

Enhancing Context-Aware Human Motion Prediction for Efficient Robot Handovers

Gerard Gómez-Izquierdo^{1†}, Javier Laplaza¹, Alberto Sanfeliu¹ and Anaís Garrell¹

Abstract—Accurate human motion prediction (HMP) is critical for seamless human-robot collaboration, particularly in handover tasks that require real-time adaptability. Despite the high accuracy of state-of-the-art models, their computational complexity limits practical deployment in real-world robotic applications. In this work, we enhance human motion forecasting for handover tasks by leveraging *siMLPe* [1], a lightweight yet powerful architecture, and introducing key improvements. Our approach, named *IntentMotion* incorporates intention-aware conditioning, task-specific loss functions, and a novel intention classifier, significantly improving motion prediction accuracy while maintaining efficiency. Experimental results demonstrate that our method reduces body loss error by over 50%, achieves 200× faster inference, and requires only 3% of the parameters compared to existing state-of-the-art HMP models in robotics. These advancements establish our framework as a highly efficient and scalable solution for real-time human-robot interaction.

I. INTRODUCTION

Human motion prediction (HMP) plays a crucial role in human-robot collaboration (HRC) by enabling robots to anticipate human movements and respond proactively. This capability is particularly important in handover tasks, where the seamless exchange of objects between humans and robots requires both accuracy and speed. The ability to predict human motion allows robots to preemptively adjust their trajectories, improving efficiency and ensuring safety. In this context, human intention (collaborative or non-collaborative) directly influences the prediction and subsequent robot response. As a result, HMP is a critical component in various real-world robotic applications, including industrial automation, assistive robotics, and service robots.

Recent advances in deep learning, including RNNs, GCNs, and Transformers, have led to highly accurate HMP by modeling spatial-temporal dynamics. However, their high computational cost limits real-time deployment in robotics. A key challenge lies in balancing accuracy with efficiency; many models favor precision over practicality and often overlook human intention, a critical factor that shapes motion.

To tackle these challenges, we adapt the lightweight *siMLPe* [1] architecture for human-robot handovers in a new approach called *IntentMotion*. It introduces task-specific loss functions, an intention-aware conditioning mechanism, and a novel intention classifier to refine motion predictions.

[†] Main author of the paper

¹Institut de Robotica y Informatica Industrial (CSIC-UPC) and Universitat Politècnica de Catalunya - BarcelonaTech (UPC), Barcelona, Spain {gerard.gomez.izquierdo}@estudiantat.upc.edu

This work was partially supported by JST Moonshot R, D grant number JPMJMS2011 and LENA (PID2022-142039NA-I00), funded by MCIN/AEI/10.13039/501100011033

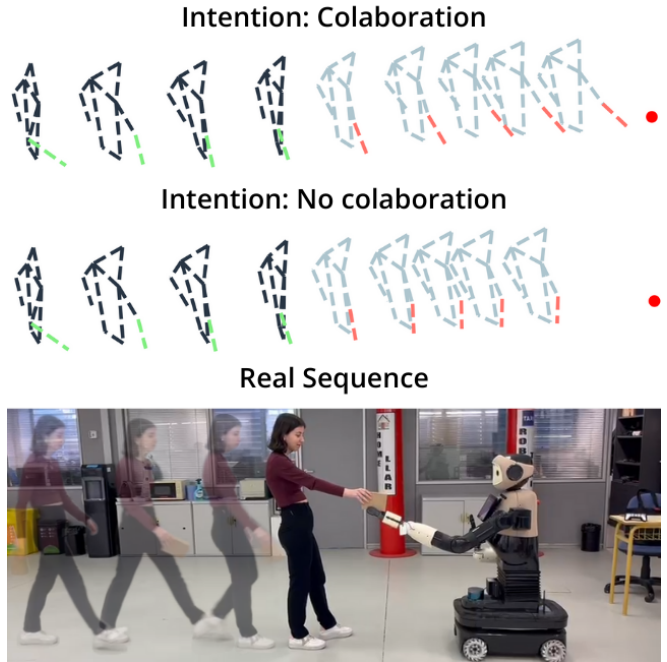


Fig. 1. **Motion prediction based on detected intention.** Both sequences share the same ground truth represented by a dark blue skeleton and a green right hand. **Top:** Prediction with collaborative intention. **Middle:** Prediction with non-collaborative intention. **Red Circle:** Represents the last frame of the robot’s end effector. **Bottom:** Real handover sequence.

