

Application of Graph Convolutions in a Lightweight Model for Skeletal Human Motion Forecasting

Luca Hermes, Barbara Hammer and Malte Schilling *

Machine Learning Group, Bielefeld University
33501 Bielefeld, Germany

Abstract. Prediction of movements is essential for successful cooperation with intelligent systems. We propose a model that integrates organized spatial information as given through the moving body’s skeletal structure. This inherent structure is exploited in our model through application of Graph Convolutions and we demonstrate how this allows leveraging the structured spatial information into competitive predictions that are based on a lightweight model that requires a comparatively small number of parameters.

1 Introduction

Human motion forecasting has many useful applications. As intelligent systems should interact with humans, it becomes necessary to predict human movements and actions. Take, as one example, a cooperative task in which a human and a robot should safely collaborate in an assembly process. Therefore, multiple techniques have been developed to tackle such tasks. Ranging from simple heuristic approaches towards learning-based approaches. Machine learning approaches for such temporal data were mostly relying on recurrent networks that were applied to vectorized representations of joints. These were using skeletal structure only implicitly. The latest advances in geometric deep learning (GDL) provide a direct way to leverage the skeletal structure for predictions via graph convolutions (GCNs). In this work, we propose a straightforward implementation of such a model. We use spatio-temporal convolutions together with a type of GCN to extract spatio-temporal features in human motion, resulting in a simple autoregressive model. We adapt dilated causal convolutions for temporal modeling as used in [1], but include local joint connectivity which leads to a lightweight spatio-temporal operation. The code will be made publicly available.¹

2 Methods

In this section, we introduce our deep geometric model for motion forecasting. The model is based on *Graph-WaveNet* [2], a spatio-temporal extension to the original *WaveNet* [1].

*This research was supported by the research training group “Dataninja” (Trustworthy AI for Seamless Problem Solving: Next Generation Intelligence Joins Robust Data Analysis) funded by the German federal state of North Rhine-Westphalia.

¹<https://github.com/LucaHermes/lightweight-motion-forecasting>

Input to the model is a time series of T consecutive skeletal states with J joints and a d_j -dimensional joint representation, $\mathbf{S} \in \mathbb{R}^{T \times J \times d_j}$ (Fig. 1 shows the model architecture). In our experiments, we represent joints as quaternions, i.e. $d_j = 4$.

First, the linear input layer (1×1) is applied to each joint and acts as a trainable embedding for the d_j -dimensional joint inputs. Secondly, N consecutive Graph-WaveNet blocks extract spatio-temporal features from the given time series. Every block produces a skip output and the sum of these outputs are, thirdly, passed to an MLP that is applied per joint.

Every Graph-WaveNet block performs a spatio-temporal convolution (ST-Conv) followed by a purely spatial graph-convolution (K-GCN), as shown in Fig. 1. Both operations use the same output dimensionality, to which we refer to as the block dimensionality d_b . A residual connection is applied which bypasses both convolutions by adding the block input to the block output. The result is fed to the subsequent block.

We use the following model configuration in our experiments: The input layer consists of 64 neurons. We use $N = 5$ blocks with $d_b = 64$ and a skip output dimensionality of 256. The two layers in the output MLP have 256 and 4 neurons, respectively. The resulting model has 4.46×10^5 trainable parameters.

Spatio-Temporal Convolution In the original Graph-WaveNet a purely temporal convolution is used to extract temporal features from the input. In contrast, we substitute the temporal convolution that acts on each joint individually with a spatio-temporal convolution that integrates information from neighboring joints. This operation is inspired by the temporal extension module (TEM) [3], which is a type of GCN placed prior to a temporal operation. Therefore, the temporal operation operates on information from the neighborhood \mathcal{N}_j .

We integrate the idea of TEM directly into the spatio-temporal convolution operation. This means, instead of applying a 1-dimensional convolution on the trajectory of a single joint, we apply a 2-dimensional spatio-temporal convolution on the trajectory of a kinematic chain of joints. The joint hierarchy is given by the kinematic tree of the skeleton. Through selecting the hip as the root joint, an ordering is introduced that extends into the leaf-joints, i.e. the head, hands and foot joints.

