

Ensemble of LSTMs and feature selection for human action prediction

Tomislav Petković, Luka Petrović, Ivan Marković and Ivan Petrović

University of Zagreb Faculty of Electrical Engineering and Computing
Laboratory for Autonomous Systems and Mobile Robotics (LAMOR)
{petkovic, luka.petrovic, ivan.markovic, ivan.petrovic}@fer.hr

Abstract. As robots are becoming more and more ubiquitous in human environments, it will be necessary for robotic systems to better understand and predict human actions. However, this is not an easy task, at times not even for us humans, but based on a relatively structured set of possible actions, appropriate cues, and the right model, this problem can be computationally tackled. In this paper, we propose to use an ensemble of long-short term memory (LSTM) networks for human action prediction. To train and evaluate models, we used the MoGaze¹ dataset – currently the most comprehensive dataset capturing poses of human joints and the human gaze. We have thoroughly analyzed the MoGaze dataset and selected a reduced set of cues for this task. Our model can predict (i) which of the labeled objects the human is going to grasp, and (ii) which of the macro locations the human is going to visit (such as table or shelf). We have exhaustively evaluated the proposed method and compared it to individual cue baselines. The results suggest that our LSTM model slightly outperforms the gaze baseline in single object picking accuracy, but achieves better accuracy in macro object prediction. Furthermore, we have also analyzed the prediction accuracy when the gaze is not used, and in this case, the LSTM model considerably outperformed the best single cue baseline.

Keywords: human action prediction, recurrent neural networks, long-short term memory networks, feature selection

1 Introduction

With the robots becoming more sophisticated, their presence in human environments increases and includes close collaboration with humans. Such applications involve proximate interaction of robots and humans, yielding novel challenges concerning system efficiency and human safety. Robots that physically share space with humans will need to understand and learn by interacting and observing humans [12]. However, this is not an easy task, at times, not even for us humans. Hence, human action and full-body motion prediction are becoming increasingly important research topics for the robotics research community [19].

¹ <https://humans-to-robots-motion.github.io/mogaze/>

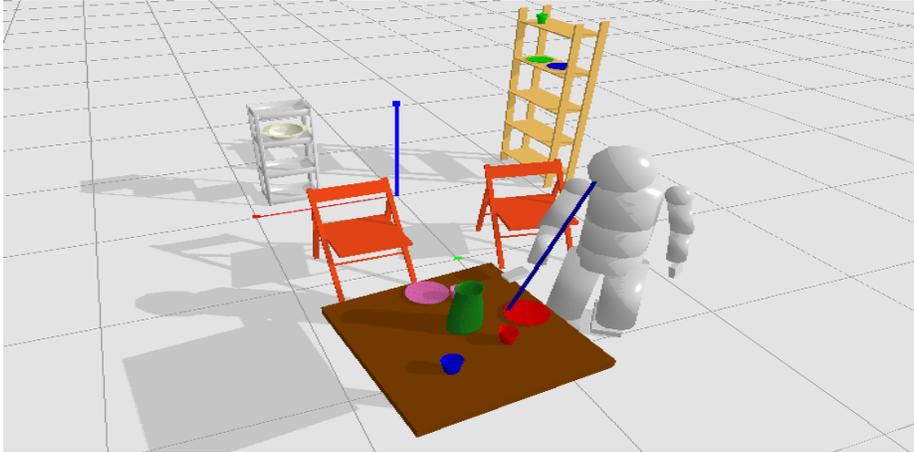


Fig. 1: Visualization of the MoGaze dataset based on PyBullet [5]. The subject is about to pick the red plate from the table which is strongly indicated by the subject’s gaze illustrated with the blue line.

In recent years, human action prediction applications ranged from robotized warehouses [2, 18] to sedentary object-picking domain [9, 20, 21] and full-body motions [13, 16, 26]. Common approaches to solving the human action prediction problem are based on hidden Markov models [10] or conditional random fields [25], which try to learn observed patterns. In works [3, 9, 12, 21] authors have indicated that the eye-gaze is a powerful predictor of human action. They tackle the problem using machine learning models such as support vector machine [9] or recurrent neural networks [12]. In [21] authors calculate the similarity between the hypothetical gaze points on the objects and the actual gaze points and use the nearest neighbor algorithm to classify the intended object.

