Early Classification of Human Motion Intent for Exoskeleton Assistance Using Kinematics

Aymeric Orhan, Duy Hoàng, Olivier Bruneau, Bastien Berret, Franck Geffard

Early Classification of Human Motion Intent for Exoskeleton Assistance Using Kinematics ⋆

Aymeric Orhan [*,**,†] **Duy Hoang** [***] **Olivier Bruneau** [*]
Bastien Berret [****] **Franck Geffard** [**]

[*] *LURPA, Mechanical Engineering Department, ENS Paris-Saclay, Université Paris-Saclay, Gif-sur-Yvette, 91190, France*
[**] *CEA, List, Université Paris-Saclay, Palaiseau, 91120, France*
[***] *Université de Versailles Saint-Quentin-en-Yvelines, Université Paris-Saclay, Versailles, 78035, France*
[****] *CIAMS, Université Paris-Saclay, Inria, Gif-sur-Yvette, 91190, France*
[†] *Corresponding author : orhanaymeric@gmail.com*

Abstract: Exoskeletons can reduce physical effort and the risk of work-related musculoskeletal disorders by providing robotic assistance in repetitive tasks. Designing exoskeletons that seamlessly assist human users without disrupting their natural movements poses a significant challenge, necessitating accurate prediction and adaptation to their motion intent. Particularly, an erroneous detection of motion intent could result in large adversarial effects where the exoskeleton resists the user's desired movement. We propose to analyze the possibility of early and accurate intent detection using only kinematic information and quantify the disturbance that a failed intention detection could entail. We evaluate different classification methods with voluntarily ambiguous experimental data on the intention underlying reaching movements towards four targets in a parasagittal plane. We show that, while recent advances in time series classification —namely using a convolutional and residual neural network— can enable earlier and more accurate intent detection, an informed adaptation of the assistance according to the classification's confidence level is necessary.

Keywords: Learning robot control; Haptic interaction; Adaptive robot control

1. INTRODUCTION

Exoskeletons have proven effective in assisting human movement across various applications, from rehabilitation Proietti et al. (2016) to workplace ergonomics Nussbaum et al. (2019). By targeting specific joints or muscles, exoskeletons can reduce strain, mitigate musculoskeletal disorders, and improve endurance. However, ensuring seamless and natural integration with the human user remains a significant challenge due to the high complexity, variability and adaptability of their movements.

To collaborate effectively with humans, robots must anticipate and predict their partners' motion intent Li et al. (2022). In other words, the robot must rapidly identify the target and predict the intended trajectory to provide effective assistance. In human-exoskeleton interaction, early and accurate intent detection is critical, as the exoskeleton must anticipate the user's goal to adapt assistance appropriately throughout the movement. Inaccurate intent detection can lead to adversarial effects, causing the exoskeleton to act towards an incorrect goal. To improve intent detection, the use of biosignal cues such as EEG as done in Huang et al. (2021) EMG, as in the work of Trigili et al. (2019), or gaze as in Dermy et al. (2019) can be exploited. However, integrating dedicated sensors to measure these biosignals increases system complexity and hinders broader exoskeleton adoption, despite advances in sensor calibration (e.g. Quesada et al. (2024) for EMGs).

To avoid complex biosignals, learning-by-demonstration (LbD) schemes use user-demonstrated movements to build a motion dictionary, enabling exoskeletons to recognize and assist these movements. In this work, we employ the learning-by-demonstration (LbD) scheme proposed in Martinez et al. (2019) and later adapted for upper-limb exoskeletons in Jamšek et al. (2021). The intent detection problem can thus be expressed as a classification problem with a partial observation of the incoming movement.

This classification task involves early recognition of a new observation's class (i.e., the goal or target) from a dictionary of learned motions, enabling accurate prediction of the human user's intended motion. This problem proves to be challenging when only kinematic information is measured through the exoskeleton's internal encoders, as shown in Dermy et al. (2019). Specifically, when demonstrated trajectories overlap significantly, identifying the movement class becomes challenging. Probabilistic methods such as Gaussian mixture models (GMMs as in Calinon and Lee (2019) or Luciani et al. (2024)), probabilistic

⋆ This work is supported in part by the French National Agency for Research (grant ANR-19-CE33-0009). This work has been submitted to the Joint MECHATRONICS 2025, ROBOTICS 2025 IFAC conferences for possible publication

flow tubes (PFTs as in Pérez-D'Arpino and Shah (2015)), probabilistic movement primitives (ProMPs as in Dermy et al. (2017)) or support vector machines (SVMs as in Cao et al. (2009)) have been employed to address this type of problem. Recent advances in time series classification using neural networks have also been applied to this problem, as done in the work of Elkholy et al. (2020). The work of Ismail Fawaz et al. (2019) identified the residual neural network (ResNet), previously applied for time series classification in Wang et al. (2017) and Geng and Luo (2018), as the most effective approach for this classification task.

