparse and unordered points or short tracklets. Powered by years of commercial development, these systems offer high temporal and spatial accuracy. Richly varied mocap data from such systems is widely used to train machine learning methods in action recognition, motion synthesis, human motion modeling, pose estimation, etc. Despite this, the largest existing mocap dataset, AMASS [28], has about 45 hours of mocap, much smaller than video datasets used in the field.

Mocap data is limited since capturing and processing it is expensive. Despite its value, there are large amounts of archival mocap in the world that have never been labeled;

this is the “dark matter” of mocap. The problem is that, to solve for the 3D body, the raw mocap point cloud (MPC) must be “labeled”; that is, the points must be assigned to physical “marker” locations on the subject’s body. This is challenging because the MPC is noisy and sparse and the labeling problem is ambiguous. Existing commercial tools, e.g. [21, 27], offer partial automation, however none provide a complete solution to automatically handle variations in marker layout, i.e. number of markers used and their rough placement on the body, variation in subject body shape and gender, and variation across capture technologies namely active vs passive markers or brands of system. These challenges typically preclude cost-effective labeling of archival data, and add to the cost of new captures by requiring manual cleanup.

Automating the mocap labeling problem has been examined by the research community [14, 16, 19]. Existing methods focus on fixing the mistakes in already labeled markers through denoising [8, 19]. Recent work formulates the problem in a matching framework, directly predicting the label assignment matrix for a fixed number of markers in a restricted setup [14]. In short, the existing methods are limited to a constrained range of motions [14], a single body shape [8, 16, 19], a certain capture scenario, a special marker layout, or require a subject-specific calibration sequence [14, 21, 27]. Other methods require high-quality real mocap marker data for training, effectively prohibiting their scalability to novel scenarios [8, 14].

To address these shortcomings we take a data-driven approach and train a neural network end-to-end with self-attention components and an optimal transport layer to predict a per-frame constrained inexact matching between mocap points and labels. Having enough “real” data for training is not feasible, therefore we opt for synthetic data. Given a marker layout, we generate synthetic mocap point clouds with realistic noise, and then train a layout-specific network that can cope with realistic variations across a whole mocap dataset. While previous works have exploited synthetic data [16, 19], they are limited in terms of body shapes, motions, marker layouts, and noise sources.

Even with a large synthetic corpus of MPC, labeling a cloud of sparse 3D points, containing outliers and missing data, is a highly ambiguous task. The key to the solution lies in the fact that the points are structured, as is their variation with articulated pose. Specifically, they are constrained by the shape and motion of the human body. Given sufficient training data, our attentional framework learns to exploit local context at different scales. Furthermore, if there were no noise, the mapping between labels and points would be one-to-one. We formulate these concepts as a unified training objective enabling end-to-end model training. Specifically, our formulation exploits a transformer architecture to capture local and global contextual information using

self-attention (Sec. 4.1). By generating synthetic mocap data with varying body shapes and poses, SOMA implicitly learns the kinematic constraints of the underlying deformable human body (Sec. 4.4). A one-to-one match between 3D points and markers, subject to missing and spurious data, is achieved by a special normalization technique (Sec. 4.2). To provide a common output framework, consistent with [28], we use MoSh [26, 28] as a post-processing step to fit SMPL-X [34] to the labeled points; this also helps deal with missing data caused by occlusion or dropped markers. The SOMA system is outlined in Fig. 3.

To generate training data, SOMA requires a rough marker layout that can be obtained by a single labeled frame, which requires minimal effort. Afterward, virtual markers are automatically placed on a SMPL-X body and animated by motions from AMASS [28]. In addition to common mocap noise models like occlusions [14, 16, 19], and ghost points [16, 19], we introduce novel terms to vary marker placement on the body surface and we copy noise from real marker data in AMASS (Sec. 4.4). We train SOMA once for each mocap dataset and apart from the one layout frame, we do not require any labeled real data. After training, given a noisy MPC *frame* as input, SOMA predicts a distribution over labels of each point, including a null label for ghost points.

We evaluate SOMA on several challenging datasets and find that we outperform the current state of the art numerically while being much more general. Additionally, we capture new MPC data using a Vicon mocap system and compare hand-labeled ground-truth to Shōgun and SOMA output. SOMA performs similarly compared with the commercial system. Finally, we apply the method on archival mocap datasets: Mixamo [10], DanceDB [4], and a previously unreleased portion of the CMU mocap dataset [11].

In summary, our main contributions are: (1) a novel neural network architecture exploiting self-attention to process sparse deformable point cloud data; (2) a system that consumes mocap point clouds directly and outputs a distribution over marker labels; (3) a novel synthetic mocap generation pipeline that generalizes to real mocap datasets; (4) a robust solution that works with archival data, different mocap technologies, poor data quality, and varying subjects and motions; (5) 220 minutes of processed mocap data in SMPL-X format, trained models, and code are released for research purposes.

2. Related Work

Learning to Process MoCap Data was first introduced by [45] in a limited scenario. More recently, [14] proposes a learning-based model for mocap labeling that directly predicts permutations of 44 input markers. The number of possible permutations is prohibitive, hence, the authors restrict them to a limited pool shown to the network during train-

ing and test time. Moreover, motions are restricted to four categories: walk, jog, jump and sit. Furthermore, [14] inherently cannot deal with ghost points. We compare directly with them and find that we are more accurate, while removing the limitations.

Solutions exist that “denoise” the possible incorrect labels of already labeled mocap markers [8, 19]. These approaches normalize the markers to a standard body size, and rely on fragile heuristics to remove ghost points, and must first compute the global orientation of the body. Our method starts one step earlier, with *unlabeled point clouds*, and outputs a fully labeled sequence, while learning to reject ghost points and dealing with varied body shapes.