Evaluated in a handover setting, *IntentMotion* predicts human motion and classifies intention (collaborative or non-collaborative) with high accuracy while significantly reducing computational cost, making it suitable for real-time HRC.

The paper is organized as follows: Section II reviews related work, Section III details our methodology, Section IV presents the experiments, Section V discusses results, and Section VI concludes the study.

II. RELATED WORK

Several studies have addressed Human Motion Prediction (HMP) in human-robot interaction to improve motion forecasting accuracy and efficiency. Liu and Wang [2] proposed a predictive framework for collaborative manufacturing, enabling robots to anticipate human motion. Zhao et al. [3] evaluated HMP models in collaborative settings, emphasizing the benefits of parameter-efficient neural networks. Mugisha et al. [4] used Gaussian Processes to predict hand motion, integrating gaze and hand data for enhanced intention detection.

Laplaza et al. [5], [6] presented a relevant approach for predicting human motion in handover tasks. Building

on Mao et al.’s architecture [7], their model incorporates human intention as a conditioning factor and was trained on a custom handover dataset. We also use this dataset to benchmark our improvements in computational efficiency and accuracy over their method.

Deep learning methods such as Recurrent Neural Networks (RNNs) have been widely applied to HMP for their ability to model temporal sequences. Encoder-decoder LSTM models [8], [9] were extended by Martinez et al. [10] with residual connections to better model motion velocities. However, RNNs often struggle with vanishing gradients and long-term dependencies.

To address these issues, Graph Convolutional Networks (GCNs) have been explored for capturing spatial dependencies. Mao et al. [7] proposed a GCN model encoding spatial-temporal features, further refined with attention mechanisms in [11] to emphasize relevant body movements. Ma et al. [12] introduced multiscale representations to improve long-term predictions.

Transformer-based models have also gained traction. Ak-san et al. [13] introduced a spatio-temporal Transformer to jointly capture spatial and temporal dynamics, while Mao et al. [11] proposed a memory-augmented variant to retrieve historical patterns. Despite their accuracy, these models are often resource-intensive and less suitable for real-time deployment.

A recent breakthrough in human motion prediction is the introduction of *siMLPe*, a lightweight yet powerful MLP-based model proposed by Guo et al. [1]. Unlike deep GCNs or Transformer-based architectures, *siMLPe* applies a Discrete Cosine Transform (DCT) to encode temporal dependencies and predicts motion using fully connected layers with layer normalization. Despite its simplicity, this model outperforms previous architectures on major benchmarks while maintaining an extremely low computational footprint. Its efficiency and accuracy make it particularly well-suited for real-time robotic applications.

The present work builds upon previous research by adapting *siMLPe* for human-robot handover tasks and integrating intention-aware conditioning to enhance motion prediction accuracy and computational efficiency.

III. METHODOLOGY

In this section, we present our approach for intention-aware human motion prediction in handover tasks. We enhance *siMLPe* [1] by introducing *IntentMotion*, our intention-aware motion prediction framework.

Here, first define the problem formulation, detailing how human motion sequences and intention labels are represented. Next, we describe the architecture of our proposed *IntentMotion* framework, which extends prior frequency-domain motion prediction models by incorporating intention as contextual information. We then introduce key modifications, including an intention embedding mechanism and novel loss functions designed to enhance the accuracy of right-hand motion. Finally, we outline the design and training

objectives of our intention classifier, which enables real-time intent recognition to dynamically refine motion forecasts.