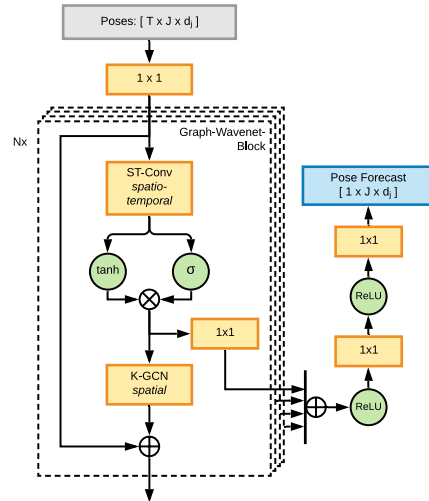


Fig. 1: Model architecture, with N consecutive spatio-temporal processing blocks, followed by two alternating ReLU and linear layers (1×1). Rounded arrows denote a channel-split.

Fig. 2 visualizes our spatio-temporal convolution. First, the parent and grandparent joint for every joint j are sampled based on the kinematic tree. Afterwards, the trajectories for these three joints are stacked and convolved using a convolutional kernel $\mathbf{W} \in \mathbb{R}^{\tau_s \times \tau_t \times d_{in} \times d_{out}}$ (blue rectangle). The kernel size is denoted as τ_s and τ_t in the spatial and temporal dimension, respectively. The applied convolution is causal and dilated in the temporal dimension following the pattern of [1, 2].

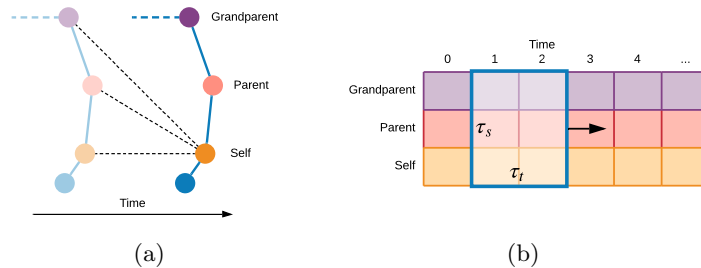


Fig. 2: (a) Sampling strategy for the ST-Conv. (b) Computation of the spatio-temporal convolution. A 2-dimensional kernel computes the convolution of the time-series from three joints. The kernel size in the spatial and the temporal domain is denoted by τ_s and τ_t , respectively.

Spatial Graph-Convolution The purely spatial convolution is computed using graph convolutions (GCNs) as defined in [4]. Instead of using the full skeletal graph, we again utilize the kinematic tree that was described above to convert the undirected skeletal graph into three directed subgraphs: In the first subgraph, an edge $e = (j, u)$ exists if and only if joint j relative to joint u is further up in the hierarchy of the kinematic tree. This subgraph retains all edges linking a joint to its immediate child joints. The second subgraph is similarly constructed, but with inverted edge direction, therefore retaining all edges linking a joint to its immediate parent joints. The third subgraph consists only of self-loops. The subgraphs are represented by three adjacency matrices \mathbf{A}_i . Note that \mathbf{A}_2 corresponds to the identity matrix. The output of this operation is computed as follows:

$$\text{K-GCN}(\mathbf{X}_{in}) = \sum_{i=0}^2 \mathbf{D}_i^{-1} \mathbf{A}_i \mathbf{X}_{in} \mathbf{W}_i$$

The inputs are given as \mathbf{X}_{in} , \mathbf{W}_i is the parameter matrix of the i^{th} GCN operation and \mathbf{D}_i denotes the degree matrix of subgraph i .

3 Results

In this section, we describe the experiments we conducted on the presented model. To quantify the performance, we perform an evaluation using protocols

from the related literature and provide qualitative results.

Dataset and Training Setup We train and evaluate the model on the Human3.6M (H3.6M) dataset [5, 6]. The dataset consists of motion capture data from seven different human subjects performing 15 different actions. Table 1 and 2 provide an overview of these actions. We follow the general preprocessing procedure in [7], downsampling the dataset by a factor of 2 to 25 Hz and mirror each trajectory across the y - z -plane.