Deep Learning on Point Cloud Data requires a solution to handle a variable number of unordered points, and a way to define a notion of “neighborhood”. To address these issues [13, 47] project 3D points into 2D images and [30, 55] rasterize the point cloud into 3D voxels to enable the use of traditional convolution operators. Han et al. [16] utilize this idea for labeling hand MPC data, by projecting them into multiple 2D depth images. Using these, they train a neural network to predict the 3D location of 19 known markers on a fixed-shape hand, assigning the label of a marker to a point closest to it in a disconnected graph matching step. In contrast, our pipeline directly works with mocap point clouds and predicts a distribution over labels for each point, end-to-end, without disconnected stages.

PointNet methods [7, 36] also process the 3D point cloud directly, while learning local features with permutation-invariant pooling operators. Further non-local networks [54] and self-attention-based [52] models can attend globally while learning to focus locally on specific regions of the input. This simple formulation enables learning robust features on sparse point clouds while being insensitive to variable numbers of points. SOMA is a novel application of this idea to mocap data. In Sec. 4.1 we demonstrate that by stacking multiple self-attention elements, SOMA can learn rich point features at multiple scales that enable robust, permutation-invariant, mocap labeling.

Inexact Graph Matching formulates the problem of finding an assignment between nodes of a model graph to data nodes, where the former occasionally has fewer elements than the latter. This is an NP-hard problem [1] in general and challenging for mocap due to occlusions and spurious data points. Such graph matching problems appear frequently in computer vision and are addressed either with engineered costs [5, 25, 49] or learned ones [6, 40]. Ghorbani et al. [14] apply this to mocap using an approximate solution with Sinkhorn normalization [2, 44] by assuming that the graph of a mocap frame is isomorphic to the graph of the labels. We relax this assumption by considering an *inexact match* between the labels and points, by opting for an *optimal transport* [53] solution.

Body Models, in the form of skeletons, are widely used to constrain the labeling problem and to solve for body pose [9, 18, 33, 37, 38, 41, 41]. Recently [24, 46] employ variants of Kalman filtering to estimate constrained configuration of a given body skeleton in real-time but are susceptible to occlusions and ghost points. The most mature commercial solution to-date is Shōgun [27], which can produce real-time labeling and skeletal solving. While this is an excellent product, out-of-the-box it works only with a Vicon-specific marker layout and requires subject-specific session calibration. Thus it is not a general solution and cannot be used on most archival mocap data or with customized marker layouts needed in many applications. MoSh [26, 28], goes beyond simple skeleton models and utilizes a realistic 3D body surface model learned from a large corpus of body scans [23, 34, 56]. It takes labeled mocap markers as well as a rough layout of their placement on the body, and solves for body parameters and the precise placement of markers on the body. We employ MoSh for post-processing auto-labeled mocap across different datasets into a unified SMPL-X representation.

3. The MoCap Labeling Problem

A mocap point cloud, MPC, is a time sequence with T frames of 3D points

$$\text{MPC} = \{P_1, \dots, P_T\}, \quad (1)$$

$$P_t = \{P_{t,1}, \dots, P_{t,n_t}\}, \quad P_{t,i} \in \mathbb{R}^3, \quad (2)$$

where $|P_t| = n_t$ for each time step $t \in \{1 : T\}$. We visualize an MPC as a chart in Fig. 2 (top), where each row holds points reconstructed by the mocap hardware, and each column represents a frame of MPC. Each point is unlabeled but these can often be locally tracked over short time intervals, illustrated by the gray bars in the figure. Note that some of these tracklets may correspond to noise or “ghost points”. For passive marker systems like Vicon [27], a point that is occluded typically appears in a new row; i.e. with a new ID. For active marker systems like PhaseSpace [22], one can have gaps in the trajectories due to occlusion. The figure shows both types.

The goal of mocap labeling is to assign each point (or tracklet) to a corresponding marker label

$$L = \{l_1, \dots, l_M, \text{null}\}, \quad (3)$$

in the marker layout as illustrated in Fig. 2 (middle), where each color is a different label. The set of marker labels include an extra null label for points that are not valid markers, hence $|L| = M + 1$. These are shown as red in the figure. Valid point labels and tracklets of them are subject to several constraints: (C_i) each point $P_{t,i}$ can be assigned to at most one label and vice versa; (C_{ii}) each point $P_{t,i}$ can be assigned to at most one tracklet; (C_{iii}) the label null is an

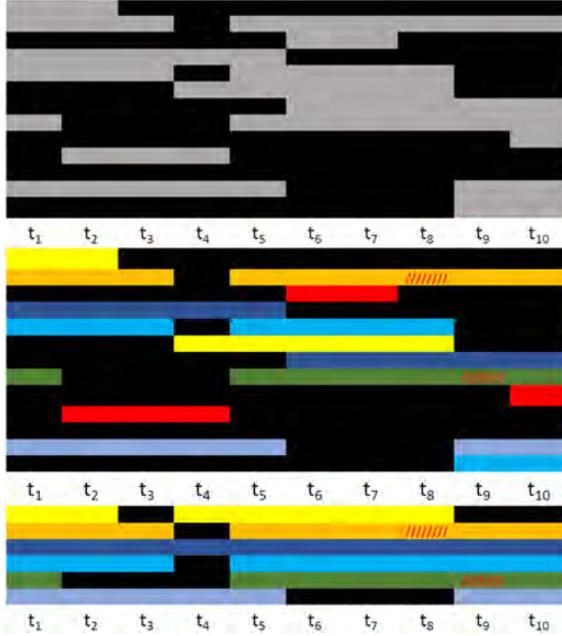


Figure 2: MoCap labeling problem. (top) Raw, unlabeled, MoCap Point Cloud (MPC). Each column represents a timestamp in a mocap sequence and each cell is a 3D point or a short tracklet (shown as a gray row). (middle) shows the MPC after labeling. Colors correspond to different labels in the marker layout. Red corresponds to ghost points (outliers). The red oblique lines show ghost points wrongly tracked as actual markers by the mocap system. (bottom) shows the final result, with the tracklets glued together to form full trajectories with only valid markers retained. Note that marker occlusion results in black (missing) sections.

exception that can be matched to more than one point and can be present in multiple tracklets in each frame.

