

Fig. 1. Given raw, unlabeled full-body motion capture point cloud (black balls) (a), RoMo generates tracklets to leverage temporal information for accurately labeling markers (colored balls) (b) and utilizes a hybrid inverse kinematics-based method (c) to solve body motions (d). RoMo demonstrates robustness against occlusion (cubes) and outliers (grey balls) and accurately solves large body motions and fine multi-fingered hand movements.

Optical motion capture (MoCap) is the "gold standard" for accurately capturing full-body motions. To make use of raw MoCap point data, the system labels the points with corresponding body part locations and solves the fullbody motions. However, MoCap data often contains mislabeling, occlusion and positional errors, requiring extensive manual correction. To alleviate this burden, we introduce RoMo, a learning-based framework for robustly labeling and solving raw optical motion capture data. In the labeling stage, RoMo employs a divide-and-conquer strategy to break down the complex full-body labeling challenge into manageable subtasks: alignment, full-body segmentation and part-specific labeling. To utilize the temporal continuity of markers, RoMo generates marker tracklets using a K-partite graph-based clustering algorithm, where markers serve as nodes, and edges are formed based on positional and feature similarities. For motion solving, to prevent error accumulation along the kinematic chain, we introduce a hybrid inverse kinematic solver that utilizes joint positions as intermediate representations and adjusts the template skeleton to match estimated joint positions. We demonstrate that RoMo achieves high labeling and solving accuracy across multiple metrics and various datasets. Extensive comparisons show that our method outperforms state-of-the-art research methods. On a real dataset, RoMo improves the F1 score of hand labeling from 0.94 to 0.98, and reduces joint position error of body motion solving by 25%. Furthermore, RoMo can be applied in scenarios where commercial systems are inadequate. The code and data for RoMo are available at https://github.com/robustmocap/RoMo.

$\label{eq:ccs} CCS \ Concepts: \bullet \ Computing \ methodologies \rightarrow Motion \ capture; \ Neural networks.$

Additional Key Words and Phrases: Character Animation, Motion Capture, Machine Learning

1 INTRODUCTION

Human pose estimation is a crucial field in computer vision and graphics. Despite the promising outcomes of marker-free, imagebased methods, marker-based motion capture (MoCap) systems are preferred for their precision and versatility in game and film industry [Taheri et al. 2020] to capture body and hand motions with high coordination. MoCap system operates by first capturing a series of 2D images from multiple perspectives of the subjects, with markers at key locations on their bodies. The system then determines the 3D positions of these markers and associates them with specific locations of body parts through a process known as marker **labeling**. Next, the system reconstructs underlying skeletal structures and body movements, which is also known as **solving**.

Despite the use of sophisticated and expensive motion capture equipment, motion capture data inevitably suffers from various issues, including mislabeling, occlusions, and positional inaccuracies. These noises necessitate extensive manual corrections to achieve high accuracy. Noises can be categorized into three main types:

- Mislabeling: These noises include label swaps, where markers are incorrectly labeled, and outliers, which occur when invalid markers are not properly excluded. Hand markers are susceptible to mislabeling due to the complexity of hand motions and their proximity.
- Occlusion: This type of noise occurs when markers become invisible to the MoCap system, typically due to actors' selfocclusion or markers entering cameras' dead zones.
- **Positional inaccuracy**: These result from markers being recorded at positions that do not accurately reflect their true locations. Positional inaccuracies can occur when correcting occlusions or due to tracking system errors.

A multitude of data-driven approaches have been developed to address these noises, focusing on cleaning the raw marker data and solving motions. Some works [Ghorbani and Black 2021; Ghorbani et al. 2019] concentrate on the labeling stage, but are limited to body markers alone. When these techniques are applied to wholebody marker data, the body and hand markers are treated equally,

115 resulting in lower accuracy for hand markers due to distinct data 116 distributions between body and hand markers, with the former 117 having larger mutual distances and lower occlusion or mislabeling 118 probabilities. Furthermore, these methods label the point cloud 119 using information of a single frame, neglecting valuable temporal information. This omission reduces labeling accuracy, especially 120 in the presence of occlusions. Another branch of research [Chen 121 122 et al. 2021; Holden 2018; Pan et al. 2023] aims to clean the noises in 123 already labeled markers during the solving stage. These approaches 124 often rely on the presumed accuracy of initial labeling or employ weak assumptions, such as acceleration thresholds [Pan et al. 2023] 125 126 and marker distances [Chen et al. 2021; Holden 2018], to address mislabels. Consequently, they are susceptible to labeling errors. 127 Moreover, all three types of noises can lead to inaccuracies in solving. 128 These methods attempt to estimate the relative joint rotations of 129 the entire body simultaneously, ignoring error accumulation along 130 131 the kinematic chain of the body, where minor errors in parent joints can significantly amplify inaccuracies in child joints' positioning. 132

To develop a MoCap labeling and solving framework that is robust 133 against the aforementioned noises, particularly mislabeling and 134 positional inaccuracies, three key challenges must be addressed. 135 136 (1) The ability to deal with the distinct data distributions of body and hand markers. Even with deeper or wider neural networks, 137 it is difficult for a single network to handle the full-body MoCap 138 139 data containing body and hand markers. (2) The exploitation of temporal continuity in marker data. A simple approach is to use the 140 marker positions to create tracklets. Yet, this strategy overlooks the 141 inter-marker relationships within the frame, leading to inaccuracies 142 during rapid movements, such as rapid scrolling and fast hand 143 144 waving. (3) The prevention of error accumulation along kinematic 145 chain. Prior methods focus on the accuracy of joints' local rotations without considering the overall positional accuracy. It is crucial to 146 devise a strategy that maintains awareness of the global positions 147 of joints to prevent error accumulation. 148

To this end, we introduce RoMo, a data-driven approach for ro-149 150 bustly labeling and solving full-body unlabeled optical MoCap data, mainly addressing two types of noises: mislabeling and positional 151 152 inaccuracy. During the labeling stage, a divide-and-conquer strategy is employed to decompose the complex full-body labeling challenge 153 into manageable subtasks: alignment, full-body marker segmenta-154 tion and part-specific labeling. This approach mitigates training 155 difficulties and enhances network performance. To make full use of 156 157 temporal continuity of markers, tracklets are generated using deep features extracted by the network. Tracklet construction is modelled 158 159 as a K-partite graph-based clustering algorithm, treating markers as nodes and edges are determined by feature similarities and po-160 sitional differences. Incorporating feature similarities extends the 161 162 receptive field of markers from nearest neighbors to all markers in 163 the point cloud, thereby improving labeling precision. In the motion solving stage, RoMo uses joint positions as intermediate representa-164 165 tions, with inverse kinematics adjusting the template skeleton to match the estimated joint positions. The utilization of global joint 166 positions prevents error accumulation along the kinematic tree and 167 enables accurate reconstruction of motions corresponding to input 168 169 markers, especially the positions of ending joints.

