

Data Augmentation of Wearable Sensor Data for Parkinson’s Disease Monitoring using Convolutional Neural Networks*

Terry T. Um
University of Waterloo
Canada
terry.t.um@gmail.com

Franz M. J. Pfister
Ludwig-Maximilians-Univ. München
Germany
fmj.pfister@me.com

Daniel Pichler
Technical University of Munich
Germany
daniel.pichler@tum.de

Satoshi Endo, Muriel Lang,
Sandra Hirche
Technical University of Munich
Germany
{s.endo,muriel.lang,hirche}@tum.de

Urban Fietzek
Schön Klinik München Schwabing
Germany
urban.fietzek@schoen-kliniken.de

Dana Kulić
University of Waterloo
Canada
dana.kulic@uwaterloo.ca

ABSTRACT

While convolutional neural networks (CNNs) have been successfully applied to many challenging classification applications, they typically require large datasets for training. When the availability of labeled data is limited, data augmentation is a critical preprocessing step for CNNs. However, data augmentation for wearable sensor data has not been deeply investigated yet.

In this paper, various data augmentation methods for wearable sensor data are proposed. The proposed methods and CNNs are applied to the classification of the motor state of Parkinson’s Disease patients, which is challenging due to small dataset size, noisy labels, and large intra-class variability. Appropriate augmentation improves the classification performance from 77.54% to 86.88%.

CCS CONCEPTS

• **Computing methodologies** → **Supervised learning by classification**; • **Applied computing** → **Consumer health**;

KEYWORDS

Data augmentation; wearable sensor; convolutional neural networks; Parkinson’s disease; health monitoring

ACM Reference Format:

Terry T. Um, Franz M. J. Pfister, Daniel Pichler, Satoshi Endo, Muriel Lang, Sandra Hirche, Urban Fietzek, and Dana Kulić. 2017. Data Augmentation of Wearable Sensor Data for Parkinson’s Disease Monitoring using Convolutional Neural Networks. In *Proceedings of 19th ACM International Conference on Multimodal Interaction (ICMI’17)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3136755.3136817>

*This work was partly supported by the EU Seventh Framework Programme FP7/2007-2013 within the ERC Starting Grant Control based on Human Models (con-humo), grant agreement No. 337654.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMI’17, November 13–17, 2017, Glasgow, UK
© 2017 Association for Computing Machinery.
ACM ISBN 978-1-4503-5543-8/17/11...\$15.00
<https://doi.org/10.1145/3136755.3136817>

1 INTRODUCTION

In recent years, convolutional neural networks (CNNs) have shown excellent performance on classification problems when large-scale labeled datasets are available (e.g. [9, 20]). However, it is challenging to apply CNNs to problems where only small labelled datasets are available. For example, collecting and labeling a large amount of medical data is often difficult. As a result, it is challenging to apply CNNs to small-scale medical data.

Data augmentation leverages limited data by transforming the existing samples to create new ones. A key challenge for data augmentation is to generate new data that maintains the correct label, which typically requires domain knowledge. However, it is not obvious how to carry out label-preserving augmentation in some domains, e.g., wearable sensor data. For example, scaling of the acceleration data may change their labels because some labels are differentiated by the intensity of motion.

In this paper, the problem of classifying the motor state of Parkinson’s disease (PD) patients is tackled using CNNs. PD motor state classification is a challenging task due to noisy labels, irrelevant motion interference, large variability over patients, and limited availability of the labelled data. In this paper, we propose data augmentation methods for wearable sensor data and successfully tackle the challenging PD classification task using CNNs.

The contributions of the paper can be summarized as follows:

- Application of CNNs to the task of PD motor state classification, using a clinician-labeled dataset of 30 PD patients (25 patient’s data are exploited) in daily-living conditions.
- A set of approaches for data augmentation of wearable sensor datasets for CNN-based classification.
- Experimental comparison of proposed data augmentation methods.

2 RELATED WORK

Most PD patients experience motor fluctuations, which are characterized by phases of bradykinesia, i.e. underscaled and slow movement, and dyskinesia, i.e. overflowing spontaneous movement [17]. Dopaminergic treatment can alleviate symptoms of bradykinesia while its over-treatment can cause dyskinesia. Thus, an accurate evaluation of a patient’s phenomenology is needed for determining the right dose of medication. Current PD motor state evaluation




Figure 1: (a) and (b) show typical examples of bradykinesia and dyskinesia in a 1 min window while (c) and (d) show atypical patterns. The blue, red, green represent X,Y,Z signals from the accelerometer, respectively.

methods rely on patient self-reports and visual observation by the clinician [17].