In this work, we propose a hybrid convolutional and residual neural network architecture to address this classification challenge. We evaluate its performance using post-hoc data from assistive trials of reaching movements in a parasagittal plane, conducted with the ABLE 7D upper-limb exoskeleton under a standard velocity-field-based flow controller. By leveraging post-hoc data, we assess performance under conditions reflective of real-world use. Our results demonstrate that this classifier outperforms traditional probabilistic approaches (e.g., GMMs, SVMs, and ProMPs) and provide a preliminary investigation into whether the performance gains can mitigate adversarial effects caused by misclassification during assistive trials.

2. MATERIALS AND METHODS

2.1 Intent Classification

In this study, we evaluated the performance of four classifiers in our application. We included GMM, SVM, and ProMPs as baseline methods, given their established use for this problem, and provide a brief overview of these approaches. In contrast, we present a detailed description of our proposed architecture, which combines ResNet and long short-term memory (LSTM) networks.

1) Probabilistic approaches

Support Vector Machines (SVM) is a supervised learning algorithm for classification and regression. In this work we use SVM with a radial basis function (RBF) kernel. Once the SVM is trained, it predicts class labels for new data points by calculating the decision boundary. The SVM then estimates posterior probabilities for each class by applying Platt scaling, which fits a sigmoid function to the SVM's decision scores. The predicted class is then assigned based on the highest posterior probability.

Gaussian Mixture Models (GMM) and **Probabilistic Movement Primitives (ProMPs)** encode the training data into one distribution per target, resulting in four distinct distributions. The likelihood of a new partial observation belonging to each distribution is computed and compared to determine the classification.

2) Proposed Neural Network: ResNet + LSTM

ResNet: Residual Convolutional Encoder. We start by applying a one-dimensional (1D) convolutional layer, followed by a series of residual blocks inspired by the ResNet architecture of He et al. (2016) adapted for time series classification in the work of Ismail Fawaz et al. (2019). Each residual block contains consecutive 1D convolutions, Batch Normalization (BN), and ReLU activation, combined with a skip connection that bypasses these convolutional transformations. This skip connection helps prevent the vanishing gradient problem, thereby ensuring smoother gradient flow through deeper layers. Together, these residual blocks enable the network to extract robust local patterns and spatiotemporal features from the input sequences.

LSTM-based Temporal Modeling. After local features have been extracted, we transpose the output feature maps to the shape $(\text{batch\_size}, \text{sequence\_length}, \text{channels})$ and pass them to an *LSTM* layer. This LSTM captures temporal dependencies over longer portions of the input sequence. Unlike simple recurrent cells, the LSTM's memory gating mechanisms (input, output, and forget gates) more effectively manage long-range dependencies. The final hidden state of the LSTM serves as a learned summary of the entire sequence, which can then be used for classification.

Residual Block. Let $\mathbf{x}$ be the input to a residual block and $\mathcal{F}(\mathbf{x})$ represent the sequence of operations in the block (`Conv1D`, BN, ReLU). A skip connection is added so that the output $\mathbf{y}$ is:

$$\mathbf{y} = \mathrm{ReLU}\big(\mathcal{F}(\mathbf{x}) + \mathbf{x}\big).$$

This formulation ensures that gradients can flow directly from later layers to earlier layers during backpropagation.

LSTM Layer. The LSTM processes the output from the convolutional encoder. At each time step t, the hidden state h_t and cell state c_t are updated based on:

$$\begin{aligned}
\mathbf{f}_t &= \sigma\big(\mathbf{W}_f\,[\mathbf{h}_{t-1}; \mathbf{x}_t] + \mathbf{b}_f\big),\\
\mathbf{i}_t &= \sigma\big(\mathbf{W}_i\,[\mathbf{h}_{t-1}; \mathbf{x}_t] + \mathbf{b}_i\big),\\
\mathbf{o}_t &= \sigma\big(\mathbf{W}_o\,[\mathbf{h}_{t-1}; \mathbf{x}_t] + \mathbf{b}_o\big),\\
\tilde{\mathbf{c}}_t &= \tanh\big(\mathbf{W}_c\,[\mathbf{h}_{t-1}; \mathbf{x}_t] + \mathbf{b}_c\big),\\
\mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} \;+\; \mathbf{i}_t \odot \tilde{\mathbf{c}}_t,\\
\mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t).
\end{aligned}$$