171

172

173

174

175

176

177

178

179

180

181

182

183

184

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

- A robust full-body MoCap labeling framework that employs divide-and-conquer strategy and integrates spatial-temporal information to precisely label body and hand markers.
- An innovative method for generating marker tracklets by solving a K-partite graph-based clustering algorithm that utilizes both markers' positions and deep features.
- A hybrid inverse kinematics MoCap solver that circumvents error accumulation along the body kinematic chain.

2 RELATED WORK

Motion capture plays a crucial role in film and game industries, imbuing characters with vitality and distinctive personalities. Accurately capturing human motions is also essential across various motion-related research domains, including action recognition [Hua et al. 2023; Yan et al. 2018], action prediction [Cao et al. 2020; Cui et al. 2021], motion synthesis [Chen et al. 2023; Tevet et al. 2023], and image-based pose estimation [Munea et al. 2020; Varol et al. 2017]. However, optical motion capture data inevitably suffers from mislabeling, occlusion, and positional inaccuracy, necessitating extensive manual corrections. Both the industry and research community have devoted significant efforts toward tackling these noises.

Within the industry, there are multiple mature commercial solutions available for real-time labeling and motion solving, such as Vicon Shogun [Vicon 2023] and OptiTrack [OptiTrack 2024]. However, these solutions are tailored to specific marker layouts and require careful calibrations for each new actor to accommodate diverse body shapes. Consequently, they do not offer a universally applicable solution for diverse scenarios requiring customized marker layouts, which are prevalent in numerous motion capture applications.

Numerous hand-crafted prior-based methods have been proposed in the research community [Aristidou et al. 2018; Baumann et al. 2011; Liu and McMillan 2006; Tautges et al. 2011; Tits et al. 2018; Wang et al. 2016; Xiao et al. 2015] to clean MoCap data. Although these methods produce satisfying results on data with specific noise patterns via careful hand-tuning, they suffer from poor generalization ability on the real world data with more complex noises. Recent advances in research community have spotlighted the utilization of neural networks. Existing neural-based strategies primarily aim at solving motions from motion capture data [Chen et al. 2021; Holden 2018; Pan et al. 2023; Pavllo et al. 2018; Perepichka et al. 2019], with a strong reliance on pre-labeled datasets. These methods often resort to some fragile heuristics to eliminate ghost points and struggle to address mislabeling caused by label swaps. Moreover, prior techniques attempt to solve the whole body joint's local rotations in a single pass, which frequently leads to cumulative errors along the body kinematic chain. In contrast, RoMo innovates by adopting global joint points as an intermediate representation and utilizing a hybrid inverse kinematics-based approach. This novel strategy ensures highly accurate solving results, markedly improving position accuracy for the terminal joints.

Few methods aim at labeling optical motion capture data using neural networks [Ghorbani and Black 2021; Ghorbani et al. 2019]. Notably, Ghorbani et al. [Ghorbani et al. 2019] introduced a novel approach using a simple feed-forward residual network that employs permutation learning for body marker labeling. However,

170 Overall, our paper makes the following contributions:



Fig. 2. RoMo's pipeline consists of three modules. **Top**: In the pre-processing stage, RoMo accepts 3D sparse unordered MoCap point clouds with varying point numbers. It then conducts point cloud alignment to eliminate the global transformations and segmentation to partition the point cloud into body and hand point clouds. Subsequently, it employs a network consisting of alternating global self-attention and local aggregation layers to extract markers' features. **Bottom Left**: In the tracklet construction stage, RoMo addresses a K-partite graph-based clustering problem to create tracklets, and assigns markers within the same tracklet to a same label. **Bottom Right**: RoMo utilizes a hybrid inverse kinematics-based method to solve the motion, which iteratively adjust the joints of template skeleton along the kinematic tree to match the estimated joint positions.

their method exhibits limitation in handling outliers. The introduction of SOMA [Ghorbani and Black 2021] marked a significant leap forward by leveraging an attention-based network architecture alongside an optimal transport layer, enhancing the accuracy of body marker labeling while efficiently managing outliers. Unlike previous methods that strive for a one-fits-all solution, RoMo employs a divide-and-conquer strategy, significantly enhancing labeling accuracy on full-body MoCap data. Furthermore, RoMo innovates by constructing tracklets through solving a K-partite problem based on proximities in marker positions and similarities in markers' deep features, thus efficiently utilizing the temporal continuity. Additionally, these methods primarily focus on global attention, overlooking the critical local interactions among adjacent markers. To bridge this gap, RoMo integrates stacked local aggregation [Wang et al. 2019] and global attention layers [Vaswani et al. 2017] to adeptly capture both local and global point features.

Tracklet is extensively studied in the field of Multiple Object Tracking (MOT) [Dai et al. 2019; Shen et al. 2018; Sheng et al. 2018; Zhang et al. 2020], and RoMo generates tracklets to exploit temporal continuity in motion. The tracklets are short trajectories with high confidence, which reduces false positive detections. The key issue is correctly identifying objects across multiple frames. Many algorithms are proposed for determining the best association, including min-cost flow [Shitrit et al. 2014], conditional random field [Choi 2015; Milan et al. 2013; Xiang et al. 2020], and multiple hypothesis tracking [Kim et al. 2015]. RoMo is the first method to use k-partite graph clustering in MoCap data processing, and it produces cuttingedge quantitative and qualitative results. Unlike previous works that process objects in videos, RoMo employs K-partite clustering on point cloud sequences with varying cardinalities.

