

Improving Human Motion Prediction Through Continual Learning

Mohammad Samin Yasar
University of Virginia
Charlottesville, VA 22904, USA
msy9an@virginia.edu

Tariq Iqbal
University of Virginia
Charlottesville, VA 22904, USA
tiqbal@virginia.edu

ABSTRACT

Human motion prediction is an essential component for enabling closer human-robot collaboration. The task of accurately predicting human motion is non-trivial. It is compounded by the variability of human motion, both at a skeletal level due to the varying size of humans and at a motion level due to individual movement's idiosyncrasies. These variables make it challenging for learning algorithms to obtain a general representation that is robust to the diverse spatio-temporal patterns of human motion. In this work, we propose a modular sequence learning approach that allows end-to-end training while also having the flexibility of being fine-tuned. Our approach relies on the diversity of training samples to first learn a robust representation, which can then be fine-tuned in a continual learning setup to predict the motion of new subjects. We evaluated the proposed approach by comparing its performance against state-of-the-art baselines. The results suggest that our approach outperforms other methods over all the evaluated temporal horizons, using a small amount of data for fine-tuning. The improved performance of our approach opens up the possibility of using continual learning for personalized and reliable motion prediction.

1 INTRODUCTION

Human motion prediction involves forecasting future human poses given past motion. For enabling efficient Human-Robot Collaboration, a crucial aspect of robot perception is real-time anticipatory modeling of human motion [15, 18, 22, 36]. Fluid tasks such as collaborative assembly, handovers, and navigating through moving crowds require combining aspects of perception, representation, and motion analysis to accurately and timely predict probable human motion [4, 17, 20, 21, 23, 24, 29]. This would enable the robot to anticipate the human pose and intent and plan accordingly around the human partner without disturbing the natural flow of the human's motion. However, accurate and timely prediction of human motion remains a non-trivial problem due to the complex and interpersonal nature of human behavior [19, 40].

To address the aperiodic and stochastic nature of human motion, prior work has framed the problem of predicting future poses like that of sequence learning, modeling the spatio-temporal aspect of human motion using Recurrent Neural Networks [2, 12, 33, 40]. These approaches aim to learn a unified representation from training samples that are expected to generalize for test data. However, generalization comes at the cost of learning individual subtleties of motion, which is crucial for human-robot collaboration. When




Figure 1: Qualitative performance of different motion prediction methods for walking on UTD-MHAD. The framework trained with Curriculum Learning (C.L.) have predictions that are closer to the ground-truth poses.

training these networks, the core assumption is that the given data points are realizations of independent and identically distributed (i.i.d) random variables. However, this assumption is often violated, e.g., when training and test data come from different distributions (dataset bias or domain shift) or the data points are highly interdependent (e.g., when the data exhibits temporal or spatial correlations) [10]. Both these cases are observed in human motion prediction, making it challenging to deploy models trained on benchmark models to the real world.

While generalization at the cost of learning individual preferences is sub-optimal, there is also a need to learn a robust representation over a diverse range of training samples. As such, training and generalizing over a benchmark dataset cannot be discarded and is, in fact, necessary as the first step to accurate motion prediction. Prior work on language modeling has demonstrated the benefit of learning a rich representation on a large training data followed by fine-tuning on a target task [5, 11, 39]. For human motion prediction, this can be posed as a continual learning problem whereby a motion prediction model acquires prior knowledge by observing a large range of human activities. This is followed by fine-tuning its parameters to accurately capture the subtleties of motion prediction for a particular individual. Such a learning setup, however, brings additional challenges to an already non-trivial problem, with prior work on continual learning demonstrating the risk of *catastrophic forgetting* [31, 34].

To address the challenges mentioned above, we propose a continual learning scheme that can improve human motion prediction accuracy while reducing the risk of catastrophic forgetting. Our framework is modular and is developed to acquire new knowledge and refine existing knowledge based on the new input. In line with

prior work on computational neuroscience, which states that the brain must carry out two complementary tasks: generalize across experiences, and retain specific episodic-like events [30, 34]; we utilize a two-phase learning scheme. Our framework aims to learn a robust representation of past observations by training on a benchmark dataset in the first phase. This is achieved by using a modular encoder-decoder architecture with adversarial regularization [40], that has state-of-the-art performance on benchmark datasets. In the second phase, we use the representation learning aspect of the framework to condition future poses and fine-tune only the decoder module on new samples in a curriculum learning setup [3]. This mitigates the problem of training from scratch while also providing performance gains, both quantitatively over short, mid, and long-term horizons and qualitatively in terms of generating motion that is perceptibly similar to the ground-truth.