where f_t, i_t, o_t represent forget, input, and output gates, respectively, and $\odot$ denotes element-wise multiplication. The final hidden state h_T after T time steps encodes the temporal information of the entire sequence. The symbols $\mathbf{W}_f, \mathbf{W}_i, \mathbf{W}_o, \mathbf{W}_c$ (and biases $\mathbf{b}_f, \mathbf{b}_i, \mathbf{b}_o, \mathbf{b}_c$) are the learnable weight matrices in the LSTM cell. Specifically, $\mathbf{W}_f$ is for the forget gate, $\mathbf{W}_i$ for the input (or update) gate, $\mathbf{W}_o$ for the output gate, and $\mathbf{W}_c$ (sometimes written as $\mathbf{W}_g$ or $\mathbf{W}_{\tilde{c}}$) for the candidate cell state.

Classification Layer. The final hidden state $\mathbf{h}_T$ is passed through a fully connected layer to produce logits for each class $y \in \{1, 2, \ldots, C\}$:

$$\hat{\mathbf{y}} \;=\; \mathbf{W}_{\text{fc}}\,\mathbf{h}_T \;+\; \mathbf{b}_{\text{fc}},$$

where C is the total number of classes ($C = 4$ in our application). We use *cross entropy loss* during training:

$$\mathcal{L} \;=\; -\sum_{c=1}^{C} y_c \,\log\Big(\mathrm{Softmax}(\hat{\mathbf{y}})_c\Big).$$

Data Preprocessing. Each input sequence is arranged as a 3D tensor $(\text{batch\_size}, \text{channels}, T)$, where T is the temporal length. Our data's multiple features (position and velocity) each treated as a separate channel.

Optimization. We train the network using stochastic gradient descent (e.g., Adam), iteratively updating the parameters to minimize the cross-entropy loss. Batch Normalization and skip connections in the residual blocks are used to stabilize training and to mitigate vanishing gradients.

2.2 Flow Controller

To obtain relevant and challenging data for our classification evaluation, we use the ABLE upper limb exoskeleton

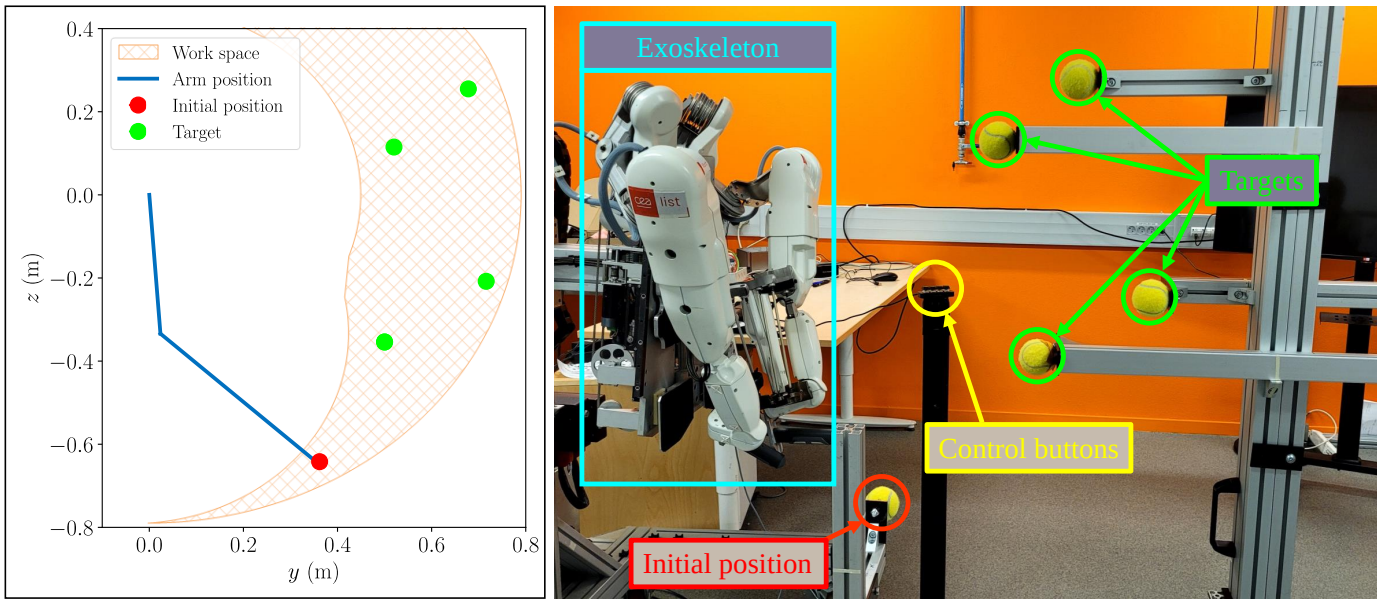

Fig. 1. Left panel: Representation of the exoskeleton's reachable space, targets, and initial position in the para-sagittal plane. Right panel: Experimental setup. Participants initiated reaching movements from the initial position toward a target, activating assistance using control buttons.