A. Problem Definition

Given a sequence of past 3D human poses, the objective is to accurately forecast the corresponding future motion. The observed human motion sequence is represented as $\mathbf{x}_{1:T} = [\mathbf{x}_1, \dots, \mathbf{x}_T] \in \mathbb{R}^{T \times C}$, where each pose $\mathbf{x}_t \in \mathbb{R}^C$ at time step t is encoded as a C -dimensional vector. Specifically, \mathbf{x}_t contains the 3D coordinates of body joints, with $C = 3 \times K$, where K denotes the total number of joints. The goal is to predict the sequence of future motion frames, formulated as $\mathbf{x}_{T+1:T+N} = [\mathbf{x}_{T+1}, \dots, \mathbf{x}_{T+N}] \in \mathbb{R}^{N \times C}$.

Additionally, each input sequence $\mathbf{X}_{1:T}$ is associated with an intention label $i \in \{0, 1\}$, where $i = 0$ corresponds to a collaborative intention and $i = 1$ to a non-collaborative intention. This intention variable encodes the expected human intent in the predicted motion frames, denoted as $\hat{\mathbf{i}}_{N+1:N+T}$, providing contextual information to enhance motion forecasting accuracy.

B. Network Overview

IntentMotion architecture, as done in [1] operates in the frequency domain by applying the Discrete Cosine Transform (DCT) to encode temporal dependencies, followed by fully connected layers, transpose operations, and layer normalization. The output is then transformed back to the original pose representation using the Inverse Discrete Cosine Transform (IDCT). Given an input sequence of past 3D human poses $\mathbf{x}_{1:T} = [\mathbf{x}_1, \dots, \mathbf{x}_T] \in \mathbb{R}^{T \times C}$ the model predicts the sequence of future poses $\mathbf{x}'_{T+1:T+N} = [\mathbf{x}'_{T+1}, \dots, \mathbf{x}'_{T+N}] \in \mathbb{R}^{N \times C}$. The prediction process can be formulated as: $\hat{\mathbf{x}}_{T+1:T+N} = \mathcal{D}^{-1}(\mathcal{F}(\mathcal{D}(\mathbf{x}_{1:T})))$, where \mathcal{F} denotes the network function, \mathcal{D} is the DCT transformation, and \mathcal{D}^{-1} represents the IDCT transformation.

C. Proposed Modifications

Figure 2 shows the architecture of our network.

1) *Incorporating Intention as Context*: To integrate intention into the model, we encode the intention label $i \in \{0, 1\}$ using an embedding layer:

$$\mathbf{e}_i = \text{Embedding}(i), \quad \mathbf{e}_i \in \mathbb{R}^{1 \times C}. \quad (1)$$

The embedded intention vector is then expanded along the temporal dimension to match the sequence length T , forming the intention matrix:

$$\mathbf{M}_i = \mathbf{e}_i \otimes \mathbf{1}_T, \quad \mathbf{M}_i \in \mathbb{R}^{T \times C}, \quad (2)$$

where \otimes denotes broadcasting across the temporal dimension.

Next, we fuse the encoded intention with the transformed motion sequence \mathbf{z}^0 , obtained after the first fully connected layer:

$$\mathbf{z}^0 = \mathcal{D}(\mathbf{x}_{1:T})\mathbf{W}_0 + \mathbf{b}_0, \quad (3)$$

where $\mathbf{W}_0 \in \mathbb{R}^{C \times C}$ and $\mathbf{b}_0 \in \mathbb{R}^C$ are learnable parameters. The fusion is achieved through element-wise addition:

$$\mathbf{U} = \mathbf{M}_i + \mathbf{z}^0. \quad (4)$$




Fig. 2. **Overview of our *siMLPe*-based approach for human motion prediction.** *FC* denotes a fully connected layer, *LN* represents layer normalization, *Trans* indicates a transpose operation, *EMB* denotes the embedding layer, and *AVG* represents the average pooling layer. *DCT* and *IDCT* correspond to the discrete cosine and inverse discrete cosine transformations, respectively. The MLP module (highlighted in pink) consists of 48 *FC* and *LN* layers repeated across the architecture. The green box represents how the intention classifier processes data before intention prediction.

