
Contrastive Self-Supervised Learning for Skeleton Representations

Nico Lingg*

Personal Robotics Lab
Imperial College London
n.lingg20@imperial.ac.uk

Miguel Sarabia, Luca Zappella, Barry Theobald

Apple
{miguelSDC, lzappella, btheobald}@apple.com

Abstract

Human skeleton point clouds are commonly used to automatically classify and predict the behaviour of others. In this paper, we use a contrastive self-supervised learning method, SimCLR, to learn representations that capture the semantics of skeleton point clouds. This work focuses on systematically evaluating the effects that different algorithmic decisions (including augmentations, dataset partitioning and backbone architecture) have on the learned skeleton representations. To pre-train the representations, we normalise six existing datasets to obtain more than 40 million skeleton frames. We evaluate the quality of the learned representations with three downstream tasks: skeleton reconstruction, motion prediction, and activity classification. Our results demonstrate the importance of 1) combining spatial and temporal augmentations, 2) including additional datasets for encoder training, and 3) and using a graph neural network as an encoder.

1 Introduction

Skeleton point clouds are commonly used to understand human behaviour, including activity classification [7, 23], gait analysis [15], and command recognition [9]; in fact, many sensors have been developed to easily capture skeleton data, whether through depth cameras or motion capture systems. However, working with skeletons has inherent challenges, such as occlusions, noisy data, missing joints, and poor skeleton fittings [19]. Self-supervised learning has successfully been used to learn representations robust to noise and missing data in computer vision [11], speech recognition [22] and natural language processing [1]. Inspired by this success, we aim to train a skeleton representation which generalises to three downstream tasks: skeleton reconstruction, motion prediction, and activity classification (see Fig. 1b).

We use SimCLR [3], a multi-view contrastive self-supervised learning algorithm, which relies on augmentations to create two views of the same input sample. SimCLR uses a loss which attracts the representations of views from the same samples while repelling the representations of views from different samples (cf. Fig. 1a). We adapt SimCLR using a graph neural network backend, ST-GCN [23], and employ eight augmentations (both spatial and temporal) to generate the views. It is worth noting that the choice of augmentations is a form of inductive bias – it determines what invariances will be built into the learned representations. In this work, we seek representations that are invariant to movement speeds, the framerate of the recording, and sensor noise.

Training with SimCLR requires making many algorithmic decisions involving network architecture, datasets, augmentations, etc. Our focus is to determine the effect of these decisions on downstream performance, hence we run exhaustive experimental ablations. Our results highlight the importance of combining both spatial and temporal augmentations, as well as the impact of using ST-GCN as an encoder with additional datasets to learn the representation. In fact, we gather six existing

*Work done during an internship at Apple.

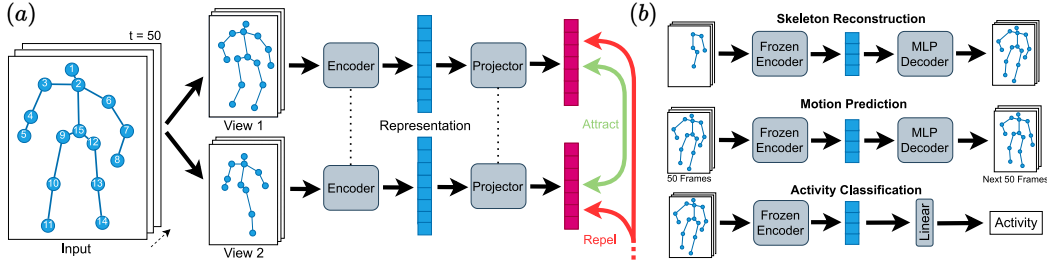


Figure 1: (a) Adapted architecture of SimCLR [3] for skeleton data. We use a custom augmentation stack (cf. Appendix A) and a ST-GCN [23] encoder. (b) Downstream tasks to evaluate the quality of the learned representation: skeleton reconstruction, motion prediction, and activity classification. Joints: 1) head, 2) neck, 3) right shoulder, 4) right elbow, 5) right wrist, 6) left shoulder, 7) left elbow, 8) left wrist, 9) right hip, 10) right knee, 11) right ankle, 12) left hip, 13) left knee, 14) left ankle, 15) torso.

skeleton datasets captured with different sensors, subjects, conditions, and labels, which results in approximately 40 million frames, each with 15 3d joint positions.

