

RESEARCH

Open Access



Automated freezing of gait assessment with marker-based motion capture and multi-stage spatial-temporal graph convolutional neural networks

Benjamin Filtjens^{1,2*} , Pieter Ginis³, Alice Nieuwboer³, Peter Slaets² and Bart Vanrumste¹

Abstract

Background: Freezing of gait (FOG) is a common and debilitating gait impairment in Parkinson's disease. Further insight into this phenomenon is hampered by the difficulty to objectively assess FOG. To meet this clinical need, this paper proposes an automated motion-capture-based FOG assessment method driven by a novel deep neural network.

Methods: Automated FOG assessment can be formulated as an action segmentation problem, where temporal models are tasked to recognize and temporally localize the FOG segments in untrimmed motion capture trials. This paper takes a closer look at the performance of state-of-the-art action segmentation models when tasked to automatically assess FOG. Furthermore, a novel deep neural network architecture is proposed that aims to better capture the spatial and temporal dependencies than the state-of-the-art baselines. The proposed network, termed multi-stage spatial-temporal graph convolutional network (MS-GCN), combines the spatial-temporal graph convolutional network (ST-GCN) and the multi-stage temporal convolutional network (MS-TCN). The ST-GCN captures the hierarchical spatial-temporal motion among the joints inherent to motion capture, while the multi-stage component reduces over-segmentation errors by refining the predictions over multiple stages. The proposed model was validated on a dataset of fourteen freezers, fourteen non-freezers, and fourteen healthy control subjects.

Results: The experiments indicate that the proposed model outperforms four state-of-the-art baselines. Moreover, FOG outcomes derived from MS-GCN predictions had an excellent ($r = 0.93$ [0.87, 0.97]) and moderately strong ($r = 0.75$ [0.55, 0.87]) linear relationship with FOG outcomes derived from manual annotations.

Conclusions: The proposed MS-GCN may provide an automated and objective alternative to labor-intensive clinician-based FOG assessment. Future work is now possible that aims to assess the generalization of MS-GCN to a larger and more varied verification cohort.

Keywords: Temporal convolutional neural networks, Graph convolutional neural networks, Freezing of gait, Parkinson's disease, MS-GCN

Background

Freezing of gait (FOG) is a common and debilitating gait impairment of Parkinson's disease (PD). Up to 80% of people with Parkinson's disease (PwPD) may develop FOG during the course of the disease [1, 2]. FOG leads

*Correspondence: benjamin.filtjens@kuleuven.be

² Department of Mechanical Engineering, Intelligent Mobile Platforms Research Group, KU Leuven, Andreas Vesaliusstraat 13, 3000 Leuven, Belgium
Full list of author information is available at the end of the article



© The Author(s) 2022. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

to sudden blocks in walking and is clinically defined as a “brief, episodic absence or marked reduction of forward progression of the feet despite the intention to walk and reach a destination” [3]. The PwPD themselves describe freezing of gait as “the feeling that their feet are glued to the ground” [4]. Freezing episodes most frequently occur while traversing under environmental constraints, during emotional stress, during cognitive overload by means of dual-tasking, and when initiating gait [5, 6]. Though, turning hesitation was found to be the most frequent trigger of FOG [7, 8]. Subjects with FOG experience more anxiety [9], have a lower quality of life [10], and are at a much higher risk of falls [11–15].

Given the severe adverse effects associated with FOG, there is a large incentive to advance novel interventions for FOG [16]. Unfortunately, the pathophysiology of FOG is complex and the development of novel treatments is severely limited by the difficulty to objectively assess FOG [17]. Due to heightened levels of attention, it is difficult to elicit FOG in the gait laboratory or clinical setting [4, 6]. Therefore, health professionals relied on subjects’ answers to subjective self-assessment questionnaires [18, 19], which may be insufficiently reliable to detect FOG severity [20]. Visual analysis of regular RGB videos has been put forward as the gold standard for rating FOG severity [20, 21]. However, the visual analysis relies on labor-intensive manual annotation by a trained clinical expert. As a result, there is a clear need for an automated and objective approach to assess FOG.

The percentage time spent frozen (%TF), defined as the cumulative duration of all FOG episodes divided by the total duration of the walking task, and the number of FOG episodes (#FOG) have been put forward as reliable outcome measures to objectively assess FOG [22]. An accurate segmentation in-time of the FOG episodes, with minimal over-segmentation errors, is required to robustly determine both outcome measures.

Several methods have been proposed for automated FOG assessment based on motion capture (MoCap) data. MoCap encodes human movement as a time series of human joint locations and orientations or their higher-order representations and is typically performed with optical or inertial measurement systems. Prior work has tackled automated FOG assessment as an action recognition problem and used a sliding-window scheme to segment a MoCap sequence into fixed partitions [23–36]. For all the samples within a partition, a single label is then predicted with methods ranging from simple thresholding methods [23, 26] to high-level temporal models driven by deep learning [27, 30, 32, 33, 36]. However, the samples within a pre-defined partition may not always share the same label. Therefore, a data-dependent heuristic is imposed to force all samples to take a

single label, most commonly by majority voting [33, 36]. Moreover, a second data-dependent heuristic is needed to define the duration of the sliding-window, which is a trade-off between expressivity, i.e., the ability to capture long-term temporal patterns, and sensitivity, i.e., the ability to identify short-duration FOG episodes. Such manually defined heuristics are unlikely to generalize across study protocols.

This study proposes to reformulate the problem of FOG annotation as an action segmentation problem. Action segmentation approaches overcome the need for manually defined heuristics by generating a prediction for each sample within a long untrimmed MoCap sequence. Several methods have been proposed to tackle action segmentation. Similar to FOG assessment, earlier studies made use of sliding-window classifiers [37, 38], which do not capture long-term temporal patterns [39]. Other approaches use temporal models such as hidden Markov models [40, 41] and recurrent neural networks [42, 43]. The state-of-the-art methods tend to use temporal convolutional neural networks (TCN), which have been shown to outperform recurrent methods [39, 44]. Dilation is frequently added to capture long-term temporal patterns by expanding the temporal receptive field of the TCN models [45]. In multi-stage temporal convolutional network (MS-TCN), the authors show that multiple stages of temporal dilated convolutions significantly reduce over-segmentation errors [46]. These action segmentation methods have historically been validated on video-based datasets [47, 48] and thus employ video-based features [49]. The human skeleton structure that is inherent to MoCap has thus not been exploited by prior work in action segmentation.

To model the structured information among the markers, this paper uses the spatial-temporal graph convolutional neural network (ST-GCN) [50] as the first stage of an MS-TCN network. ST-GCN applies spatial graph convolutions on the human skeleton graph at each time step and applies dilated temporal convolutions on the temporal edges that connect the same markers across consecutive time steps. The proposed model, termed multi-stage spatial-temporal graph convolutional neural network (MS-GCN), thus extends MS-TCN to skeleton-based data for enhanced action segmentation within MoCap sequences.

The MS-GCN was tasked to recognize and localize FOG segments in a MoCap sequence. The predicted segments were quantitatively and qualitatively assessed versus the agreed-upon annotations by two clinical-expert raters. From the predicted segments, two clinically relevant FOG outcomes, the %TF and #FOG, were computed and statistically validated. To the best of our knowledge, the proposed MS-GCN is a novel neural

network architecture for skeleton-based action segmentation in general and FOG segmentation in particular. The benefit of MS-GCN for FOG assessment is four-fold: (1) It exploits ST-GCN to model the structured information inherent to MoCap. (2) It allows modeling of long-term temporal context to capture the complex dynamics that precede and succeed FOG. (3) It can operate on high temporal resolutions for fine-grained FOG segmentation with precise temporal boundaries. (4) To accomplish (2) and (3) with minimal over-segmentation errors, MS-GCN utilizes multiple stages of refinements.