to assist reaching motions in a para-sagittal plane. We use a state-of-the-art Flow Controller (FC) proposed by Jamšek et al. (2021) for this purpose. This FC predicts upcoming trajectory using ProMPs and modulates an assistance around this prediction. In this section we briefly introduce this FC and the trajectory prediction method but we refer the reader to the original paper for a complete presentation of this controller.

ProMP Encoding and Prediction. ProMPs Paraschos et al. (2013) are employed to encode both demonstrated and artificial trajectories into a probabilistic framework. The trajectories are represented as weight vectors associated with Gaussian basis functions. Observations of partial movements are used to condition the learned distribution, enabling real-time prediction of the remaining trajectory. The prediction is updated continuously at 20 Hz during the assistive task, ensuring responsiveness to user movements.

Flow Controller. The FC consists of a velocity-field-based controller, which creates a flow field that is shaped based on a reference path (the prediction obtained with ProMPs inference). The controller proposes a compromise between a corrective and assistive effort based on the distance to the reference path. The assistive force is modulated according to the current hand velocity.

2.3 Experimental setup

We conducted an experiment with 11 participants. The experimental design and protocol are outlined below.

Task The experiment consisted of two blocks. In the first block, participants performed 15 reaching movements toward four targets (T_1, T_2, T_3, T_4) positioned in front of them, as shown in Fig. 1. These demonstrations trained the ProMPs for each target to predict upcoming movements. The remaining block involved the same reaching movements assisted by a flow controller Jamšek et al. (2021), conducted in a random order.

Participants Eleven healthy, right-handed adults (3 female; mean age: 29.91 ± 2.39 years; height: 1.75 ± 0.1 m; weight: 67.09 ± 12.27 kg) participated. All were free of neurological disorders or injuries. Written informed consent was obtained in accordance with the Helsinki declaration, and the study was approved by the local ethics committee.

Data Collection The reaching task was performed using the ABLE upper-limb exoskeleton Garrec et al. (2008), which features seven active joints replicating human shoulder, elbow, and wrist motion. Movements were measured via the exoskeleton's encoders at 1 kHz.

3. RESULTS

In this section we first compare the result of our proposed classifier to other classic classifiers of the literature. Then, we analyze whether the performance of this classifier could be enough for early motion detection using kinematic data only.

3.1 Accuracy comparison of our Neural Network with GMMs, SVMs and ProMPs.

To investigate the effects of earliness on classification accuracy for each of these classifiers, we evaluate their performance after 10%, 15% and 20% of the motion is demonstrated. Each method is trained with the position and velocity trajectories for each target. The mean and standard deviation of the accuracy for each method is computed with a 10-fold cross validation across 30 movements per target per participant (1320 movements in total, 330 per target). Movements lasted on average 0.7 ± 0.2, 1.0 ± 0.2, 0.9 ± 0.2 and 1.2 ± 0.2 s for Target 1 through 4 respectively. Results can be found in Table 1. We validated an accuracy of 1 for all classes before the end of the movement (at 50% for Targets 1 and 2 and at 85% for Targets 3 and 4). This result is consistent with the Bhattacharrya coefficient shown in Fig. 2 as they show that the overlapping of the distributions end around these percentages of the trajectories.