3 METHOD

RoMo takes as input a sequence of unlabeled motion capture point clouds, denoted as $\{P^1, ..., P^T\}$, where each P^i represents a sparse and unordered point cloud at time i, containing n^i points. The cardinalities of these clouds nⁱ changes over time due to marker occlusion or ghost points. Initially, RoMo assigns each point to a corresponding marker label $\{L_1, ..., L_N, null\}$, with L_i signifying a valid label and null indicating outliers. The labeled MoCap data is then structured as $M \in \mathbb{R}^{T \times \breve{N} \times 3}$, representing the ordered positions of markers, alongside $O \in \{0, 1\}^{T \times N}$, which marks the visibility of markers. Subsequently, RoMo solves the motions, $R \in \mathbb{R}^{T \times (3+K \times 9)}$, incorporating both the global translation of the body and the local rotations of each joint, as well as the underlying template skeleton $S \in \mathbb{R}^{K \times 3}$, representing the joints' offsets relative to their parent joints. An overview of RoMo's pipeline is presented in Fig. 2. In the following part of this section, we will first explain the labeling process, and then describe the solving framework.

Initially, marker features are extracted using a neural network comprised of alternating global self-attention [Vaswani et al. 2017] and local aggregation [Wang et al. 2019] layers to capture both global and local features f (top of Fig. 2). To simplify training, before feature extraction, RoMo removes the global orientation of markers and segments the point cloud into three parts: body, left hand, and right hand. The alignment and segmentation employ networks that share a similar structure with the feature extraction network. After the

feature extraction, RoMo performs Sinkhorn normalization [Adams and Zemel 2011] to transform the features f to labeling confidences c which satisfies a relaxed one-to-one correspondence.

346 To harness the temporal information, tracklets are constructed 347 (bottom left of Fig. 2), and markers within the same tracklet are assigned identical labels. A tracklet depicts the motion of a single 348 marker through 3D trajectories over a short period of time. The 349 350 construction of tracklets is formulated as a K-partite graph-based 351 clustering problem [Zhang et al. 2020]. Initially, a graph G is created, treating markers as nodes and forming edges based on similarities 352 between markers, which include both positional and feature sim-353 ilarities. Subsequently, a clustering for \mathcal{G} is sought using greedy 354 algorithm. For the label association, RoMo calculates the tracklet's 355 356 confidence using the L-q norm [Ghorbani et al. 2019] of markers' 357 confidence, and selects label with the highest confidence.

In the solving stage, motions are solved frame-by-frame (bottom 358 right of Fig. 2). The joint position, denoted as $J \in \mathbb{R}^{T \times K \times 3}$, serves as 359 an intermediate representation, with inverse kinematics applied to 360 361 estimate joint rotations. The rotations are decomposed into twist and swing components, $R = R_{sw}R_{tw}$, where the swing component's axis 362 is perpendicular to the bone direction, and the twist component's 363 axis is parallel with the bone direction. Utilizing the labeled mark-364 ers M^i at frame *i*, RoMo estimates the joint positions $J^i \in \mathbb{R}^{K \times 3}$. 365 This allows the calculation of a closed-form solution for the swing 366 367 rotation component R_{sw}^i by applying Rodrigues formula [Askey 2005] on the relative joint positions with child joints. Additionally, 368 RoMo estimates the twist angle R_{tw}^i , yielding the complete joint 369 370 rotations. This process is iteratively performed along the kinematic chain to sequentially estimate joint rotations. The solving network 371 372 also estimates the joint offsets of each frame S^{i} . To maintain bone 373 length consistency throughout the motion sequence, the overall joint offset is determined by averaging estimated offsets across the 374 motion sequence: $S_i = \text{mean}(S_i^1, ..., S_i^T)$. 375

3.1 Point Cloud Pre-processing

376

377

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395

396

397

398

399

Unlike previous methods [Ghorbani and Black 2021; Ghorbani et al. 2019] that preserve body orientations, RoMo begins by normalizing the global body orientation, which aligns point clouds to accelerate network convergence and enhance the accuracy. An orthogonal transformation matrix is estimated through a smaller network and then directly applied to the input point cloud. This auxiliary network mirrors the architecture of the feature extraction network, employing similar components. The network is trained based on the rotation of the root joint, with regularization applied to ensure the output matrix *A* closely approximates an orthogonal matrix:

$$\mathcal{L}_{align} = ||A^{i} - \hat{R}_{root}^{i}||_{2} + \lambda_{reg}||I - A^{i}A^{i^{I}}||_{2}.$$
(1)

Subsequently, the point cloud is segmented. Given that body and hand markers exhibit distinct data distributions, employing a unified network for feature extraction from the entire body's point cloud is overly complex. To simplify the learning process, a divide-and-conquer strategy is implemented, wherein the wholebody point cloud is first divided into body,left hand and right hand parts, followed by feature extraction. This segmentation process utilizes a neural network with a structure identical to that of the alignment network, and it is trained using cross-entropy loss.

Following the transformer-based point cloud processing networks [Lu et al. 2022], RoMo's network architecture incorporates two types of layers: self-attention layers and local aggregation layers. Self-attention layers initiate by projecting the input features into key, query, and value vectors, and then employ the first two to calculate marker attention. The resultant layer output is derived from the product of attention weights and the value vector. Relying solely on the self-attention layers [Ghorbani and Black 2021] overlooks local spatial information, which is critical for point cloud processing. To address this, RoMo integrates local aggregation layers [Wang et al. 2019] to capture the neighboring details. This process involves identifying the k-nearest neighbors within the feature space of points, and utilizing feature differences to aggregate relative information regarding neighboring markers. Such data are subsequently aggregated and processed through a multi-layer perceptron (MLP) to produce the local aggregation layer's output.

The confidence of labels to markers is designed to satisfy a relaxed one-to-one correspondence: each point can be mapped to at most one label and vise versa, with the exception of the *null* label, which can be assigned to multiple points. To enforce this constraint, RoMo employs log-domain optimal transport [Peyré and Cuturi 2020; Sarlin et al. 2020], utilizing the Sinkhorn-Knopp algorithm iteratively for optimization [Adams and Zemel 2011]. Following [Ghorbani and Black 2021], RoMo constructs the initial confidence matrix $C_{init}^i \in [0, 1]^{n^i \times N}$, using the extracted marker features f, and modifies it by adding an extra row and column to account for unmatched points and labels, resulting in the augmented confidence matrix $C_{aug}^i \in [0, 1]^{(n^i+1)\times(N+1)}$. Sinkhorn normalization is then iteratively applied to alternatively normalize the rows and columns, ensuring they sum to 1. After k iterations, we obtain the final normalized confidence matrix C^i .