4. SOMA

4.1. Self-Attention on MoCap Point Clouds

The SOMA system pipeline is summarized in Fig. 3. The input to SOMA is a single frame of sparse, unordered, points, the cardinality of which varies with each timestamp due to occlusions and ghost points. To process such data, we exploit multiple layers of self-attention [52], with a multi-head formulation, concatenated via residual operations [17, 52]. Multiple layers increase the capacity of the model and enable the network to have a local and a global view of the point cloud, which helps disambiguate points.

We define self-attention span as the average of attention weights over random sequences picked from our validation dataset. Figure 4 visualizes the attention placed on markers at the first and last self-attention layers; the intensity of the red color correlates with the amount of attention. Note that the points are shown on a body in a canonical pose but the

actual mocap point cloud data is in many poses. Deeper layers focus attention on geodesically near body regions (wrist: upper and lower arm) or regions that are highly correlated (left foot: right foot and left knee), indicating that the network has figured out the spatial structure of the body and correlations between parts even though the observed data points are non-linearly transformed in Euclidean space by articulation. In Sup. Mat., we provide further computational details and demonstrate the self-attention span as a function of network depth. Also, we present a model-selection experiment to choose the optimum number of layers.

4.2. Constrained Point Labeling

In the final stage of the architecture, SOMA predicts a non-square score matrix $S \in \mathbb{R}^{n_t \times M}$. To satisfy the constraints C_i and C_{ii} , we employ a log-domain, stable implementation of optimal transport [40] described by [35]. The optimal transport layer depends on iterative Sinkhorn normalization [2, 43, 44], which constrains rows and columns to sum to 1 for *available* points and labels. To deal with missing markers and ghost points, following [40], we introduce *dustbins* by appending an extra last row and column to the score matrix. These can be assigned to multiple *unmatched* points and labels, hence respectively summing to n_t and M . After normalization, we reach the augmented assignment matrix, $A' \in [0, 1]^{(n_t+1) \times |L|}$, from which we drop the appended row, for unmatched labels, yielding the final normalized assignment matrix $A \in [0, 1]^{n_t \times |L|}$.

While Ghorbani et al. [14] use a similar score normalization approach, their method cannot handle unmatched cases, which is critical to handle real mocap point cloud data, in its raw form.

4.3. Solving for the Body

Once mocap points are labeled, we “solve” for the articulated body that lies behind the motion. Typical mocap solvers [19, 21, 27] fit a skeletal model to the labeled markers. Instead, here we fit an entire articulated 3D body mesh to markers using MoSh [26, 28]. This technique gives an animated body with a skeletal structure so nothing is lost over traditional methods while yielding a full 3D body model, consistent with other recent mocap datasets [28, 48]. Here we fit the SMPL-X body model, which provides forward compatibility for datasets with hands and face captures. For more details on MoSh we refer the reader to the original paper [26, 28].

4.4. Synthetic MoCap Generation

Human Body Model. To synthetically produce realistic mocap training data with ground truth labels, we leverage a gender-neutral, state of the art statistical body model, SMPL-X [34], that uses vertex-based linear blend skinning with learned corrective blend shapes to output the global

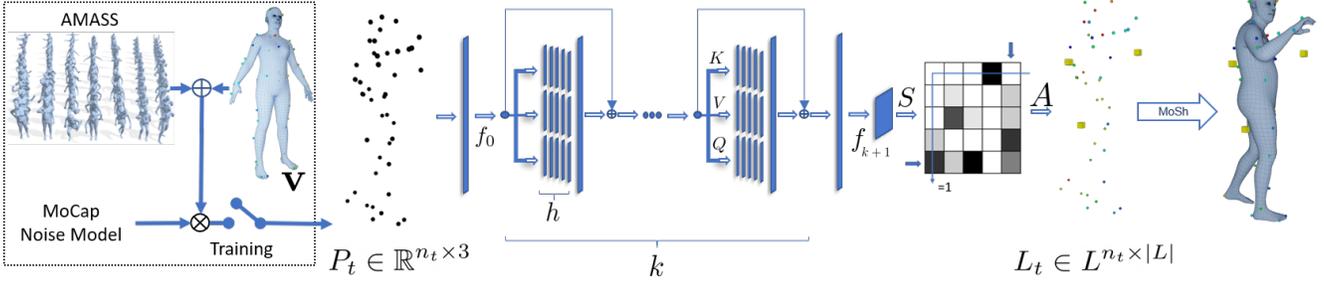


Figure 3: We train SOMA solely with synthetic data, Sec. 4.4. At runtime, SOMA receives an unprocessed 3D sparse mocap point cloud, P_t , with a varying number of points. These are median centered and passed through the pipeline, consisting of self-attention layers, Sec. 4.1, and a final normalization to encourage bijective label-point correspondence, Sec. 4.2. The network outputs labels, L_t , assigned to each point, that correspond to markers in the training marker layout, \mathbf{v} , with an additional null label. Finally, a 3D body is fit to the labeled points using MoSh, Sec. 4.3. The dimensionality of the features are $\{K, V, Q\} \in \mathbb{R}^{n_t \times \frac{d_{model}}{h}}$, $f_0 \in \mathbb{R}^{n_t \times d_{model}}$, $f_{k+1} \in \mathbb{R}^{n_t \times 256}$, $A \in \mathbb{R}^{n_t \times |L|}$.