This enriched representation \mathbf{U} encodes both the historical motion and the expected human intention, enabling the model to generate predictions that are contextually aware of the intention label.

2) *Objective Function*: The objective function of our model is designed to improve the predictive performance by incorporating additional loss terms that explicitly account for the movement dynamics of the right hand during handover tasks. Our total loss function, \mathcal{L}_h , extends the loss formulation of *siMLPe* [1], which originally includes the terms \mathcal{L}_{re} and \mathcal{L}_v , by adding three additional terms: \mathcal{L}_c , \mathcal{L}_{rer} , and \mathcal{L}_{vr} . The final objective function is given by:

$$\mathcal{L}_h = \mathcal{L}_{re} + \mathcal{L}_v + \mathcal{L}_c + \mathcal{L}_{rer} + \mathcal{L}_{vr}. \quad (5)$$

The standard reconstruction loss \mathcal{L}_{re} and velocity loss \mathcal{L}_v from [1] are defined as:

$$\mathcal{L}_{re} = \|\mathbf{x}'_{T+1:T+N} - \mathbf{x}_{T+1:T+N}\|_2^2, \quad (6)$$

$$\mathcal{L}_v = \|\mathbf{v}'_{T+1:T+N} - \mathbf{v}_{T+1:T+N}\|_2^2, \quad (7)$$

where \mathbf{x}' and \mathbf{x} denote the predicted and ground-truth motion sequences, respectively, and \mathbf{v} represents the velocity, computed as $\mathbf{v}_t = \mathbf{x}_{t+1} - \mathbf{x}_t$.

To overcome the limitation of *siMLPe*, which tends to predict static right-hand poses, we introduce an additional loss term \mathcal{L}_c that ensures the right hand moves towards the robot end effector in collaborative scenarios. This term consists of two sub-components: \mathcal{L}_r and \mathcal{L}_b :

$$\mathcal{L}_c = 0.05\mathcal{L}_r + 0.95\mathcal{L}_b. \quad (8)$$

Here, \mathcal{L}_r minimizes the Euclidean distance between the last predicted right hand position $\mathbf{x}'_{rh,T+N}$ and the robot end effector \mathbf{ree}_{T+N} :

$$\mathcal{L}_r = \|\mathbf{x}'_{rh,T+N} - \mathbf{ree}_{T+N}\|_2^2. \quad (9)$$

Additionally, \mathcal{L}_b penalizes deviations from the natural movement constraints of the right hand by minimizing the

difference between the right hand and K other body joints in the predicted and ground-truth sequences:

$$\mathcal{L}_b = \left\| \frac{1}{K} \sum_{k=1}^K \|\mathbf{x}'_{rh} - \mathbf{x}'_k\|_2 - \frac{1}{K} \sum_{k=1}^K \|\mathbf{x}_{rh} - \mathbf{x}_k\|_2 \right\|_2^2. \quad (10)$$

This term enforces consistency in the natural movement patterns of the right hand relative to the rest of the body.

Apart from \mathcal{L}_c , to further guide the prediction of right-hand motion, we introduce \mathcal{L}_{rer} and \mathcal{L}_{vr} , which emphasize minimizing positional and velocity errors of the right hand. The error for the right hand, \mathcal{L}_{rer} , is formulated as:

$$\mathcal{L}_{rer} = \|\mathbf{x}'_{rh,T+1:T+N} - \mathbf{x}_{rh,T+1:T+N}\|_2^2. \quad (11)$$

Similarly, the velocity consistency loss \mathcal{L}_{vr} is given by:

$$\mathcal{L}_{vr} = \|\mathbf{v}'_{rh,T+1:T+N} - \mathbf{v}_{rh,T+1:T+N}\|_2^2, \quad (12)$$

where \mathbf{v}'_{rh} and \mathbf{v}_{rh} represent the predicted and ground-truth velocities of the right hand, respectively.

It is important to note that the loss term \mathcal{L}_c is applied selectively, being computed only for sequences labeled with collaborative intent. This ensures that the model learns context-specific behavior without introducing biases in non-collaborative scenarios.