2 PROBLEM FORMULATION

Formally defined, human motion prediction is the problem of predicting the future human pose over a horizon, given their past pose and any additional contextual information. In this paper, we assume that there is only one agent in the scene. For any particular scenario, the input to our model is the past or *observed* trajectory frames, spanning time $t = 1$ to τ ; $\mathbf{X} = \{x_1, \dots, x_\tau\}$. Each frame $x_t \in \mathbb{R}^N$ denotes the N -dimensional body pose. N depends on the number of joints in the skeleton, J and the dimension of the joints D , where $N = J \times D$. The expected output of the model is the future trajectory frames over horizon H , i.e. the ground truth pose over the horizon $t = \tau + 1$ to $\tau + H$: $\mathbf{Y} = \{y_{\tau+1}, \dots, y_{\tau+H}\}$.

Our first objective is to learn the underlying representation which would allow the model to generate feasible and accurate human poses $\hat{\mathbf{Y}} = \{\hat{y}_{\tau+1}, \dots, \hat{y}_{\tau+H}\}$. We assume that future human pose is conditioned on the past observed or generated poses and predict each frame in an auto-regressive manner as formulated below:

$$p_\theta(\hat{\mathbf{Y}}) = \prod_{\delta=\tau+1}^{\tau+H} p_\theta(\hat{y}_\delta | \hat{y}_{\tau, \delta-1}, x_{1:\tau}) \quad (1)$$

where the joint distribution is parameterized by θ .

Next, we use these learned parameters to fine-tune for a specific agent who was not observed during the training phase, using a continual learning setup. Instead of updating all the model parameters, we update a specific module, say the decoder module, with corresponding parameters θ^* . We formulate this as follows, similar to prior work in continual learning [26]:

$$\log p(\theta^* | D) = \log p(D_B | \theta) + \log p(\theta | D_A) - \log p(D_B) \quad (2)$$

where D_A represents the first phase's training data, which involves learning a representation from the large data distribution. D_B represents the second phase's training data, whereby we aim to learn the parameters for a specific human. $\log p(\theta | D_A)$ embeds all the prior information learned during the training phase.

3 CONTINUAL LEARNING FOR HUMAN MOTION PREDICTION

The collective goal of our approach is to accurately predict human motion while being flexible to parameter or architectural updates, given new data. Our overall framework is comprised of an encoder




Figure 2: Motion prediction architecture

and decoder, trained end-to-end with adversarial regularization on the latent variables, building on top of our prior work [40]. The encoder aims to learn a rich representation over past trajectories, which the decoder can use to condition its prediction. To improve model stability and robustness of the latent space, we use adversarial regularization through discriminators. This acts as a regularizer during training and can improve the network's stability during parameter updates over new data. We will first describe the overall model for motion prediction and then discuss its flexibility for fine-tuning on a particular agent.

3.1 Overall model for motion prediction

Motion Encoder: The encoder learns a representation over the high-dimensional observed trajectory, projecting the input to a low-dimensional latent space. To obtain a rich and more robust representation over the past trajectories, we extract the past velocity and acceleration features along with the provided positional values, in line with prior work on motion prediction [40]. The velocity and acceleration features are first and second-order derivatives of the position values for each skeleton joint.

For encoding spatio-temporal representation from the position, velocity, and acceleration data, we employ Recurrent Neural Networks, in particular Gated Recurrent Units (GRU). We use unidirectional GRUs, as we wish to predict human motion in real-time. For each stream, the stream-specific GRU aims to extract the spatio-temporal representation that summarizes the input sequence, with the operation formulated as follows:

$$h_{s,t} = GRU(h_{s,t-1}, x_{s,t}, \phi_s) \quad (3)$$

where s represents the specific stream: position, velocity or acceleration, $x_{s,t}$ denotes the input to the GRU at t , $h_{s,t-1}$ corresponds to the past hidden state and ϕ_s represent the parameters of the GRU. The output from each GRU represents disparate information corresponding to the past trajectory and needs to be fused adaptively. As such, we use a multi-head self-attention mechanism [38] which is tasked to disentangle and extract relevant stream-specific representation.

$$\begin{aligned} h_t &= \text{Concat}(h_{pos,t}; h_{vel,t}; h_{acc,t}); \\ h_{att,t} &= \text{Attention}(h_t; \phi_{att}) \end{aligned} \quad (4)$$

where $h_{att,t}$ is the output of the attention mechanism, and ϕ_{att} represents the parameters. The output $h_{att,t}$ is used to obtain the latent representation, which is tasked to characterize the observed trajectory.