position of $|V| = 10,475$ vertices:

$$\text{SMPL-X}(\theta_b, \beta, \gamma) : \mathbb{R}^{|\theta_b| \times |\beta| \times |\gamma|} \rightarrow \mathbb{R}^{3N}. \quad (4)$$

Here $\theta_b \in \mathbb{R}^{3(J+1)}$ is the axis-angle representation of the body pose where $J = 21$ is the number of body joints of an underlying skeleton in addition to a root joint for global rotation. We use $\beta \in \mathbb{R}^{10}$ and $\gamma \in \mathbb{R}^3$ to respectively parameterize body shape and the global translation. Compared to the original SMPL-X notation, here we discard parameters that control facial expressions, face and hand poses; i.e. respectively ψ, θ_f, θ_h . We build on SMPL-X to enable extension of SOMA to datasets with face and hand markers but SMPL-X can be converted to SMPL if needed. For more details we refer the reader to [34].

MoCap Noise Model. Various noise sources can influence mocap data, namely: subject body shape, motion, marker layout and the exact placement of the markers on body, occlusions, ghost points, mocap hardware intrinsics, and more. To learn a robust model, we exploit AMASS [28] that we refit with a neutral gender SMPL-X body model and sub-sample to a unified 30 Hz. To be robust to subject body shape we generate AMASS motions for 3664 bodies from the CAESAR dataset [39]. Specifically, for training we take parameters from the following mocap sub-datasets of AMASS: CMU [11], Transitions [28] and Pose Prior [3]. For validation we use HumanEva [42], ACCAD [12], and TotalCapture [51].

Given a marker layout of a target dataset, \mathbf{v} , as a vector of length M in which the index corresponds to the marker label and the entry to a vertex on the SMPL-X body mesh, together with a vector of marker-body distances \mathbf{d} we can place virtual markers, $X \in \mathbb{R}^{M \times 3}$ on the body:

$$X = \text{SMPL-X}(\theta_b, \beta, \gamma) |_{\mathbf{v}} + \mathbf{N}|_{\mathbf{v}} \odot \mathbf{d}. \quad (5)$$

Here $\mathbf{N} \in \mathbb{R}^{V \times 3}$ is a matrix of vertex normals and $|_{\mathbf{v}}$ picks the vector of elements (vertices or normals) corresponding to vertices defined by the marker layout.

With this, we produce a library of mocap frames and corrupt them with various controllable noise sources. Specifically, to generate a noisy layout, we randomly sample a vertex in the 1-ring neighborhood of the original vertex specified by the marker layout, effectively producing a different marker placement, $\tilde{\mathbf{v}}$, for each data point. Instead of normalizing the global body orientation, common to previous methods [8, 14, 16, 19], we add a random rotation $r \in [0, 2\pi]$ to the global root orientation of every body frame to augment this value. Further, we copy the noise for each label from the real AMASS mocap markers to help generalize to mocap hardware differences. We create a database of the differences between the simulated and actual markers of AMASS and draw random samples from this noise model to add to the synthetic marker positions.

Furthermore, we append ghost points to the generated mocap frame, by drawing random samples from a 3D Gaussian distribution with mean and standard deviation equal to the median and standard deviation of the marker positions, respectively. Moreover, to simulate marker occlusions we take random samples from a uniform distribution representing the index of the markers and occlude selected markers by replacing their value with zero. The number of added ghost points and occlusions in each frame can also be subject to randomness.

At test time, to mimic broken trajectories of passive mocap systems, we randomly choose a trajectory and break it at random timestamps. To break a trajectory we copy marker values at the onset of the breakage and create a new trajectory whose previous values up-to-the breakage are zero and the rest are replaced by the marker of interest. The original marker trajectory after breakage is replaced by zeros.

Finally, at train and test times we randomly permute the markers to create an unordered set of 3D mocap points. In contrast to [14], the permutations are random and not limited to a specific set of permutations.

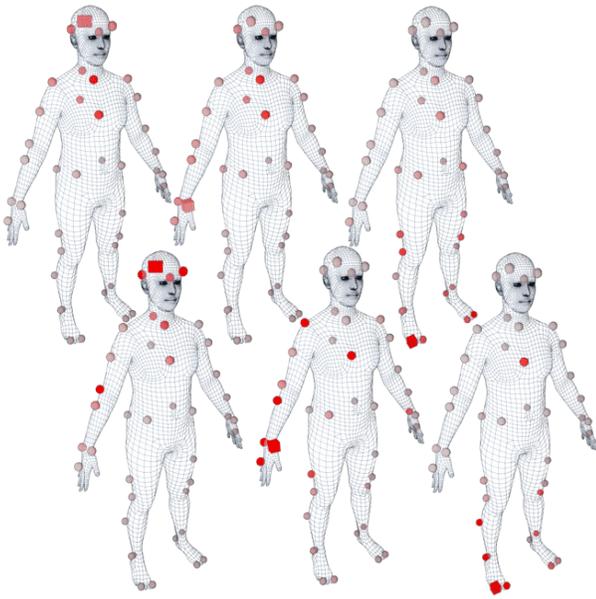


Figure 4: Attention span for different markers on the body in a canonical pose. The cube shows the marker of interest and color intensity depicts the average value of attention across frames of 50 randomly selected sequences. Each column shows a different marker. At the first layer (top) we see wider attention compared to the deepest layer (bottom).