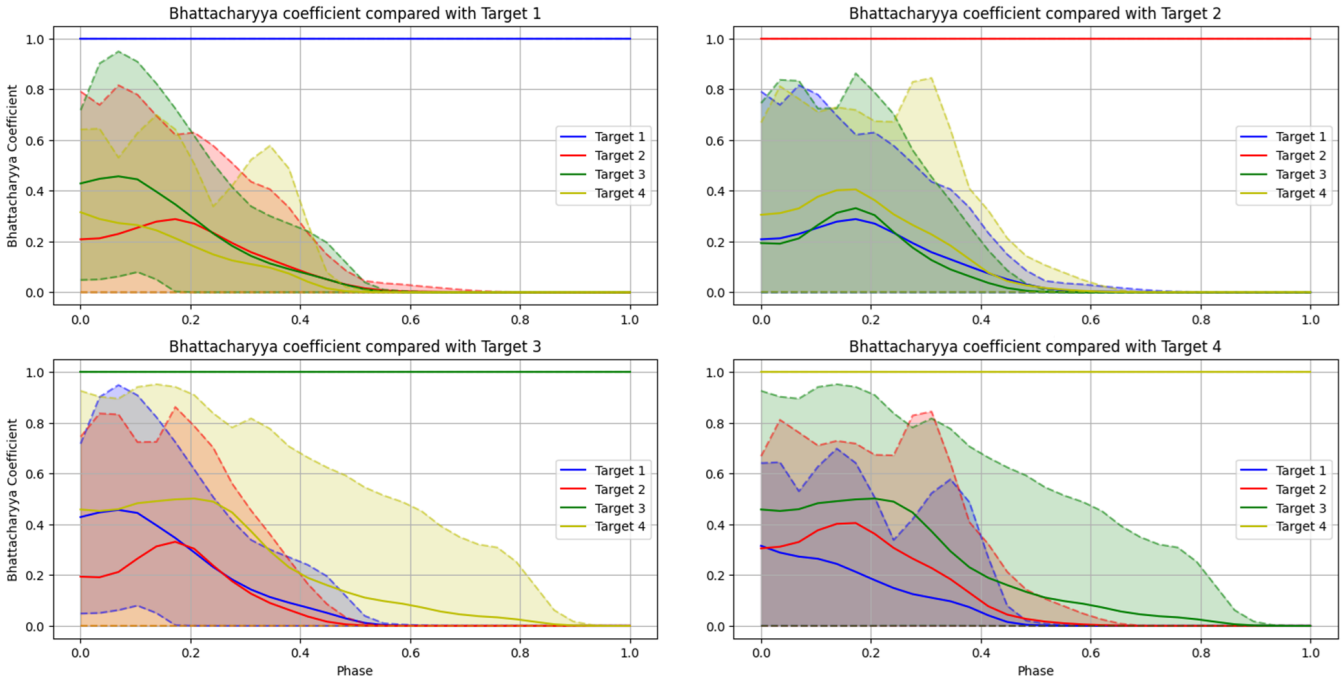

Fig. 2. Bhattacharrya coefficient for each target movement sets. Pairwise computation of the Battacharya coefficient is plotted for each target. Mean value in plain line, mean and max value of the coefficient represented in dashed lines.

Table 1. Classifiers accuracy comparisons at different early phase of the movement

Accuracy at	10%	15%	20%
GMMs	0.49 ±0.08	0.49 ±0.07	0.66 ±0.06
SVMs	0.35 ±0.02	0.35 ±0.01	0.42 ±0.04
ProMPs	0.37 ±0.02	0.35 ±0.03	0.48 ±0.11
ResNet + LSTM	**0.67 ±0.07**	**0.69 ±0.05**	**0.76 ±0.07**

3.2 Distribution Distance Analysis

As per design of our experimental setup, an ambiguous region of the reaching movement is present at the movement start as the starting position is the same for each motion. This distance between each movement distribution classified per target can be represented using the Bhattacharya coefficient. This coefficient can help quantify how much two distributions overlap each other, which can directly inform on the ambiguity of a classification problem. If the coefficient is null, the distributions are entirely distinct and classification is immediate, whereas if the coefficient is equal to 1, the classification is impossible at this stage. Using the formulation of the ProMPs, we encoded the kinematic results of our experimental results into Gaussian distributions. The resulting Bhattacharya coefficient is plotted along the phase for each target in Fig. 2. These distributions are 4-dimensional multivariate Gaussians and include the position and velocity result as represented by the ProMP framework.

3.3 Assistance Error Estimation

One could argue that, if the classification is erroneous, its impact could be limited by the fact that the assistance is likely similar at this phase as trajectories align. To evaluate this argument we calculate the average norm and direction error that classification to an erroneous target would entail throughout the assistive trials. The mean of these errors are represented for all assistive trials for target 3 in Fig. 3.

3.4 Classification Results

Finally, we evaluate the per-class accuracy of our classifier, along with its per-class false positive rate Fig. 4. We employed 10-fold cross validation across all assistive trials. The reaching motions toward Target 1 are more easily classified, as supported by the lower Battacharrya coefficient and represented in Fig. 2 when pairwise comparisons is achieved with target 1 distributions.

In summary, the comparison results demonstrated that our classifier consistently outperformed classical probabilistic methods for the considered data set. However, even after 30% of the trajectory, the classifier's accuracy and false positive rate revealed persistent misclassifications. We quantified the potential adversarial effects of these errors by computing the average direction and force errors that would result from misclassification (see Fig. 3). While the errors were relatively low during the early and ambiguous stages of movement, they could not be disregarded. Specifically, a mean direction error of 20° in these regions indicates that the classifier's performance would be insufficient to fully prevent adversarial effects.