Methods

Dataset

Two existing MoCap datasets [51, 52] were included for analysis. The first dataset [51], includes forty-two subjects. Twenty-eight of the subjects were diagnosed with PD by a movement disorders neurologist. Fourteen of the PwPD were classified as freezers based on the first question of the New Freezing of Gait Questionnaire (NFOG-Q): “Did you experience “freezing episodes” over the past month?” [19]. The remaining fourteen subjects were age-matched healthy controls. The second dataset [52], includes seventeen PwPD and FOG, as classified by the NFOG-Q. The subjects underwent a gait assessment at baseline and after twelve months follow-up. Five subjects only underwent baseline assessment and four subjects dropped out during the follow-up. The clinical characteristics are presented in Table 1.

Protocol

Both datasets were recorded with a Vicon 3D motion analysis system recording at a sampling frequency of 100 Hz. Retro-reflective markers were placed on anatomical landmarks according to the full-body or lower-limb

plug-in-gait model [53, 54]. Both datasets featured a nearly identical standardized gait assessment protocol, where two retro-reflective markers placed 0.5 m from each other indicated where subjects either had to walk straight ahead, turn 360°left, or turn 360°right. For dataset 1, the subjects were additionally instructed to turn 180°left and turn 180°right. The experimental conditions were offered randomly and performed with or without a verbal cognitive dual-task [55, 56]. All gait assessments were conducted during the off-state of the subjects’ medication cycle, i.e., after an overnight withdrawal of their normal medication intake. The experimental conditions are visualized in Fig. 1.

For dataset 1, two clinical experts, blinded for NFOG-Q score, annotated all FOG episodes by visual inspection of the knee-angle data (flexion-extension) in combination with the MoCap 3D images. For dataset 2, the FOG episodes were annotated by one of the authors (BF) based on visual inspection of the MoCap 3D images. To ensure that the results were unbiased, the FOG trials of dataset 2 were used to enrich the training dataset and not for the evaluation of the model. For both datasets, the onset of FOG was determined at the heel strike event prior to delayed knee flexion. The termination of FOG was determined at the foot-off event that is succeeded by at least two consecutive movement cycles [51].

FOG segmentation

Marker-based optical MoCap describes the 3D movement of optical markers in time, where each marker represents the 3D coordinates of the corresponding anatomical landmark. The duration of a MoCap trial can vary substantially due to high inter-and intra-subject variability. The goal is to segment a FOG episode in time, given a variable-length MoCap trial. The MoCap trial can be represented as $X \in \mathbb{R}^{N \times T \times C_{in}}$, where N specifies the number of optical markers, T the number of samples, and C_{in} the feature dimension. Each MoCap trial X is associated with a ground truth label vector $Y_{exp}^{T \times l}$, where the label l represents the manual annotation of FOG and functional gait (FG) by the clinical experts. A deep neural network segments a FOG episode in time by learning a function $f : X \rightarrow Y$ that transforms a given input sequence $X = x_0, \dots, x_T$ into an output sequence $\hat{Y} = \hat{y}_0, \dots, \hat{y}_T$ that closely resembles the manual annotations Y_{exp} .

From the 3D marker coordinates, the marker displacement between two consecutive samples was computed as $X(n, t + 1, :) - X(n, t, :)$. The two markers on the femur and tibia, which were wand markers in dataset 1 and thus placed away from the primary axis, were excluded. The heel marker was excluded due to close proximity with the ankle marker. The reduced marker configuration consists

Table 1 Subject characteristics

	Dataset 1			Dataset 2
	Controls	Non-freezers	Freezers	Freezers
Age	65 ± 6.8	67 ± 7.4	69 ± 7.4	67 ± 9.3
PD duration		7.8 ± 4.8	9.0 ± 4.8	10 ± 6.3
MMSE [81]	29 ± 1.3	29 ± 1.2	28 ± 1.1	28 ± 1.3
UPDRS III [82]		34 ± 9.9	38 ± 14	39 ± 12
H&Y [83]		2.4 ± 0.3	2.5 ± 0.5	2.4 ± 0.5

The subject characteristics of the fourteen healthy control subjects (controls), fourteen PwPD and without FOG (non-freezers), fourteen PwPD and FOG (freezers) of dataset 1, and seventeen PwPD and FOG (freezers) of dataset 2 at the baseline assessment. All characteristics are given in terms of mean ± standard deviation. For dataset 1, the characteristics were measured during the ON-phase of the medication cycle, while for dataset 2 the characteristics were measured while OFF medication

of nine optical markers: the marker in the middle of the left and right posterior superior iliac spine, the markers on the left and right anterior superior iliac spine, the markers on the left and right lateral femoral condyle, the markers on the left and right lateral malleolus, and the markers on the left and right second metatarsal head. As a result, an input sequence $X \in \mathbb{R}^{N \times T \times C_{in}}$ is composed of nine optical markers (N), variable duration (T), and with the feature dimension (C_{in}) composed of the 3D displacement of each marker.

MS-GCN

The proposed multi-stage graph convolutional neural network (MS-GCN), generalizes the multi-stage temporal convolutional neural network (MS-TCN) [46] to graph-based data. A visual overview of the model architecture is provided in Fig. 2.

Formally, MS-GCN features a prediction generation stage of several ST-GCN blocks, which generates an initial prediction $Y \in \mathbb{R}^{T \times l}$. The first layer of the prediction generation stage is a batch normalization (BN) layer that normalizes the inputs and accelerates training [57]. The normalized input is passed through a 1×1 convolutional layer that adjusts the input dimension C_{in} to the number of filters C in the network, formalized as:

$$f_{adj} = W_1 * f_{in} + b, \tag{1}$$

where $f_{adj} \in \mathbb{R}^{T \times N \times C}$ is the adjusted feature map, $f_{in} \in \mathbb{R}^{T \times N \times C_{in}}$ the input MoCap sequence, $b \in \mathbb{R}^C$ the bias term, $*$ the convolution operator, $W_1 \in \mathbb{R}^{1 \times 1 \times C_{in} \times C}$ the weights of the 1×1 convolution filter with C_{in} input feature channels and C equal to the number of feature channels in the network.

The adjusted input is passed through several blocks of ST-GCN [50]. Each ST-GCN first applies a graph convolution, formalized as:

$$f_{gcn} = \sum_p A_p f_{adj} W_p M_p, \tag{2}$$

where $f_{adj} \in \mathbb{R}^{T \times N \times C}$ is the adjusted input feature map, $f_{gcn} \in \mathbb{R}^{T \times N \times C}$ the output feature map of the spatial graph convolution, and W_p the $1 \times 1 \times C \times C$ weight matrix. The matrix $A_p \in \{0, 1\}^{N \times N}$ is the adjacency matrix, which represents the spatial connection between the joints. The graph is partitioned into three subsets based on the spatial partitioning strategy [50]. The matrix M_p is a learnable $N \times N$ attention mask that indicates the importance of each node and its spatial partitions.

Next, after passing through a BN layer and ReLU non-linearity, the ST-GCN block performs a dilated temporal convolution [45]. The dilated temporal convolution

is, in turn, passed through a BN layer and ReLU non-linearity, and lastly, a residual connection is added between the activation map and the input. This process is formalized as:

$$f_{out} = \delta(BN(W * _d f_{gcn} + b)) + f_{adj}, \tag{3}$$

where $f_{out} \in \mathbb{R}^{T \times N \times C}$ is the output feature map, $b \in \mathbb{R}^C$ the bias term, $*_d$ the dilated convolution operator, $W \in \mathbb{R}^{k \times 1 \times C \times C}$ the weights of the dilated convolution filter with kernel size k . The output feature map is passed through a spatial pooling layer that aggregates the spatial features among the N joints.

Lastly, the aggregated feature map is passed through a 1×1 convolution and a softmax activation function to get the probabilities for the l output classes for each sample in-time, formalized as:

$$\hat{y}_t = \zeta(W_1 * f_{out} + b), \tag{4}$$

where \hat{y}_t are the class probabilities at time t , f_{out} the output of the pooled ST-GCN block at time t , $b \in \mathbb{R}^l$ the bias term, $*$ the convolution operator, ζ the softmax function, $W_1 \in \mathbb{R}^{1 \times C \times l}$ the weights of the 1×1 convolution filter with C input channels and l output classes.