2 Background

Our work is at the intersection of both multi-view self-supervised learning, and skeleton representation learning; therefore we will review the relevant background of both fields.

Multi-View Self-Supervised Learning. Multi-view self-supervised learning is a form of representation learning that uses augmentations to generate different views of the input. The views are then used to contrast between the representations of views from the same input and representations of views from other inputs. In recent years, various multi-view self-supervised learning frameworks have been proposed, including CPC [18], MoCO v3 [4], DINO [2] and SimCLR [3]. These approaches have shown impressive results in transferability to smaller datasets like CIFAR 10, where DINO achieves an accuracy of 99.6%. SimCLR, in particular, has been shown to be flexible, stable to train and does not need extra logic to mine negatives samples [3]. For these reasons, we make use of SimCLR in this work.

Skeleton Representation Learning. There are several related approaches to learning representations of skeleton point-clouds. In particular, [21] introduced a custom pretext task based on motion consistency and motion continuity as well as an ST-GCN encoder to learn skeleton representations. For this work, we reuse ST-GCN as our representation encoder. More recently, [5] made use of a transformer architecture and pretext tasks inspired by BERT [8] to achieve the state-of-the-art in activity classification accuracy. Additionally, CP-STN [24] combines contrastive learning with pretext tasks (e.g., masked sequences prediction) to extract spatial and temporal information to learn fully discriminative representations. Other works, like CrossSCLR [13], further attempt to improve contrastive learning by using extra hard contrastive pairs stored in a memory bank. Unlike these approaches, which use custom pretext tasks to learn their representations, we take advantage of SimCLR to train the representation encoder.

3 Learning a Skeleton Representation with SimCLR

We aggregate six skeleton datasets captured with different sensors, subjects, conditions, and labels (refer to Table 1 for an overview). We hypothesise that the combined diversity of the data leads to better representations (as shown in Section 4). Our custom pre-processing pipeline² consists of the following five steps: 1) transform the datasets joint representation into 3d global coordinates, 2) reduce the joint-space to 15 fundamental joints commonly provided by skeleton sensing systems, 3) scale each skeleton to a height of 2m, i.e. have the z-axis extend from -1.0 to 1.0, 4) translate the skeleton so that the torso is at the centre-of-coordinates, and 5) rotate each frame around the x- and z-axes so that the skeleton faces the camera. We do not alter frame-rate, nor the skeleton structure, so each skeleton retains its scaled bone lengths.

²Code available at: <https://github.com/apple/ml-skeleton-preprocessing>

Table 1: Datasets used in this work. Training set of datasets marked with * are used to pre-train skeleton representation with SimCLR, those marked with \diamond are used for downstream task training.

DATASET	NUMBER OF TRAIN FRAMES	NUMBER OF TEST FRAMES	UNIQUE HUMAN SUBJECTS	FRAMES PER SECOND	RECORDING DEVICE
NTU RGB+D 60* \diamond [20]	3,328,560 (8%)	1,456,493 (92%)	40 (22%)	30	Depth Camera
NTU RGB+D 120* [14]	4,324,668 (11%)	–	66 (37%)	30	Depth Camera
TRINITY SPEECH-GESTURE* [10]	4,579,806 (11%)	–	1 (1%)	120	MoCap System
TALKING WITH HANDS* [12]	25,680,106 (64%)	–	50 (28%)	90	MoCap System
DANCEDB* [16]	1,679,209 (4%)	–	17 (9%)	120	MoCap System
HDM05 \diamond [17]	492,340 (1%)	123,085 (8%)	5 (3%)	120	MoCap System
TOTAL	40,084,689	1,579,578	179	–	

We pre-train the skeleton representation by feeding two augmented views of a 50-frame skeleton sequence into a graph encoder (ST-GCN [23]) which outputs a vector of 128 elements. The representation is then projected through a 3-layer multi-layer perceptron (MLP). Lastly, we use the NT-Xent loss [3] to attract projections from the same skeleton sequences and repel projections from different inputs. We pre-train SimCLR for 500 epochs with the datasets described above and then keep only the trained encoder for downstream tasks. We experimented with a larger number of epochs for pre-training, but found no improvement in downstream task performance. Pre-training and training hyperparameters are fixed across all experiments and are detailed in Appendix B.