4. DISCUSSION

In this paper, we leveraged recent advances in time series classification for early intent detection of human reaching trajectories. We analyzed the applicability of our approach using post-hoc data of repetitive trials with a state-of-the-art LbD assistive control scheme (the Flow Controller of Martinez et al. (2019) and Jamšek et al. (2021)) implemented in the ABLE upper-limb exoskeleton. Our results

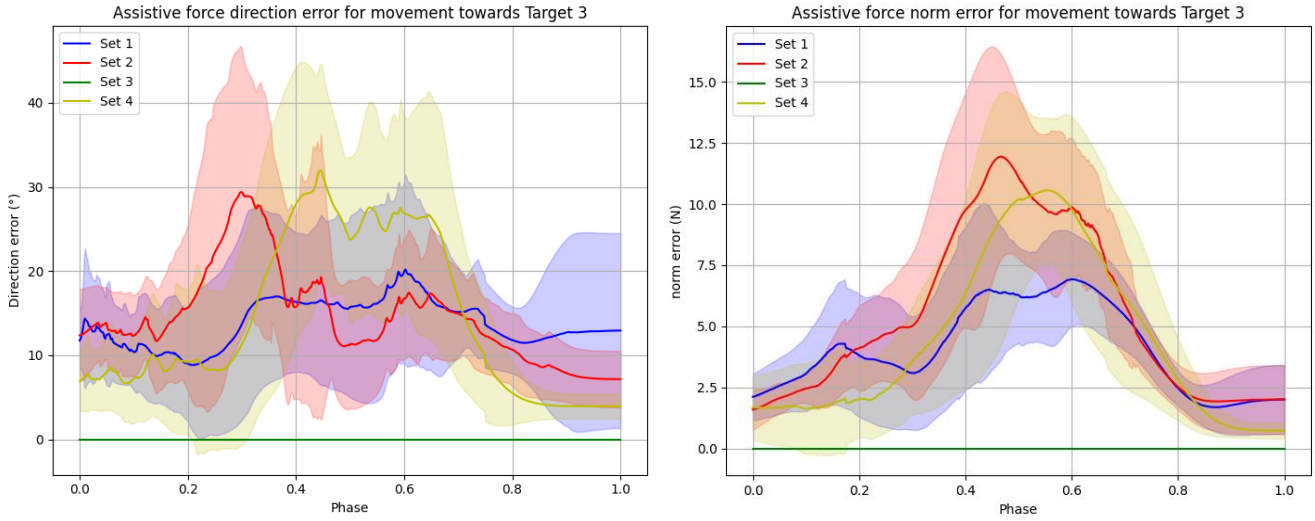

Fig. 3. Norm and direction error applied by the assistive force in the event of a misclassification of the movements across all assistive trials for target 3. Mean in plain line and standard deviation represented by the shaded area

