

Wearable-based Parkinson’s Disease Severity Monitoring using Deep Learning

Jann Goschenhofer^{1,2} (✉), Franz MJ Pfister^{1,2}, Kamer Ali Yuksel², Bernd Bischl¹, Urban Fietzek^{3,4}, and Janek Thomas¹

¹ Dept. of Statistics, Ludwig-Maximilians University, Munich, Germany
jann.goschenhofer@stat.uni-muenchen.de

² ConnectedLife GmbH, Munich, Germany

³ Dept. of Neurology, Ludwig-Maximilians University, Munich, Germany

⁴ Dept. of Neurology and Clinical Neurophysiology, Schoen Clinic Schwabing, Munich, Germany

Abstract. One major challenge in the medication of Parkinson’s disease is that the severity of the disease, reflected in the patients’ motor state, cannot be measured using accessible biomarkers. Therefore, we develop and examine a variety of statistical models to detect the motor state of such patients based on sensor data from a wearable device. We find that deep learning models consistently outperform a classical machine learning model applied on hand-crafted features in this time series classification task. Furthermore, our results suggest that treating this problem as a regression instead of an ordinal regression or a classification task is most appropriate. For consistent model evaluation and training, we adopt the leave-one-subject-out validation scheme to the training of deep learning models. We also employ a class-weighting scheme to successfully mitigate the problem of high multi-class imbalances in this domain. In addition, we propose a customized performance measure that reflects the requirements of the involved medical staff on the model. To solve the problem of limited availability of high quality training data, we propose a transfer learning technique which helps to improve model performance substantially. Our results suggest that deep learning techniques offer a high potential to autonomously detect motor states of patients with Parkinson’s disease.

Keywords: Motor State Detection · Sensor Data · Time Series Classification · Deep Learning · Personalized Medicine · Transfer Learning

1 Introduction

Parkinson’s disease (PD) is one of the most common diseases of the elderly and the second most common neurodegenerative disease in general after Alzheimer’s [38]. Two million Europeans are affected and 1% of the population over the age of 60 in industrial nations are estimated to suffer from PD [36, 1]. Fortunately, the disease can be managed by applying the correct personalized dosage and schedule of medication, which has to be continuously adapted regarding the

progress of this neurodegenerative disease. Crucial for the optimal medication is knowledge about the current latent motor state of the patients, which can not yet be measured effortlessly, autonomously and continuously. The motoric capabilities of the patients are distinguishable into three different motor states which can vary substantially over the course of a day within hours. The most prominent symptom is the tremor but the disease defining symptom is the loss of amplitude and slowness of movement, also referred as bradykinesia [35]. In contrast to bradykinesia, an overpresence of dopaminergic medication can make affected patients execute involuntary excessive movement patterns which may remind an untrained observer of a bizarre dance. This hyperkinetic motor state is termed dyskinesia [40]. In a very basic approximation, people with Parkinson’s disease (PwP) can be in three motor states: 1) the bradykinetic state (OFF), 2) a state without apparent symptoms (ON), and 3) the dyskinetic state (DYS) [31]. If the true motor state of PwP was known at all times, the medication dose could be optimized in such a way, that the patient has an improved chance to spend the entirety of his waking day in the ON state. An example for such a closed-loop approach can be found in Diabetes therapy, where the blood sugar level serves as a biomarker for the disease severity. Patients suffering from Diabetes can continuously measure their blood sugar level and apply the individual, correct medication dose of insulin in order to balance the disease. Analogously, an inexpensive, autonomous and precise method to assess the motor state might allow for major improvements in personalized, individual medication of PwP.

Advancements in both wearable devices equipped with motion sensors and statistical modeling tools accelerated the scientific community in researching solutions for motor state detection of PwP since the early 2000s. In 1993, Ghika et al. did pioneering work in this field by proposing a first computer-based system for tremor measurement [14]. A comprehensive overview on the use of machine learning and wearable devices in a variety of PD related problems was recently provided by Ahlrichs et al. [1]. A variety of studies compare machine learning approaches applied on hand-crafted features with deep learning techniques where the latter show the strongest performance [25, 38, 26, 27, 40, 20, 24, 9, 41]. In the present setting, a leave-one-subject-out (LOSO) validation is necessary to yield unbiased performance estimates of the models [37]. Thus, it is surprising that only a subset of the reviewed literature deploys a valid LOSO validation scheme [25, 24, 41, 9, 40]. It is noteworthy that one work proposes modeling approaches with a continuous response [26], while the rest of the literature tackles this problem as a classification task to distinguish between the different motor states. Amongst the deep learning approaches, it is surprising that none of the related investigations describe their method to tune the optimal amount of training epochs for the model, which is not a trivial problem as discussed in Section 3.3. A structured overview on the related literature is given in Table 1.