To obtain the different views of our skeleton sequences, we use eight augmentations from two categories: spatial augmentations (joint dropout, joint jitter, axis mirroring, random scaling, frame dropout), and temporal augmentations (speed-up, slow-down). We describe the augmentations in detail in Appendix A. The purpose of these augmentations is to learn a skeleton representation invariant to noise in the skeleton joints and the actual speed of the human movements.

4 Ablations on Skeleton Representations

We analyse the effect of different algorithmic decisions on the quality of the learned representations by measuring the impact of these changes on the performance of the following downstream tasks:

Skeleton Reconstruction: The input is a sequence of skeleton frames from which 80% of randomly selected frames have been dropped. The target is the uncorrupted skeleton sequence. We use a 3-layer MLP decoder, a mean-squared error (MSE) loss, and evaluate against the NTU-60 and HDM05 datasets. We report the test mean absolute joint error (MAE) in millimetres for a two-metre tall normalised skeleton. We introduce the MAE metric to have a consistent measure across test datasets.

Motion Prediction: The target is to predict the next 50 frames of motion given the representation of the previous 50 frames of data. As above, we use the 3-layer MLP, MSE loss, evaluate on NTU-60 and HDM05, and report the test MAE in millimetres. We use MAE for consistency reasons (e.g. [23] and [6]).

Activity Classification: We use NTU-60 [20] labels, a linear classifier, and cross-entropy loss to classify motion segments. We use 50 frames from one skeleton to train the linear classifier rather than the standard 300 frames from two skeletons commonly used in the literature. Moreover, some of the activities in NTU-60 require joints not present in our representation, making it challenging to draw comparisons against existing approaches which take advantage of all the joints.

The results of our ablations are presented in Table 2. Note that we report the performance from the final epoch. Our baseline models, which consist of a linear classifier/decoder trained for 200 epochs on top of the frozen representation, achieve an MAE of between 50mm and 60mm in most skeleton reconstruction and motion prediction tasks (except in motion prediction for HDM05 where the MAE is 23.4mm) and 45.8% accuracy in the 60-class activity classifier task³.

We remark that while fine-tuning the representation yields an improvement in activity classification, it actually hurts performance on skeleton reconstruction and motion prediction. This is due to the training sets needing more samples to fully fine-tune the downstream tasks models. Our results also

³For context, [23] and [6] achieve 81.5% and 90% accuracy respectively. Still, a direct comparison cannot be made due to the reduced joint space and temporal receptive field of our models

Table 2: Effect of various ablations for training skeleton representations with SimCLR on downstream tasks. Each cell denotes mean performance and standard deviation across 3 different random seeds. For each row in the table, we pre-train 3 different encoder networks/linear classifiers using SimCLR and train each network on the downstream task with a frozen backbone. Baseline configuration has all augmentations (except random rotation), is trained with the full training set size, and was pre-trained with all five skeleton datasets. We highlighted changes over 8% in bold.