4.5. Implementation Details

Loss. The total loss for training SOMA is formulated as, $\mathbb{L} = c_l \mathbb{L}_A + c_{reg} \mathbb{L}_{reg}$, where:

$$\mathbb{L}_A = \frac{-1}{\sum_{i,j} G'_{i,j}} \sum_{i,j} W_{i,j} \cdot G'_{i,j} \cdot \log(A'_{i,j}), \quad (6)$$

$$\mathbb{L}_{reg} = \|\phi\|_2^2. \quad (7)$$

A' is the augmented assignment matrix, and G' is its ground-truth version. W is a matrix to down-weight the influence of the over-represented class, i.e. the null label, by the reciprocal of its occurrence frequency. \mathbb{L}_{reg} is L_2 regularization on the model parameters. For further architecture details, please see the Sup. Mat..

Using SOMA. The automatic labeling pipeline starts with a labeled mocap frame that can roughly resemble the marker layout of the target dataset. If the dataset has significant variations in marker layout or many displaced or removed markers, one labeled frame per each major variation is needed. We then train one model for the whole dataset. After training with synthetic data produced for the target marker layout, we apply SOMA on mocap sequences in a *per-frame* mode; i.e. each frame is processed independently. On a GPU, auto-labeling runs at $52 \pm 12Hz$ in non-batched mode and, for a batch of 30 frames runtime is $1500 \pm 400Hz$. In cases where the mocap hardware pro-

Train \ Test	B		B+C		B+G		B+C+G	
	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1
B	97.93	97.37	83.55	81.11	86.54	85.97	73.95	71.93
B+C	97.25	96.31	97.21	96.05	95.37	95.29	93.79	92.77
B+G	98.06	97.19	96.14	94.03	97.87	97.65	95.27	93.62
B+C+G	95.74	94.44	95.56	93.91	95.73	95.22	95.50	94.44

Table 1: Per-frame labeling with synthetic noise model evaluated on real mocap markers of HDM05 with added noise. “B, C, G” respectively stand for Base, Occlusion, and Ghost points. We report average accuracy and F1 scores.

vides tracklets of points, we assign the most frequent label for a tracklet to all of the member points; we call this *tracklet labeling*. For detailed examples of using SOMA, including a general model to facilitate labeling the initial frame, i.e. “label priming”, see Sup. Mat.

5. Experiments

Evaluation Datasets. We evaluate SOMA quantitatively on various mocap datasets with real marker data and synthetic noise; namely: BMLrub [50], BMLmovi [15], and KIT [29]. The individual datasets offer various marker layouts with different marker density, subject shape variation, body pose, and recording systems. We take original marker data from their respective public access points and further corrupt the data with controlled noise, namely marker occlusion, ghost points, and per-frame random shuffling (Sec. 4.4). For per-frame experiments, broken trajectory noise is not used. We also collect a new “SOMA dataset”, which we use for direct comparison with Vicon.

To avoid overfitting hyper-parameters to test datasets, we utilize a separate dataset for model selection and validation experiments; namely HDM05 [32], containing 215 sequences, across 4 subjects, on average using 40 markers.

Evaluation Metrics. Primarily, we report mean and standard deviation of per-frame accuracy and F1 score in percentages. Accuracy is the proportion of correctly predicted labels over all labels, and F1 score is regarded as the harmonic-average of the precision and recall:

$$F1 = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}, \quad (8)$$

where recall is the proportion of correct predicted labels over actual labels and precision is regarded as the proportion of actual correct labels over predicted labels. The final F1 score for a mocap sequence is the average of the per-frame F1 scores.

5.1. Effectiveness of the MoCap Noise Generation

Here we train and test SOMA with various amounts of synthetic noise. The training data is synthetically produced for the layout of HDM05 as described in Sec. 4.4. We

Method	Number of Exact Per-Frame Occlusions						
	0	1	2	3	4	5	5+G
Holzreiter et al. [20]	88.16	79.00	72.42	67.16	61.13	52.10	—
Maycock et al. [31]	83.19	79.35	76.44	74.91	71.17	65.83	—
Ghorbani et al. [14]	97.11	96.56	96.13	95.87	95.75	94.90	—
SOMA-Real	99.08	98.97	98.85	98.68	98.48	98.22	98.29
SOMA-Synthetic	99.16	98.92	98.54	98.17	97.61	97.07	95.13
SOMA*	98.38	98.28	98.17	98.03	97.86	97.66	97.56

Table 2: Comparing SOMA with prior work on the same data. We train SOMA in three different ways: once with *real* data; once with *synthetic* markers placed on bodies of the same motions obtained from AMASS; and ultimately once with the data produced in Sec. 4.4, designated with *. Refer to Sup. Mat. for standard deviations.

test on original markers of HDM05 corrupted with synthetic noise. B stands for no noise, $B+C$ for up-to 5 markers occluded per-frame, $B+G$ for up-to 3 ghost points per-frame, and $B+C+G$ for the full treatment of the occlusion and ghost point noise. Table 1 shows different training and testing scenarios. In general, matching the model training noise to the true noise produces the best results but training with the full noise model ($B+C+G$) gives competitive performance and is a good choice when the noise level is unknown. Training with more noise improves robustness.

5.2. Comparison with Previous Work

We compare SOMA to prior work in Tab. 2 under the same conditions. Specifically, we use train and test data splits of the BMLrub dataset explained by [14]. The test happens on real markers with synthetic noise. We train SOMA once with real marker data and once by synthetic markers produced by motions of the same split. Additionally, we train SOMA with the full synthetic data outlined in Sec. 4.4. The performance of other competing methods is as reported by [14]. All versions of SOMA outperform previous methods. The model (SOMA*) trained with synthetic markers and varied training poses from AMASS is competitive with the models trained on limited real data or synthetic marker data with dataset-specific motions. This is likely due to the rich variation in our noise generation pipeline. In contrast to prior work, SOMA is robust to increasing occlusion and can process ghost points without extra heuristics; i.e. last column in Tab. 2.