Contributions This paper closes the main literature gaps in machine learning based monitoring of PD: the optimal problem setting for this task is discussed, a customized performance measure is introduced and a valid LOSO validation

Table 1. Overview on results from the literature on Motor State detection for PwP. In the method column, the MLP refers to a Multi-layer Perceptron, FE to manual feature extraction, SVM to a Support Vector Machine, RF to a Random Forest and LSTM for Long-short-term-memory network. In the label column, the names of the class labels are depicted. From this column, one can infer that only two authors used continuous labels and thus regression models for their task. Generally, a comparison of the reviewed approaches is difficult due to high variation in the data sets, methods and evaluation criteria.

Author	Method	Validation	Subjects	Sensors	Position	Setting	Labels	Results
[25]	FE, SVM	LOSO	19	6	wrist, ankle	lab	ON, OFF	Acc.: 90.5%
[41]	CNN	LOSO	30	1	wrist	free	OFF, ON, DYS	Acc.: 63.1%
[38]	FE, SVM	Holdout Patients	20	1	belt	lab	ON, OFF	Acc.: 94.0%
[24]	LSTM	LOSO	12	1	ankle	free	ON, OFF	Acc.: 73.9 %
	FE, SVM	LOSO	12	1	ankle	free	ON, OFF	Acc.: 65.7 %
[9]	CNN	LOSO	10	2	wrist	lab	ON, OFF	Acc.: 90.9%
[19]	FE, MLP	Leave-one-day-out	34	2	wrist	free	OFF, ON, DYS, Sleep	F1: 55%
[20]	FE, MLP	7-fold CV	34	2	wrist	lab	OFF, ON, DYS, Sleep	F1: 76%
[27]	FE, MLP	Train set	23	6	trunk, wrist, leg	lab	ON, OFF	F1: 97%
[40]	FE, MLP	LOSO	29	6	wrist, leg, chest, waist	lab	DYS Y/N	Acc.: 84.3%
[26]	FE, MLP	80/20 Split	13	6	trunk, wrist, leg	free	Continuous	Acc.: 77%
[30]	FE, RF	LOSO	20	1	wrist	lab	ON, OFF	AUC: 0.73
	FE, RF	LOSO	20	1	wrist	lab	Tremor Y/N	AUC: 0.79
Our approach ^a	CNN	LOSO	28	1	wrist	free	Continuous	MAE: 0.77
	CNN	LOSO	28	1	wrist	free	9-class	±1 Acc.: 86.95%

^a Performance measures are detailed in Section 5

strategy is applied to compare time series classification (TSC) deep learning and classical machine learning approaches. Furthermore, the application of a transfer learning strategy in this domain is investigated.

This paper is structured as follows: The used data sets are described in Section 2. In Section 3, peculiarities of the problem as well as the transfer learning strategy are discussed. Furthermore, in Section 4 model architectures and problem settings are proposed and their results are discussed in Section 5.

2 Data

Data was collected from PwP to model the relation between raw movement sensor data and motor states. The acceleration and rotation of patient’s wrists was measured via inertial measurement units (IMUs) integrated in the Microsoft band 2 fitness tracker [32] with a standard frequency of 62.5Hz. The wrist was chosen as sensor location as it is the most comfortable location for a wearable device to be used in the patients’ daily lives and was shown to be sufficient for the detection of Parkinson-related symptoms [30, 7]. The raw sensor data was downsampled to a frequency of 20Hz as PD related patterns do not exceed this frequency [20]. A standard procedure in human activity recognition is the segmentation of continuous sensor data streams into smaller windows. As the data in this study was annotated by a medical doctor on a minute-level, the window length was set to one minute. To increase the amount of training data, the windows were segmented with an overlap of 80% which is in line with related literature [44, 9, 19]. To neutralize any direction-specific information, the L_2 -norms of the accelerometer and gyroscope measurements are used as model input, leading to two time series per window. Finally, the data was normalized to a $[0, 1]$ range via quantile transformation.