In the last few years, multiple datasets concerning motion and action prediction have become publicly available but, to the best of our knowledge, none of them couple these two problems. Examples of purely motion prediction datasets are: KITTI [7], ETH [17] and UCY [15]. We encourage the reader to examine Table 2 in [19] for a detailed listing. These datasets offer enough diverse data to train and test human motion prediction models focused on answering the question “*Where is the human going to be during the next N steps?*”, but they are not adequately labeled with the context which would help to answer “*What is (the goal of) the observed human motion?*”. On the other hand, datasets tailored for models focused on the second question, like the CMU’s motion capture database [1] and G3D [4] that excel in action diversity, but they are focused on distinguishing between different actions (jumping, punching, sitting), do not incorporate complicated motion patterns, and usually are not long enough for a long term human motion prediction problem. The MoGaze [12] dataset (Fig. 1) positions itself as an excellent blend of the aforementioned datasets because all the recorded motion has a labeled purpose (an object picking). Its subset has already been used by the authors for human motion prediction problems based

on RNNs and trajectory optimization [13,14]. Therein, they used Euclidean distance of the right hand to each object as an action prediction signal, improving their original motion prediction result.

In this paper, we rely on findings in [12,13,21] and propose a human action prediction model based on the ensemble of recurrent neural networks, namely long-short term memory (LSTM) networks. To train and evaluate our model, we used and thoroughly analyzed the MoGaze dataset and selected a reduced set of features for human action prediction. Our model can predict (i) which of the labeled objects the human is going to grasp, and (ii) which of the macro locations the human is going to visit (such as table or shelf). Experimental results are focused on comparative performance with respect to individual cue accuracy, especially the gaze which is a powerful action predictor, and demonstrate the effectiveness of the proposed model.

2 Proposed Human Action Prediction

We design the proposed action prediction model as a general model for full-body motion that successfully captures relations between input cues and picked objects, but without learning specific relations between objects in a dataset. Another important aspect that needs to be captured by an action prediction model are long-term dependencies, since goal inferring cues appear usually much earlier than the actual picking action. For example, imagine a human who is given the task of picking a specific object. First, they would look towards the object and start walking to it. While walking, the gaze of the human would not be fixed only on the object, but would also wander around the scene, especially if there are obstacles to be negotiated. Given that, a well-designed human action prediction model should take into account that the gaze becomes fixed at the object early in the sequence and can wonder thereafter. In other words, to successfully infer the goal, the model should be able to remember the most important past cue values, e.g. early gaze fixation at the object, as well as capture local tendencies, such as human approaching the object. In order to achieve that, in the ensuing section we propose an ensemble of long short-term memory (LSTM) networks that also take distances of human joints towards each of the goals. However, each joint in the input feature set adds to the complexity and the network parameter number, which not only increases the run-time, but can also impede the training process. Consequently, we further introduce a feature selection method based on signal correlations and individual effectiveness to act as an action prediction cue.

2.1 Ensembles of LSTMs for Action Prediction

A recurrent neural network (RNN) is a feed forward network derivate that allows previous network outputs to be used as its inputs with the introduction of a hidden state. By memorizing past states, RNN is able to process sequences of inputs [6]. An LSTM [8] is an RNN extension that takes advantage of its