Latent Representation: The latent representation aims to capture relevant spatial and temporal semantics from the observed data, which can then be used to condition motion prediction. The latent representation is comprised of a continuous random variable

and a categorical random variable. The motivation behind using both continuous and categorical variables is to jointly model the continuous aspect of human motion, such as the spatial semantics of a particular activity and the discrete characteristics of human motion such as the class activity or segment.

To obtain the continuous latent variable z_t , the output from the self-attention module is passed through a linear layer. In the case of the categorical latent variable c_t , the output from the self-attention module is passed through a linear layer followed by a softmax layer.

$$\begin{aligned} z_t &= \text{Linear}(h_{att,t}) \\ h_{c,t} &= \text{Linear}(h_{att,t}) \\ c_t &= \text{softmax}(h_{c,t}) \end{aligned} \quad (5)$$

Adversarial Regularization: To enforce a prior on the latent space, we use adversarial learning, similar to the Adversarial Autoencoders (AAE) [32] framework. This serves the purpose of a regularizer as there is a modification to the overall objective function: the objective function now consists of a reconstruction loss and an adversarial loss. We reason that this helps improve the stability of the overall framework for continual learning, as the parameters are updated based on two competing objectives: the reconstruction loss and the discriminator loss.

Similar to the GAN [13] and AAE [32] setups, the encoder aims to confuse the discriminators by trying to ensure that its output is similar to the aggregated prior. The discriminators are trained to distinguish the true samples generated using a given prior, from the latent space output of the encoder, thus establishing a min-max adversarial game between the networks [13, 32].

We use two discriminators, one for the continuous latent variable and the other for the categorical latent variable. The discriminators compute the probability that a point z_t or c_t is a sample from the prior distribution that we are trying to model (positive samples), or from the latent space (negative sample). We use a Gaussian prior for continuous latent variables and a uniform distribution prior for categorical latent variables.

Decoder: The decoder uses the latent representation and the past generated pose to predict the future pose for each time step. It is auto-regressive, i.e., it uses the output of the previous timestep to predict the current pose and has only one stream: position as the expected output is future joint positions of the human.

The input to the decoder is the latent representation: z_t and c_t and the past generated pose, or the seed pose at time t if it is predicting the first time-step, $t+1$. This is then passed to an attention mechanism that allows the decoder to adaptively condition its output on the latent variables that provide long-term information over the observed frames and the immediate generated frame. The output from this attention mechanism is next passed to a GRU cell, similar to the one at the encoder. This is followed by a Structured Prediction Layer (SPL) [2], which predicts each joint hierarchically following a skeleton tree, thus allowing the decoder to enforce structural prior on its final output. The operations at the decoder are formulated as follows:

$$\begin{aligned} p_t &= \text{Concat}(z_t, c_t, h_{dec,t-1}) \\ p_{att,t} &= \text{Attention}(p_t, \phi_{att}) \\ h_{dec,t} &= \text{GRU}(S_{t-1}, p_{att,t}, \phi_{pos}) \\ S_{t+1} &= \gamma(h_{dec,t}) \end{aligned} \quad (6)$$

3.2 Curriculum learning for the decoder:

The encoder-decoder architecture with adversarial regularization is trained to convergence on the training set. This training is followed by providing the overall architecture with unseen but small samples of motion data. This aims to relax the i.i.d assumption of the training procedure as our framework now has access to limited motion samples of the agent that it is trying to model.

Our choice of continual learning scheme is the curriculum learning setup [3], whereby we first train the network on a comparatively simpler task of representation learning, followed by a relatively difficult task of fine-tuning its parameters for a specific human subject. Our implementation is based on findings in connectionist models [35, 37], in particular self-organizing maps which reduce the levels of functional plasticity (i.e., ability to acquire knowledge in neural networks) through a two-phase training of the topographic neural map [27, 28]. The first phase is the organization phase, where the neural network is trained with a high learning rate and large spatial neighborhood size, allowing the network to reach an initial rough topological organization. The second phase is referred to as the tuning phase, where the learning rate and the neighborhood size are iteratively reduced for fine-tuning [34]. We aim to adopt these findings to a sequence learning framework.