We consider the machine learning problem of the feature space $\mathcal{X} \subset \mathbb{R}^p$, with $p = 1200 \cdot 2$, a target space \mathcal{Y} described below and a performance measure $\mathcal{P} : \mathcal{Y} \times f(\mathcal{X}) \rightarrow \mathbb{R}$ measuring the prediction quality of a model $f : \mathcal{X} \rightarrow \mathcal{Y}$, trained on the data set $\mathcal{D} = \{(x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)})\}$ where a tuple $(x^{(i)}, y^{(i)}) \in \mathcal{X} \times \mathcal{Y}$, $i = 1, \dots, n$ refers to a single labeled one minute window with a frequency of 20Hz.

The disease severity \mathcal{Y} is measured on a combined version of the UPDRS [16] and the mAIMS scale [29]. The UPDRS scale is based on a diagnostic questionnaire for physicians to rate the severity of the bradykinesia of PwP on a scale with 0 representing the ON state to 4, the severely bradykinetic state. The mAIMS scale is analogue to the UPDRS, but in contrast used for the clinical evaluation of dyskinetic symptoms. Both scales were combined and the UPDRS scale was flipped to cover the whole disease spectrum. The resulting label scale takes values in $\mathcal{Y} = \{-4, \dots, 4\}$ where $y^{(i)} = -4$ means a patient is in a severely bradykinetic state, $y^{(i)} = 0$ is assigned to a patient in the ON state and $y^{(i)} = 4$ resembles a severely dyskinetic motor state. The sensor data was labeled by a medical doctor who shadowed the PwP during one day in a free living setting. Thus, the rater monitored each patient, equipped with an IMU, while they

performed regular daily activities and the rater clinically evaluated the patients’ motor state at each minute. In total, 9356 windows were extracted from the data of 28 PwP. By applying the above described preprocessing steps, the amount of windows was increased to 45944.

3 Challenges

3.1 Class Imbalance

The labeled data set suffers from high label imbalance towards the center of the scale as shown in Figure 1. Thus, machine learning models will be biased towards predicting the majority classes [21].

A straightforward way of dealing with this problem is to reweight the loss contribution of different training data samples. This way, the algorithm incurs heavier loss for errors on samples from minority classes than for those of majority classes, putting more focus on the minority classes during training. The weights for the classes $j \in \mathcal{Y} = \{-4, \dots, 4\}$ are calculated as follows:

$$c_j = \frac{n}{n_j}; \quad \tilde{c}_j = |\mathcal{Y}| \cdot \frac{c_j}{\sum_{j \in \mathcal{Y}} c_j} \quad (1)$$

where $|\mathcal{Y}|$ describes the amount of classes, n is the total amount of samples, n_j the amount of samples for class j and thus c_j is the inverse relative frequency of class j in the data. Further, the weights $c_j, j \in \mathcal{Y}$ are normalized such that the sum of the weights is equal to the amount of classes. The individual weight of one sample is referred to as $\omega^{(i)}$ which is the normalized weight \tilde{c}_j associated with the label $y^{(i)}$ of this sample i such that $y^{(i)} = j$.

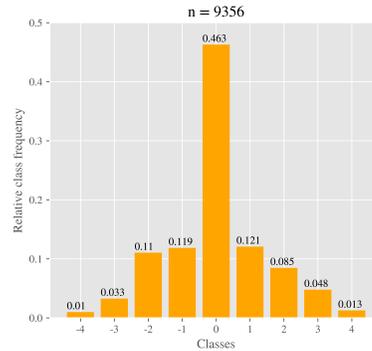


Fig. 1. Label distribution of the data which is highly centered around $y = 0$.

3.2 Custom Performance Measure

It is crucial for the practical application of the final model to select an adequate performance measure which reflects the practical requirements on the model. Based on discussions with involved medical doctors, we found that larger errors should be penalized heavier which implies a quadratic error. Additionally, errors in the wrong direction of the scale, e.g. $\hat{y}^{(i)} = -1, y^{(i)} = 1$, should have a higher negative impact than errors with the same absolute distance in the correct direction, e.g. $\hat{y}^{(i)} = 3, y^{(i)} = 1$. The rationale behind this is that an exaggerated diagnostic evaluation which follows the true pathological scenario harms the

patient less than an opposing one. Furthermore, the cost of predicting the wrong pathological direction increases with the severity of the disease: diagnostic errors weigh heavier on patients with strong symptoms compared to patients that are only mildly affected by the disease. In summary, three main requirements on the custom performance measure were identified: non-linearity, asymmetry and not being translation invariant.