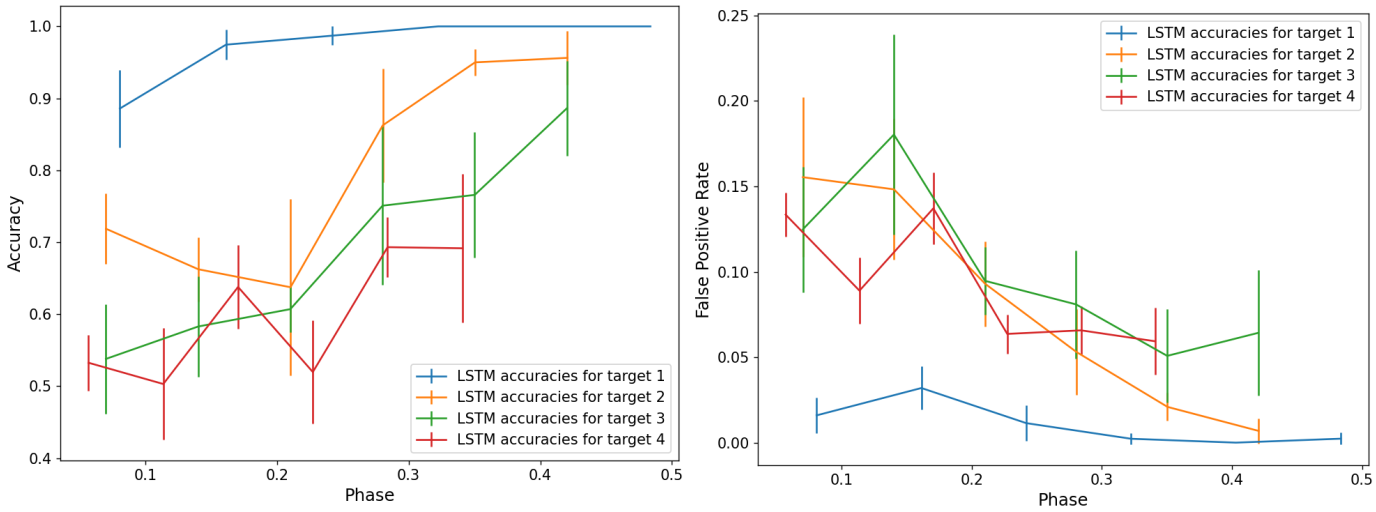

Fig. 4. LSTM classifier per class accuracy and false positive ratio at different phase. Classification occurs at a fixed timestep to reflect real life usage of the classifier (phase is not obtainable at the moment of the classification as the precise duration of the movement is still uncertain), leading to the differences in the phase axis between LSTM prediction accuracy observed here. Mean time for classification of the observation was measured at $0.571ms$ (13th Gen Intel(R) Core(TM) i7-13620H at 2.40 GHz), indicating real time usability at $1kHz$.

demonstrate that the proposed hybrid architecture, combining residual and convolutional neural networks, significantly outperforms traditional probabilistic methods in addressing the problem.

In the following, we estimated the adversarial resistive effects that our classifier could entail by calculating the assistance error in direction and norm in the case of erroneous classification. Our findings suggest that the classification accuracy reached with this new method cannot guarantee preventing adversarial effects of erroneous intent detections, supporting the current consensus on the necessity of exploiting biosignals to address this issue. Overall, our approach could improve human-exoskeleton interaction for assistance schemes that cannot rely on biosignals. However, further experimental tests using this classifier during robot-assisted trials are still necessary to thoroughly evaluate the criticality of the adversarial effects that remain despite our classification improvements.

In future works, the proposed classifier could be used to modulate the assistive scheme thanks to an evaluation of the confidence level of the classification as done for a library of GMMs in the work of Raiola et al. (2015) where each GMM contributes to the definition of the assistance proportionally to their likelihood. Other recent advances could also complement our approach to help foster a broader adoption of assistive exoskeletons. The work of Orhan et al. (2024) proposes to bypass the training of the LbD schemes by using artificially generated trajectories.

Similarly to this approach our classifier could be trained with these artificial demonstrations in future works. If the trajectories are impacted by the interaction, their data-augmentation strategy could be applied after a few assistive trials to reduce the necessary data for training our proposed classifier. Incorporating biosignal information, such as gaze as in Dermy et al. (2017), to complement and improve the accuracy of kinematic-based classifier is also a promising avenue for future work.

In conclusion, our approach holds significant potential to enhance the detection and classification of human movements in LbD schemes, particularly for repetitive tasks, offering a robust foundation for more efficient and accurate performance.

REFERENCES

Calinon, S. and Lee, D. (2019). Learning Control. In A. Goswami and P. Vadakkepat (eds.), *Humanoid Robotics: A Reference*, 1261–1312. Springer Netherlands, Dordrecht. doi:10.1007/978-94-007-6046-2_68.

Cao, D., Masoud, O.T., Boley, D., and Papanikolopoulos, N. (2009). Human motion recognition using support vector machines. *Computer Vision and Image Understanding*, 113(10), 1064–1075. doi:10.1016/j.cviu.2009.06.002.

Dermy, O., Charpillet, F., and Ivaldi, S. (2019). Multimodal Intention Prediction with Probabilistic Movement Primitives. In F. Ficuciello, F. Ruggiero, and A. Finzi (eds.), *Human Friendly Robotics*, 181–196. Springer International Publishing, Cham. doi:10.1007/978-3-319-89327-3_14.

Dermy, O., Paraschos, A., Ewerton, M., Peters, J., Charpillet, F., and Ivaldi, S. (2017). Prediction of Intention during Interaction with iCub with Probabilistic Movement Primitives. *Frontiers in Robotics and AI*, 4. doi:10.3389/frobt.2017.00045.

Elkholy, H.A., Azar, A.T., Magd, A., Marzouk, H., and Ammar, H.H. (2020). Classifying Upper Limb Activities Using Deep Neural Networks. In A.E. Hassanien, A.T. Azar, T. Gaber, D. Oliva, and F.M. Tolba (eds.), *Proceedings of the International Conference on Artificial Intelligence and Computer Vision (AICV2020)*, 268–282. Springer International Publishing, Cham. doi:10.1007/978-3-030-44289-7_26.