Following prior work on developmental and curriculum learning [3, 14], we fine-tuned the architecture on the new data. We adopt techniques that will allow us to retain previous knowledge and avoid catastrophic forgetting during fine-tuning. In particular, we rely on *discriminative fine-tuning* [16], whereby we fine-tune only the decoder network at a different learning rate while freezing the encoder and the discriminator networks.

Fine-tuning the decoder: In line with equation 5, the input to the model is the sequence of observed poses: $\mathbf{X} = \{x_1, \dots, x_T\}$, with the output being of the encoder being z_T and c_T . However, instead of imposing a prior on the latent space and training the encoder-decoder end-to-end, we only update the decoder's parameters. We also use a lower learning rate and rely on a small number of training samples to improve model stability and reduce the likelihood of catastrophic forgetting.

We leverage the representation learning capability of the framework that it attained when training on a large and diverse dataset. The encoder network is tasked to provide a representation summarizing the past observation that is used by the decoder to condition its prediction, similar to equation 6. The pre-trained weights from the training set are used to initialize the overall architecture and act as prior knowledge. The decoder weights are updated based on the reconstruction loss on the new data.

4 EXPERIMENTAL SETUP

4.1 Dataset

We evaluated the performance of our approach on the widely used human-activity dataset: UTD-MHAD [9]. The dataset comprises 27 action classes covering activities from hand gestures to training exercises and daily activities, thus providing relevant activities for human-robot collaboration. Each activity was performed by 8 different subjects, with each subject repeating the activity 4 times. In our experiments, we use only Skeleton data for predicting human motion, following previous work in this domain [2, 8, 12, 25, 33], and

Table 1: MSE (in cm²) comparison of fine-tuning vs no fine-tuning on UTD-MHAD for different test subjects (Lower is better)

Frames	Subject 2						Subject 4						Subject 6					
	2	4	8	10	13	15	2	4	8	10	13	15	2	4	8	10	13	15
Zero-Velocity	11.20	27.37	66.85	86.23	112.39	127.64	13.11	32.33	77.22	97.80	123.99	138.53	10.37	25.56	64.62	85.29	115.33	135.17
Our method without Curriculum-Learning	5.41	14.75	33.87	41.89	51.33	56.12	7.78	18.26	41.66	51.46	62.67	68.1	6.99	16.18	38.51	49.40	63.7	71.68
Our method with Curriculum-Learning	7.62	16.17	32.35	38.63	45.72	49.13	6.98	14.79	30.29	36.51	43.81	47.61	6.51	13.44	27.49	33.66	41.74	46.41

considered each of the 20 provided joints. For all experiments, the model predicted output for the next 15 frames, using observation over the past 15 frames.

4.2 Generalized Representation Learning

We used the cross-subject evaluation scheme, training and validating on odd-numbered subjects for the first phase, thus providing the framework with a large training sample and maximizing the likelihood of encountering diverse demonstrations. To evaluate the performance, we hold out a section of the data for the validation set and early stopping. This reduces the likelihood of overfitting on the training data while also provide the mechanism for stopping the training procedure.

4.3 Curriculum Learning for a specific subject

Having learned a generalized representation, the second phase involved training the framework in a curriculum learning setup. Here, the experiments are conducted on a particular held-out even-numbered subject. We fine-tuned only the decoder using a reduced learning rate, with the encoder weights initialized from the first phase. As each subject has 4 trials, we trained on one trial and tested on the other 3 trials.

4.4 State-of-the-art method and baseline

For evaluating the efficacy of our curriculum learning setup, we compared against a non-curriculum learning framework, and the zero-velocity baseline [33]. The first benchmark [40] is comprised of an encoder-decoder framework, with adversarial regularization, but with no provision for curriculum learning. The zero-velocity baseline assumes that all the future predictions are identical to the last observed pose and is challenging to outperform for short-term prediction [2, 33]. It also allows us to gauge the movement dynamics, with a lower MSE for zero-velocity suggesting less movement and vice-versa for higher MSE.