5.3. Performance Robustness

Performance on Various MoCap Datasets could vary due to variations in the marker density, mocap quality, subject shape and motions. To assess such cases we take three full scale mocap datasets and corrupt the markers with synthetic noise including up to 50 broken trajectories that best mimic the situation with a realistic unlabeled mocap scenario. Additionally we evaluate tracklet labeling explained in Sec. 4.5. Table 3 shows consistent high performance of SOMA across the datasets. When tracklets are available,

as they often are, tracklet labeling improves performance across the board.

Performance on Subsets or Supersets of a Specific Marker Layout could vary since this introduces “structured” noise. A superset marker layout is the set of all labels in a dataset, which may combine multiple marker layouts. A *base model* trained on a superset marker layout and tested on subsets would be subject to structured occlusions, while a model trained on subset and tested on the superset *base mocap* would see structured ghost points. These situations commonly happen across a dataset when trial coordinators improvise on the planned marker layout for a special take with additional or fewer markers. Alternatively, markers often fall off during the trial. To quantitatively determine the range of performance variance we take the marker layout of the validation dataset, HDM05, and omit markers in progressive steps; Tab. 4. The model, trained on subset layout and tested on base markers (superset), shows greater deterioration in performance than the base model trained on the superset and tested on reduced marker sets.

5.4. Ablation Studies

Table 5 shows effect of various components in the final performance of SOMA on the validation dataset, HDM05. The self-attention layers and the novel random marker placement noise play the most significant role in overall performance of the model. The optimal transport layer marginally improves accuracy of the model compared to the Log-Softmax normalization.

5.5. Application

Comparison with a Commercial Tool. To compare directly with Vicon’s Shōgun auto-labeling tool, we record a new “SOMA” mocap dataset with 2 subjects, performing 11 motion types, including dance, clap, kick, etc., using a Vicon system with 54 infrared “Vantage 16” [27] cameras operating at 120 Hz. In total, we record 69 motions and intentionally use a marker layout preferred by Vicon. We manually label this dataset and treat these labels as ground truth. We process the reconstructed mocap point cloud data with both SOMA (using tracklet labeling) and with Shōgun. The results in Tab. 6 show that SOMA achieves sub-millimeter accuracy and similar performance compared with the proprietary tool while not requiring subject calibration data. In Sup. Mat. we present further details of this dataset.

Processing Real MoCap Datasets with different capture technologies and unknown subject calibration is presented in Tab. 7. For each dataset, SOMA is trained on the marker superset using only synthetic data. SOMA effectively enables running MoSh on mocap point cloud data to extract realistic bodies. The results are not perfect and we manually remove sequences that do not pass a subjective quality bar (see Sup. Mat. for examples). Table 7 indicates

Datasets	Per-Frame		Tracklet		#	#	#	#
	Acc.	F1	Acc.	F1	Markers	Motions	Frames	Subjects
BMLrub [50]	98.15 ± 2.78	97.75 ± 3.23	98.77 ± 1.58	98.65 ± 1.89	41	3013	3757725	111
KIT[29]	94.97 ± 2.42	95.51 ± 2.65	95.46 ± 1.87	97.10 ± 2.00	53	3884	3504524	48
BMLmovi[15]	95.90 ± 4.65	95.12 ± 5.26	97.33 ± 2.29	96.87 ± 2.60	67	1863	1255447	89

Table 3: Performance of SOMA on real marker data of various datasets with large variation in number of subjects, body pose, markers, and hardware specifics. We corrupt the real marker data with additional noise, and forget the labels, turning it into a raw MPC before passing through SOMA pipeline.

# Markers Removed	3		5		6		12	
	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1
Base Model	95.35 ± 6.52	94.40 ± 7.37	94.08 ± 7.04	91.75 ± 8.32	93.41 ± 7.84	90.73 ± 9.33	88.25 ± 14.04	8.80 ± 16.07
Base MoCap	91.78 ± 10.13	90.68 ± 10.90	90.73 ± 9.38	90.12 ± 9.96	91.89 ± 8.97	91.00 ± 9.71	87.46 ± 10.77	86.78 ± 12.00

Table 4: Robustness to variations in marker layout. First row: A *base model* is trained with full marker layout (superset) and tested per-frame on real markers from the validation set (HDM05) with omitted markers (subset). Second row: One model is trained per varied layout (subset) and tested on *base mocap* markers (superset).

Version	Accuracy	F1
Base	95.50 ± 5.33	94.66 ± 6.03
- AMASS Noise Model	94.73 ± 5.52	93.73 ± 6.25
- CAESAR bodies	95.21 ± 6.83	94.31 ± 7.57
- Log-Softmax Instead of Sinkhorn	91.51 ± 10.69	90.10 ± 11.47
- Random Marker Placement	89.41 ± 8.06	87.78 ± 8.85
- Transformer	11.36 ± 6.54	7.54 ± 6.22

Table 5: Ablation study of SOMA components on the HDM05 dataset. The numbers reflect the contribution of each component in overall per-frame performance of SOMA. We take the full base model and remove one component at a time.

	$V2V_{mm}^{mean}$	$V2V_{mm}^{median}$	Acc.	F1
Shōgun	0.00 ± 0.11	0.00	100.0 ± 0.00	100.0 ± 0.00
SOMA	0.08 ± 2.09	0.00	99.94 ± 0.47	99.92 ± 0.64

Table 6: SOMA vs Shōgun. On a manually labeled dataset with passive markers, we compare SOMA against a commercial tool for labeling performance and surface reconstruction.

	Type	# Points	# Subjects	Minutes	Success Ratio
CMU-II [11]	P	40-255	41	116.30	80.0%
DanceDB [4]	A	38	20	203.38	81.26%
Mixamo [10]	A	38-96	29	195.37	78.31%
SOMA	P	53-140	2	18.27	100.00%
Total				533.32	

Table 7: Processing uncleaned, unlabeled mocap datasets with SOMA. Input to the pipeline are mocap sequences with possibly varying number of points; SOMA labels the points as markers and afterwards MoSh is applied on the data to solve for the body surface. P and A stand for passive and active marker systems respectively.

the percentage of successful minutes of mocap. Failures are typically due to poor mocap quality. Note that the SOMA dataset is very high quality with many cameras and, here, the success rate is 100%. For sample renders refer to the accompanying video.