We split the data by subject into a train (subjects 1, 7, 8, 9, 11), validation (subject 6) and test set (subject 5). The models are trained using a batch size of 16 trajectories over 3000 epochs on the mean absolute error in the quaternion space. An epoch consists of five samples from every sequence. A sequence is a single trial of one subject, performing a single action. Each sample contains successive frames from a single downsampled or mirrored version of the trajectories. The model is trained using the Adam optimizer with an initial learning rate of $\eta = 0.001$ that is decayed exponentially by a factor of $\alpha = 0.999$ after every epoch. We use a seed trajectory length of 32 frames ($= 1.28sec$) to condition the model, because this perfectly fits into the receptive field. The target sequence has a length of 10 frames ($= 400ms$) and was generated autoregressively.

Model Evaluation We show qualitative and quantitative results for the proposed model. For our model evaluation, we use the non-mirrored, but downsampled dataset. For the quantitative evaluation, we follow the standard evaluation protocol of [8]. This protocol first constructs an evaluation set that consists of four random samples from every sequence in the test set, i.e. the trials of subject 5.² An error metric quantifies the performance of the model on this evaluation set. The metric used by [8] is the Euclidean distance between the predicted and target rotations converted to Euler angles. The following equation summarizes this metric:

$$d(\mathbf{x}, \mathbf{y}) = \frac{1}{T} \sum_{t=0}^T \sqrt{\sum_j^J \sum_d^d (x_{t,j,d} - y_{t,j,d})^2},$$

where T is the number of time-steps, J is the number of joints and the sum over d accumulates the error in the x , y and z dimension of the given Euler angles. The final results correspond to the average taken over the four samples. Following [9], the results for a running average over 2 and 4 frames (Run. avg. 2/4) and a zero-velocity-model are also documented as baselines. The zero-velocity model returns the first observed frame as the prediction for all successive frames.

Table 1 shows the results under the standard protocol for four actions. The running average and zero-velocity baselines are included, as well as results for multiple SoTA models. Table 2 lists additional results for the remaining 12 actions (we include results from other models when available for comparison). Our

²We were able to draw the same sequences as [8] using the *RandomState* implementation of the *random* package of *NumPy* (v. 1.19.2) with a seed of $s = 1234567890$.

forecasting model generally shows competitive results that in some cases, e.g. *eating*, *directions*, and *walkingtogether*, outperform the referenced approaches. The original Graph-WaveNet architecture converges to a zero-velocity model.

Table 1: Quantitative results on action forecasting under the standard protocol of [8]. We further specify the input and output mode of some models (*input/output*), where either velocities (*vel.*) or absolute angles (*abs.*) are used.

Milliseconds	Walking				Eating				Smoking				Discussion				No. of Parameters
	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400	
Run.-avg. 4	0.64	0.87	1.07	1.20	0.40	0.59	0.77	0.88	0.37	0.58	1.03	1.02	0.60	0.90	1.11	1.15	-
Run.-avg. 2	0.48	0.74	1.02	1.17	0.32	0.52	0.74	0.87	0.30	0.52	0.99	0.97	0.41	0.74	0.99	1.09	-
Zero-velocity	0.39	0.68	0.99	1.15	0.27	0.48	0.73	0.86	0.26	0.48	0.97	0.95	0.31	0.67	0.94	1.04	-
LSTM-3LR [8]	0.77	1.00	1.29	1.47	0.89	1.09	1.35	1.46	1.34	1.65	2.04	2.16	1.88	2.12	2.25	2.23	1.48×10^7
GRU <i>sup.</i> [9]	0.28	0.49	0.72	0.81	0.23	0.39	0.62	0.76	0.33	0.61	1.05	1.15	0.31	0.68	1.01	1.09	3.37×10^6
QuaterNet GRU <i>abs./vel.</i> [7]	0.21	0.34	0.56	0.62	0.20	0.35	0.58	0.70	0.25	0.47	0.93	0.90	0.26	0.60	0.85	0.93	9.5×10^6
QuaterNet CNN <i>abs./vel.</i> [7]	0.25	0.40	0.62	0.70	0.22	0.36	0.58	0.71	0.26	0.49	0.94	0.90	0.30	0.66	0.93	1.00	8.8×10^6
DCT-GCN <i>short-term</i> [10]	0.18	0.31	0.49	0.56	0.16	0.29	0.50	0.62	0.22	0.41	0.86	0.80	0.20	0.51	0.77	0.85	2.6×10^6
DMGNN [11]	0.18	0.31	0.49	0.58	0.17	0.30	0.49	0.59	0.21	0.39	0.81	0.77	0.26	0.65	0.92	0.99	6.26×10^7
Ours <i>abs./vel.</i>	0.23	0.37	0.61	0.69	0.18	0.31	0.54	0.66	0.23	0.46	0.93	0.90	0.31	0.70	0.97	1.07	4.46×10^5
Ours <i>vel./vel.</i>	0.19	0.34	0.57	0.63	0.16	0.29	0.50	0.60	0.22	0.41	0.85	0.81	0.22	0.57	0.84	0.93	4.46×10^5