	SKELETON RECONSTRUCTION		MOTION PREDICTION		ACTIVITY CLASS.
	MEAN ABSOLUTE ERROR (mm)		MEAN ABSOLUTE ERROR (mm)		ACCURACY (%)
	<i>Lower is better</i>		<i>Lower is better</i>		<i>Higher is better</i>
	NTU-60	HDM05	NTU-60	HDM05	NTU-60
BASILINE	62.0 (± 2.3)	49.7 (± 1.5)	52.1 (± 0.6)	23.4 (± 1.0)	45.8 (± 1.0)
FINETUNING	65.7 (± 1.7), $\Delta+3.7$	55.0 (± 2.9), $\Delta+5.3$	56.5 (± 1.2), $\Delta+4.4$	44.7 (± 3.2), $\Delta+21.3$	64.7 (± 1.0), $\Delta+18.9$
30 FPS	60.8 (± 0.7), $\Delta-1.2$	53.8 (± 1.6), $\Delta+4.1$	52.8 (± 0.6), $\Delta+0.7$	37.1 (± 0.6), $\Delta+13.7$	45.1 (± 0.1), $\Delta-0.7$
MLP BACKBONE	64.6 (± 2.3), $\Delta+2.6$	58.8 (± 3.3), $\Delta+9.1$	59.0 (± 2.3), $\Delta+6.9$	29.7 (± 1.5), $\Delta+6.3$	29.3 (± 0.9), $\Delta-16.5$
DOWNSTREAM TASK TRAINING SET SIZE					
15%	59.5 (± 0.9), $\Delta-2.5$	52.0 (± 0.8), $\Delta+2.3$	58.3 (± 1.0), $\Delta+6.2$	23.7 (± 0.9), $\Delta+0.3$	39.1 (± 1.7), $\Delta-6.7$
10%	58.2 (± 0.6), $\Delta-3.8$	52.6 (± 1.2), $\Delta+2.9$	59.4 (± 0.7), $\Delta+7.3$	54.3 (± 2.3), $\Delta+30.9$	37.1 (± 1.0), $\Delta-8.7$
5%	57.9 (± 0.5), $\Delta-4.1$	73.1 (± 1.3), $\Delta+23.4$	76.7 (± 4.0), $\Delta+24.6$	73.5 (± 6.6), $\Delta+50.1$	33.3 (± 0.4), $\Delta-12.5$
SIMCLR PRE-TRAINING DATASETS					
W/O TRINITY	60.4 (± 0.4), $\Delta-1.6$	49.6 (± 0.9), $\Delta-0.1$	51.9 (± 0.6), $\Delta-0.2$	23.6 (± 0.8), $\Delta+0.2$	44.5 (± 1.5), $\Delta-1.3$
W/O TWH	60.0 (± 0.2), $\Delta-2.0$	50.3 (± 0.7), $\Delta+0.6$	52.4 (± 0.4), $\Delta+0.3$	23.3 (± 0.4), $\Delta-0.1$	44.9 (± 0.5), $\Delta-0.9$
W/O DANCEDB	61.3 (± 0.6), $\Delta-0.7$	51.2 (± 0.8), $\Delta+1.5$	54.0 (± 1.7), $\Delta+1.9$	24.5 (± 0.8), $\Delta+1.1$	45.5 (± 1.4), $\Delta-0.3$
W/O NTU-60 & 120	64.2 (± 0.3), $\Delta+2.2$	50.9 (± 0.6), $\Delta+1.2$	54.9 (± 0.5), $\Delta+2.8$	24.7 (± 0.5), $\Delta+1.3$	40.5 (± 0.2), $\Delta-5.3$
NTU-60 ONLY	62.4 (± 0.7), $\Delta+0.4$	53.2 (± 0.6), $\Delta+3.5$	53.3 (± 0.3), $\Delta+1.2$	26.4 (± 0.9), $\Delta+3.0$	39.2 (± 0.2), $\Delta-6.6$
AUGMENTATIONS					
W/O AXIS MIRRORING	59.8 (± 1.2), $\Delta-2.2$	48.2 (± 1.0), $\Delta-1.5$	51.5 (± 0.9), $\Delta-0.6$	22.9 (± 0.4), $\Delta-0.5$	44.5 (± 0.5), $\Delta-1.3$
W/O RANDOM SCALING	62.3 (± 0.5), $\Delta+0.3$	51.8 (± 0.5), $\Delta+2.1$	53.8 (± 0.2), $\Delta+1.7$	26.9 (± 0.3), $\Delta+3.5$	36.5 (± 0.5), $\Delta-9.3$
W/O JOINT JITTER	62.1 (± 1.1), $\Delta+0.1$	51.4 (± 0.7), $\Delta+1.7$	52.9 (± 1.1), $\Delta+0.8$	24.1 (± 0.9), $\Delta+0.7$	43.4 (± 0.9), $\Delta-2.4$
W/O SLOW DOWN	61.9 (± 0.4), $\Delta-0.1$	51.5 (± 1.3), $\Delta+1.8$	53.3 (± 0.5), $\Delta+1.2$	25.6 (± 1.5), $\Delta+2.2$	43.3 (± 0.6), $\Delta-2.5$
W/O SPEED UP	60.9 (± 0.5), $\Delta-1.1$	50.2 (± 0.6), $\Delta+0.5$	52.2 (± 0.4), $\Delta+0.1$	23.5 (± 0.7), $\Delta+0.1$	44.1 (± 1.5), $\Delta-1.7$
W/O FRAME DROPOUT	60.6 (± 0.6), $\Delta-1.4$	49.6 (± 0.6), $\Delta-0.1$	51.9 (± 1.4), $\Delta-0.2$	23.7 (± 0.8), $\Delta+0.3$	44.9 (± 0.7), $\Delta-0.9$
W/O JOINT DROPOUT	68.8 (± 1.4), $\Delta+6.8$	56.7 (± 2.6), $\Delta+7.0$	51.8 (± 1.1), $\Delta-0.3$	25.3 (± 0.8), $\Delta+1.9$	42.1 (± 0.9), $\Delta-3.7$
W/ RANDOM ROTATION	63.2 (± 0.6), $\Delta+1.2$	52.7 (± 1.2), $\Delta+3.0$	53.9 (± 0.8), $\Delta+1.8$	24.9 (± 0.7), $\Delta+1.5$	48.1 (± 0.3), $\Delta+2.3$
SPATIAL ONLY	62.5 (± 0.9), $\Delta+0.5$	51.5 (± 1.7), $\Delta+1.8$	54.2 (± 2.0), $\Delta+2.1$	25.2 (± 0.8), $\Delta+1.8$	42.0 (± 0.9), $\Delta-3.8$
TEMPORAL ONLY	71.6 (± 0.2), $\Delta+9.6$	58.8 (± 3.4), $\Delta+9.1$	52.8 (± 0.3), $\Delta+0.7$	28.4 (± 2.0), $\Delta+5.0$	23.0 (± 0.1), $\Delta-22.8$