Researchers have proposed automating the evaluation with wearable sensors (e.g. [5, 19]). However, most approaches to date have been limited to standardized motor tasks in clinical settings [18]. To enable automated evaluation of PD motor states which covers a wide range of PD symptoms across patients, a large amount of wearable sensor data in daily-living conditions is needed [4]. Deep learning (DL) approaches [15] provide a promising methodology to deal with the large variability of PD data [5, 8, 13]. Given the difficulty in collecting such large datasets, data augmentation is needed [6].

Data augmentation is an indispensable preprocessing step for achieving peak performance in DL approaches (e.g. [9, 12]). For augmenting time-series data, Le Guennec et al. [14] used window slicing and window warping methods, which extracts multiple small-size windows from a single window and lengthens/shortens a part of the window data, respectively. Unlike data augmentations for image [2] and speech recognition [3], however, data augmentation for wearable sensor data has not been systematically investigated yet to the best of our knowledge. In this paper, we propose various data augmentation methods that enable the classification of PD motor states from wearable data and evaluate them using CNN.

3 PD MOTOR STATE CLASSIFICATION

3.1 Challenges in PD Data

We consider two frequent PD motor states: bradykinesia, which is characterized by decreased movement speed and may be accompanied by tremor, and dyskinesia, which is characterized by involuntary extremity movements. Figure 1 illustrates exemplar one minute data windows of both motor states, from a single accelerometer worn on the wrist of PD patients. Bradykinesia data typically appear as constant signals indicating less movement (Fig 1(a)) while dyskinesia data consist of fluctuating movements (Fig 1(b)).

However, there are a significant number of examples that deviate from the stereotypical expressions. For example, bradykinesia accompanied by tremor can show fluctuating signals which look like a dyskinesia state (Fig 1(c)). On the other hand, dyskinesia with voluntary suppression can show constant signals which look like a bradykinesia state (Fig 1(d)).

There are several factors that can cause an apparent disagreement between the observed data pattern and the expert label. First, if the body of the patient indicates, e.g., a dyskinesia state, but the hand which wears the wearable sensor does not move because the patient is, e.g., holding a chair for suppressing the symptom, the assigned label based on the overall body expression will be mismatched with the recorded data from the wearable device. Also, the expert rater typically rates the symptoms for a fixed length window, but arbitrary segmentation into fixed length windows may not result in single motor state windows. Furthermore, the interference of voluntary movements, e.g., waving the hand, can make bradykinesia states look like dyskinesia, and, e.g., voluntary rest, appear like bradykinesia. Finally, bradykinesia accompanied by tremor can also make it difficult to distinguish between bradykinesia and dyskinesia.

The factors described above introduce noisy labels, and lead to large intra-class variability and significant overlap between two classes. As a result, it makes the PD motor state classification more challenging, particularly given a small amount of data.

3.2 Data Augmentation Methods for Wearable Sensor Data

Data augmentation can be viewed as an injection of prior knowledge about the invariant properties of the data against certain transformations. Augmented data can cover unexplored input space, prevent overfitting, and improve the generalization ability of a DL model [6]. In image recognition, it is well-known that minor changes due to jittering, scaling, cropping, warping and rotating do not alter the data labels because they are likely to happen in real world observations. However, label-preserving transformations for wearable sensor data are not obvious and intuitively recognizable (Fig 2).

One factor that can introduce label-invariant variability of wearable sensor data are differences in sensor placement between participants. For example, an upside-down placement of the sensor can invert the sign of the sensor readings without changing the labels. Therefore, augmentation by applying arbitrary **rotations (Rot)** to the existing data can be used as a way of simulating different sensor placements.

Another factor that can introduce variability is the temporal location of activity events, e.g., tremor, in the window. Since the fixed size window segmentation is arbitrary, the location of the observed symptom in the window does not have any meaning. Thus, we may augment data by perturbing the location of the windows or events.

Permutation (Perm) is a simple way to randomly perturb the temporal location of within-window events. To perturb the location of the data in a single window, we first slice the data into N same-length segments, with N ranging from 1 to 5, and randomly permute the segments to create a new window. **Time-warping (TimeW)** is another way to perturb the temporal location. By smoothly distorting the time intervals between samples, the temporal locations of the samples can be changed using time-warping.