Inspired by econometric forecasting [8], the following asymmetric performance measure, which satisfies the first two previous requirements, is introduced:

$$P_{\alpha}(\mathcal{D}, f) = \frac{1}{|\mathcal{D}|} \sum_{x^{(i)}, y^{(i)} \in \mathcal{D}} \left[\alpha + \text{sign}(y^{(i)} - f(x^{(i)})) \right]^2 \left(f(x^{(i)}) - y^{(i)} \right)^2 \quad (2)$$

where $\alpha \in [-1, 1]$ controls the asymmetry such that:

$$\alpha \begin{cases} \in [-1, 0[, & \text{penalization of underestimation,} \\ = 0, & \text{symmetric loss,} \\ \in]0, 1], & \text{penalization of overestimation.} \end{cases} \quad (3)$$

This performance measure is the squared error multiplied by a factor that depends on the parameter α and on the over- or underestimation of the true label via the *sign* function. As motivated in the third requirement, the asymmetry should depend on the true label values. Therefore, y is connected with α by introducing α^* such that $\alpha = \frac{y^{(i)}}{4} \alpha^*$ where $y^{(i)} \in \mathcal{Y} = \{-4, \dots, 4\}$, hence $\alpha^* \in [0, 1]$. The constant denominator 4 is used to link α and α^* in such a way that the sign of α that governs the direction of the asymmetric penalization is controlled by the true labels y . This leads to the formalization:

$$P_{\alpha^*}(\mathcal{D}, \hat{f}) = \frac{1}{|\mathcal{D}|} \sum_{x^{(i)}, y^{(i)} \in \mathcal{D}} \left[\frac{y^{(i)}}{4} \alpha^* + \text{sign}(y^{(i)} - \hat{f}(x^{(i)})) \right]^2 \left(\hat{f}(x^{(i)}) - y^{(i)} \right)^2 \quad (4)$$

The parameter $\alpha^* = 0.25$ was set based on the feedback of the involved medical experts⁵. The model will be heavily penalized for the overestimation of negative labels and for the underestimation of positive labels. For instance, the performance measure for $y^{(i)} = 2$ and prediction $\hat{y}^{(i)} = 1$ is higher (1.265) than for $\hat{y}^{(i)} = 3$ (0.765). The asymmetry of the measure is reciprocally connected to the magnitude of the label y in both, the negative as well as the positive direction, e.g. for $y^{(i)} = 1$ it is more symmetric than for $y^{(i)} = 3$. Furthermore, P_{α^*} collapses to a regular quadratic error for $y^{(i)} = 0$. The behavior of the measure is further illustrated in Figure 2.

3.3 Leave-One-Subject-Out Validation

Proposed models are expected to perform well on data from patients not seen before. Using regular cross validation (CV) strategies, subject-specific information

⁵ Feedback was collected by comparing multiple cost matrices as shown in Figure 3.

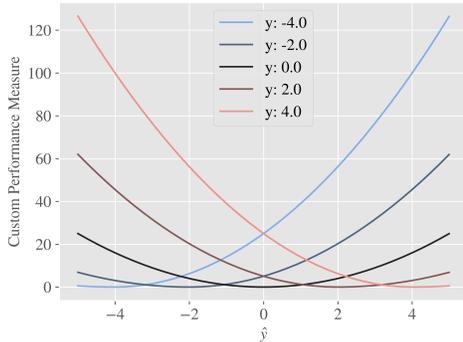


Fig. 2. Behavior of the performance measure $P_{\alpha^*=0.25}$ on the y-axis for different labels y and the corresponding predictions \hat{y} on the x-axis.