4.5 Evaluation Metric

We evaluated the performance of all models using the Mean Squared Error (MSE), which is the l_2 distance between the ground-truth and the predicted poses at each timestep, averaged over the number of joints and sequence length, in line with prior work [1, 6, 7, 40]. The MSE is calculated as:

$$\mathcal{L}(\mathcal{X}, \hat{\mathcal{X}}) = \frac{1}{T \cdot K} \sum_{t=1}^T \sum_{i=1}^K (x_t^i - \hat{x}_t^i)^2 \quad (7)$$

where, T and K are the total number of frame and joints respectively.

5 RESULTS AND DISCUSSION

Results: We present the results of all approaches on the UTD-MHAD on table 1. We report the performance of all approaches at distinct frame intervals to circumvent the problem of frame drops during data collection and subsequent evaluation [40]. Our frame intervals aim to evaluate all models on short (2 & 4), mid (8 & 10), and long-term motion prediction (13 & 15). Table I depicts the performance of all approaches with respect to the test subjects. The results in Table I suggest that fine-tuning the framework allows it to outperforms all other methods and the zero-velocity baseline for short, mid, and long-term prediction.

Discussion: Our proposed approach outperformed the prior state-of-the-art approach and baseline both quantitatively by having lower MSE and qualitatively in terms of generating motion closer to the ground-truth pose (see Fig. 1). This shows the benefit of the curriculum learning approach while also suggesting that our overall framework is robust to catastrophic forgetting. The performance gain is especially significant over the mid and long-term, as the decoder is trained to learn the spatial-temporal movement pattern of a specific subject and can generate the future pose with higher accuracy.

Using a curriculum-learning setup, albeit on a small training sample, allows the framework to capture individual human motion subtleties, as seen by the lower MSE, particularly over the mid and long-term horizons. The approach is particularly useful when there is significant movement over the given horizon, as seen for Subjects 4 and 6 (table 1), who have higher MSE loss on the zero-velocity baseline. For Subject 2, there is overall less movement as seen by the zero-velocity MSE loss, and hence the performance gain is not significant over the mid and long-term and even worse over the short-term. Overall, the results are particularly promising as we did not fine-tune the encoder, instead only focusing on the decoder. Further improvement can be attained by fine-tuning the encoder.

6 CONCLUSION

In this work, we present a curriculum learning approach that opens the possibility of continual learning for human motion prediction, especially if the model is especially *deployed in the wild*. Our framework first learns a general representation over diverse training samples before fine-tuning on a target human subject. Our experiments suggest the feasibility of curriculum learning with performance gains over non-curriculum learning approaches. Future work will focus on fine-tuning to activities and domains that were not observed in training in a zero-shot learning setup.

7 ACKNOWLEDGEMENT

This work is supported by the CCAM Innovation Award. The authors thank Tim Bakker, Tomonari Furukawa, and Qing Chang for their support.