The feature extraction network's training incorporates a weighted negative likelihood loss: $\mathcal{L}_{label} = -\frac{1}{\sum \hat{C}^i} \sum W \hat{C}^i \log(C^i)$, with \hat{C}^i as the ground truth confidence matrix and W serving as a weight matrix to mitigate the impact of the *null* label.

3.2 Tracklet Generation and Labeling

The previous section focused solely on information within a single frame, overlooking the temporal continuity of marker motions. To harness temporal information in sequential MoCap data, RoMo generates tracklets by leveraging the marker positions and marker features extracted by neural network. While some commercial software, such as Vicon [Vicon 2023], automatically generates tracklets, these are not universally available, notably in certain archival Mo-Cap datasets like GRAB [Taheri et al. 2020] or limited hardware. Additionally, these methods do not utilize deep features of markers, limiting the precision of tracklet construction.

Drawing inspiration from [Zhang et al. 2020], RoMo treats tracklet construction as a K-partite graph-based clustering algorithm. After construction, RoMo assigns the markers in the same tracklet with the same label, determined by the L-q norm of confidence of the labels computed in the preceding subsection.

407

408

409

400

401

402

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

432

433

434

435

436

437

438

439

440

441

442

443

444

445

446

447

448

449

450

451

452

453

454

455

following three components:

- Node set \mathcal{V} represents the markers in the point cloud, where v_m^i denotes the m-th marker in the point cloud P^i .
- Edge set & comprises edges e^{ij}_{mn} indicates temporal connections, which connect any two nodes (vⁱ_m, v^j_n) that represent markers from different frame,
- markers from different frame,
 Weight set W: the weight w^{ij}_{mn} associated with edge e^{ij}_{mn} quantifies the similarity between markers vⁱ_m and v^j_n, incorporating both spatial proximity and feature resemblance to measure marker similarity.

$$w_{mn}^{ij} = w_{pos,mn}^{ij} + \lambda_{fet} w_{fet,mn}^{ij},$$

$$w_{pos,mn}^{ij} = ||p_m^i, p_n^j||_2,$$

$$w_{fet,mn}^{ij} = 1 - \text{cosine_similarity}(f_m^i, f_n^j),$$
(2)

where $w_{pos,mn}$ represents the Euclidean distance between marker positions, while $w_{fet,mn}$ quantifies the cosine distance between the deep features of the corresponding nodes, with λ_{fet} serving as a predefined weighting factor of feature similarity.

To identify a clustering solution for the graph \mathcal{G} , RoMo seeks to minimize the total sum of all weights, as outlined below:

$$\arg \min_{\substack{e_{jn}^{ij} \\ m,n}} \sum_{\substack{e_{mn}^{ij} \\ m,n}} e_{mn}^{ij} w_{mn}^{ij};$$
s.t.
$$\sum_{\substack{v_n^{j} \in \mathcal{V}^{i}, i \neq j \\ mn}} e_{mn}^{ij} \leq 1; \sum_{\substack{v_n^{j} \in \mathcal{V}^{i}, i = j \\ mn}} e_{mn}^{ij} = 0;$$

$$e_{mn}^{ij} = 0, w_{pos,mn}^{ij} > th_{pos}; e_{mn}^{ji} = 0, w_{fet,mn}^{ij} > th_{fet};$$

$$e_{mn}^{ij} \in \{0,1\}; e_{mn}^{ij} = e_{nm}^{ji};$$
(3)

where \mathcal{V}^i refers to the node set of the i-th frame, while th_{pos} and th_{fet} are thresholds for position and feature differences, respectively. The first condition ensures that any given node v_m^i connects to at most one node in the node set of a different frame \mathcal{V}^j . The edge value e_{mn}^{ij} is restricted to either 0 or 1, resulting in integer solutions for this optimization problem. To expedite computation, a greedy algorithm is employed. In the resulting solution, a collection of cliques is obtained, which are considered as tracklets.

For the association of tracklets with marker labels, a straightforward approach is to assign the most frequently occurring label from single-frame labeling, yet this method overlooks scenarios involving labels that are prevalent but have low confidence. To circumvent this issue, RoMo's tracklet labeling method calculates tracklet confidence from markers' confidence. Specifically, for a tracklet $Tr = \{p_{l}^{i}, ..., p_{m}^{j}, ..., p_{n}^{k}\}$, the confidence of it being assigned to marker label L_{i} is calculated as $C_{L_{i}} = (\sum_{p_{l}^{i} \in Tr} |c_{p_{l}^{i},L_{j}}|^{q})^{\frac{1}{q}}$, representing the L - q norm of the confidence for label L_{i} of markers within Tr. For q = 0, this score mimics a voting mechanism, whereas for q = 1, it serves as the sum of confidence of label L_{i} .

3.3 Hybrid Inverse Kinematics-based Solving



Fig. 3. a) An illustration of rotation decomposition. b) RoMo utilizes global joint positions and avoids error accumulation along kinematic chain.

Directly estimating local joint rotations in one pass may cause error accumulation along the kinematic chain, where minor rotational inaccuracies in parent joints can lead to significant positional discrepancies in child joints, as shown in the right of Fig. 3. This process is further challenged by high non-linearity, as rotations belong to the 3D rotation group $R \in SO(3)$ and marker positions belong to the 3D position group $M \in R^3$, which complicates network training due to its inherent difficulty [Kanazawa et al. 2018; Kolotouros et al. 2019]. Inspired by [Li et al. 2021], RoMo leverages 3D keypoints as an intermediate representation and uses an inverse kinematics-based solver. This approach simplifies network training and mitigates error accumulation. Compared to [Li et al. 2021], RoMo employs a different feature extractor for the task of optical motion capture solving. Next, we dive into details of the solver, the times in notations are omitted for simplicity.