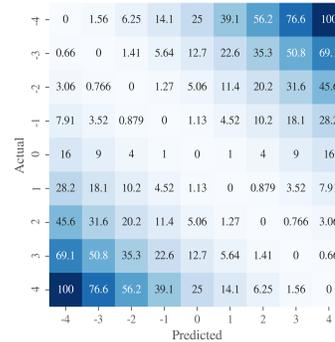


Fig. 3. Cost factors resulting from $P_{\alpha^*=0.25}$ that are associated with each combination of actual and predicted values.

could be exploited resulting in an overly optimistic estimate of the generalization performance [37]. Consequently, a leave-one-subject-out (LOSO) validation scheme is often applied in settings where much data are gathered from few subjects [2, 12, 9]. Thereby, a model is trained on all except one subject and then tested on the left out subject, yielding an unbiased performance estimate. This is repeated for each individual subject and all resulting estimates are averaged.

The usage of early stopping [17] requires the introduction of a tuning step to determine the optimal amount of training epochs e^* in each of the LOSO folds, which in turn requires a second inner split of the data set. In a setting with unlimited computational resources, one would run a proper LOSO validation in the inner folds, determine e^* , train the model on the whole data except the left out subject and evaluate the trained model on that subject. With a total amount of 28 patients, this would result in the training of $28 \cdot 27 = 756$ models for the validation of one specific architecture. As a cheaper solution, the first 80% one minute windows per patient are used for training and the last 20% for early stopping.

3.4 Transfer Learning

One of the most important requirements for the successful training of deep neural networks with strong generalization performance is the availability of a large amount of train data. Next to strong regularization and data set augmentation, one prominent method to fight overfitting and improve the model’s generalization performance is transfer learning [43]. A model architecture is first trained on source task \mathcal{D}_A . The learned knowledge, manifested in the model’s weights, is used to initialize a model that should be trained on the target task \mathcal{D}_B . The model is then fine-tuned on \mathcal{D}_B which often leads to faster model convergence and, dependent on the similarity of the tasks, to an improvement in model performance. Though TSC is still an emerging topic in the deep learning community,

first investigations into the adoption of transfer learning to time series data have been made [11].

As a source task for the motor state detection, we train the model to classify between one-minute windows that were either gathered from PwP or from healthy patients. Therefore, we use a weakly labeled data set that contains 70175 one-minute windows of sensor data along with the binary target if the corresponding patient suffers from Parkinson’s disease or not. Among those patients, 50% were healthy and 50% suffered from PD. The proposed deep learning models were trained on this task and their weights were used for initialization of the models which were then fine-tuned on the actual data as described in Section 5.

4 Problem Setting and Models

4.1 Problem Setting

As explained in Section 2, the target was measured on a discrete scale $y \in \mathcal{Y} = \{-4, \dots, 4\}$ where $y = -4$ represents severe bradykinesia, $y = 0$ the ON state and $y = 4$ severe dyskinesia. This gives rise to the question whether the problem should be modeled as a classification, an ordinal regression or a regression task. The majority of previous research in this domain treats the problem as binary sub-problems with the goal to just detect whether the PwP experience symptoms, regardless of their severity. The granular labeling scheme used here follows an ordinal structure. For instance, a patient with $y = -4$ suffers from more severe bradykinesia than one with $y = -3$. In contrast, simple multi-class classification treats all class labels as if they were unordered. A simple way of including this ordinal information is to treat the labels as if they were on a metric scale and apply standard regression methods. However, this implies a linear relationship between the levels of the labels. For example, a change in the motor state from $y = -4$ to $y = -3$, $\delta_{-4,-3}$, could have a different meaning than $\delta_{-2,-1}$, though they would be equivalent on a metric scale. The formally correct framing of such problems is ordinal regression which takes into account the ordered structure of the target but does not make the strong linearity assumption [18]. This model class is methodologically located at the intersection of classification and metric regression. All three problem settings are compared in Section 5.

4.2 Models

Random Forest A Random Forest [3] was trained on manually extracted features from the raw sensor data, similar to related literature [20, 9, 24, 38]. From each sample window of both signal norms, a total of 34 features such as mean, variance and energy were extracted (a complete list can be found in the digital Appendix). This is a standard procedure in TSC [6, 4]. The Random Forest was specifically chosen as a machine learning baseline due to its low dependency on hyperparameter settings and its strong performance in general.

FCN The Fully Convolutional Net (FCN) was introduced as a strong baseline deep learning architecture for TSC [42]. The implementation resembles that of Wang et al. except that the final layer consists of $|\mathcal{Y}| = 9$ or 1 neuron(s) for classification and regression, respectively.