REFERENCES

- [1] Vida Adeli, Ehsan Adeli, Ian Reid, Juan Carlos Nieves, and S Hamid Rezatofighi. 2020. Socially and Contextually Aware Human Motion and Pose Forecasting. *IEEE Robotics and Automation Letters* (2020).
- [2] Emre Aksan, Manuel Kaufmann, and Otmar Hilliges. 2019. Structured prediction helps 3d human motion modelling. In *Proceedings of the IEEE International Conference on Computer Vision*. IEEE, 7144–7153.
- [3] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*. 41–48.
- [4] Abhijat Biswas, Henny Admoni, and Aaron Steinfeld. 2019. Fast on-board 3d torso pose recovery and forecasting. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*. IEEE.
- [5] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).
- [6] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellstrom. 2017. Deep representation learning for human motion prediction and classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, 6158–6166.
- [7] Judith Butepage, Ali Ghadirzadeh, Özge Öztimur Karadag, Márten Björkman, and Danica Kragic. 2019. Imitating by generating: deep generative models for imitation of interactive tasks. *arXiv preprint arXiv:1910.06031* (2019).
- [8] Judith Butepage, Hedvig Kjellström, and Danica Kragic. 2018. Anticipating many futures: Online human motion prediction and generation for human-robot interaction. In *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 1–9.
- [9] Chen Chen, Roozbeh Jafari, and Nasser Kehtarnavaz. 2015. UTD-MHAD: A multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor. In *2015 IEEE International conference on image processing (ICIP)*. IEEE, 168–172.
- [10] Trevor Darrell, Marius Kloft, Massimiliano Pontil, Gunnar Rätsch, and Erik Rodner. 2015. Machine learning with interdependent and non-identically distributed data (dagstuhl seminar 15152). In *Dagstuhl Reports*, Vol. 5. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [12] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. 2015. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*. IEEE, 4346–4354.
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [14] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature* 538, 7626 (2016), 471–476.
- [15] Guy Hoffman. 2019. Evaluating fluency in human–robot collaboration. *IEEE THMS* (2019).
- [16] Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146* (2018).
- [17] T. Iqbal, S. Li, C. Fourie, B. Hayes, and J. A. Shah. 2019. Fast Online Segmentation of Activities from Partial Trajectories. In *2019 ICRA*. 5019–5025. <https://doi.org/10.1109/ICRA.2019.8794054>
- [18] T. Iqbal, S. Rack, and L. D. Riek. 2016. Movement Coordination in Human-Robot Teams: A Dynamical Systems Approach. *IEEE Transactions on Robotics* 32, 4 (2016), 909–919.
- [19] Tariq Iqbal and Laurel D. Riek. 2014. Role Distribution in Synchronous Human-Robot Joint Action. In *Proc. of IEEE Symp. on Robot and Human Interactive Communication, Towards a Framework for Joint Action*.
- [20] T. Iqbal and L. D. Riek. 2016. A Method for Automatic Detection of Psychomotor Entrainment. *IEEE Transactions on Affective Computing* 7, 1 (2016), 3–16.
- [21] Tariq Iqbal and Laurel D. Riek. 2017. Coordination Dynamics in Multi-human Multi-robot Teams. *IEEE Robotics and Automation Letters (RA-L)* (2017).
- [22] T. Iqbal and L. D. Riek. 2017. Human Robot Teaming: Approaches from Joint Action and Dynamical Systems. *Humanoid Robotics: A Reference*, Springer (2017).
- [23] Md Mofijul Islam and Tariq Iqbal. 2020. HAMLET: A Hierarchical Multimodal Attention-based Human Activity Recognition Algorithm. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [24] Md Mofijul Islam and Tariq Iqbal. 2021. Multi-GAT: A Graphical Attention-based Hierarchical Multimodal Representation Learning Approach for Human Activity Recognition. In *IEEE RA-L*.
- [25] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. 2016. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, 5308–5317.
- [26] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* 114, 13 (2017), 3521–3526.
- [27] Teuvo Kohonen. 1982. Self-organized formation of topologically correct feature maps. *Biological cybernetics* 43, 1 (1982), 59–69.
- [28] Teuvo Kohonen. 1997. Exploration of very large databases by self-organizing maps. In *Proceedings of international conference on neural networks (icnn'97)*, Vol. 1. IEEE, PL1–PL6.
- [29] Alyssa Kubota, Tariq Iqbal, Julie A Shah, and Laurel D Riek. 2019. Activity recognition in manufacturing: The roles of motion capture and sEMG+ inertial wearables in detecting fine vs. gross motion. In *2019 ICRA*. IEEE, 6533–6539.
- [30] Dharshan Kumaran, Demis Hassabis, and James L McClelland. 2016. What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in cognitive sciences* 20, 7 (2016), 512–534.
- [31] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. 2020. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information fusion* 58 (2020), 52–68.
- [32] Alyreza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644* (2015).
- [33] Julieta Martinez, Michael J Black, and Javier Romero. 2017. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2891–2900.
- [34] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks* 113 (2019), 54–71.
- [35] Fiona M Richardson and Michael SC Thomas. 2008. Critical periods and catastrophic interference effects in the development of self-organizing feature maps. *Developmental science* 11, 3 (2008), 371–389.
- [36] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. 2020. Human motion trajectory prediction: A survey. *IJRR* (2020).
- [37] Michael SC Thomas and Mark H Johnson. 2006. The computational modeling of sensitive periods. *Developmental Psychobiology* 48, 4 (2006), 337.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [39] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237* (2019).
- [40] Mohammad Samin Yasar and Tariq Iqbal. 2021. A Scalable Approach to Predict Multi-Agent Motion for Human-Robot Collaboration. In *IEEE RA-L*.