Next, the initial prediction is passed through one or more refinement stages. The first layer of the refinement stage is a 1×1 convolutional layer that adjusts the input dimension l to the number of filters C in the network, formalized as:

$$f_{adj} = W_1 * f_{in} + b, \tag{5}$$

where $f_{adj} \in \mathbb{R}^{T \times C}$ is the adjusted feature map, $f_{in} \in \mathbb{R}^{T \times l}$ the softmax probabilities of the previous stage, $b \in \mathbb{R}^C$ the bias term, $*$ the convolution operator, $W_1 \in \mathbb{R}^{1 \times l \times C}$ the weights of the 1×1 convolution filter with l input feature channels and C equal to the number of feature channels in the network.

The adjusted input is passed through ten blocks of TCN. Each TCN block applies a dilated temporal convolution [45], BN, ReLU non-linear activation, and a residual connection between the activation map and the input. Formally, this process is defined as:

$$f_{out} = \delta(BN(W * _d f_{adj} + b)) + f_{adj}, \tag{6}$$

where $f_{out} \in \mathbb{R}^{T \times C}$ is the output feature map, $b \in \mathbb{R}^C$ the bias term, $*_d$ the dilated convolution operator, $W \in \mathbb{R}^{k \times C \times C}$ the weights of the dilated convolution filter with kernel size k , and δ the ReLU function.

Lastly, the feature map is passed through a 1×1 convolution and a softmax activation function to get the probabilities for the l output classes for each sample in-time, formalized as:

$$\hat{y}_t = \zeta(W_1 * f_{out} + b), \tag{7}$$

where \hat{y}_t are the class probabilities at time t , f_{out} the output of the last TCN block at time t , $b \in \mathbb{R}^l$ the bias term, $*$ the convolution operator, ζ the softmax function, $W_1 \in \mathbb{R}^{1 \times C \times l}$ the weights of the 1×1 convolution filter with C input channels and l output classes.

Model comparison

To put the MS-GCN results into context, four strong DL baselines were included. Specifically, the state-of-the-art in skeleton-based action recognition, spatial-temporal graph convolutional network (ST-GCN) [50]. The state-of-the-art in action segmentation, multi-stage temporal convolutional neural network (MS-TCN) [46]. Two commonly used sequence to sequence models in human movement analysis [58, 59], a bidirectional long short term memory-based network (LSTM) [60], and a temporal convolutional neural network-based network (TCN) [39].

Implementation details

To train the models, this paper used the same loss as MS-TCN which utilized a combination of a classification loss (cross-entropy) and smoothing loss (mean squared error) for each stage. The combined loss is defined as:

$$L = L_{cls} + \lambda L_{T-MSE}, \tag{8}$$

where the hyperparameter λ controls the contribution of each loss function. The classification loss L_{cls} is the cross entropy loss:

$$L_{cls} = \frac{1}{T} \sum_t -y_{t,l} \log(\hat{y}_{t,l}). \tag{9}$$

The smoothing loss L_{T-MSE} is a truncated mean squared error of the sample-wise log-probabilities:

$$L_{T-MSE} = \frac{1}{TC} \sum_{t,c} \tilde{\Delta}_{t,c}^2 \tag{10}$$

$$\tilde{\Delta}_t = \begin{cases} \Delta_{t,c} & : \Delta_{t,c} \leq \tau, \\ \tau & : \text{otherwise,} \end{cases}$$

$$\Delta_{t,l} = |\log(\hat{y}_{t,l}) - \log(\hat{y}_{t-1,l})|,$$

In each loss function, T are the number of samples and $\hat{y}_{t,l}$ is the probability of FOG or FG at sample t . To train the entire network, the sum of the losses over all stages is minimized:

$$L = \sum_s L_s \tag{11}$$

To allow an unbiased comparison, the model and optimizer hyperparameters were selected according to

MS-TCN [46]. Specifically, the multi-stage models had 1 prediction generation stage and 4 refinement stages. Each stage had 10 layers of 64 filters that applied graph and/or dilated temporal convolutions with kernel size 3 and ReLU activations. The temporal convolutions were acausal, i.e., they could take into account both past and future input features, with a dilation factor that doubled at each layer, i.e., 1, 2, 4, ..., 512. The single-stage models, i.e., ST-GCN and TCN, used the same configuration but without refinement stages. The Bi-LSTM used a configuration that is conventional in human movement analysis, with two forward LSTM layers and two backward LSTM layers, each with 64 cells [59, 61]. For the loss function, τ was set to 4 and λ was set to 0.15. All experiments used the Adam optimizer [62] with a learning rate of 0.0005. All models were trained for 100 epochs with a batch size of 16.

For the temporal models, i.e., LSTM, TCN, and MS-TCN, the input is reshaped into their accepted formats. Specifically, the data is shaped into $T \times C_{in} * N$, i.e., the spatial feature dimension N is thus collapsed.

The LSTM was additionally evaluated as an action recognition model. For this evaluation, the MoCap sequences were partitioned into two-second windows and majority voting was used to force all samples to take a single label. These settings are commonly used in FOG recognition [33, 36]. The last hidden LSTM state, which constitutes a compressed representation of the entire sequence, was fed to a feed-forward network to generate a single label for the sequence. To localize the FOG episodes during evaluation, predictions for each sample were made by sliding the two-second partition in steps of one. This setting enables an objective comparison with the proposed action segmentation approaches as predictions are made at a temporal frequency of 100 Hz for both action detection schemes.


Evaluation

For dataset 1, FOG was provoked for ten of the fourteen freezers during the test period, with seven subjects freezing within the visibility of the MoCap system. For dataset 2, eight of the seventeen freezers froze within the visibility of the MoCap system. The training dataset consists of the FOG and non-FOG trials of the seven subjects who froze in front of the MoCap system of dataset 1, enriched with the FOG trials of the eight subjects who froze in front of the MoCap system of dataset 2. Only the FOG trials of dataset 2 were considered to balance out the number of FOG and FG trials. Only the subjects of dataset 1 were considered for evaluation, as motivated in the procedure. Detailed dataset characteristics are provided in Table 2.

Table 2 Dataset characteristics

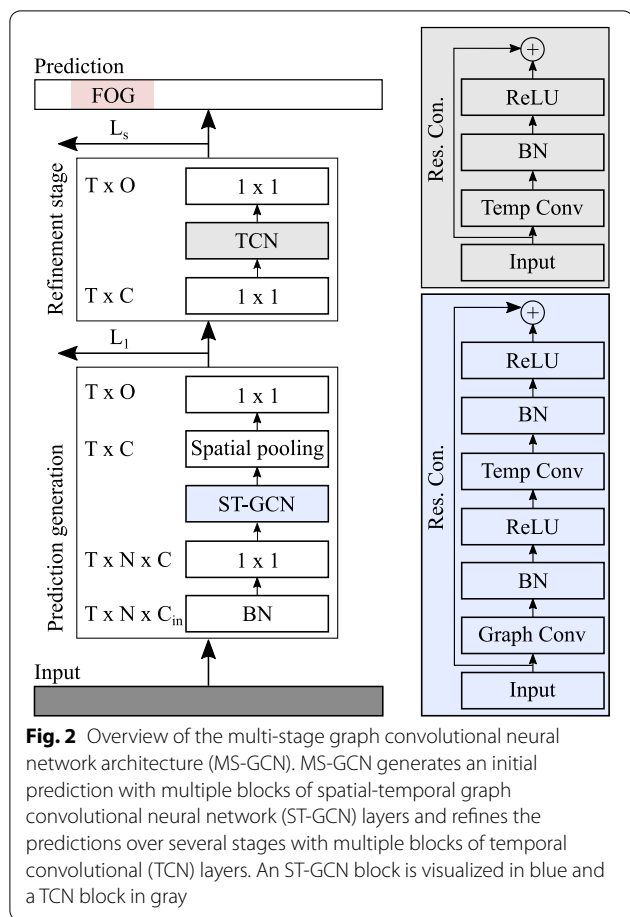
ID	#Trials	#FOG trials	#FOG	%TF	Total duration	Avg duration (\pm SD)
S1	27	3	9	33.9	1.05	0.35 (\pm 0.28)
S2	22	11	13	12.3	4.01	0.36 (\pm 0.12)
S3	27	4	5	6.89	0.49	0.12 (\pm 0.02)
S4	21	9	18	36.7	2.45	0.27 (\pm 0.19)
S5	24	1	3	36.1	0.29	0.29
S6	7	5	7	14.4	0.85	0.17 (\pm 0.08)
S7	31	1	1	20.1	0.08	0.08
D2	68	68	134	28.4	24.4	0.36 (\pm 0.18)