FCN Inception Inception modules led to substantial performance increases in computer vision and are motivated by the observation that the kernel size of the convolutional layers are often chosen rather arbitrarily by the deep learning practitioner [39]. The rationale is to give the model the opportunity to choose from different kernel sizes for each convolutional block and distribute the amount of propagated information amongst the different kernels. One inception module consists of branches with kernel sizes 1, 5, 7 and 13 respectively and a depth of 64 each, plus one additional max-pooling branch with a kernel size of 3, followed by a convolution block with depth 64 and a kernel size 1. The final FCN Inception architecture essentially follows the original FCN with simple convolutional layers being replaced by 1D inception modules.

FCN ResNet Similar to the inception modules, the introduction of residual learning has met with great enthusiasm in the deep learning community [22]. The main advantage of such Residual Networks (ResNet) over regular CNNs is the usage of skip-connections between subsequent layers. These allow the information to flow around layers and skip them in case they do not contribute to the model performance, which makes it possible to train much deeper networks. Unlike inception modules, this model class was already adapted for TSC and proven to be a strong competitor for the original FCN [42]. The FCN ResNet was shown to outperform the standard FCN especially in multivariate TSC problems [10]. Others argue that the ResNet is prone to overfitting and thus found it to perform worse than the FCN [42]. For the comparison in Section 5, three residual modules are stacked where each of the modules is identical to the standard FCN in order to provide comparability among architectures. The module depths were chosen as proposed by Wang et al. [42].

FCN Broad Pathologically, the disease severity changes rather slowly over time. Thus, it can be hypothesized that additional input information and a broader view on the data could be beneficial for the model. This model is referred to as FCN Broad and includes the following extension: the raw input data from the previous sample window x_{t-1} and the following sample window x_{t+1} are padded to the initial sample window x_t , which results in a channel depth of 6 for the input layer.

FCN Multioutput A broad variety of techniques for ordered regression exist [23, 13, 5, 33]. As a neural network based approach for ordered regression is required, a simple architecture is to create a single CNN, which is trained jointly on a variety of binary ranking-based sub-tasks [33]. A key element to allow the network to exploit the ordinal structure in the data is a rank-based transformation of labels. The categorical labels $y \in \mathcal{Y}$ are transformed into $K = |\mathcal{Y}| - 1$ rank-based labels by:

$$y_k^{(i)} = \begin{cases} 1, & \text{if } y^{(i)} > r_k \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

where r_k is the rank for the k -th sub-problem for $k \in \{1, \dots, K\}$. Following this label transformation, a multi-output CNN architecture is proposed where each of the K outputs refers to one binary ranking-based sub-task. These are optimized jointly on a single CNN corpus. Thus, the sub-task k is trained on a binary classification problem minimizing the binary cross entropy loss. The total model output consists of K probability outputs for each input sample. In order to train the CNN jointly on those sub-tasks, the individual losses are combined to one cumulative loss:

$$L^{\text{ranks}}(y^{(i)}, f(x^{(i)})) = \sum_{k=1}^K L_k^b(y_k^{(i)}, \hat{y}_k^{(i)}) \quad (6)$$

where L_k^b is the binary cross-entropy loss for sub-task output $\hat{y}_k^{(i)}$. For inference, the K outputs are summed up such that $\hat{y}^{(i)} = \sum_{k=1}^K \hat{y}_k^{(i)} - 4$, where the scalar 4 is subtracted from the sum over all probability outputs to map the predictions back to the initial label scale, yielding a continuous output.

FCN Ordinal A second ordinal regression model was created by training a regular FCN with an additional distance-based weighting factor in the multi-class cross entropy loss L^m :

$$L^{\text{ordinal}}(y^{(i)}, f(x^{(i)})) = |y^{(i)} - \hat{y}^{(i)}| \cdot L^m(y^{(i)}, \hat{y}^{(i)}) \quad (7)$$

This way, the model is forced to learn the inherent ordinal structure of the data as it is penalized higher for predictions that are very distant to the true labels.