show that using a normalised frame-rate across all datasets of 30 fps does not improve performance. Replacing the ST-GCN backbone with an MLP decreases the performance across most tasks, most notably with a drop of 16.5% in activity classification.

The results confirm that task performance improves with the number of downstream training samples. Interestingly, we observed a 6.6% decrease in accuracy when pre-training only with NTU-60. This shows the importance of high diversity in pre-training data. However, removing a single pre-training dataset does not have a sizeable effect on downstream performance. Likewise, removing individual augmentations does not have a particular effect, except for random scaling (e.g. the performance drops by 9.3% for activity classification) and joint dropout (e.g. the MAE on skeleton reconstruction increases by 6.8mm for NTU-60 and 7mm for HDM05). However, when we remove *all* spatial augmentations or *all* temporal augmentations simultaneously, we observe a drop in performance across most tasks (e.g. with only temporal augmentations, the MAE on skeleton reconstruction increases by 9.6mm on NTU-60 and decreases by 22.8% in accuracy for activity classification).

5 Conclusion and Further Work

In this work, we showed that both spatial and temporal augmentations are needed to train skeleton representations. However, the specific type of augmentation within each category is less critical. Our results indicate that more data can help the model learn generic representations, even if the dataset were collected for different purposes. Finally, we demonstrated that a graph-based encoder (ST-GCN [23]) produces higher quality representations.

In the future, we would like to explore how other modalities, such as audio, can further enrich skeleton representations.

Acknowledgments

We would like to thank Zak Aldeneh, Jason Ramapuram, Arno Blass, Xavier Suau, Federico Danieli, Andy Keller, Borja Rodríguez, Katherine Metcalf, Nick Apostoloff, Russ Webb, Josh Susskind, and the anonymous reviewers for their helpful and knowledgeable feedback on this article. Additionally, we are grateful to Jason Ramapuram and Dan Busbridge for their excellent scalable implementation of SimCLR.