6. Conclusion

SOMA addresses the problem of robustly labeling of raw mocap point cloud sequences of human bodies in motion, subject to noise and variations across subjects, motions, marker placement, marker density, mocap quality, and capture technology. SOMA solves this problem using several innovations including a novel self-attention mechanism and a matching component that deals with outliers and missing data. We train SOMA end-to-end on synthetic data using several techniques to add realistic noise that enable generalization to real data. We extensively validate the performance of SOMA showing that it is more accurate than previous research methods and comparable in accuracy to a commercial system while being significantly more flexible. SOMA is also freely available for research purposes.

Limitations and Future Work. SOMA performs per-frame MPC labeling and, hence, does not exploit temporal information. A temporal model could potentially improve accuracy. As with any learning-based method, SOMA may be limited in generalizing to new motions outside the training data. Using AMASS, however, the variability of the training data is large and we did not observe problems with generalization. By exploiting the full SMPL-X body model in synthetic data generation pipeline we plan to extend the method to label hand and face markers. Relying on feed-forward components, SOMA is extremely fast and coupled with a suitable mocap solver could potentially recover bodies in real-time from mocap point clouds.

Acknowledgments: We thank Senya Polikovsky, Markus Höschle, Galina Henz (GH), and Tobias Bauch for the mocap facility. We thank Alex Valisu, Elisha Denham, Leyre Sánchez Viñuela, Felipe Mattioni and Jakob Reinhardt for mocap cleaning. We thank GH, and Tsvetelina Alexiadis for trial coordination. We thank Benjamin Pellkofer and Jonathan Williams for website developments. **Disclosure:** <https://files.is.tue.mpg.de/black/CoI/ICCV2021.txt>

References

- [1] Mohammad Abdulkader Abdulrahim. *Parallel Algorithms for Labeled Graph Matching*. PhD thesis, Colorado School of Mines, USA, 1998. AAI0599838. [3](#)
- [2] Ryan Prescott Adams and Richard S. Zemel. Ranking via Sinkhorn Propagation, 2011. [3](#), [4](#)
- [3] Ijaz Akhter and Michael J. Black. Pose-conditioned joint angle limits for 3D human pose reconstruction. In *CVPR*, pages 1446–1455, 2015. [5](#)
- [4] Andreas Aristidou, Efstathios Stavrakis, Margarita Papaefthimiou, George Papagiannakis, and Yiorgos Chrysanthou. Style-based motion analysis for dance composition. *The Visual Computer*, 34:1725–1737, 2018. [2](#), [8](#)
- [5] Alexander C. Berg, Tamara L. Berg, and Jitendra Malik. Shape matching and object recognition using low distortion correspondences. In *CVPR*, volume 1, pages 26–33, 2005. [3](#)
- [6] Tibério S. Caetano, Julian J. McAuley, Li Cheng, Quoc V. Le, and Alex J. Smola. Learning graph matching. In *ICCV*, pages 1–8, 2007. [3](#)
- [7] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, pages 77–85, 2017. [3](#)
- [8] Kang Chen, Yupan Wang, Song-Hai Zhang, Sen-Zhe Xu, Weidong Zhang, and Shi-Min Hu. MoCap-Solver: A neural solver for optical motion capture data. *ACM Transactions on Graphics (TOG)*, 40(4), 2021. [2](#), [3](#), [5](#)
- [9] Renato Contini. Body segment parameters, part II. *Artificial Limbs*, 16(1):1–19, 1972. [3](#)
- [10] Adobe Mixamo MoCap Dataset, 2019. [2](#), [8](#)
- [11] Carnegie Mellon University (CMU) MoCap Dataset, 2019. [2](#), [5](#), [8](#)
- [12] Advanced Computing Center for the Arts and Design (AC-CAD) MoCap Dataset, 2019. [5](#)
- [13] Liuhao Ge, Hui Liang, Junsong Yuan, and Daniel Thalmann. Robust 3D hand pose estimation in single depth images: From single-view CNN to multi-view CNNs. In *CVPR*, pages 3593–3601, 2016. [3](#)
- [14] Saeed Ghorbani, Ali Etemad, and Nikolaus F. Troje. Autolabelling of markers in optical motion capture by permutation learning. In *Advances in Computer Graphics*, pages 167–178, Cham, 2019. Springer International Publishing. [2](#), [3](#), [4](#), [5](#), [7](#)
- [15] Saeed Ghorbani, Kimia Mahdavian, Anne Thaler, Konrad Kording, Douglas James Cook, Gunnar Blohm, and Nikolaus F. Troje. MoVi: A large multi-purpose human motion and video dataset. *PLOS ONE*, 16(6):1–15, 2021. [6](#), [8](#)
- [16] Shangchen Han, Beibei Liu, Robert Wang, Yuting Ye, Christopher D. Twigg, and Kenrick Kin. Online optical marker-based hand tracking with deep labels. *ACM Transactions on Graphics (TOG)*, 37(4):1–10, 2018. [2](#), [3](#), [5](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [4](#)
- [18] Lorna Herda, Pascal Fua, Ralf Plänkers, Ronan Boulic, and Daniel Thalmann. Using skeleton-based tracking to increase the reliability of optical motion capture. *Human Movement Science*, 20(3):313–341, 2001. [3](#)
- [19] Daniel Holden. Robust solving of optical motion capture data by denoising. *ACM Transactions on Graphics (TOG)*, 37(4):165:1–165:12, 2018. [2](#), [3](#), [4](#), [5](#)
- [20] Stefan Holzreiter. Autolabeling 3d tracks using neural networks. *Clinical Biomechanics*, 20(1):1–8, 2005. [7](#)
- [21] NaturalPoint, Inc. Optitrack motion capture systems, 2019. [2](#), [4](#)
- [22] PhaseSpace, Inc., 2019. [3](#)
- [23] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3D deformation model for tracking faces, hands, and bodies. In *CVPR*, pages 8320–8329, 2018. [3](#)
- [24] Vladimir Joukov, Jonathan F. S. Lin, Kevin Westermann, and Dana Kulić. Real-time unlabeled marker pose estimation via constrained extended kalman filter. In *Proceedings of the International Symposium on Experimental Robotics*, pages 762–771, Cham, 2020. Springer International Publishing. [3](#)
- [25] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *ICCV*, page 1482–1489, USA, 2005. IEEE Computer Society. [3](#)
- [26] Matthew Loper, Naureen Mahmood, and Michael J. Black. MoSh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics (TOG)*, 33(6):1–13, 2014. [1](#), [2](#), [3](#), [4](#)
- [27] Vicon Motion Systems, Ltd. Motion Capture Systems, 2019. [2](#), [3](#), [4](#), [7](#)
- [28] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *ICCV*, pages 5441–5450, 2019. [1](#), [2](#), [3](#), [4](#), [5](#)
- [29] Christian Mandery, Ömer Terlemez, Martin Do, Nikolaus Vahrenkamp, and Tamim Asfour. The KIT whole-body human motion database. In *ICAR*, pages 329–336, 2015. [6](#), [8](#)
- [30] Daniel Maturana and Sebastian Scherer. VoxNet: A 3D convolutional neural network for real-time object recognition. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928, 2015. [3](#)
- [31] Jonathan Maycock, Tobias Rohlig, Matthias Schroder, Mario Botsch, and Helge Ritter. Fully automatic optical motion tracking using an inverse kinematics approach. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 461–466, 2015. [7](#)
- [32] Michael Clausen Bernhard Eberhardt Björn Krüger Andreas Weber Meinard Müller, Tido Röder. Documentation mocap database HDM05. Technical Report CG-2007-2, Universität Bonn, 2007. [6](#)
- [33] Johannes Meyer, Markus Kuderer, Jörg Müller, and Wolfram Burgard. Online marker labeling for fully automatic skeleton tracking in optical motion capture. In *ICRA*, pages 5652–5657, 2014. [3](#)
- [34] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, pages 10967–10977, 2019. [2](#), [3](#), [4](#), [5](#)
- [35] Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. [4](#)
- [36] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on