The key idea is to estimate 3D joint positions rather than the complete rotations. With the estimated 3D joint positions J, a solution for the swing rotation R_{sw} , whose axis is perpendicular to bone direction, is obtainable. The twist rotation R_{tw} , whose axis is parallel to the direction, however, requires the assistance of a neural network. An illustration of rotation decomposition is shown in the left of Fig. 3. The full rotation is the combination of the swing and twist rotation $R = R_{sw}R_{tw}$. Although this method still relies on the neural network for rotation estimation, it significantly reduces learning complexity. Unlike the three degree of freedoms (DoF) associated with rotation, the twist angle is constrained to a single DoF. Additionally, due to the physical constraints of the human body, the twist angle exhibits a limited range of variation [Li et al. 2021]. Only the joints on limbs exhibit a wide range of twist angles, constituting a small proportion of all the joints. We present a detailed analysis of joints' twist angle's distribution in the supplementary material.

The hybrid inverse kinematics process is performed iteratively along the kinematic tree. In our experiments, we discovered that using only the network on the point cloud to estimate the root orientation can lead to an error of about 2 degrees. As a result, we use Singular Value Decomposition (SVD) along with the estimated joint positions to improve the root orientation solving accuracy. For the estimation of rotations of subsequent children joints, we calculate the swing rotation R_{sw} from the estimated joint positions. A straightforward approach is to employ the relative positions derived from the estimated positions of child joints:

$$\vec{j}_{i} = J_{ch(i)} - J_{i}, \quad \vec{t}_{i} = T_{ch(i)} - T_{i},$$

$$\cos\alpha_{i} = \frac{\vec{j}_{i} \cdot \vec{t}_{i}}{||\vec{j}_{i}|| \; ||\vec{t}_{i}||}, \quad \sin\alpha_{i} = \frac{||\vec{j}_{i} \times \vec{t}_{i}||}{||\vec{j}_{i}|| \; ||\vec{t}_{i}||},$$

$$\vec{n}_{i} = \frac{\vec{j}_{i} \times \vec{t}_{i}}{||\vec{j}_{i} \times \vec{t}_{i}||},$$
(4)

$$R_{sw,i} = I + \sin\alpha_i [\vec{n}_i]_{\times} + (1 - \cos\alpha_i) [\vec{n}_i]_{\times}^2,$$

where \vec{j}_i, \vec{t}_i are the estimated and template relative joint positions, respectively. The rest is the Rodrigues formula, where I is the identity matrix and $[\vec{n}_i]_{\times}$ is the skew symmetric matrix of \vec{n}_i .

However, this approach does not mitigate error accumulation as it relies solely on the relative positions of child joints, disregarding errors introduced by parent joints. To account for such discrepancies, RoMo employs forward kinematics on the template skeleton with the current estimated joint rotations to obtain the global joint positions J'_i , which can also reflect the error introduced by estimated parent joints' offsets and rotations. The relative joint position is then calculated as $\vec{J}'_i = J_{ch(i)} - J'_i$. We replace the first equation in Equation 4 using \vec{J}'_i . This process is illustrated in the right of Fig. 3.

RoMo's motion solver uses the heterogeneous graph network outlined in [Pan et al. 2023] as its backbone. This network considers markers and joints as distinct nodes in a heterogeneous graph and conducts graph convolution operations to extract their local and global features to solve motions. The motion solving network estimates joint positions P, twist rotations R_{tw} , and bone lengths S. For continuity, the twist rotation is represented by a two-dimensional vector $[\cos(R_{tw}), \sin(R_{tw})]$, derived from the trigonometric functions of the angle. Direct computation of bone lengths from joint positions was explored but resulted in lower accuracy. The network's training is based on a set of loss functions specified below:

$$\mathcal{L}_{solving} = \lambda_{pos} ||J - \hat{J}||_2 + \lambda_{tw} ||R_{tw} - \hat{R}_{tw}||_2 + \lambda_{skel} ||S - \hat{S}||_2,$$
(5)

where \hat{J} , \hat{R}_{tw} , \hat{S} are ground truth joint positions, twist rotations and joint offsets, respectively. λ_{pos} , λ_{tw} , λ_{skel} are weight factors.

4 EXPERIMENTS

We evaluate RoMo quantitatively on several full-body MoCap datasets. The Production and Front-waist dataset are captured from a game studio and cleaned by hand, with different marker layouts. We take the raw markers as input and clean markers as ground truth. The quality of these datasets is exceptionally high, as they are captured in a professional setting with a large number of cameras. The occlusion probabilities for body and hand markers are 0.5% and 6%, respectively. The GRAB dataset [Taheri et al. 2020] captures people interacting with everyday objects. We simulate noises by adding occlusions, where different markers on different body parts have diverse occlusion probabilities, which are propositional to the real data. To emulate data captured under in limited conditions, we set the overall occlusion rate at 15%. The supplementary material contains additional information about our dataset, network architectures, hyper-parameters and implementation details.

4.1 Point Cloud Labeling

For marker labeling, we compare RoMo with two neural-based approaches: [Ghorbani et al. 2019] and SOMA [Ghorbani and Black 2021]. We report two metrics: F1 score and accuracy in percentages. F1 score is the harmonic-average of the precision and recall score, and accuracy is the proportion of correctly predicted labels over all labels. Since the body and hand markers have distinct data distributions, we separately display their metrics to clearly show the methods' performance on diverse body parts.

We present the quantitative result on the left of Table 1. Thanks to the high data quality, the accuracy of body marker labeling in the Production and Front-waist datasets exceeds 98%, with RoMo marginally outperforming previous methods. However, the accuracy for hand data in comparing methods drops dramatically to about 94%, which can be attributed to two main factors. Firstly, the distinct data distributions between body and hand data make it challenging for a single network to effectively address both types. Secondly, hand data are have higher occlusion probabilities and closer marker proximities. Frame-by-frame labeling failed to handle such complex data. Our method maintains high labeling accuracy through a divide-and-conquer strategy and the use of a tracklet-based labeling schema. We further present a qualitative comparison in Fig. 4. RoMo demonstrates robustness to occlusions and outliers and achieves the highest accuracy compared to other methods, particularly under complex occlusions. Additionally, the utilization of tracklets significantly boosts labeling performance. We present additional qualitative comparisons on Fig. 6.