Overview of the number of motion capture trials (#Trials), number of FOG trials (#FOG trials), number of FOG episodes (#FOG), percentage time spent frozen (%TF), total duration of the FOG trials (in minutes), and average duration of the FOG trials (\pm standard deviation (SD)) (in minutes). For dataset 1, the characteristics are given per subject. For dataset 2 (D2), which was only used to enrich the training dataset and not for model evaluation, a single summary is provided



The evaluation dataset was partitioned according to a leave-one-subject-out cross-validation approach. This cross-validation approach repeatedly splits the data according to the number of subjects in the dataset. One subject is selected for evaluation, while the other subjects are used to train the model. This procedure is repeated until all subjects have been used for evaluation. This approach mirrors the clinically relevant scenario of FOG assessment in newly recruited subjects [63], where the model is tasked to assess FOG in unseen subjects.

From a machine learning perspective, action segmentation papers tend to use sample-wise metrics, such as accuracy, precision, and recall. However, sample-wise metrics do not heavily penalize over-segmentation errors. As a result, methods with significant qualitative differences, as was observed between the single-stage ST-GCN and MS-GCN, can still achieve similar performance on the sample-wise metrics. In 2016 Lea et al. [39] proposed a segment-wise F1-score to address those drawbacks. To compute the segment-wise F1-score, action segments



are first classified as true positive (TP), false positive (FP), or false negative (FN) by comparing the intersection over union (IoU) to a pre-determined threshold, as visualized in Fig. 3. The segment-wise F1-score has several advantages for FOG segmentation. (1) It penalizes over and under-segmentation errors, which would result in an inaccurate #FOG severity outcome. (2) It allows for

minor temporal shifts, which may have been caused by annotator variability and do not impact the FOG severity outcomes. (3) It is not impacted by the variability in FOG duration, since it is dependent on the number of FOG episodes and not on their duration.


This paper also reports a sample-wise metric. More specifically, the sample-wise Matthews correlation coefficient (MCC), defined as [64]:

$$MCC = \frac{TP * TN - FP * FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (12)$$

A perfect MCC score is equal to one hundred, whereas minus one hundred is the worst value. An MCC score of zero is reached when the model always picks the majority class. The MCC can thus be considered a balanced measure, i.e., correct FOG and FG classification are of equal importance. The discrepancy between sample-wise MCC and the segment-wise F1 score allows assessment of potential over and under-segmentation errors. Conclusions were based on the segment-wise F1-score at high IoU overlap.

For the model validation, the entirety of dataset 1 was used, i.e., MoCap trials without FOG and MoCap trials with FOG, of the seven subjects who froze during the protocol. The machine learning metrics were used to evaluate MS-GCN with respect to the four strong baselines. While a high number of trials without FOG can inflate the metrics, correct classification of FOG and non FOG segments are, however, of equal importance for assessing FOG severity and thus also for assessing the performance of a machine learning model. To further assess potential false-positive scoring, an additional analysis was performed on trials without FOG of the healthy controls, non-freezers, and freezers that did not freeze during the protocol.

From a clinical perspective, FOG severity is typically assessed in terms of percentage time-frozen (%TF) and



number of detected FOG episodes (#FOG) [22]. The %TF quantifies the duration of FOG relative to the trial duration, and is defined as:

$$\%TF = \left(\frac{1}{T} \sum_t y_{FOG} \right) * 100, \quad (13)$$

where T are the number of samples in a MoCap trial and y_{FOG} are the FOG samples predicted by the model or the samples annotated by the clinical experts. To evaluate the goodness of fit, the linear relationship between observations by the clinical experts and the model predictions was assessed. The strength of the linear relationship was classified according to [65]: ≥ 0.8 : strong, $0.6-0.8$: moderately strong, $0.3-0.5$: fair, and < 0.3 : poor. The correlation describes the linear relationship between the experts' observations and the model predictions but ignores bias in predictions. Therefore, a linear regression analysis was performed to evaluate whether the linear association between the expert annotations and model predictions was statistically significant. The significance level for all tests was set at 0.05. For the FOG severity statistical analysis, only the trials with FOG were considered, as trials without FOG would inflate the reliability scores.

Results

Model comparison

All models were trained using a leave-one-subject-out cross-validation approach. The metrics were summarized in terms of the mean \pm standard deviation (SD) of the seven subjects that froze during the protocol, where the SD aims to capture the variability across different subjects. According to the results shown in Table 3, the ST-GCN-based models outperform the TCN and LSTM-based models on the MCC metric. This result confirms the notion that explicitly modeling the spatial hierarchy within the skeleton-based data results in a better representation [50]. Moreover, the multi-stage refinements improve the F1 score at all evaluated overlapping thresholds, the metric that penalizes over-segmentation errors, while the sample-wise MCC remains mostly consistent across stages. This result confirms the notion that multi-stage refinements can reduce the number of over-segmentation errors and improve neural network models for fine-grained activity segmentation [46]. Additionally, the results suggest that the sliding window scheme is ill-suited for fine-grained FOG annotation at high temporal frequencies.

MS-GCN detailed results

This section provides an in-depth analysis of the performance of the MS-GCN model. According to the results shown in Table 4, the model correctly detects 52 of 56


FOG episodes. A detection was considered as a TP if at least one sample overlapped with the ground-truth episode. Thus, without imposing a constraint on how much the predicted segment should overlap with the ground-truth segment, as is the case when computing the segment-wise F1 score. The model proved robust, with only six episodes incorrectly detected in a trial that the experts did not label as FOG. In terms of the clinical metrics, the model provides an accurate assessment of #FOG and %TF for five of the seven subjects. For S2 the model overestimates FOG severity, while for S3 the model underestimates FOG severity.

One FOG segmentation trial for each of the seven subjects is visualized in Fig. 4. The sample-wise MCC and segment-wise F1@50 for each trial are included for comparison. A near-perfect FOG segmentation can be observed for the trials of S1, S4, S5, and S7. For the two chosen trials of S3 and S6, the model did not detect two of the sub-0.5-second FOG episodes. For S2, it is evident that the model overestimates the number of FOG episodes.

A quantitative assessment of the MS-GCN predictions for the fourteen healthy control subjects (controls), fourteen non-freezers (non-freezers), and the seven freezers that did not freeze during the protocol (freezers-) further demonstrates the robustness of the MS-GCN. The results are summarized in Table 5. According to Table 5, no false-positive FOG segments were predicted.

Automated FOG assessment: statistical analysis

The clinical experts observed at least one FOG episode in 35 MoCap trials of dataset 1. The number of detected FOG episodes (#FOG) per trial varied from 1 to 7 amounting to 56 FOG episodes, while the percentage time-frozen (%TF) varied from 4.2 to 75. For the %TF, the model predictions had a very strong linear relationship with the experts observations, with a correlation value [95% confidence interval (CI)] of $r = 0.93$ [0.87, 0.97]. For the #FOG, the model predictions had a moderately strong linear relationship with the experts' observations, with a correlation value [95% CI] of $r = 0.75$ [0.55, 0.87]. A linear regression analysis was performed to evaluate whether the linear association between the experts' annotations and model predictions was statistically significant. For the %TF, the intercept [95% CI] was -1.79 [$-6.8, 3.3$] and the slope [95% CI] was 0.96 [0.83, 1.1]. For the #FOG, the intercept [95% CI] was 0.36 [$-0.22, 0.94$] and the slope [95% CI] was 0.73 [0.52, 0.92]. Given that the 95 % CIs of the slopes exclude zero, the linear association between the model predictions and expert observations was statistically significant (at the 0.05 level) for both FOG severity outcomes. The linear relationship is visualized in Fig. 5.