D. Proposed Intention Classifier

To enable automatic conditioning of the motion predictor with the human's intention in real-world scenarios, we propose an intention classifier that infers intention labels directly from body keypoint trajectories. This classifier is designed to be lightweight and efficient, drawing inspiration from [1] while adapting its structure for classification tasks. Figure 2 also shows the architecture of the classifier.

1) *Architecture*: The classifier adopts a similar backbone to *IntentMotion* but omits the intention embedding module. Instead, the output feature map from the last fully connected layer is aggregated across the temporal dimension via average pooling. Let $\mathbf{F} \in \mathbb{R}^{T \times C}$ denote the feature representation

extracted from an input sequence of T frames, where C is the feature dimension. The pooled feature is computed as:

$$\bar{\mathbf{F}} = \frac{1}{T} \sum_{t=1}^T \mathbf{F}_t. \quad (13)$$

This pooled feature is subsequently mapped to class logits through a series of fully connected layers with nonlinear activations. Formally, the classifier computes:

$$\mathbf{z} = g(\bar{\mathbf{F}}) \in \mathbb{R}^M, \quad (14)$$

where $g(\cdot)$ represents the mapping function composed of two fully connected layers, and M is the number of intention classes (in our experiments, $M = 2$). The predicted intention is then given by:

$$\hat{y} = \operatorname{argmax}(\mathbf{z}). \quad (15)$$

2) *Objective Function*: The training objective for the classifier modifies the loss formulation of the motion predictor. In addition to the \mathcal{L}_{re} loss and the velocity loss \mathcal{L}_v already defined in [1], we incorporate a cross-entropy loss \mathcal{L}_{ce} to supervise the intention predictions. The overall objective is formulated as:

$$\mathcal{L}_{class} = \mathcal{L}_{re} + \mathcal{L}_v + \mathcal{L}_{ce}. \quad (16)$$

Here, the cross-entropy loss is defined by:

$$\mathcal{L}_{ce} = - \sum_{i=1}^M y_i \log \left(\frac{\exp(z_i)}{\sum_{j=1}^M \exp(z_j)} \right), \quad (17)$$

where y_i is the ground-truth label for class i and z_i denotes the corresponding logit. This formulation ensures that the classifier performs classification capturing the spatio-temporal characteristics of the motion to provide reliable intention cues to condition the motion predictions.

IV. EXPERIMENTATION

A. Dataset

We adopt a dataset originally detailed by Laplaza et al. [5] to train and evaluate our approach. In this dataset, human-robot handover interactions are captured in controlled experiments using the anthropomorphic robot IVO. During each trial, a human volunteer, acting as the giver, transfers a 10 cm cylindrical object to the robot, which functions as the receiver by aligning its motion to successfully grasp the object. All interactions commence from a 6-meter separation and conclude upon completion of the transfer.

Model	Body L_2 (m)	$\leq 0.35\text{m}$ (%)	$\leq 0.40\text{m}$ (%)	Right Hand L_2 (m)
RNN [10]	0.739	3.49	11.62	0.677
Hist. Rep. Its. [7]	0.403	34.13	37.14	0.188
Laplaza et al. [5]	0.355	32.15	35.71	0.151
siMLPe [1]	0.177	96.95	100.00	0.217

TABLE I

COMPARISON OF ORIGINAL siMLPE WITH PREVIOUS MODELS

The dataset encompasses multiple scenarios: one without obstacles, another with a single obstacle, and a final setup featuring three obstacles. Additionally, various approach paths were defined to capture both lateral and straight trajectories, as shown in Figure 3. Volunteers executed each trajectory under three behavioral conditions: a natural execution, a gesture-modulated variant, and an adversarial variant. For our purposes, the original four intention labels are consolidated into two classes. Specifically, both natural and gesture-modulated behaviors are treated as *collaboration*, while neutral and adversarial behaviors are grouped as *no collaboration*.