Table 2: Quantitative results for the remaining actions; Continuation of table 1.

Milliseconds	Directions				Greeting				Phoning				Posing			
	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400
GRU <i>sup.</i>	0.26	0.47	0.72	0.84	0.75	1.17	1.74	1.83	0.23	0.43	0.69	0.82	0.36	0.71	1.22	1.48
DCT-GCN <i>st</i>	0.26	0.45	0.71	0.79	0.36	0.60	0.95	1.13	0.53	1.02	1.35	1.48	0.19	0.44	1.01	1.24
DMGNN	0.25	0.44	0.65	0.71	0.36	0.61	0.94	1.12	0.52	0.97	1.29	1.43	0.20	0.46	1.06	1.34
Ours <i>abs./vel.</i>	0.32	0.47	0.68	0.80	0.42	0.72	1.14	1.36	0.54	1.00	1.34	1.47	0.27	0.55	1.05	1.27
Ours <i>vel./vel.</i>	0.24	0.43	0.77	0.81	0.35	0.61	1.01	1.20	0.53	1.00	1.28	1.40	0.26	0.51	1.08	1.32
Milliseconds	Purchases				Sitting				Sittingdown				Takingphoto			
	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400
GRU <i>sup.</i>	0.51	0.97	1.07	1.16	0.41	1.05	1.49	1.63	0.39	0.81	1.40	1.62	0.24	0.51	0.90	1.05
DCT-GCN <i>st</i>	0.43	0.65	1.05	1.13	0.29	0.45	0.80	0.97	0.30	0.61	0.90	1.00	0.14	0.34	0.58	0.70
DMGNN	0.41	0.61	1.05	1.14	0.26	0.42	0.76	0.97	0.32	0.65	0.93	1.05	0.15	0.34	0.58	0.71
Ours <i>abs./vel.</i>	0.56	0.75	1.03	1.15	0.31	0.50	0.91	1.12	0.33	0.65	0.96	1.09	0.19	0.42	0.73	0.93
Ours <i>vel./vel.</i>	0.42	0.61	1.08	1.15	0.30	0.49	0.90	1.09	0.29	0.65	0.97	1.08	0.15	0.34	0.58	0.72
Milliseconds	Waiting				Walkingdog				Walkingtogether				Average			
	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400
GRU <i>sup.</i>	0.28	0.53	1.02	1.14	0.56	0.91	1.26	1.40	0.31	0.58	0.87	0.91	0.36	0.67	1.02	1.15
DCT-GCN <i>st</i>	0.23	0.50	0.91	1.14	0.46	0.79	1.12	1.29	0.15	0.34	0.52	0.57	0.27	0.51	0.83	0.95
DMGNN	0.22	0.49	0.88	1.10	0.42	0.72	1.16	1.34	0.15	0.33	0.50	0.57	0.27	0.52	0.83	0.95
Ours <i>abs./vel.</i>	0.25	0.51	0.93	1.10	0.46	0.79	1.16	1.32	0.17	0.37	0.58	0.65	0.32	0.57	0.90	1.04
Ours <i>vel./vel.</i>	0.21	0.51	0.97	1.17	0.43	0.78	1.10	1.24	0.15	0.32	0.50	0.54	0.27	0.52	0.87	0.98