Fig. 4. A qualitative comparison with other methods on hand labeling, where each column represents a timestamp in a MoCap sequence. In cases of outliers, we either randomly place them in positions corresponding to occluded markers or discard them if there isn't enough space.

Table 2 shows the effect of various components of RoMo on the validation split of the Production dataset. The point cloud segmentation and tracklet-based labeling play the most significant role in the overall performance of marker labeling. Other components marginally improves the performance of the model.

SIGGRPAH ASIA'24, December 3-6, 2024, Tokyo

			[Ghorl	bani et	so	МА	Ro	Mo		Vic	on /	[Ho	lden	MoC	n-Solver	Local	MoCan	Ro	Мо	742
			al. 2	019]		11111	100	1110		Mos	Sh++	20	18]	widea	ip solver	Local	mocap	100.	110	743
	Due des stiens	F1 ↑	99.69	92.28	99.85	94.58	99.94	98.62	MPJRE↓	4.21	1.37	2.54	0.80	1.89	0.58	1.22	0.45	1.09	0.42	715
-	FIGUACTION	Acc. ↑	99.59	93.02	99.87	96.21	99.96	98.87	MPJPE↓	1.75	0.49	1.02	0.27	0.89	0.20	0.61	0.15	0.43	0.14	744
	Front Waist	F1 ↑	97.37	87.29	98.26	90.24	99.46	97.86	MPJRE↓	3.93	1.44	2.79	1.02	2.25	0.70	1.75	0.60	1.55	0.62	745
		Acc. ↑	98.59	92.18	99.28	93.87	99.68	98.23	MPJPE↓	1.66	0.30	1.11	0.32	0.99	0.25	0.82	0.21	0.69	0.21	746
	GRAB	F1 ↑	99.65	94.48	99.70	96.92	99.65	98.47	MPJRE ↓	3.32	6.37	2.69	2.21	2.15	1.72	1.83	1.59	1.70	1.61	747
		Acc. ↑	99.72	95.57	99.68	97.27	99.69	98.69	MPJPE↓	1.79	0.79	1.21	0.27	1.13	0.21	0.98	0.18	0.83	0.19	748
Table 1. Comparison with other methods on labeling (left) and solving (right). Cells with white background display metrics for body, those with gray										ay	749									
	hackground rei	present m	etrics fo	r hand	The F1	and accu	racy are	multiplie	d by 100. Th	e units	1.37 2.54 0.80 1.89 0.58 1.22 0.45 1.09 0.42 0.49 1.02 0.27 0.89 0.20 0.61 0.15 0.43 0.14 1.44 2.79 1.02 2.25 0.70 1.75 0.60 1.55 0.62 745 0.30 1.11 0.32 0.99 0.25 0.82 0.21 0.69 0.21 746 6.37 2.69 2.21 2.15 1.72 1.83 1.59 1.70 1.61 747 0.79 1.21 0.27 1.13 0.21 0.98 0.18 0.83 0.19 748 white background display metrics for body, those with gray 749 749 749 749 0.79 1.21 0.27 1.13 0.21 0.98 0.18 0.83 0.19 748 white background display metrics for body, those with gray 749 749 749 749									

background represent metrics for hand. The F1 and accuracy are multiplied by 100. The units of MPJRE and MPJPE are in degree and centimeter, respectively.

	F	1	Accuracy			
Base	98	.62	98.87			
- Local aggregation layers	98.59	-0.03	98.32	-0.55		
- Point cloud segmentation	95.20	-4.74	95.31	-3.56		
- Global transformation removal	98.10	-0.52	98.23	-0.64		
- Tracklet-based labeling	96.21	-2.41	96.79	-2.08		
- Feature similarity in edge weights	98.01	-0.61	98.38	-0.49		

Table 2. Ablation study of RoMo's components on the Production dataset's hand markers. We take the full model as the baseline and remove one component at a time.

4.2 Motion Solving

For motion solving, we compare RoMo with three neural-based approaches [Chen et al. 2021; Holden 2018; Pan et al. 2023]. Before solving, we apply the state-of-the-art occlusion fixing algorithm [Pan et al. 2023] to fill the missing markers, which first uses the Euclidean distance matrix optimization algorithm [Zhou et al. 2020] and then employ a bidirectional long short time memory network (BiLSTM) [Schuster and Paliwal 1997] to optimize the occluded markers' potisions. To ensure a fair comparison, the body and hand markers are aligned using the same strategy and trained using separate networks for neural approaches. Additionaly, we compare our method with two optimization-based methods: Vicon [Vicon 2023] for Production and Front-waist dataset, and Mosh++ [Loper et al. 2014] for the GRAB dataset. We use two quantitative metrics: mean per joint rotation error (MPJRE) and mean per joint rotation error (MPJPE). The former reflects the angle differences of the joints and latter represents the global joint position difference.

The quantitative results are displayed on the right in Table 1. For body motion, RoMo's MPJRE and MPJPE are approximately 15% and 25% lower, respectively, than those achieved by the best comparison methods. RoMo outperforms other neural-based methods for two primary reasons. Firstly, it employs joint positions as intermediate representations, simplifying the learning process for neural net-works. Secondly, RoMo leverages the inverse kinematics process to utilize the global joint positions, effectively avoiding the error accumulation along the kinematic tree. Other methods [Chen et al. 2021; Holden 2018; Pan et al. 2023] directly estimates the relative joint rotation of the whole body simultaneously. They are fragile to error accumulation, resulting in large discrepancy on ending joints. RoMo's solving accuracy of hand motions are comparable with the state-of-the-art methods, since the hand bone length of are relatively limited, and does not amplify parental joints' rotation errors.

We present the qualitative results of body motion solving in Fig. 5 and highlight RoMo's improvements. RoMo demonstrates better solving accuracy, especially on ending joints such as wrists and ankles. We present additional qualitative comparison on Fig. 7.

To better demonstrate RoMo's robustness against positional errors, we introduce random jitters to the input marker positions and record the resulting joint position errors. We compare RoMo with a state-of-the-art solution method [Pan et al. 2023]. Additionally, we validated the use of global joint positions J', with the results presented in Table 3. Methods that rely solely on joints' relative rotations [Pan et al. 2023] or relative positions tend to suffer from errors accumulation along the kinematic chain, leading to increased errors at higher jitter intensities. In contrast, RoMo, by using global joint positions, exhibits greater robustness to such noise.