B. Training Details

Both *IntentMotion* and *IntentMotion Classifier* are trained using sequences of 50 frames as input, with the subsequent 10 frames (corresponding to 1 second of data) serving as the target output. To extend the prediction horizon to 2.5 seconds, the models are deployed in an autoregressive way. In the case of the intention classifier, the final predicted intention is determined by taking the mode across three consecutive blocks of 10 predicted frames.

Models are trained for 5000 epochs using the Adam optimizer combined with a cosine annealing scheduler, which adjusts the learning rate from 10^{-2} down to 10^{-5} . A batch size of 256 is used during training. The feature dimension is defined as $C = 3 \times K$, where $K = 9$ represents the number of upper-body keypoints extracted from the handover data.

In addition, data augmentation is employed by randomly inverting the motion sequence during training, similar to the procedure in [1]. An extensive ablation study was conducted to fine-tune all model hyperparameters, thereby ensuring optimal performance across both tasks.

C. Evaluation Details

To evaluate *IntentMotion*, we employ metrics analogous to those by Laplaza et al. [5]. The primary evaluation metric is the mean L2 distance between the predicted and ground truth upper-body 3D joint positions, computed for each frame.

To facilitate a more detailed analysis, we report the percentage of predicted frames with errors below predefined thresholds of 0.35m and 0.40m, metrics that directly enable comparison with previous approaches (see Table I). In our ablation study shown in Table III, we further introduce thresholds of 0.30m and 0.20m, given that our best model nearly achieves 100% accuracy under the 0.35m and 0.40m criteria.

Moreover, since the original *siMLPe* [1] architecture tends to produce static predictions for the right-hand, which is critical in handover tasks, we incorporate a dedicated L2

Model	Mean inference time (ms)	# Parameters
Laplaza [5] + int.	3,680.6	6,113,782
IntentMotion	18.5	126,558
IntentMotion Classifier	18.7	265,032

TABLE II

COMPARISON OF INFERENCE TIME AND PARAMETERS




Fig. 3. **The figure presents three top-view scenarios from the dataset.** The human (right figure) and the robot (left figure) interact in different environments, with obstacles depicted as red squares. Dashed lines indicate human movement, while the robot follows corresponding target points.

error metric for the right-hand joint to quantify improvements in its motion prediction (see Table III).

For the *IntentMotion Classifier*, our evaluation is centered on the macro F1-score. This metric is well-suited to handle the inherent class imbalance in our dataset, where collaborative samples predominate over non-collaborative ones.

To robustly assess both models, as done in [5] we adopt a leave-one-out cross-validation strategy over 10 subjects, where each subject is sequentially held out for testing while the remaining subjects form the training set. This approach ensures a comprehensive evaluation of model generalization across different individuals.

Finally, we evaluate the computational efficiency of our models by measuring the average inference time (in milliseconds over 100 inferences) and by comparing the number of model parameters with those in [5] (see Table II).

V. RESULTS AND DISCUSSION

This section presents a comprehensive evaluation of our proposed *IntentMotion* framework, analyzing both quantitative performance metrics and qualitative insights. We first compare our model against existing baselines in terms of computational efficiency and motion prediction accuracy, emphasizing the improvements introduced by intention conditioning. We then conduct ablation studies to assess the contributions of individual loss functions and intention-aware mechanisms. Additionally, we evaluate the performance of the *IntentMotion Classifier* in recognizing human intent, highlighting its impact on motion prediction. Finally, qualitative results provide visual evidence of how our approach refines human-robot handover interactions, demonstrating its ability to adapt motion predictions based on inferred intention.

A. Quantitative Metrics

As shown in Table II, our proposed model achieves a substantial reduction in both inference time and parameter count compared to the baseline in [5], requiring only 3% of the parameters and achieving 200× faster inference. This gain in computational efficiency is particularly valuable in robotic scenarios, where multiple specialized models may need to operate in real time to support robust human-robot collaboration.

Moreover, as shown in Table I our approach significantly lowers the L_2 error for the entire body, improving from 0.355m to 0.165m relative to [5]. Such accuracy is essential for safe and precise interactions in collaborative tasks, where

even minor positional inaccuracies can undermine human comfort or safety.