Discussion

Existing approaches treat automatic FOG assessment as an action recognition task and employ a sliding-window scheme to localize the FOG segments within a MoCap sequence. Such approaches require manually defined heuristics that may not generalize across study protocols. For instance, the most common FOG recognition scheme uses two-second partitions with majority voting to force all labels within a partition to a single label [33, 36]. Yet, such settings would induce a bias on the ground-truth annotations as sub-second episodes would never be the majority label. For the present dataset, this bias would neglect all the FOG episodes of S3. While shorter partitions could overcome this issue, they would restrict the amount of temporal context exposed to the model.

To address these issues, this paper reformulated FOG assessment as an action segmentation task. Action segmentation frameworks overcome the need for fixed

partitioning by generating a prediction for each sample. Therefore, these frameworks rely only on the observations and their assumed model and not on manual heuristics that are unlikely to generalize across study protocols. As predictions vary at a high temporal frequency, action segmentation is inherently more challenging than recognition. To address this task, a novel neural network architecture, entitled MS-GCN, was proposed. MS-GCN extends MS-TCN [46], the state-of-the-art model in action segmentation, to graph-based input data that is inherent to MoCap.

MS-GCN was quantitatively compared with four strong deep learning baselines. The comparison confirmed the notions that: (1) the multi-stage refinements reduce over-segmentation errors, and (2) the graph convolutions give a better representation of skeleton-based data than regular temporal convolutions. As a result, MS-GCN showed state-of-the-art FOG segmentation

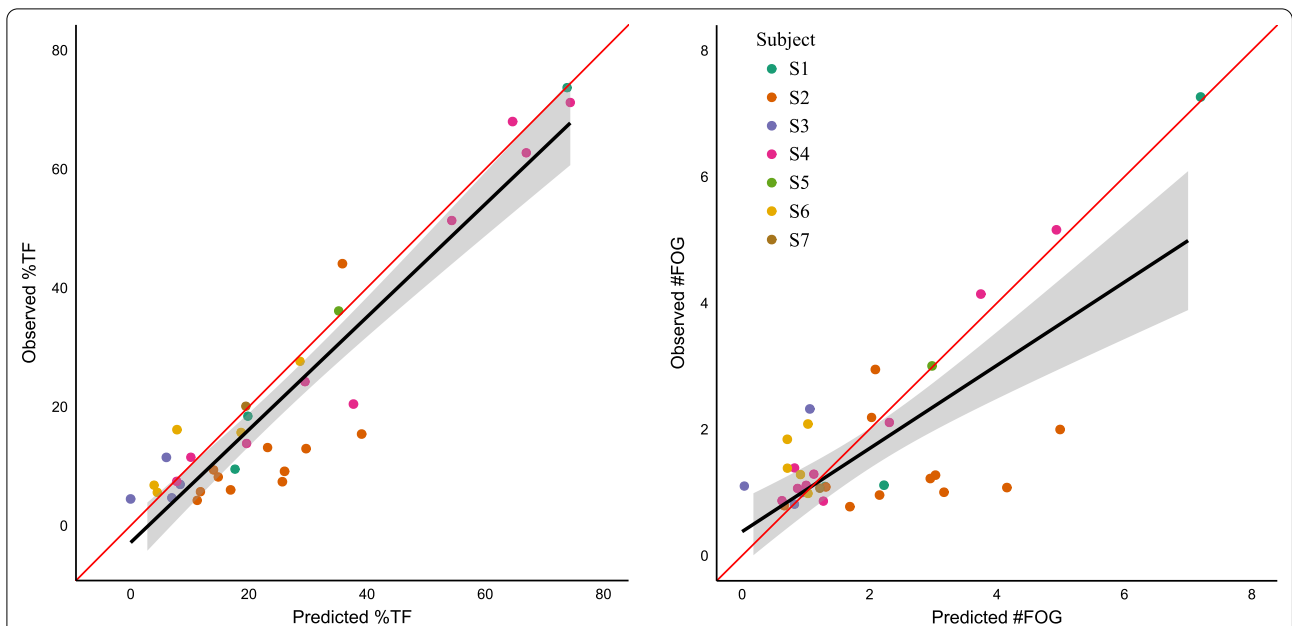


Fig. 5 Assessing the performance of the MS-GCN (6 stages) for automated FOG assessment. More specifically, the performance to measure the percentage time-frozen (%TF) (left) and the number of FOG episodes (#FOG) (right) during a standardized protocol. The ideal regression line with a slope of one and an intercept of zero is visualized in red. All results were derived from the test set, i.e., subjects that the model had never seen. Observe the overestimation of %TF and #FOG for S2

Table 3 Model comparison results

Model	F1@10	F1@25	F1@50	F1@75	MCC
Bi-LSTM†	25.9 ± 8.40	21.8 ± 9.03	15.0 ± 5.60	11.9 ± 6.26	62.4 ± 23.2
Bi-LSTM	63.7 ± 21.7	63.2 ± 22.0	50.8 ± 25.4	40.9 ± 28.4	78.8 ± 21.1
TCN	45.4 ± 16.8	42.7 ± 18.6	35.8 ± 14.8	27.0 ± 16.6	81.1 ± 12.9
ST-GCN	53.2 ± 21.2	51.5 ± 21.7	46.7 ± 22.5	37.6 ± 26.6	83.0 ± 11.5
MS-TCN	68.2 ± 29.4	66.8 ± 29.3	60.2 ± 30.5	54.9 ± 33.1	77.3 ± 22.2
MS-GCN	77.8 ± 15.3	77.8 ± 15.3	74.2 ± 21.0	57.0 ± 30.1	82.7 ± 15.5

Overview of the FOG segmentation performance in terms of the segment-wise F1@50 and sample-wise MCC for MS-GCN and the four strong baselines. The †denotes the sliding window FOG detection scheme. The best score is denoted in bold. All results were derived from the test set, i.e., subjects that the model had never seen

Table 4 Detailed MS-GCN results

ID	F1@50	MCC	#TP	#FP	#FOG	%TF
S1	87.5	95.7	9 / 9	0 / 24	10 / 9	37.1 / 33.9
S2	31.6	60.3	13 / 13	3 / 11	24 / 13	22.6 / 12.3
S3	60.0	59.8	3 / 5	2 / 23	3 / 5	5.34 / 6.89
S4	71.1	87.2	18 / 18	0 / 12	18 / 18	40.6 / 36.7
S5	85.7	96.9	3 / 3	1 / 23	3 / 3	35.2 / 36.1
S6	83.3	80.7	5 / 7	0 / 2	5 / 7	12.7 / 14.4
S7	100	98.5	1 / 1	0 / 30	1 / 1	19.5 / 20.1
	74.2	82.7	52 / 56	6 / 125	64 / 56	24.7 / 22.9

Detailed overview of the FOG assessment performance of the proposed MS-GCN model for each subject. The fourth column depicts the number of true positive FOG detections (TP) with respect to the number of FOG episodes. The fifth column depicts the number of false-positive (FP) FOG detections with respect to the number of trials that did not contain FOG. The sixth and seventh columns depict the #FOG and %TF computed from the model annotated segmentations with respect to those computed from the expert annotated segmentations. All results were derived from the test set, i.e., subjects that the model had never seen

Table 5 MS-GCN robustness

Subjects	FP
Controls (k = 404)	0
Non-freezers (k = 423)	0
Freezers- (k = 195)	0

Overview of MS-GCN's robustness to false-positive FOG detections on the MoCap trials of the 14 healthy controls, 14 non-freezers, and 7 freezers that did not freeze in front of the cameras during the protocol. The letter k denotes the number of MoCap trials for each group

performance. Two common outcome measures to assess FOG, the %TF and #FOG [22], were computed and statistically assessed. MS-GCN showed a very strong ($r = 0.93$) and moderately strong ($r = 0.75$) linear relationship with the experts' observations for %TF and #FOG, respectively. For context, the intraclass correlation coefficient between independent assessors was reported to be 0.87 [66] and 0.73 [22] for %TF and 0.63 [22] for #FOG.