References

- [1] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *CoRR*, abs/2005.14165, 2020. URL <https://arxiv.org/abs/2005.14165>.
- [2] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *CoRR*, abs/2104.14294, 2021. URL <https://arxiv.org/abs/2104.14294>.
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 1597–1607. PMLR, 2020. URL <https://proceedings.mlr.press/v119/chen20j.html>.
- [4] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *CoRR*, abs/2104.02057, 2021. URL <https://arxiv.org/abs/2104.02057>.
- [5] Yuxiao Chen, Long Zhao, Jianbo Yuan, Yu Tian, Zhaoyang Xia, Shijie Geng, Ligong Han, and Dimitris N. Metaxas. Hierarchically self-supervised transformer for human skeleton representation learning, 2022.
- [6] Yuxiao Chen, Long Zhao, Jianbo Yuan, Yu Tian, Zhaoyang Xia, Shijie Geng, Ligong Han, and Dimitris N. Metaxas. Hierarchically self-supervised transformer for human skeleton representation learning, 2022. URL <https://arxiv.org/abs/2207.09644>.
- [7] Yuxin Chen, Ziqi Zhang, Chunfeng Yuan, Bing Li, Ying Deng, and Weiming Hu. Channel-wise topology refinement graph convolution for skeleton-based action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13359–13368, 2021.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*, abs/1810.04805, 2019.
- [9] Sean Ryan Fanello, Ilaria Gori, Giorgio Metta, and Francesca Odone. Keep it simple and sparse: Real-time action recognition. *Journal of Machine Learning Research*, 14(44):2617–2640, 2013. URL <http://jmlr.org/papers/v14/fanello13a.html>.
- [10] Ylva Ferstl and Rachel McDonnell. Iva: Investigating the use of recurrent motion modelling for speech gesture generation. In *IVA '18 Proceedings of the 18th International Conference on Intelligent Virtual Agents*, Nov 2018. URL <https://trinityspeechgesture.scss.tcd.ie>.
- [11] Adnan Khan, Sarah AlBarri, and Muhammad Arslan Manzoor. Contrastive self-supervised learning: A survey on different architectures. In *2022 2nd International Conference on Artificial Intelligence (ICAI)*, pp. 1–6, 2022. doi: 10.1109/ICAI55435.2022.9773725.
- [12] Gilwoo Lee, Zhiwei Deng, Shugao Ma, Takaaki Shiratori, Siddhartha Srinivasa, and Yaser Sheikh. Talking with hands 16.2m: A large-scale dataset of synchronized body-finger motion and audio for conversational motion analysis and synthesis. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 763–772, 2019. doi: 10.1109/ICCV.2019.00085.

- [13] Li Linguo, Wang Minsi, Ni Bingbing, Wang Hang, Yang Jiancheng, and Zhang Wenjun. 3d human action representation learning via cross-view consistency pursuit. In *CVPR*, 2021.
- [14] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+d 120: A large-scale benchmark for 3d human activity understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(10):2684–2701, 2020.
- [15] Rashmi M. and Ram Mohana Reddy Guddeti. Human identification system using 3d skeleton-based gait features and lstm model. *Journal of Visual Communication and Image Representation*, 82:103416, 2022. ISSN 1047-3203. doi: <https://doi.org/10.1016/j.jvcir.2021.103416>. URL <https://www.sciencedirect.com/science/article/pii/S1047320321002807>.
- [16] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision*, pp. 5442–5451, October 2019.
- [17] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber. Documentation mocap database hdm05. Technical Report CG-2007-2, Universität Bonn, June 2007.
- [18] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2018. URL <https://arxiv.org/abs/1807.03748>.
- [19] Pierre Plantard, Edouard Auvinet, Anne-Sophie Le Pierres, and Franck Multon. Pose estimation with a kinect for ergonomic studies: Evaluation of the accuracy using a virtual mannequin. *Sensors*, 15(1):1785–1803, 2015. doi: 10.3390/s150101785.
- [20] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1010–1019, 2016.
- [21] Yukun Su, Guosheng Lin, and Qingyao Wu. Self-supervised 3d skeleton action representation learning with motion consistency and continuity. In *International Conference on Computer Vision*, pp. 13308–13318. IEEE/CVF, 2021. doi: 10.1109/ICCV48922.2021.01308.
- [22] Luyu Wang, Pauline Luc, Yan Wu, Adrià Recasens, Lucas Smaira, Andrew Brock, Andrew Jaegle, Jean-Baptiste Alayrac, Sander Dieleman, Joao Carreira, and Aäron van den Oord. Towards learning universal audio representations. In *International Conference on Acoustics, Speech and Signal Processing*, pp. 4593–4597. IEEE, 2022. doi: 10.1109/ICASSP43922.2022.9746790.
- [23] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018.
- [24] Yiwen Zhan, Yuchen Chen, Pengfei Ren, Haifeng Sun, Jingyu Wang, Qi Qi, and Jianxin Liao. Spatial temporal enhanced contrastive and pretext learning for skeleton-based action representation. In Vineeth N. Balasubramanian and Ivor Tsang (eds.), *Proceedings of The 13th Asian Conference on Machine Learning*, volume 157 of *Proceedings of Machine Learning Research*, pp. 534–547. PMLR, 17–19 Nov 2021.