The ablation results in Table III further highlight the benefits of incorporating intention conditioning. When the ground-truth intention label is provided, every performance metric exceeds that of the finetuned *IntentMotion* baseline, demonstrating how context awareness enhances predictive fidelity and reinforcing the importance of a reliable intention classifier.

Because handover tasks require highly accurate right-hand positioning, we additionally measure the L_2 distance for the right-hand joint. This evaluation underscores the limitations of the original *siMLPe* framework, which often produces static right-hand poses. By contrast, our model, through new loss terms and refined training, avoids static right hand predictions and reduces the joint-level error from 0.217m to 0.195m (see Table III).

Table III also illustrates how the introduction of \mathcal{L}_c , \mathcal{L}_{rer} , and \mathcal{L}_{vr} yields better motion predictions. Although \mathcal{L}_c alone do not provide quantitative gains, its combination with the right-hand-focused losses (\mathcal{L}_{rer} and \mathcal{L}_{vr}) substantially improves both body-level and right-hand accuracy.

The next section offers a deeper examination of how these individual losses specifically enhance the quality of right-hand trajectories.

Finally, we assess *IntentMotion Classifier*'s performance. As reported in Table IV, incorporating the cross-entropy loss \mathcal{L}_{ce} alongside \mathcal{L}_{re} and \mathcal{L}_v proposed by [1] notably boosts classification f1-score. This improvement likely stems from the model's ability to capture key motion characteristics, such as velocity and joint positioning, that are crucial for recognizing human intention. Furthermore, we evaluate two data-processing strategies for the classifier input and find that averaging the temporal dimension leads to slightly higher performance than flattening data.

B. Qualitative Results

In addition to the improvements quantified in our ablation studies, we emphasize that standard metrics may not fully capture the nuances of model performance. Therefore, we include visualizations to elucidate how our architectural modifications and changes in the objective function impact the generated predictions.

As illustrated in Figure 1, our model dynamically adjusts its predictions based on the inferred intention. When a collaborative intention is detected, the predicted motion

Model	Body L_2 (m)	$\leq 0.20\text{m}$ (%)	$\leq 0.30\text{m}$ (%)	Right Hand L_2 (m)
siMLPe [1]	0.177	64.53	93.48	0.217
IntentMotion FT	0.171	69.11	95.27	0.215
IntentMotion FT + Int	0.168	71.66	97.68	0.200
IntentMotion FT + Int + \mathcal{L}_c	0.175	65.94	96.58	0.237
IntentMotion FT + Int + $\mathcal{L}_c + \mathcal{L}_{vr}$	0.172	64.32	95.30	0.205
IntentMotion FT + Int + $\mathcal{L}_c + \mathcal{L}_{vr} + \mathcal{L}_{rer}$	0.165	70.66	98.36	0.195

TABLE III
INTENTMOTION RESULTS ON LEAVE-ONE-OUT EVALUATION

Model	F1 macro	Accuracy (%)
Classifier input flat. + \mathcal{L}_{ce}	0.417	63.94
Classifier input flat. + $\mathcal{L}_{ce} + \mathcal{L}_{re} + \mathcal{L}_v$	0.633	79.40
Classifier input avg. + $\mathcal{L}_{ce} + \mathcal{L}_{re} + \mathcal{L}_v$	0.654	81.10

TABLE IV
INTENTION CLASSIFIER RESULTS ON LEAVE-ONE-OUT EVALUATION

depicts a prototypical handover scenario, characterized by the human approaching the robot and elevating the right hand to deliver the object. Conversely, when a non-collaborative intention is forced from the same input data, the prediction reflects a more reserved behavior, where the body halts its forward progression earlier and the right hand remains in a lower position. These visualizations clearly demonstrate that our architecture effectively incorporates intention as contextual information, leading to different predictions even when provided with the same input motion data.