A benefit of MS-GCN is that it is not strictly limited to marker-based MoCap data. The MS-GCN architecture naturally extends to other graph-based input data, such as single- or multi-camera markerless pose estimation [67, 68], and FOG assessment protocols that employ multiple on-body sensors [24, 25]. Both technologies are receiving increased attention due to the potential to assess FOG not only in the lab but also in an at-home environment and thereby better capture daily-life FOG severity. Furthermore, up until now, deep learning-based gait assessment [58, 61, 69, 70] did not yet exploit the inherent graph-structured data. The established improvement in FOG assessment by this research might, therefore, signify further improvements in deep learning-based gait assessment in general.

Several limitations are present. The first and most prominent limitation is the lack of variety in the standardized FOG-provoking protocol. FOG is characterized by several apparent subtypes, such as turning and destination hesitation, and gait initiation [7]. While turning was found to be the most prominent [7, 8], it should still be established whether MS-GCN can generalize to other FOG subtypes under different FOG provoking protocols. For now, practitioners are advised to closely follow the experimental protocol used in this study when employing MS-GCN. The second limitation is the small sample size. While MS-GCN was evaluated based on the clinically relevant use-case scenario of FOG assessment in newly recruited subjects, the sample size of the dataset is relatively small compared to the deep learning literature. The third limitation is based on the observation that FOG assessment in the clinic and lab is prone to two shortcomings. (1) FOG can be challenging to elicit in the lab due to elevated levels of attention [4, 6], despite

providing adequate FOG provoking circumstances [51, 71]. (2) Research has demonstrated that FOG severity in the lab is not necessarily representative of FOG severity in daily life [4, 72]. Future work should therefore establish whether the proposed method can generalize to tackle automated FOG assessment with on-body sensors or markerless MoCap captured in less constrained environments. Fourth, due to the opaqueness inherent to deep learning, clinicians have historically distrusted DNNs [73]. However, prior case studies [74, 75], have demonstrated that interpretability techniques are able to visualize what features the model has learned [76–78], which can aid the clinician in determining whether the assessment was based on credible features.

Conclusions

FOG is a debilitating motor impairment of PD. Unfortunately, our understanding of this phenomenon is hampered by the difficulty of objectively assessing FOG. To tackle this problem, this paper proposed a novel deep neural network architecture. The proposed architecture, termed MS-GCN, was quantitatively validated versus the expert clinical opinion of two independent raters. In conclusion, it can be established that MS-GCN demonstrates state-of-the-art FOG assessment performance. Furthermore, future work is now possible that aims to assess the generalization of MS-GCN to other graph-based input data, such as markerless MoCap or multiple on-body sensor configurations, and to other FOG subtypes captured under less constrained protocols. Such work is important to increase our understanding of this debilitating phenomenon during everyday life.

Abbreviations

FOG: Freezing of gait; PD: Parkinson's Disease; PwPD: People with Parkinson's Disease; %TF: Percentage time spent frozen; #FOG: Number of FOG episodes; MoCap: Motion capture; TCN: Temporal convolutional neural network; MS-TCN: Multi-stage temporal convolutional neural network; GCN: Graph convolutional neural networks; ST-GCN: Spatial-temporal graph convolutional neural network; MS-GCN: Multi-stage spatial-temporal graph convolutional neural network; NFOG-Q: New freezing of gait questionnaire; H [MYAMPY]: Hoehn and Yahr; MMSE: Mini-mental state examination; UPDRS: Unified Parkinson's Disease Rating Scale; SD: Standard deviation; D2: Dataset 2; FG: Functional gait; TP: True positive; TN: True negative; FP: False positive; FN: False negative; MCC: Matthews correlation coefficient; CI: Confidence interval; BTK: Biomechanical toolkit.

Acknowledgements

We thank the employees of the gait laboratory for technical support during data collection.

Author contributions

Study design by BF, PG, AN, PS, and BV. Data analysis by BF. Design and implementation of the neural network architecture by BF. Statistics by BF and BV. Subject recruitment, data collection, and data preparation by AN. The first draft of the manuscript was written by BF and all authors commented on subsequent revisions. The final manuscript was read and approved by all authors.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Availability of data and materials

The input set was imported and labeled using Python version 2.7.12 with Biomechanical Toolkit (BTK) version 0.3 [79]. The MS-GCN architecture was implemented in Pytorch version 1.2 [80] by adopting the public code repositories of MS-TCN [46] and ST-GCN [50]. All models were trained on an NVIDIA Tesla K80 GPU using Python version 3.6.8. The datasets analyzed during the current study are not publicly available due to restrictions on sharing subject health information.

Declarations

Ethics approval and consent to participate

The study was approved by the local ethics committee of the University Hospital Leuven and all subjects gave written informed consent.

Consent for publication

Not applicable.

Competing interests

The authors declare that there is no conflict of interest regarding the publication of this article.

Author details

¹Department of Electrical Engineering (ESAT), eMedia Research Lab/STADIUS, KU Leuven, Andreas Vesaliusstraat 13, 3000 Leuven, Belgium. ²Department of Mechanical Engineering, Intelligent Mobile Platforms Research Group, KU Leuven, Andreas Vesaliusstraat 13, 3000 Leuven, Belgium. ³Department of Rehabilitation Sciences, Research Group for Neurorehabilitation (eNRGy), KU Leuven, Tervuursevest 101, 3001 Heverlee, Belgium.