A Augmentations

Our version of SimCLR for skeleton representations makes use of the following augmentations:

- Axis mirroring:** Multiply by -1 all coordinates in the x, y or z axis.
- Random scale:** Randomly scale the dimension of a joint for all frames in the input.
- Joint Jitter:** Add gaussian noise (with mean zero and variance 0.02) to each joint and frame.
- Slow down:** Linearly interpolate between two frames and select random continuous subset.
- Speed up:** Skip 1 or 2 frames when preparing input representations. Pad with zeros.
- Frame Dropout:** Completely zeroing a frame on the input.
- Joint Dropout:** Randomly zero some joints from the input.
- Random rotation:** Randomly rotate the skeleton in XYZ for all frames in the input.

Visualisations of the augmentations can be seen in figure 2. The augmentations are applied to the input with the probabilities and strengths specified in table 3. They are always applied in the order specified above.

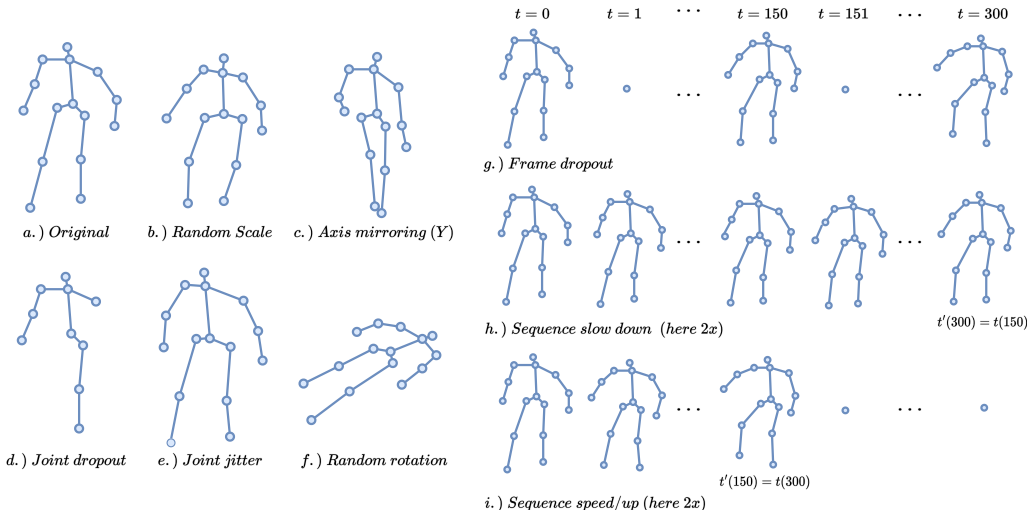


Figure 2: The augmentations used to train the SSL model with SimCLR. We use spatial (fig. b-g) and temporal (fig. h-i) augmentations.

B Training hyper-parameters

Table 3: Augmentations hyper-parameters for pre-training with probabilities and strengths

Augmentation	Probability	Strength
Axis mirroring	0.5	-
Random Scale	0.7	$\mu: 0, \sigma^2: 0.02$
Joint Jitter	0.5	$\mu: 0, \sigma^2: 0.02$
Slow down	0.5	0.5x-0.9x
Speed up	0.5	1.2-2x
Frame Dropout	0.5	0.5
Joint Dropout	0.5	0.5
Random rotation	0.6	$\pm\pi$

Table 4: Hyper-parameter used in pretraining and downstream tasks

	Pretraining	Downstream Task
Batch Size	8096	1024
Training Epochs	500	200
Optimizer	Lars Momentum	Adam
Learning Rate	0.01	0.01
LR Warmup	10	-
LR Update Schedule	cosine	-
Backbone Feature Size	128	-
Head Latent Size	256	-
Head Output Size	128	-
Head Output Size	128	-

Table 5: Network Parameter used in pretraining and downstream tasks

	Trainable Parameters
SimCLR	
Encoder (ST-GCN)	0.968458M
Encoder (MLP)	1.351168M
Frozen	
Activity Classification	0.00774M
Motion Prediction Frozen	2.905182M
Skeleton Reconstruction Frozen	2.905182M
Fine-Tune	
Activity Classification	0.843206M
Motion Prediction Frozen	3.740648M
Skeleton Reconstruction Frozen	3.740648M