Moreover, the introduction of the loss terms \mathcal{L}_c , \mathcal{L}_{rer} , and \mathcal{L}_{vr} specifically enhances the modeling of right-hand dynamics. Visual comparisons indicate that \mathcal{L}_c helps the model better distinguish between intention-driven outputs. Additionally, the combined use of \mathcal{L}_c , \mathcal{L}_{rer} , and \mathcal{L}_{vr} results in significantly smoother and more realistic right-hand trajectories. In the absence of \mathcal{L}_{rer} and \mathcal{L}_{vr} , the predicted right-hand movement appears more abrupt, underscoring the efficacy of including them in this task.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we proposed a lightweight and efficient approach for human motion prediction in human-robot handovers, demonstrating that *IntentMotion* achieves state-of-the-art performance while drastically improving real-time feasibility. By integrating human intention as contextual information, we enhanced both prediction accuracy and quality, allowing for more natural and adaptive human-robot interactions. Furthermore, we introduced task-specific loss terms designed to capture handover-specific motion patterns, reinforcing the model’s ability to generalize effectively across various scenarios.

A key contribution of this study is the development of an intention classifier within the *IntentMotion* framework, enabling automatic intent recognition and linking human intention to motion prediction. This enhances forecasting accuracy and facilitates seamless collaboration by allowing robots to anticipate human actions more effectively.

Our findings emphasize the potential of lightweight, context-aware architectures for human-robot collaboration. By reducing computational complexity while maintaining accuracy, our approach supports real-time deployment in industrial automation, assistive robotics, and service robots. Future research will refine intention recognition, explore hybrid architectures integrating diverse contextual cues, and extend the method to broader HRI applications beyond handover tasks, advancing intelligent and human-centric robotic systems.

REFERENCES

- [1] W. Guo, Y. Du, X. Shen, V. Lepetit, X. Alameda-Pineda, and F. Moreno-Noguer, “Back to mlp: A simple baseline for human motion prediction,” in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2023, pp. 4809–4819.
- [2] H. Liu and L. Wang, “Human motion prediction for human-robot collaboration,” *Journal of Manufacturing Systems*, vol. 44, pp. 287–294, 2017.
- [3] W. Zhao, L. Sun, C. Liu, and M. Tomizuka, “Experimental evaluation of human motion prediction toward safe and efficient human robot collaboration,” in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 4349–4354.
- [4] S. Mugisha, V. K. Guda, C. Chevallereau, D. Chablat, and M. Zoppi, “Motion prediction with gaussian processes for safe human-robot interaction in virtual environments,” *IEEE Access*, 2024.
- [5] F. M.-N. Javier Laplaza and A. Sanfeliu, “Context and intention aware 3d human body motion prediction using an attention deep learning model in handover tasks,” in *IEEE*, 2022, p. 4743–4748.
- [6] J. Laplaza, A. Garrell, F. Moreno-Noguer, and A. Sanfeliu, “Context and intention for 3d human motion prediction: experimentation and user study in handover tasks,” in *2022 31st IEEE international conference on robot and human interactive communication (RO-MAN)*. IEEE, 2022, pp. 630–635.
- [7] W. Mao, M. Liu, and M. Salzmann, “History repeats itself: Human motion prediction via motion attention,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*. Springer, 2020, pp. 474–489.
- [8] K. Fragkiadaki, S. Levine, P. Felsen, and J. Malik, “Recurrent network models for human dynamics,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4346–4354.
- [9] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik, “Human pose estimation with iterative error feedback,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4733–4742.
- [10] J. Martinez, M. J. Black, and J. Romero, “On human motion prediction using recurrent neural networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2891–2900.
- [11] W. Mao, M. Liu, M. Salzmann, and H. Li, “Multi-level motion attention for human motion prediction,” *International journal of computer vision*, vol. 129, no. 9, pp. 2513–2535, 2021.
- [12] T. Ma, Y. Nie, C. Long, Q. Zhang, and G. Li, “Progressively generating better initial guesses towards next stages for high-quality human motion prediction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6437–6446.
- [13] E. Aksan, M. Kaufmann, P. Cao, and O. Hilliges, “A spatio-temporal transformer for 3d human motion prediction,” in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 565–574.