Received: 30 June 2021 Accepted: 10 May 2022

Published online: 21 May 2022

References

- Perez-Lloret S, Negre-Pages L, Damier P, Delval A, Derkinderen P, Destée A, Meissner WG, Schelosky L, Tison F, Rascol O. Prevalence, determinants, and effect on quality of life of freezing of gait in Parkinson disease. *JAMA Neurol.* 2014;71(7):884–90.
- Hely MA, Reid WGJ, Adena MA, Halliday GM, Morris JGL. The Sydney multicenter study of Parkinson's disease: the inevitability of dementia at 20 years. *Mov Disord.* 2008;23(6):837–44.
- Nutt JG, Bloem BR, Giladi N, Hallett M, Horak FB, Nieuwboer A. Freezing of gait: moving forward on a mysterious clinical phenomenon. *Lancet Neurol.* 2011;10(8):734–44.
- Snijders AH, Nijkrake MJ, Bakker M, Munneke M, Wind C, Bloem BR. Clinimetrics of freezing of gait. *Mov Disord.* 2008;23(Suppl 2):468–74.
- Nonnekes J, Snijders AH, Nutt JG, Deuschl G, Giladi N, Bloem BR. Freezing of gait: a practical approach to management. *Lancet Neurol.* 2015;14(7):768–78.
- Okuma Y. Practical approach to freezing of gait in Parkinson's disease. *Pract Neurol.* 2014;14(4):222–30.
- Schaafsma JD, Balash Y, Gurevich T, Bartels AL, Hausdorff JM, Giladi N. Characterization of freezing of gait subtypes and the response of each to levodopa in Parkinson's disease. *Eur J Neurol.* 2003;10(4):391–8.
- Giladi N, Balash J, Hausdorff JM. Gait disturbances in Parkinson's disease. In: Mizuno Y, Fisher A, Hanin I, editors. *Mapping the Progress of Alzheimer's and Parkinson's Disease*. Boston: Springer; 2002. p. 329–35.
- Giladi N, Hausdorff JM. The role of mental function in the pathogenesis of freezing of gait in Parkinson's disease. *J Neurol Sci.* 2006;248(1–2):173–6.
- Moore O, Kreitler S, Ehrenfeld M, Giladi N. Quality of life and gender identity in Parkinson's disease. *J Neural Transm.* 2005;112(11):1511–22.
- Bloem BR, Hausdorff JM, Visser JE, Giladi N. Falls and freezing of gait in Parkinson's disease: a review of two interconnected, episodic phenomena. *Mov Disord.* 2004;19(8):871–84.
- Grimbergen YAM, Munneke M, Bloem BR. Falls in Parkinson's disease. *Curr Opin Neurol.* 2004;17(4):405–15.
- Gray P, Hildebrand K. Fall risk factors in Parkinson's disease. *J Neurosci Nurs.* 2000;32(4):222–8.
- Rudzińska M, Bukowczan S, Stożek J, Zajdel K, Mirek E, Chwata W, Wójcik-Pędziwiatr M, Banaszkiwicz K, Szczudlik A. Causes and consequences of falls in Parkinson disease patients in a prospective study. *Neurol Neurochir Pol.* 2013;47(5):423–30.
- Pellicioni PHS, Menant JC, Latt MD, Lord SR. Falls in Parkinson's disease subtypes: risk factors, locations and circumstances. *Int J Environ Res Public Health.* 2019;16(12):2216.
- Gilat M, Lúcia Silva de Lima A, Bloem BR, Shine JM, Nonnekes J, Lewis SJG. Freezing of gait: promising avenues for future treatment. *Parkinsonism Relat Disord.* 2018;52:7–16.
- Mancini M, Bloem BR, Horak FB, Lewis SJG, Nieuwboer A, Nonnekes J. Clinical and methodological challenges for assessing freezing of gait: future perspectives. *Mov Disord.* 2019;34(6):783–90.
- Giladi N, Shabtai H, Simon ES, Biran S, Tal J, Korczyn AD. Construction of freezing of gait questionnaire for patients with parkinsonism. *Parkinsonism Relat Disord.* 2000;6(3):165–70.
- Nieuwboer A, Rochester L, Herman T, Vandenberghe W, Emil GE, Thomaes T, Giladi N. Reliability of the new freezing of gait questionnaire: agreement between patients with Parkinson's disease and their carers. *Gait Posture.* 2009;30(4):459–63.
- Shine JM, Moore ST, Bolitho SJ, Morris TR, Dilda V, Naismith SL, Lewis SJG. Assessing the utility of freezing of gait questionnaires in Parkinson's disease. *Parkinsonism Relat Disord.* 2012;18(1):25–9.
- Gilat M. How to annotate freezing of gait from video: a standardized method using Open-Source software. *J Parkinsons Dis.* 2019;9(4):821–4.
- Morris TR, Cho C, Dilda V, Shine JM, Naismith SL, Lewis SJG, Moore ST. A comparison of clinical and objective measures of freezing of gait in Parkinson's disease. *Parkinsonism Relat Disord.* 2012;18(5):572–7.
- Moore ST, MacDougall HG, Ondo WG. Ambulatory monitoring of freezing of gait in Parkinson's disease. *J Neurosci Methods.* 2008;167(2):340–8.
- Moore ST, Yungher DA, Morris TR, Dilda V, MacDougall HG, Shine JM, Naismith SL, Lewis SJG. Autonomous identification of freezing of gait in Parkinson's disease from lower-body segmental accelerometry. *J Neuroeng Rehabil.* 2013;10:19.
- Popovic MB, Djuric-Jovicic M, Radovanovic S, Petrovic I, Kostic V. A simple method to assess freezing of gait in Parkinson's disease patients. *Braz J Med Biol Res.* 2010;43(9):883–9.
- Delval A, Snijders AH, Weerdesteijn V, Duysens JE, Defebvre L, Giladi N, Bloem BR. Objective detection of subtle freezing of gait episodes in Parkinson's disease. *Mov Disord.* 2010;25(11):1684–93.
- Hu K, Wang Z, Mei S, Ehgoetz Martens KA, Yao T, Lewis SJG, Feng DD. Vision-based freezing of gait detection with anatomic directed graph representation. *IEEE J Biomed Health Inform.* 2020;24(4):1215–25.
- Ahlich C, Samà A, Lawo M, Cabestany J, Rodríguez-Martín D, Pérez-López C, Sweeney D, Quinlan LR, Laighin G, Counihan T, Browne P, Hadas L, Vainstein G, Costa A, Annicchiarico R, Alcaine S, Mestre B, Quispe P, Bayes À, Rodríguez-Molinero A. Detecting freezing of gait with a triaxial accelerometer in Parkinson's disease patients. *Med Biol Eng Comput.* 2016;54(1):223–33.
- Rodríguez-Martín D, Samà A, Pérez-López C, Català A, Moreno Arostegui JM, Cabestany J, Bayes À, Alcaine S, Mestre B, Prats A, Crespo MC, Counihan T, Browne P, Quinlan LR, ÓLaighin G, Sweeney D, Lewy H, Azuri J, Vainstein G, Annicchiarico R, Costa A, Rodríguez-Molinero A. Home detection of freezing of gait using support vector machines through a single waist-worn triaxial accelerometer. *PLoS ONE.* 2017;12(2):0171764.
- Masiála S, Huijbers W, Atzmueller M. Feature-Set-Engineering for detecting freezing of gait in Parkinson's disease using deep recurrent neural networks. pre-print 2019. [arXiv:1909.03428](https://arxiv.org/abs/1909.03428).
- Tahafchi P, Molina R, Roper JA, Sowalsky K, Hass CJ, Gunduz A, Okun MS, Judy JW. Freezing-of-Gait detection using temporal, spatial, and physiological features with a support-vector-machine classifier. In: 2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 2867–2870; 2017.