Small changes in magnitude may preserve the labels, depending on the target task. **Scaling (Scale)** changes the magnitude of the data in a window by multiplying by a random scalar, while




Figure 2: Various data augmentations that are used in the experiments: jittering, scaling, rotating, permuting, magnitude-warping, time-warping methods. Combinations of various data augmentations can also be applied.

magnitude-warping (MagW) changes the magnitude of each sample by convolving the data window with a smooth curve varying around one. In addition, **jittering (Jitter)** is also considered as a way of simulating additive sensor noise. These data augmentation methods may increase robustness against multiplicative and additive noise and improve performance.

Lastly, **cropping (Crop)**, which is similar to image cropping or window slicing in [14], is applied for diminishing the dependency on event locations. Note that cropping can capture an event-free region, which might change the label. Also, note that cropping with random locations over epochs will eventually converge to a sliding window method with arbitrary stride sizes.

In a nutshell, jittering, cropping, rotating, permuting, magnitude-warping and time-warping methods are applied for augmenting wearable sensor data. In the next section, the performance of PD motor state classification with the proposed data augmentation methods is evaluated using CNNs.

4 EXPERIMENTS

4.1 Data Preparation

A dataset of 30 patients' motor states was collected using Microsoft Band 2 [1] in daily-living conditions without requesting specific motor tasks¹. The 30 PD patients are 67 ± 10 years old, median Hoehn & Yahr stage 2, average disease duration 11 ± 5 years, and MoCA points 26 ± 3 . Among them, 25 patient's data are used for this research and each one minute interval is labeled by a clinical expert. The data are collected at a frequency of 62.5Hz and resampled to 120Hz to deal with sampling irregularities. The first 58-seconds of data (6960 samples) from each one minute window is used to make same-length instances.

Similar to previous works (e.g. [19], [7], [8], [5]) acceleration data only are used for the PD motor state classification. Also, *no-symptom* data are removed to simplify the problem and focus on characterizing data augmentation methods. Data collected during walking, laying and eating activities are also removed due to limited observation of movement during these activities. Note that no other preprocessing, e.g., data normalization or smoothing, is applied because they may confound the data label and subsequent results.

The resulting dataset consists of 3530 min (58.8 hours) of bradykinesia and dyskinesia data. For cross-validation, the 25 PD patients




Figure 3: 7-layer CNN with a global average pooling (GAP). The 7-layer CNN consists of 16-32-64-64-64-64-64 feature maps which reduce the size of the inputs to 2319×3 , 772×3 , 385×3 , 193×3 , 97×3 , 49×3 , 48×1 , respectively.

are divided into five subject groups. The performance of PD motor state classification is reported in Section 4.3 using the average values of 5-fold cross-validation results.

4.2 The CNN Architecture

In this research, CNNs are used for PD motor state classification. CNNs are more suitable for small-scale datasets than long short-term memories (LSTMs) [10] because CNNs generally use a smaller number of parameters compared to fully-connected LSTMs. Deep and sparse 7-layer CNNs (Figure 3) are employed to capture the large variability of the small-scale PD data.

A convolutional layer, a batch normalization layer [11], and an activation layer using rectified units (ReLU) form a single convolutional layer of the 7-layer CNN. With strided convolutions using 4×1 , 4×1 , 3×1 , 3×3 , 2×3 , 2×3 , 2×3 convolution filters, the sizes of the inputs are reduced from 6960×3 to 48×1 over layers (Figure 3). Note that XYZ signals of the accelerometer are convolved in layers 4,5,6 and 7 to capture inter-vector-component features. For reducing the number of parameters for small-scale datasets, a global averaging pooling (GAP) layer [16] is applied at the end instead of fully-connected layers.

4.3 Results

Classification of PD motor states is performed using the CNN with the various data augmentation methods. For baseline results, a support vector machine (SVM) with an RBF kernel is applied to 540 dimensional statistical features: mean, variance, skewness, kurtosis,

¹The study was approved by the ethics committee of Technical University of Munich (Az. 234/16 S).

Table 1: The results of PD motor state classification with various data augmentation methods. R,P,T,M represent *Rot*, *Perm*, *TimeW*, *MagW*, respectively.