- point sets in a metric space. In *NIPS*, volume 30. Curran Associates, Inc., 2017. 3
- [37] Maurice Ringer and Joan Lasenby. Multiple hypothesis tracking for automatic optical motion capture. In *ECCV*, pages 524–536, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg. 3
- [38] Maurice Ringer and Joan Lasenby. A procedure for automatically estimating model parameters in optical motion capture. *Image and Vision Computing*, 22(10):843–850, 2004. British Machine Vision Conference (BMVC). 3
- [39] Kathleen Robinette, Sherri Blackwell, Hein A. M. Daanen, Mark Boehmer, Scott Fleming, Tina Brill, David Hoferlin, and Dennis Burnsides. Civilian american and european surface anthropometry resource (CAESAR) final report. Technical Report AFRL-HE-WP-TR-2002-0169, US Air Force Research Laboratory, 2002. 5
- [40] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *CVPR*, pages 4937–4946, 2020. 3, 4
- [41] Tobias Schubert, Alexis Gkoggidis, Tonio Ball, and Wolfram Burgard. Automatic initialization for skeleton tracking in optical motion capture. In *ICRA*, pages 734–739, 2015. 3
- [42] Leonid Sigal, Alexandru O. Balan, and Michael J. Black. HumanEva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *IJCV*, 87(4):4–27, 2010. 5
- [43] Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35(2):876–879, 1964. 4
- [44] Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967. 3, 4
- [45] Yang Song, Luis Goncalves, and Pietro Perona. Unsupervised learning of human motion. *TPAMI*, 25(7):814–827, 2003. 2
- [46] Jannik Steinbring, Christian Mandery, Florian Pfaff, Florian Faion, Tamim Asfour, and Uwe D. Hanebeck. Real-time whole-body human motion tracking based on unlabeled markers. In *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 583–590, 2016. 3
- [47] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik G. Learned-Miller. Multi-view convolutional neural networks for 3D shape recognition. In *ICCV*, 2015. 3
- [48] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *ECCV*, pages 581–600, Cham, 2020. Springer International Publishing. 4
- [49] Lorenzo Torresani, Vladimir Kolmogorov, and Carsten Rother. Feature correspondence via graph matching: Models and global optimization. In *ECCV*, pages 596–609, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. 3
- [50] Nikolaus F. Troje. Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. *Journal of Vision*, 2(5):2–2, 2002. 6, 8
- [51] Matthew Trumble, Andrew Gilbert, Charles Malleison, Adrian Hilton, and John Collomosse. Total capture: 3D human pose estimation fusing video and inertial sensors. In *BMVC*, pages 14.1–14.13, 2017. 5
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, undefinedukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, page 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. 3, 4
- [53] Cédric Villani. *Optimal transport – Old and New*, volume 338. Springer Science and Business Media, 2008. 3
- [54] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018. 3
- [55] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *CVPR*, pages 1912–1920, 2015. 3
- [56] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T. Freeman, Rahul Sukthankar, and Cristian Sminchisescu. GHUM & GHUML: Generative 3d human shape and articulated pose models. In *CVPR*, pages 6183–6192, 2020. 3