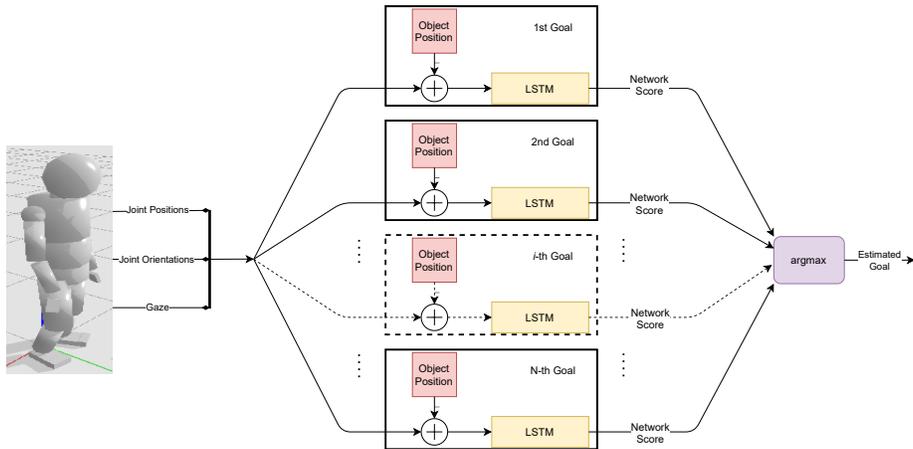


Fig. 2: Proposed human action prediction pipeline. We evaluate observed sequences with respect to each goal using an ensemble of identical LSTM networks as classifiers. The network with the largest score is considered to be associated with the current prediction of the subject’s goal.

gates, thus enabling it to capture long-term dependencies by partially solving the vanishing gradient problem. They have been used in a plethora of sequence regression and classification problems, such as trajectory prediction [27], speech emotion recognition [28] and action prediction [22]. As previously stated, one of the fundamental requirements for our model is to avoid learning relations between objects specific to a dataset and to be able to remember the most important past cue values. For example, in the MoGaze dataset the objects are placed on three macro locations, two shelves and a table, which do not move during the experiments. If we give the model, e.g. distances to all the goals as an input, the model could implicitly learn relations between those macro locations that would not hold for other datasets. Furthermore, our model should also be capable of handling a varying number of goals.

Given that, we have decided to train only a single LSTM model per picking object. The selected model has an input layer, an LSTM layer with 20 hidden units, a fully connected layer, a softmax layer, and a classification layer. Every 10 frames it takes the whole history of selected input features labeled as $\{0, 1\}$ and outputs a score between $[0, 1]$ – network’s confidence in the estimated action. Joint positions are used to calculate Euclidean distances towards each of the 10 goals in the dataset, and gaze and orientation vectors are used to calculate the distance between them and the vector pointing towards the position of each goal. All of the features are normalized based on the average value on the training set. During runtime, we evaluate all of the selected features for each of the N goals and send them as inputs to the ensemble of N LSTM networks, where N is the number of objects – for the MoGaze dataset $N = 10$. We compare the output of each of the networks and select the goal whose network has the maximum score. An illustration of this process is shown in Fig. 2.

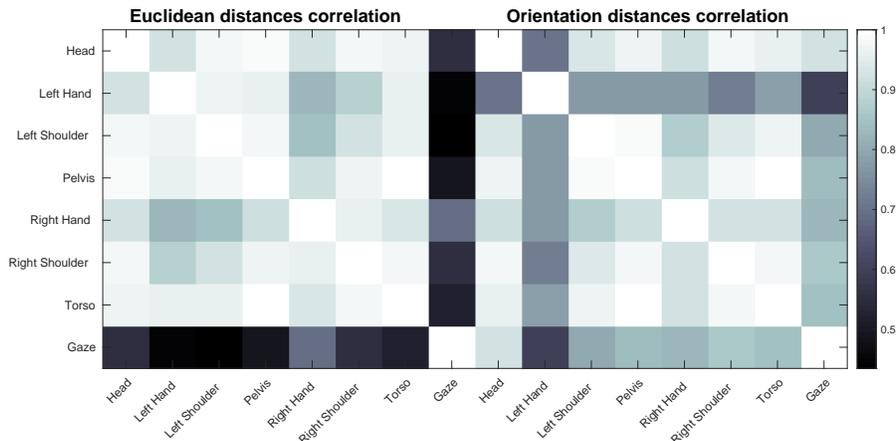


Fig. 3: Correlations of various dataset features. Both orientations and Euclidean distances of hands correlate poorly with the other joints distances.

The main contributions of the current work lie in feature selection for human action prediction and demonstration that the problem at hand can be successfully approached using LSTMs based on selected features. Because of that, hyperparameters and the layout of the proposed LSTM network have not been further scrutinized and we plan to analyze possible improvements and alternative architectures in future work.