Fig. 3 visualizes predicted and true individual quaternion dimensions. Overall, the predictions are very similar to ground truth, but finer details are smoothed out. This is clearly visible in the trajectory of the right knee joint performing the *walkingdog* action. However, the model is still able to correctly predict the phase timing and thus model major movement features even in the long-term future ($> 400ms$).

4 Conclusions

We combine the well-established causal dilated convolutions from WaveNet with geometric deep learning principles resulting in a lightweight autoregressive model.

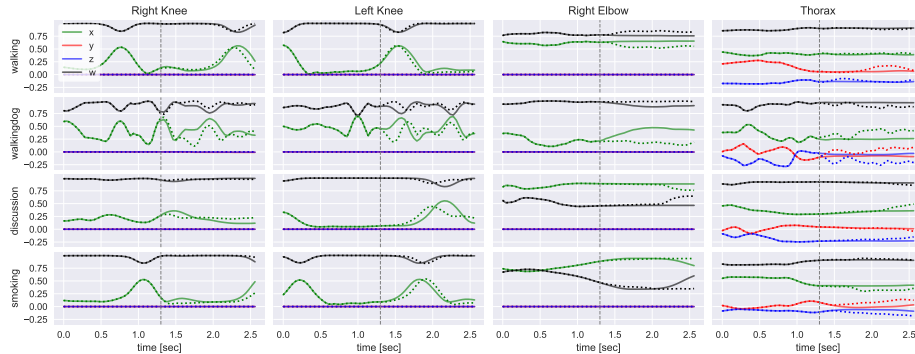


Fig. 3: Ground truth (dotted lines) and prediction (solid lines) for four different joints (columns) and four different actions (rows) of the second trial from subject 5. Each plot shows 32 seed frames (= 1.28sec) and 32 target frames.

With about 4.46×10^5 parameters, our model is an order of magnitude less complex compared to current models (cf. Table 1). Nonetheless, it shows competitive results compared to current approaches when evaluated on the H3.6M dataset for skeletal human motion forecasting. Furthermore, we show that this model is able to predict the phase timing even in a long-term forecasting setup.

References

- [1] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu. Wavenet: A generative model for raw audio. *CoRR*, abs/1609.03499, 2016.
- [2] Z. Wu, S. Pan, G. Long, J. Jiang, and C. Zhang. Graph wavenet for deep spatial-temporal graph modeling. *CoRR*, abs/1906.00121, 2019.
- [3] Y. Obinata and T. Yamamoto. Temporal extension module for skeleton-based action recognition, 2020.
- [4] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907, 2016.
- [5] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014.
- [6] C. Ionescu, F. Li, and C. Sminchisescu. Latent structured models for human pose estimation. In *International Conference on Computer Vision*, 2011.
- [7] D. Pavllo, C. Feichtenhofer, M. Auli, and D. Grangier. Modeling human motion with quaternion-based neural networks. *CoRR*, abs/1901.07677, 2019.
- [8] K. Fragkiadaki, S. Levine, and J. Malik. Recurrent network models for kinematic tracking. *CoRR*, abs/1508.00271, 2015.
- [9] Julieta Martinez, Michael J. Black, and Javier Romero. On human motion prediction using recurrent neural networks. *CoRR*, abs/1705.02445, 2017.
- [10] W. Mao, M. Liu, M. Salzmann, and H. Li. Learning trajectory dependencies for human motion prediction. *CoRR*, abs/1908.05436, 2019.
- [11] M. Li, S. Chen, Y. Zhao, Y. Zhang, Y. Wang, and Q. Tian. Dynamic multiscale graph neural networks for 3d skeleton-based human motion prediction. *arXiv:2003.08802 [cs.CV]*, 2020.