	0	± 0.2 cm	$\pm 0.5 cm$	± 1.0 cm
LocalMoCap [Pan et al. 2023]	0.61	0.72	0.98	1.62
RoMo (w/o global joint position)	0.47	0.68	0.83	1.49
RoMo	0.43	0.62	0.76	1.25

Table 3. The MPJPE under various jitter intensities on the body data of Production dataset.



Fig. 5. Qualitative comparisons of solved motions. To compare with the ground truth, we render the skeletons and overlay the ground truth's skeleton onto those generated by solving methods. Positions with significant differences are indicated with red boxes.

CONCLUSION, LIMITATIONS AND FUTURE WORK

We introduce RoMo, a framework designed for robust labeling and solving of optical motion capture data of full-body, addressing two primary types of noise: mislabeling and positional errors. RoMo

856

857

858

859

860

861

862

863

864

865

866

867

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

905

906

907

908

909

799 utilizes innovative techniques such as a tracklet generation algo-800 rithm based on positional and feature similarities, coupled with 801 an inverse kinematics-based motion capture solver to avoid error 802 accumulation. Our evaluations on diverse and complex benchmarks 803 demonstrate RoMo's superior performance compared to previous methods, providing precise labeling of body and hand markers and 804 robustly solving full-body motions. Notably, RoMo achieves accu-805 racy comparable to commercial systems but offers the advantages 806 807 of being cost-free and highly flexible.

RoMo does have its limitations. Currently, it still splits the label-808 ing and solving stages, since the input of labeling stage is unordered 809 markers while the solving relies on labeled marker. Merging the 810 two stages, that is, solving motions directly from unordered point 811 cloud data, could diminish the reliance on heuristic assumptions 812 813 and improve the labeling and solving network's generalization capabilities. Furthermore, integrating motion information into the la-814 815 beling stage could potentially enhance label accuracy. Additionally, 816 RoMo does not account for variations in marker layouts. Training directly on a layout superset encompassing all markers of diverse 817 layouts [Ghorbani and Black 2021] substantially reduces labeling 818 accuracy. Exploring solutions to accommodate marker layout varia-819 820 tions within a single network without compromising performance presents a promising direction to augment RoMo's flexibility. 821

REFERENCES

822

823

824

825

826

827

831

836

855

- Ryan Prescott Adams and Richard S. Zemel. 2011. Ranking via Sinkhorn Propagation. arXiv:1106.1925 [stat.ML]
- Andreas Aristidou, Daniel Cohen-Or, Jessica K Hodgins, and Ariel Shamir. 2018. Selfsimilarity Analysis for Motion Capture Cleaning. In Computer graphics forum (CGF), Vol. 37, 297-309.
- 828 Richard Askey. 2005. The 1839 paper on permutations: its relation to the Rodrigues 829 formula and further developments. Mathematics and Social Utopias in France: Olinde 830 Rodrigues and His Times 28 (2005), 105-118.
- Jan Baumann, Björn Krüger, Arno Zinke, and Andreas Weber. 2011. Data-Driven Completion of Motion Capture Data. VRIPHYS 2011 - 8th Workshop on Virtual 832 Reality Interactions and Physical Simulations, 111–118.
- Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. 833 2020. Long-term human motion prediction with scene context. In ECCV 2020. 834 387 - 404.835
 - Kang Chen, Yupan Wang, Song-Hai Zhang, Sen-Zhe Xu, Weidong Zhang, and Shi-Min Hu. 2021. MoCap-Solver: A Neural Solver for Optical Motion Capture Data. ACM Transactions on Graphics (TOG) 40, 4, Article 84 (2021), 11 pages.
- 837 Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. 2023. 838 Executing your Commands via Motion Diffusion in Latent Space. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 839 18000-18010.
- 840 Wongun Choi. 2015. Near-online multi-target tracking with aggregated local flow 841 descriptor. In Proceedings of the IEEE International Conference on Computer Vision. 3029-3037 842
- Qiongjie Cui, Huaijiang Sun, Yue Kong, and Xiaoning Sun. 2021. Deep Human Dynamics 843 Prior. In Proceedings of the 29th ACM International Conference on Multimedia (MM). 4371-4379. 844
- Peng Dai, Xue Wang, Weihang Zhang, and Junfeng Chen. 2019. Instance Segmentation 845 Enabled Hybrid Data Association and Discriminative Hashing for Online Multi-846 Object Tracking. IEEE Transactions on Multimedia 21, 7 (2019), 1709-1723.
- Nima Ghorbani and Michael J Black. 2021. Soma: Solving optical marker-based mocap 847 automatically. In Proceedings of the IEEE/CVF International Conference on Computer 848 Vision (ICCV), 11117-11126.
- Saeed Ghorbani, Ali Etemad, and Nikolaus F Troje. 2019. Auto-labelling of markers in 849 optical motion capture by permutation learning. In Advances in Computer Graphics: 850 36th Computer Graphics International Conference, CGI 2019. 167-178
- 851 Daniel Holden. 2018. Robust solving of optical motion capture data by denoising. ACM Transactions on Graphics (TOG) 37, 4 (2018), 1-12. 852
- Yilei Hua, Wenhan Wu, Ce Zheng, Aidong Lu, Mengyuan Liu, Chen Chen, and Shiqian 853 Wu. 2023. Part Aware Contrastive Learning for Self-Supervised Action Recognition. 854 International Joint Conference on Artificial Intelligence (IJCAI) (2023).

- Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. 2018. Endto-end recovery of human shape and pose. In Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR). 7122-7131.
- Chanho Kim, Fuxin Li, Arridhana Ciptadi, and James M Rehg. 2015. Multiple hypothesis tracking revisited. In Proceedings of the IEEE International Conference on Computer Vision. 4696-4704.
- Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. 2019. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In Proceedings of the IEEE/CVF international conference on computer vision (ICCV). 2252-2261.
- Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, and Cewu Lu. 2021. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR). 3383-3393.
- Guodong Liu and Leonard McMillan. 2006. Estimation of Missing Markers in Human Motion Capture. The Visual Computer (TVC) 22 (2006), 721-728.
- Matthew M. Loper, Naureen Mahmood, and Michael J. Black. 2014. MoSh: Motion and Shape Capture from Sparse Markers. ACM Transactions on Graphics (TOG) 33, 6 (2014), 220:1-220:13.
- Dening Lu, Qian Xie, Mingqiang Wei, Kyle Gao, Linlin Xu, and Jonathan Li. 2022. Transformers in 3d point clouds: A survey. arXiv preprint arXiv:2205.07417 (2022).
- Anton Milan, Konrad Schindler, and Stefan Roth. 2013. Detection-and trajectorylevel exclusion in multiple object tracking. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 3682-3689.
- Tewodros Legesse Munea, Yalew Zelalem Jembre, Halefom Tekle Weldegebriel, Longbiao Chen, Chenxi Huang, and Chenhui Yang. 2020. The progress of human pose estimation: A survey and taxonomy of models applied in 2D human pose estimation. IEEE Access 8 (2020), 133330-133348.
- OptiTrack, 2024. Motive.
- Xiaoyu Pan, Bowen Zheng, Xinwei Jiang, Guanglong Xu, Xianli Gu, Jingxiang Li, Qilong Kou, He Wang, Tianjia Shao, Kun Zhou, and Xiaogang Jin. 2023. A Locality-based Neural Solver for Optical Motion Capture. SIGGRAPH Asia 2023 Conference Papers, Article 117 (2023), 10 pages.
- Dario Pavllo, Thibault Porssut, Bruno Herbelin, and Ronan Boulic. 2018. Real-time Marker-based Finger Tracking with Neural Networks. In 2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR). IEEE, 651-652.

Maksym Perepichka, Daniel Holden, Sudhir P Mudur, and Tiberiu Popa. 2019. Robust Marker Trajectory Repair for MoCap using Kinematic Reference. In Proceedings of the 12th ACM SIGGRAPH Conference on Motion, Interaction and Games (MIG). 1-10.

Gabriel Peyré and Marco Cuturi. 2020. Computational Optimal Transport. arXiv:1803.00567 [stat.ML]

- Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. 2020. SuperGlue: Learning Feature Matching with Graph Neural Networks. arXiv:1911.11763 [cs.CV]
- Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. IEEE transactions on Signal Processing 45, 11 (1997), 2673-2681.
- Jianbing Shen, Zhiyuan Liang, Jianhong Liu, Hanqiu Sun, Ling Shao, and Dacheng Tao. 2018. Multiobject tracking by submodular optimization. IEEE Transactions on Cybernetics 49, 6 (2018), 1990-2001.
- Hao Sheng, Jiahui Chen, Yang Zhang, Wei Ke, Zhang Xiong, and Jingyi Yu. 2018. Iterative multiple hypothesis tracking with tracklet-level association. IEEE Transactions on Circuits and Systems for Video Technology 29, 12 (2018), 3660-3672.
- Horesh Ben Shitrit, Jérôme Berclaz, François Fleuret, and Pascal Fua. 2014. Multi-Commodity Network Flow for Tracking Multiple People. IEEE Transactions on Pattern Analysis and Machine Intelligence 36, 8 (2014), 1614-1627.
- Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. 2020. GRAB: A Dataset of Whole-Body Human Grasping of Objects. In European Conference on Computer Vision (ECCV). 581-600.
- Jochen Tautges, Arno Zinke, Björn Krüger, Jan Baumann, Andreas Weber, Thomas Helten, Meinard Müller, Hans-Peter Seidel, and Bernd Eberhardt. 2011. Motion Reconstruction Using Sparse Accelerometer Data. ACM Transactions on Graphics (TOG) 30, 3, Article 18 (2011), 12 pages.
- Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. 2023. Human Motion Diffusion Model. In The Eleventh International Conference on Learning Representations (ICLR).
- Mickaël Tits, Joëlle Tilmanne, and Thierry Dutoit. 2018. Robust and Automatic Motioncapture Data Recovery using Soft Skeleton Constraints and Model Averaging. PloS one 13 (2018), 1-21.
- Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. 2017. Learning from Synthetic Humans. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 109 - 117.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems (NIPS) 30 (2017). Vicon. 2023. Vicon Shogun.
- 910 911 912

- Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. 2019. Dynamic graph cnn for learning on point clouds. ACM Transactions on Graphics (TOG) 38, 5 (2019), 1–12.
 Zhoa Wang, Shuang Liu, Rongoiang Qian Tao Jiang Xiaosong Yang, and Jian L Zhang.
- ⁹¹⁵ Zhao Wang, Shuang Liu, Rongqiang Qian, Tao Jiang, Xiaosong Yang, and Jian J Zhang.
 ⁹¹⁶ 2016. Human Motion Data Refinement Unitizing Structural Sparsity and Spatialtemporal information. In 2016 IEEE 13th International Conference on Signal Processing (ICSP). 975–982.
- 918
 Jun Xiang, Guohan Xu, Chao Ma, and Jianhua Hou. 2020. End-to-end learning deep CRF

 919
 models for multi-object tracking deep CRF models. IEEE Transactions on Circuits

 920
 and Systems for Video Technology 31, 1 (2020), 275–288.
- Jun Xiao, Yinfu Feng, Mingming Ji, Xiaosong Yang, Jian J Zhang, and Yueting Zhuang. 2015. Sparse Motion Bases Selection for Human Motion Denoising. *Signal Processing* 110 (2015), 108–122.
- Sijie Yan, Yuanjun Xiong, and Dahua Lin. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI conference* on artificial intelligence (AAAI), Vol. 32.
- Yang Zhang, Hao Sheng, Yubin Wu, Shuai Wang, Weifeng Lyu, Wei Ke, and Zhang Xiong. 2020. Long-term tracking with deep tracklet association. *IEEE Transactions* on Image Processing (TIP) 29 (2020), 6694–6706.
- Shenglong Zhou, Naihua Xiu, and Hou-Duo Qi. 2020. Robust Euclidean embedding via EDM optimization. *Mathematical Programming Computation* 12 (2020), 337–387.

- - ...

SIGGRPAH ASIA'24, December 3-6, 2024, Tokyo

Anon. Submission Id: 497





Fig. 7. Additional qualitative comparisons of solved motions.