2.2 Feature Selection

The LSTM networks excel at capturing long-term dependencies of input signals with the cost of adding a substantial amount of parameters to the model. Because of that, we propose a method for feature selection to reduce the input size and enable a quicker performance of the model. We will demonstrate proposed feature selection on the MoGaze dataset which has recorded positions and orientations of multiple joints. Our analysis will be focused on the following joints: head, torso, pelvis, both hands and both shoulders. We selected these joints based on the intuition that other joints, such as knees and elbows, do not contain any additional information for inferring which object the human is about to pick. Following the preliminary results published in [12], wherein the authors asserted that gaze is the most powerful predictor of human action, we compared correlations of our selected signals and the gaze, and the results can be seen in Fig. 3.

In the MoGaze dataset participants were advised to use only one hand for performing tasks [12], resulting in heavily favored use of the right hand. As can be seen in Fig. 3, the position and orientation distances of the left hand correlate poorly with other joints; however, pelvis, shoulders, head and torso orientation distances correlate reasonably high with each other. Such high correlations suggest that only one of the features could be selected as an input to the network.

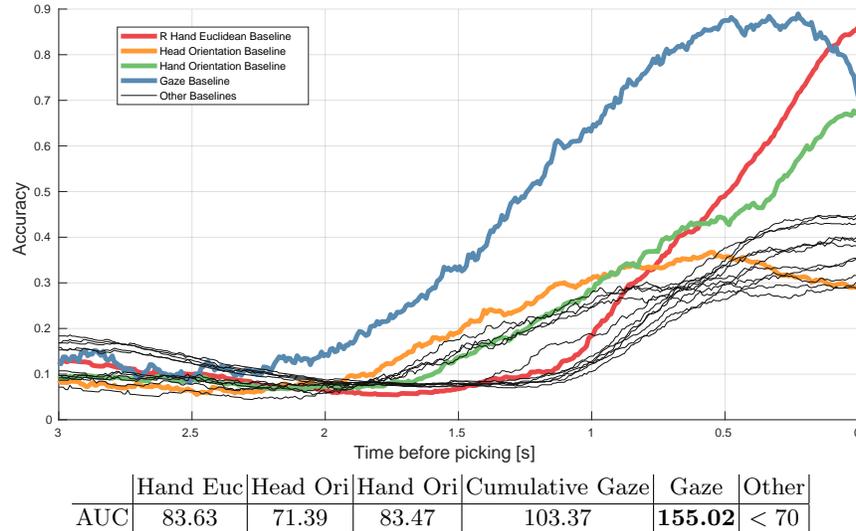


Fig. 4: Comparison of AUCs for all baselines. Larger AUC value indicates that the baseline is a better predictor of action. The gaze baseline clearly outperforms other baselines until the last few moments when the proximity of the reaching hand becomes a better indicator of the picked object.

Furthermore, we also calculated the cumulative gaze feature which at frame t takes the value of the sum of the number of frames that the gaze towards the goal was under the threshold of 0.1 m. Our motivation in constructing this feature lies in the intuition that during the action, human fixes gaze for a prolonged period of time on the goal it wants to visit, and it would be useful to record the entire history of that signal. In order to determine which of the aforementioned cues could act as a good action predictor, we compared the individual accuracy of all of the selected features. These single cue predictor results are called *baselines*, since any other more complex model should obtain better results. We followed the accuracy evaluation for a time span of 3 seconds, with the grasp of the object occurring at the third second, as proposed in [12].

Fig. 4 presents the calculated area under the accuracy curve (AUC) for all of the baselines. We decided to use the AUC score evaluation from the computer vision applications such as face alignment [24], where the authors calculate the area under the accumulated error distribution curve, which resembles our accuracy curves. We find the AUC score as a reasonable scalar value containing information about the accuracy during the whole prediction period. In the MoGaze dataset, information is sampled at 120 Hz, which means that maximum AUC for the three second sequence is 360, while the expected AUC of a random goal guesser for 10 goals would be 36. This gives us border values for accuracy evaluation. Based on the AUC scores shown in Fig. 4, we selected the following reduced feature set: gaze, cumulative gaze, head orientation, hand position and hand orientation.