Garrec, P., Friconneau, J., Measson, Y., and Perrot, Y. (2008). ABLE, an innovative transparent exoskeleton for the upper-limb. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1483–1488. IEEE, Nice. doi:10.1109/IROS.2008.4651012.

Geng, Y. and Luo, X. (2018). Cost-Sensitive Convolution based Neural Networks for Imbalanced Time-Series Classification. doi:10.48550/arXiv.1801.04396.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. doi:10.1109/CVPR.2016.90.

Huang, C., Xiao, Y., and Xu, G. (2021). Predicting Human Intention-Behavior Through EEG Signal Analysis Using Multi-Scale CNN. *IEEE/ACM transactions on computational biology and bioinformatics*, 18(5), 1722–1729. doi:10.1109/TCBB.2020.3039834.

Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., and Muller, P.A. (2019). Deep learning for time series classification: A review. *Data Mining and Knowledge Discovery*, 33(4), 917–963. doi:10.1007/s10618-019-00619-1.

Jamšek, M., Kunavar, T., Bobek, U., Rueckert, E., and Babič, J. (2021). Predictive Exoskeleton Control for Arm-Motion Augmentation Based on Probabilistic Movement Primitives Combined With a Flow Controller. *IEEE Robotics and Automation Letters*, 6(3), 4417–4424. doi:10.1109/LRA.2021.3068892.

Li, Y., Sena, A., Wang, Z., Xing, X., Babic, J., van Asseldonk, E., and Burdet, E. (2022). A review on interaction control for contact robots through intent detection. *Progress in Biomedical Engineering*, 4. doi:10.1088/2516-1091/ac8193.

Luciani, B., Costante, S., Braghin, F., Pedrocchi, A., and Gandolla, M. (2024). Imitation Learning Using Gaussian Mixture Models and Dynamic Movement Primitives for Rehabilitation Exoskeletons: A Comparison. In *2024 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, 122–127. doi:10.1109/BioRob60516.2024.10719741.

Martinez, A., Lawson, B., Durrough, C., and Goldfarb, M. (2019). A Velocity-Field-Based Controller for Assisting Leg Movement During Walking With a Bilateral Hip and Knee Lower Limb Exoskeleton. *IEEE Transactions on Robotics*, 35(2), 307–316. doi:10.1109/TRO.2018.2883819.

Nussbaum, M.A., Lowe, B.D., de Looze, M., Harris-Adamson, C., and Smets, M. (2019). An Introduction to the Special Issue on Occupational Exoskeletons. *IISE Transactions on Occupational Ergonomics and Human Factors*, 7(3-4), 153–162. doi:10.1080/24725838.2019.1709695.

Orhan, A., Verdel, D., Bruneau, O., Geffard, F., and Berret, B. (2024). Combining Model-based and Data-based approaches for online predictions of human trajectories. In *IEEE RAS EMBS 10th International Conference on Biomedical Robotics and Biomechatronics (BioRob 2024)*. Heidelberg (Germany).

Paraschos, A., Daniel, C., Peters, J.R., and Neumann, G. (2013). Probabilistic Movement Primitives. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.

Pérez-D'Arpino, C. and Shah, J.A. (2015). Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 6175–6182. doi:10.1109/ICRA.2015.7140066.

Proietti, T., Crocher, V., Roby-Brami, A., and Jarrasse, N. (2016). Upper-Limb Robotic Exoskeletons for Neurorehabilitation: A Review on Control Strategies. *IEEE reviews in biomedical engineering*, 9, 4–14. doi:10.1109/RBME.2016.2552201.

Quesada, L., Verdel, D., Bruneau, O., Berret, B., Amorim, M.A., and Vignais, N. (2024). *EMG-to-torque Models for Exoskeleton Assistance: A Framework for the Evaluation of in Situ Calibration.* doi:10.1101/2024.01.11.575155.

Raiola, G., Lamy, X., and Stulp, F. (2015). Co-manipulation with multiple probabilistic virtual guides. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 7–13. doi:10.1109/IROS.2015.7353107.

Trigili, E., Grazi, L., Crea, S., Accogli, A., Carpaneto, J., Micera, S., Vitiello, N., and Panarese, A. (2019). Detection of movement onset using EMG signals for upper-limb exoskeletons in reaching tasks. *Journal of NeuroEngineering and Rehabilitation*, 16(1), 45. doi:10.1186/s12984-019-0512-1.

Wang, Z., Yan, W., and Oates, T. (2017). Time series classification from scratch with deep neural networks: A strong baseline. In *2017 International Joint Conference on Neural Networks (IJCNN)*, 1578–1585. doi:10.1109/IJCNN.2017.7966039.