	SVM	CNN	Jitter	Scale	Crop	Rot	Perm
Train	98.82	99.92	99.78	99.84	65.77	100.0	99.33
Test	70.72	77.54	77.52	79.46	73.58	82.62	81.16
	MagW	TimeW	P,T	R,P	R,T	R,P,T	R,P,T,M
Train	100.0	94.67	96.63	99.08	94.70	94.43	94.20
Test	79.33	82.00	81.75	86.76	85.01	86.88	85.60




Figure 4: Training curves for *CNN*, *Rot*, *Perm*, *TimeW*, *Rot+Perm* and *Rot+Perm+TimeW* methods. The curves of *Rot+Perm+TimeW* shows slow training improvement and a better generalization performance.

and maximum values are extracted from 1 min data using 5 and 10-sec sliding windows. Also, a CNN is applied to raw 1 min data without data augmentation for baseline comparison. All experiments except for the SVM are performed for 400 epochs and the median values from the last 10 epoch results are used for averaging the 5-fold cross-validation results.

Different random parameter values are applied for data augmentation. For jittering, a standard deviation (STD) value is sampled from a Gaussian distribution with 0.03 STD, and 1 min of Gaussian noise is generated using the sampled STD value. For scaling, a random scalar is sampled from a Gaussian distribution with a mean of 1 and 0.1 STD. For rotation, an arbitrary rotation matrix is generated for each instance. For permutation, a random integer N is determined by rounding a positive value sampled from a Gaussian distribution with 5.0 STD. For magnitude-warping and time-warping, random sinusoidal curves are generated using arbitrary amplitude, frequency, and phase values. The implemented code for the proposed data augmentation methods is available online: <https://github.com/terryum/Data-Augmentation-For-Wearable-Sensor-Data>

The main results are presented in Table 1. Jittering fails to improve the performance of PD motor state classification because it introduces rapid fluctuations which look similar to dyskinesia. Cropping also fails because it drops the information of 2/3 window samples, which could be a critical loss given the small dataset. Cropping of an event-free region also hinders the learning process and can be a cause of the poor performance. Scaling and magnitude-warping also fail because changing of the intensity of the signal may alter the labels.

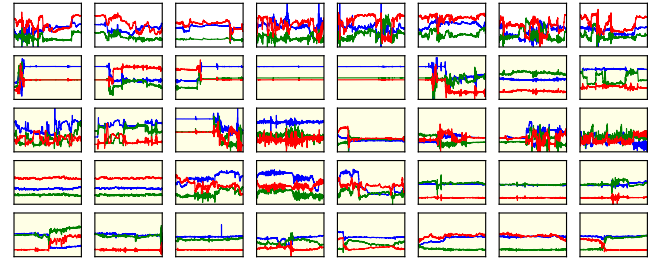


Figure 5: Randomly selected 40 incorrect mispredictions from the Fold-1 results of *Rot+Perm+TimeW* experiment. Fluctuating signals from bradykinesia (white) and constant signals from dyskinesia (yellow) are often misclassified.

On the other hand, rotation, permutation, and time-warping methods improve the performance of PD motor state classification. The best performance among the single data augmentation methods is achieved by rotation. Permutation and time-warping also provide performance improvements by perturbing the temporal locations of samples. These results indicate that the major sources of variability are different sensor placements between participants and event locations in an arbitrarily segmented window. The proposed rotation, permutation, and time-warping methods effectively compensate the unnecessary variations and improve the performance by 3.6-5.1% accuracy.

Combinations of various data augmentation methods show better performance than that of a single data augmentation method. The combinations of *Rot+Perm* and *Rot+TimeW* show better performance than the baseline of *CNN* by 7.5-9.2%. The best performance is achieved by *Rot+Perm+TimeW* with 86.88% accuracy. These results indicate that rotation can be used to alleviate sensor pose variability while either permutation or time-warping can be employed for addressing the variability of temporal locations of events in a window.

Training curves of the experiments are depicted in Fig 4. The *Rot+Perm+TimeW* curve shows slow training improvement and a better generalization performance than others thanks to the regularization effect provided by the data augmentation. Some of the failed predictions are presented in Fig 5. From the figure, it can be observed that CNNs often misclassify fluctuating bradykinesia and constant dyskinesia data, which can be considered as seemingly-noisy labels as described in Section 3.1.

5 CONCLUSION

In this paper, an automatic classification algorithm for PD motor state monitoring is developed based on wearable sensor data. PD motor state classification is a challenging task because of large inter-class variability, noisy labels, interference by irrelevant motion signals and limited data availability. The challenging PD task is successfully tackled using a 7-layer CNN and the proposed data augmentation methods. The combination of rotational and permutational data augmentation methods improves the baseline performance of 77.52% accuracy to 86.88%. Systematic experiments with various data augmentation methods provide a direction towards a general approach for augmentation for wearable sensor data.