3 Experimental Results

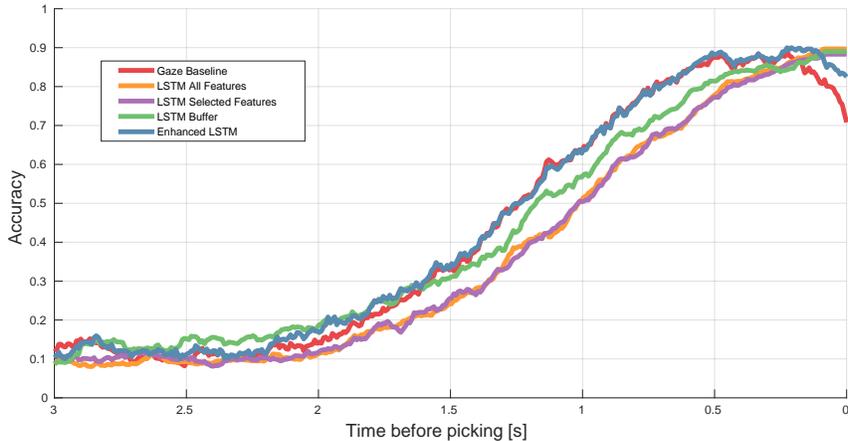
In this section we describe in detail the MoGaze dataset that we used for training and evaluation. Furthermore, we discuss the results and present several versions of the LSTM models that we designed for human action prediction.

3.1 Dataset and Training

The MoGaze dataset includes a total of 180 minutes of motion capture data with 1627 pick and place actions [12]. For one participant (participant no. 3) the eye-tracker device did not work so we excluded this session leaving a total of 1435 picking segments. Each group of segments is preceded by the instructions to the participant, e.g. *"Set the table for 2 persons"*, but we do not use this information. Each segment consists of multiple frames before the actual picking happens and a labeled of the object that is eventually picked. Frames in which an object is being carried are discarded. It might be beneficial to additionally discard the parts of segments before the instruction has been given as well as those involving moving of chairs; however, we decided to leave the dataset intact for easier future comparison. In order to prepare the dataset for model training, we split the data into training and testing datasets. The training dataset consisted of sessions with subjects 1-2 and 4-5 (a total of 853 segments), while the testing dataset contained sessions with subjects 6-7 (a total of 582 segments). We have performed feature selection as in Section 2.2 and trained the network using MATLAB with Adam optimization [11] training with 5 epochs and batch size of 5.

3.2 Experimental Results

For accuracy evaluation, we used AUC for the last three seconds of each experiment. Since we want our network to be applicable in real time applications, all the inputs are recorded each 10 steps (12Hz) and network outputs are considered to be valid until the next iteration. First, we trained and tested the proposed network using the whole signal history of all recorded joints' orientations and distances as inputs, including the gaze (we dubbed this version simply as *LSTM*). With the AUC of 134.24, the LSTM model performed worse than the gaze baseline (155.02) during majority of experiments and only succeeded to beat it in the last few moments. On top of that, average prediction time was too slow to work with the aimed frequency of 12 Hz. Therefore, we further tried to improve the results and execution time by reducing the input feature set as discussed in Section 2.2 (dubbed *LSTM Select*). While it slightly improved execution time, the AUC remained unchanged for the LSTM Select model. In order to reduce the run time, we decided to reduce the complexity by using a buffer to send only the last 20 frames to the network (dubbed *LSTM Buff*). The LSTM Buff model had a satisfying run-time and achieved AUC of 151.98; however, its score was still lower than that of the gaze baseline. As tuning the network hyperparameters mentioned at the beginning of this section did not result in any considerable improvements, we decided to further analyze the inputs and the



	Gaze	LSTM	LSTM Select	LSTM Buff	Enhanced LSTM
AUC	155.02	134.24	134.43	151.98	159.01

Fig. 5: Proposed models’ performances. Eye-gaze has proven to be strongest indicator of human action on this dataset.

outputs of the proposed model to see where accuracy could be increased and the results are shown in Fig. 6. For example, in Fig. 6a), one can see that our model has clearly distinguished which object is the goal after step 250, while Euclidean distance is not the lowest at any point. Also, the gaze signal corresponding with the picked goal has a spike around frame 325. This did not have major effect on the output of the proposed model. In Fig. 6b), the proposed model predicted high probability of the picked object very early but only succeeded to isolate it from nearby objects after step 700. Gaze and Euclidean distance alone could not make such distinction. In contrast, in Fig. 6c) and Fig. 6d), proposed model behaved poorly. One possible explanation could be that the subject looked at other objects without moving towards any of the goals, thus not giving the model enough information to conclude the exact goal until the end of the segment.