32. Camps J, Samà A, Martín M, Rodríguez-Martín D, Pérez-López C, Alcaine S, Mestre B, Prats A, Crespo MC, Cabestany J, Bayés À, Català A. Deep learning for detecting freezing of gait episodes in parkinson's disease based on accelerometers. In: *Advances in Computational Intelligence*, 2017;pp. 344–355. Springer.
33. Sigcha L, Costa N, Pavón I, Costa S, Arezes P, López JM, De Arcas G. Deep learning approaches for detecting freezing of gait in Parkinson's disease patients through On-Body acceleration sensors. *Sensors*. 2020;20(7):1895.
34. Mancini M, Priest KC, Nutt JG, Horak FB. Quantifying freezing of gait in Parkinson's disease during the instrumented timed up and go test. *Conf Proc IEEE Eng Med Biol Soc*. 2012;2012:1198–201.
35. Mancini M, Shah VV, Stuart S, Curtze C, Horak FB, Safarpour D, Nutt JG. Measuring freezing of gait during daily-life: an open-source, wearable sensors approach. *J Neuroeng Rehabil*. 2021;18(1):1.
36. O'Day J, Lee M, Seagers K, Hoffman S, Jih-Schiff A, Kidziński Ł, Delp S, Bronte-Stewart H. Assessing inertial measurement unit locations for freezing of gait detection and patient preference. 2021.
37. Rohrbach M, Amin S, Andriluka M, Schiele B. A database for fine grained activity detection of cooking activities. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1194–1201 2012.
38. Ni B, Yang X, Gao S. Progressively parsing interactional objects for fine grained action detection. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1020–1028 2016.
39. Lea C, Flynn MD, Vidal R, Reiter A, Hager GD. Temporal convolutional networks for action segmentation and detection. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1003–1012, 2017. <https://doi.org/10.1109/CVPR.2017.113>.
40. Kuehne H, Gall J, Serre T. An end-to-end generative framework for video segmentation and recognition. *IEEE Workshop on Applications of Computer Vision (WACV)*, 2015. [arXiv:1509.01947](https://arxiv.org/abs/1509.01947).
41. Tang K, Fei-Fei L, Koller D. Learning latent temporal structure for complex event detection. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1250–1257, 2012.
42. Singh B, Marks TK, Jones M, Tuzel O, Shao M. A multi-stream bi-directional recurrent neural network for Fine-Grained action detection. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1961–1970, 2016.
43. Huang D-A, Fei-Fei L, Nibbles JC. Connectionist temporal modeling for weakly supervised action labeling. In: Leibe B, Matas J, Sebe N, Welling M, editors. *Computer Vision—ECCV 2016*. Cham: Springer; 2016. p. 137–53.
44. Bai S, Zico Kolter J, Koltun V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. pre-print, 2018. [arXiv:1803.01271](https://arxiv.org/abs/1803.01271).
45. Yu F, Koltun V. Multi-Scale context aggregation by dilated convolutions. pre-print, 2015. [arXiv:1511.07122](https://arxiv.org/abs/1511.07122).
46. Farha YA, Gall J. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3570–3579, 2019. <https://doi.org/10.1109/CVPR.2019.00369>.
47. Fathi A, Ren X, Rehg JM. Learning to recognize objects in egocentric activities. In: *CVPR 2011*, pp. 3281–3288, 2011.
48. Stein S, McKenna SJ. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In: *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. UbiComp '13, pp. 729–738. Association for Computing Machinery, New York, NY, USA 2013.
49. Carreira J, Zisserman A. Quo vadis, action recognition? a new model and the kinetics dataset. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4724–4733, 2017. <https://doi.org/10.1109/CVPR.2017.502>.
50. Yan S, Xiong Y, Lin D. Spatial temporal graph convolutional networks for skeleton-based action recognition. In: *AAAI 2018*.
51. Spildooren J, Vercruyse S, Desloovere K, Vandenbergh W, Kerckhofs E, Nieuwboer A. Freezing of gait in Parkinson's disease: the impact of dual-tasking and turning. *Mov Disord*. 2010;25(15):2563–70.
52. Vervoort G, Bengevoord A, Strouwen C, Bekkers EMJ, Heremans E, Vandenbergh W, Nieuwboer A. Progression of postural control and gait deficits in Parkinson's disease and freezing of gait: a longitudinal study. *Parkinsonism Relat Disord*. 2016;28:73–9.
53. Kadaba MP, Ramakrishnan HK, Wootten ME. Measurement of lower extremity kinematics during level walking. *J Orthop Res*. 1990;8(3):383–92.
54. Davis RB, Öunpuu S, Tyburski D, Gage JR. A gait analysis data collection and reduction technique. *Hum Mov Sci*. 1991;10(5):575–87.
55. Canning CG, Ada L, Johnson JJ, McWhirter S. Walking capacity in mild to moderate Parkinson's disease. *Arch Phys Med Rehabil*. 2006;87(3):371–5.
56. Bowen A, Wenman R, Mickelborough J, Foster J, Hill E, Tallis R. Dual-task effects of talking while walking on velocity and balance following a stroke. *Age Ageing*. 2001;30(4):319–23.
57. Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift 2015. [arXiv:1502.03167](https://arxiv.org/abs/1502.03167).
58. Filtjens B, Nieuwboer A, D'cruz N, Spildooren J, Slaets P, Vanrumste B. A data-driven approach for detecting gait events during turning in people with Parkinson's disease and freezing of gait. *Gait Posture*. 2020;80:130–6.
59. Matsushita Y, Tran DT, Yamazoe H, Lee J-H. Recent use of deep learning techniques in clinical applications based on gait: a survey. *J Comput Design Eng*. 2021;8(6):1499–532.
60. Graves A, Schmidhuber J. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Netw*. 2005;18(5–6):602–10.
61. Kidziński Ł, Delp S, Schwartz M. Automatic real-time gait event detection in children using deep neural networks. *PLoS ONE*. 2019;14(1):0211466.
62. Kingma DP, Ba J. Adam: a method for stochastic optimization. pre-print 2014 [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
63. Saeb S, Lonini L, Jayaraman A, Mohr DC, Kording KP. The need to approximate the use-case in clinical machine learning. *Gigascience*. 2017;6(5):1–9.
64. Matthews BW. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochim Biophys Acta*. 1975;405(2):442–51.
65. Chan YH. *Biostatistics 104: correlational analysis*. Singapore Med J. 2003;44(12):614–9.
66. Walton CC, Mowszowski L, Gilat M, Hall JM, O'Callaghan C, Muller AJ, Georgiades M, Szeto JYY, Ehgoetz Martens KA, Shine JM, Naismith SL, Lewis SJG. Cognitive training for freezing of gait in Parkinson's disease: a randomized controlled trial. *NPJ Parkinsons Dis*. 2018;4:15.
67. Cao Z, Hidalgo G, Simon T, Wei S-E, Sheikh Y. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE Trans Pattern Anal Mach Intell*. 2021;43(1):172–86. <https://doi.org/10.1109/TPAMI.2019.2929257>.
68. Mathis A, Mamidanna P, Cury KM, Abe T, Murthy VN, Mathis MW, Bethge M. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci*. 2018;21(9):1281–9.
69. Kidziński Ł, Yang B, Hicks JL, Rajagopal A, Delp SL, Schwartz MH. Deep neural networks enable quantitative movement analysis using single-camera videos. *Nat Commun*. 2020;11(1):4054.
70. Lempereur M, Rousseau F, Rémy-Nérès O, Pons C, Houx L, Quellec G, Brochard S. A new deep learning-based method for the detection of gait events in children with gait disorders: proof-of-concept and concurrent validity. *J Biomech*. 2020;98: 109490.
71. Nieuwboer A, Dom R, De Weerd T, Desloovere K, Fieuws S, Broens-Kaucsik E. Abnormalities of the spatiotemporal characteristics of gait at the onset of freezing in Parkinson's disease. *Mov Disord*. 2001;16(6):1066–75.
72. Rahman S, Griffin HJ, Quinn NP, Jahanshahi M. The factors that induce or overcome freezing of gait in Parkinson's disease. *Behav Neurol*. 2008;19(3):127–36.
73. Barredo Arrieta A, Diaz-Rodríguez N, Del Ser J, Bennetot A, Tabik S, Barbedo A, Garcia S, Gil-Lopez S, Molina D, Benjamins R, Chatila R, Herrera F. Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf Fusion*. 2020;58:82–115.
74. Horst F, Lapuschkin S, Samek W, Müller K-R, Schöllhorn WI. Explaining the unique nature of individual gait patterns with deep learning. *Sci Rep*. 2019;9(1):2391.
75. Filtjens B, Ginis P, Nieuwboer A, Afzal MR, Spildooren J, Vanrumste B, Slaets P. Modelling and identification of characteristic kinematic features preceding freezing of gait with convolutional neural networks and layer-wise relevance propagation. *BMC Med Inform Decis Mak*. 2021;21(1):341.
76. Bach S, Binder A, Montavon G, Klauschen F, Müller K-R, Samek W. On pixel-wise explanations for non-linear classifier decisions by Layer-Wise relevance propagation. *PLoS ONE*. 2015;10(7):0130140.

77. Sundararajan M, Taly A, Yan Q. Axiomatic attribution for deep networks. In: Proceedings of the 34th International Conference on Machine Learning—Volume 70. ICML'17, pp. 3319–3328. JMLR.org, 2017.
78. Shrikumar A, Greenside P, Kundaje A. Learning important features through propagating activation differences. In: Precup, D., Teh, Y.W. (eds.) Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 70, pp. 3145–3153. PMLR, International Convention Centre, Sydney, Australia 2017. <http://proceedings.mlr.press/v70/shrikumar17a.html>.
79. Barre A, Armand S. Biomechanical ToolKit: open-source framework to visualize and process biomechanical data. *Comput Methods Programs Biomed.* 2014;114(1):80–7.
80. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Kopf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, Chintala S. Pytorch: An imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 32. Curran Associates, Inc., 2019. <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>.
81. Folstein MF, Folstein SE, McHugh PR. "mini-mental state": A practical method for grading the cognitive state of patients for the clinician. *J Psychiatr Res.* 1975;12(3):189–98.
82. ...Goetz CG, Tilley BC, Shaftman SR, Stebbins GT, Fahn S, Martinez-Martin P, Poewe W, Sampaio C, Stern MB, Dodel R, Dubois B, Holloway R, Jankovic J, Kulisevsky J, Lang AE, Lees A, Leurgans S, LeWitt PA, Nyenhuis D, Olanow CW, Rascol O, Schrag A, Teresi JA, van Hilten JJ, LaPelle N. Movement Disorder Society UPDRS Revision Task Force: movement disorder society-sponsored revision of the unified parkinson's disease rating scale (MDS-UPDRS): scale presentation and clinimetric testing results. *Mov Disord.* 2008;23(15):2129–70.
83. Hoehn MM, Yahr MD. Parkinsonism: onset, progression and mortality. *Neurology.* 1967;17(5):427–42.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