REFERENCES

- [1] Microsoft Band 2. 2015. <https://www.microsoft.com/microsoft-band/>
- [2] Ken Chatfield, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Return of the Devil in the Details: Delving Deep into Convolutional Nets. *CoRR* abs/1405.3531 (2014).
- [3] Xiaodong Cui, Vaibhava Goel, and Brian Kingsbury. 2015. Data Augmentation for Deep Neural Network Acoustic Modeling. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.* 23, 9 (Sept. 2015), 1469–1477.
- [4] Silvia Del Din, Alan Godfrey, Claudia Mazzà, Sue Lord, and Lynn Rochester. 2016. Free-living monitoring of Parkinson's disease: Lessons from the field. *Movement Disorders* 31, 9 (2016), 1293–1313.
- [5] Bjoern M. Eskofier, Sunghoon I. Lee, Jean-Francois Daneault, Fatemeh N. Golabchi, Gabriela Ferreira-Carvalho, Gloria Vergara-Diaz, Stefano Sapienza, Gianluca Costante, Jochen Klucken, Thomas Kautz, and Paolo Bonato. 2016. Recent machine learning advancements in sensor-based mobility analysis: Deep learning for Parkinson's disease assessment. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 655–658.
- [6] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.
- [7] Robert I. Griffiths, Katya Kotschet, Sian Arfon, Zheng Ming Xu, William Johnson, John Drago, Andrew Evans, Peter Kempster, Sanjay Raghav, and Malcolm K. Horne. 2012. Automated assessment of bradykinesia and dyskinesia in Parkinson's disease. *Journal of Parkinson's disease* 2, 1 (2012), 47–55.
- [8] Nils Y. Hammerla, James M. Fisher, Peter Andras, Lynn Rochester, Richard Walker, and Thomas Plotz. 2015. PD Disease State Assessment in Naturalistic Environments Using Deep Learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, 1742–1748.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. *CoRR* abs/1512.03385 (2015). <http://arxiv.org/abs/1512.03385>
- [10] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (Nov 1997), 1735–1780.
- [11] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *ArXiv e-prints* (Feb. 2015). arXiv:cs.LG/1502.03167
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 1097–1105.
- [13] Ken J. Kubota, Jason A. Chen, and Max A. Little. 2016. Machine learning for large-scale wearable sensor data in Parkinson's disease: Concepts, promises, pitfalls, and futures. *Movement Disorders* 31, 9 (2016), 1314–1326.
- [14] Arthur Le Guennec, Simon Malinowski, and Romain Tavenard. 2016. Data Augmentation for Time Series Classification using Convolutional Neural Networks. In *ECML/PKDD Workshop on Advanced Analytics and Learning on Temporal Data*. Riva Del Garda, Italy. <https://halshs.archives-ouvertes.fr/halshs-01357973>
- [15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (28 May 2015), 436–444. *Insight*.
- [16] Min Lin, Qiang Chen, and Shuicheng Yan. 2013. Network In Network. *ArXiv e-prints* (Dec. 2013). arXiv:1312.4400
- [17] João Massano and Kailash P. Bhatia. 2012. Clinical Approach to Parkinson's Disease: Features, Diagnosis, and Principles of Management. *Cold Spring Harb Perspect Med* 2, 6 (Jun 2012), a008870.
- [18] Christiana Ossig, Angelo Antonini, Carsten Buhmann, Joseph Classen, Ilona Csoti, Björn Falkenburger, Michael Schwarz, Jürgen Winkler, and Alexander Storch. 2016. Wearable sensor-based objective assessment of motor symptoms in Parkinson's disease. *Journal of Neural Transmission* 123, 1 (2016), 57–64.
- [19] Shyamal Patel, Konrad Lorincz, Richard Hughes, Nancy Huggins, John Growdon, David Standaert, Metin Akay, Jennifer Dy, Matt Welsh, and Paolo Bonato. 2009. Monitoring Motor Fluctuations in Patients With Parkinson's Disease Using Wearable Sensors. *IEEE Transactions on Information Technology in Biomedicine* 13, 6 (Nov 2009), 864–873.
- [20] Terry Taewoong Um, Vahid Babakeshizadeh, and Dana Kulić. 2017. Exercise Motion Classification from Large-Scale Wearable Sensor Data Using Convolutional Neural Networks. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1–6.