Given the previous analysis, we can see that the gaze baseline acts as the strongest predictor of the object that the human is going to pick, and furthermore, it can distinguish the actual goal among the nearby objects with pinpoint accuracy – something that our network model struggles with. The fact that the gaze is such a strong indicator of action is not surprising; indeed, visual fixation is necessary for object identification and comes after the brain automatically and in parallel gathers basic features, such as colors, shape, and motion [23]. Given that, we have enhanced our model with a few simple conditions. If the highest LSTM score is larger than the *score threshold*, we consider that object to be the goal. Otherwise, we check if the subject is looking directly at any of the objects, by comparing the minimum gaze distance to the objects with the *gaze threshold*. If the subject is looking at the object, we output that goal, otherwise we select the highest LSTM score. This enhancement led to 159.01 AUC,

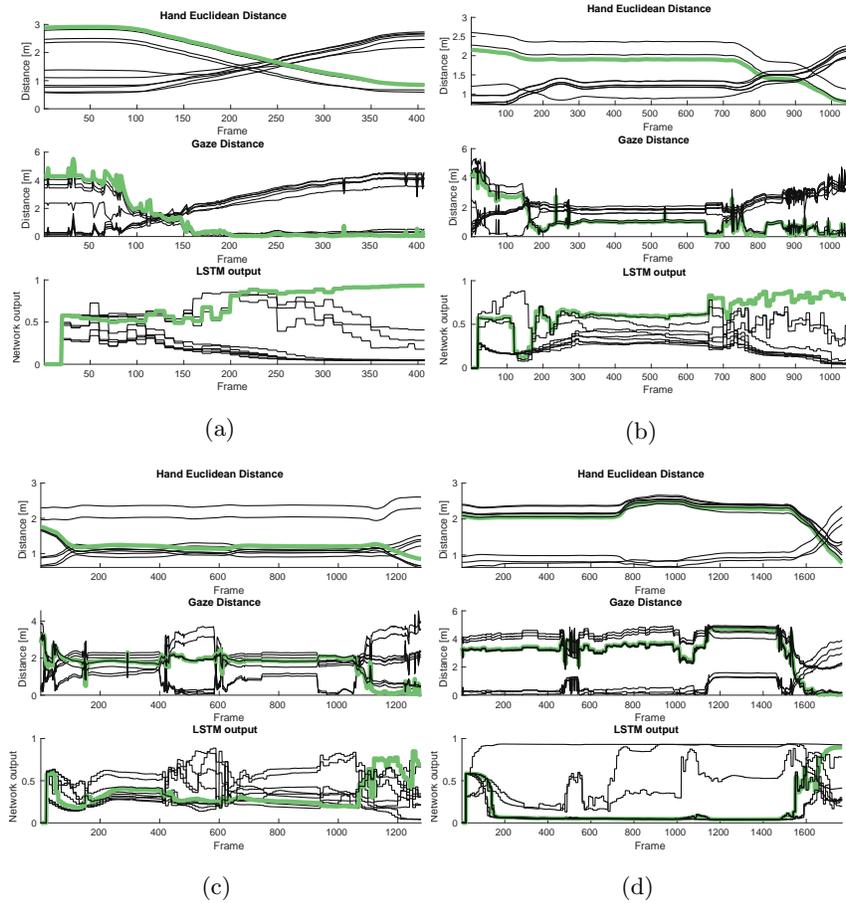


Fig. 6: Examples of selected model inputs and outputs. The feature corresponding with the object being picked is colored green while the black lines denote features associated with other objects.

slightly improving over the gaze baseline by 2.6%. We used the score threshold of 0.49 and gaze threshold of 0.2, which have been selected using grid search on the training set with the AUC as the target function. At the time of writing this paper, the only method for human action prediction on the MoGaze dataset was published in [12]. The method is based on RNNs but the authors do not provide implementation details rendering a direct comparison difficult. However, by qualitatively comparing the accuracy figures, we can assert that the proposed approach seems to yield more accurate action prediction results.

Although our LSTM models underperformed or showed only a slightly better result than the gaze baseline, we have observed that the LSTMs are better at identifying the macro objects from which the human is going to pick a specific object. Namely, all objects in the Mogaze dataset are placed on three specific macro locations: the table or two shelves. By analyzing the estimated macro

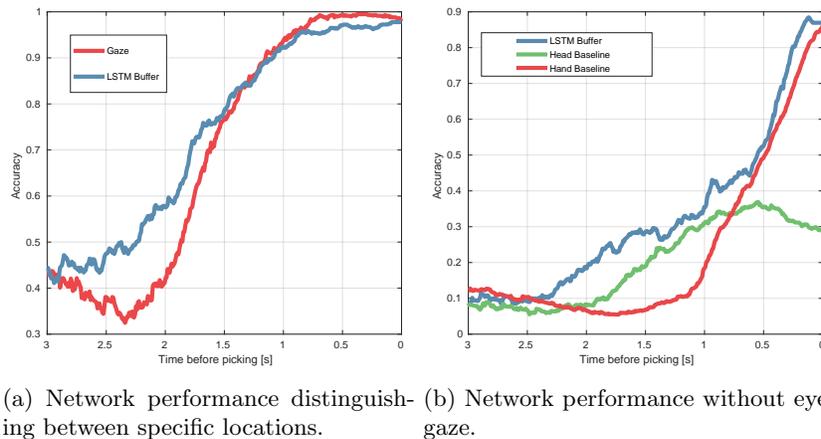


Fig. 7: The proposed action prediction model has better performance than any baseline if the gaze feature is not available or when one only needs to determine the macro location of the object being picked.

location accuracy, the LSTM Buff model scored 266.31 AUC, while the gaze baseline scored 251.31 AUC, as shown in Fig. 7a). A possible explanation could be that test subjects had to naturally look around while moving towards the goal, preventing them from keeping the gaze fixed on the correct macro location. Our model, on the contrary, managed to successfully leverage the motion cues towards the goal and obtained a higher score in identifying the macro locations.

Since gaze is the most powerful predictor of human action, it begs the question of what kind of performance could be achieved if such a cue was not available? Such a question could be further motivated since presently gaze tracking is done using a cumbersome apparatus that can impede subject’s efficiency and comfort in real-life applications. Thus it would be interesting to test what performances could be achieved if the gaze information is not available. We tested the behavior of the proposed LSTM Buff giving only hand and head positions and orientations as input. Because of the reduced number of features, we have increased the number of hidden units to 40 for this model. Our model successfully combined those signals and scored AUC of 118.60 which is considerably higher than the baselines (83.63 and 71.39), as can be seen in Fig. 7b).

4 Conclusion and Future Work

In this paper, we have proposed a novel human action prediction model based on the ensemble of the LSTM. We have thoroughly analyzed the dataset and selected appropriate cues for human action prediction. The LSTM networks receive distances and orientations of selected human joints towards each of the 10 goals as recorded. The proposed framework is trained to classify whether the observed sequence belongs to an object that the human is going to pick or not.

During the runtime, we compare the network outputs for each object and select the one with the highest score as our prediction of the human action. We have exhaustively evaluated the proposed method and the experiments demonstrated its effectiveness in comparison to single cue baselines.

For future work, we plan to design a gaze estimation model that would increase the accuracy of human action prediction on scenarios where gaze is unavailable and exploit this model to create a decision-making system for human-robot shared environments.

Acknowledgment

This research has been supported by the European Regional Development Fund under the grant KK.01.1.1.01.0009 (DATACROSS).

References

1. Carnegie mellon motion capture database, <http://mocap.cs.cmu.edu>.
2. Human intention recognition for human aware planning in integrated warehouse systems. pages 1–6, 2020.
3. T. Bader, M. Vogelgesang, and E. Klaus. Multimodal integration of natural gaze behavior for intention recognition during object manipulation. *Proceedings of the 2009 international conference on Multimodal interfaces - ICMI-MLMI '09*, page 199, 2009.
4. V. Bloom, D. Makris, and V. Argyriou. G3d: A gaming action dataset and real time action recognition evaluation framework. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2012.
5. E. Coumans and Y. Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. 2016.
6. S. Dupond. A thorough review on the current advance of neural network structures. *Annual Reviews in Control*, 14:200–230, 2019.
7. A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012.
8. S. Hochreiter and J. Schmidhuber. Lstm can solve hard long time lag problems. In *Advances in neural information processing systems*, pages 473–479, 1997.
9. C.-M. Huang, S. Andrist, A. Sauppé, and B. Mutlu. Using gaze patterns to predict task intent in collaboration. *Frontiers in Psychology*, 6(July):1–12, 2015.
10. R. Kelley, A. Tavakkoli, C. King, M. Nicolescu, M. Nicolescu, and G. Bebis. Understanding human intentions via hidden markov models in autonomous mobile robots. *Proceedings of the 3rd international conference on Human robot interaction - HRI '08*, page 367, 2008.
11. D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
12. P. Kratzer, S. Bihlmaier, N. Balachandra Midlagajni, R. Prakash, M. Toussaint, and J. Mainprice. Mogaze: A dataset of full-body motions that includes workspace geometry and eye-gaze. *IEEE Robotics and Automation Letters (RAL)*, 2020.

13. P. Kratzer, N. B. Midlagajni, M. Toussaint, and J. Mainprice. Anticipating human intention for full-body motion prediction in object grasping and placing tasks. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1157–1163. IEEE, 2020.
14. P. Kratzer, M. Toussaint, and J. Mainprice. Prediction of human full-body movements with motion optimization and recurrent neural networks. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.
15. A. Lerner, Y. Chrysanthou, and D. Lischinski. Crowds by example. In *Computer graphics forum*, volume 26, pages 655–664. Wiley Online Library, 2007.
16. Y. Li, C. Lan, J. Xing, W. Zeng, C. Yuan, and J. Liu. Online human action detection using joint classification-regression recurrent neural networks. In *European Conference on Computer Vision*, pages 203–220. Springer, 2016.
17. S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *IEEE International Conference on Computer Vision (ICCV)*, pages 261–268, 2009.
18. T. Petković, D. Puljiz, I. Marković, and B. Hein. Human intention estimation based on hidden markov model motion validation for safe flexible robotized warehouses. *Robotics and Computer-Integrated Manufacturing*, 57:182–196, 2019.
19. A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.
20. P. Schydlo, M. Rakovic, L. Jamone, and J. Santos-Victor. Anticipation in human-robot cooperation: A recurrent neural network approach for multiple action sequences prediction. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–6. IEEE, 2018.
21. L. Shi, C. Copot, and S. Vanlanduit. What are you looking at? detecting human intention in gaze based human-robot interaction. *arXiv preprint arXiv:1909.07953*, 2019.
22. C. Si, W. Chen, W. Wang, L. Wang, and T. Tan. An attention enhanced graph convolutional lstm network for skeleton-based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1227–1236, 2019.
23. A. M. Treisman and G. Gelade. A Feature-integration Theory of Attention. *Cognitive psychology*, 12(1):97–136, 1980.
24. G. Trigeorgis, M. A. Snape, Patrick and Nicolaou, E. Antonakos, and S. Zafeiriou. Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4177–4187, 2016.
25. S. B. Wang, A. Quattoni, L.-P. Morency, D. Demirdjian, and T. Darrell. Hidden conditional random fields for gesture recognition. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 2, pages 1521–1527. IEEE, 2006.
26. P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng. View adaptive recurrent neural networks for high performance human action recognition from skeleton data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2117–2126, 2017.
27. P. Zhang, W. Ouyang, P. Zhang, J. Xue, and N. Zheng. Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12085–12094, 2019.
28. J. Zhao, X. Mao, and L. Chen. Speech emotion recognition using deep 1d & 2d cnn lstm networks. *Biomedical Signal Processing and Control*, 47:312–323, 2019.