5 Results

The models described in Section 4 were implemented in pytorch [34]. Model weights were initialized by Xavier-uniform initialization [15] and ADAM [28] (learning rate = 0.00005, $\beta_1 = 0.9$, $\beta_2 = 0.99$) was used for training with a weight decay of 10^{-6} . The performances of the models were compared in a LOSO evaluation as discussed in Section 3.3, using the performance measure $P_{\alpha^*=0.25}$ as introduced in Section 3.2. Finally, the sequence of motor state predictions is smoothed via a Gaussian filter whose μ and σ parameters were optimized using the same LOSO scheme that was used for model training. The results are summarized in Table 2. An additional majority voting model which constantly predicts $\hat{y} = 0$ is added as a naive baseline.

The FCN was applied in all three problem settings. From Table 2, one can observe that regression performs better than ordered regression and classification.

Table 2. Results for different models in multiple problem settings, measured using the performance measure introduced in Section 3.2 evaluated by LOSO validation. Additional commonly used performance measures are shown for completeness where the MAE is reported in a class-weighted (MAE w.) and a regular version and Acc. ± 1 refers to accuracy relaxed by one class level.

Frame	Model	$P_{\alpha^*=0.25}$	F1	Acc.	Acc. ± 1	MAE w.	MAE
Baseline	Majority vote	2.900	0.293	0.702	0.463	0.661	0.960
Classification	FCN	0.800	0.366	0.809	0.340	0.312	0.890
	Random Forest	1.542	0.394	0.802	0.459	0.465	0.802
Ordinal	FCN	0.752	0.321	0.767	0.302	0.311	0.985
	Multioutput FCN	0.922	0.361	0.820	0.352	0.344	0.873
Regression	FCN	0.635	0.346	0.843	0.338	0.293	0.836
	FCN Inception	0.726	0.380	0.841	0.370	0.304	0.842
	FCN ResNet	0.841	0.334	0.809	0.309	0.336	0.924
	FCN Broad	0.673	0.347	0.835	0.339	0.294	0.852
	Random Forest	1.310	0.411	0.848	0.436	0.423	0.760

Similar results were obtained for the Random Forest baseline, where regression is superior to classification. It seems that the simple assumption of linearity between labels does not have a derogatory effect and a simpler model architecture as well as training process is of larger importance.

The comparison of the deep learning models with the Random Forest offers another interesting finding. For both, regression and classification, all deep learning models outperform the classic machine learning models. This finding justifies the focus on deep learning approaches and is in line with previous research discussed in the Introduction.

Niu et al. [33] claim that the Multioutput CNN architecture outperforms regular regression models in ordinal regression tasks. This can not be supported by the current results as the Multioutput FCN shows weaker performance than each of the deep learning architectures in the regression frame.

Looking at the results from the regression setting, one can observe that the simple FCN manages to outperform all more complex architectures as well as the Random Forest baseline. This could be explained by the increased complexity of these models: the FCN consists of 283,145 weights, while the FCN Inception contains 514,809 and the FCN ResNet 512,385 weights. This problem is aggravated by the limited amount of training data.

As shown in Table 3, the transfer learning approach consistently improved the performance of all tested FCN architectures. This strategy also helped to further push the best achieved performance by the regression FCN, making it the overall best performing model. Transfer learning has the biggest effect on the performance of the Multioutput FCN, which indicates that this model requires

a higher amount of training data. This is reasonable as it is arguably the most complex model considered. Further increasing the amount of training data might improve these complex models even more.

Table 3. Performance of the transfer learning approaches compared to their non-pretrained counterparts. Transfer learning consistently improves model performances. Additional commonly used measures are shown for the pretrained models only where the MAE is reported in a class-weighted (MAE w.) and a regular version and ± 1 Acc. refers to accuracy relaxed by one class level.

Frame	Model	$P_{\alpha^*=0.25}$			F1	Acc.	Acc. ± 1	MAE	MAE w.
		regular	transfer	Gain					
Classification	FCN	0.800	0.771	0.029	0.375	0.361	0.813	0.318	0.897
Ordinal	FCN	0.752	0.616	0.136	0.350	0.326	0.802	0.295	0.921
	Multioutput FCN	0.922	0.657	0.265	0.367	0.360	0.829	0.301	0.857
Regression	FCN	0.635	0.600	0.035	0.407	0.388	0.870	0.273	0.772

Some resulting predictions⁶ from the best performing model are illustrated in Figure 5 and a confusion matrix of the model predictions is shown in Figure 4. It is noteworthy that despite the class weighting scheme and the transfer learning efforts, the final model fails in correctly predicting the most extreme class labels.



Fig. 4. Row-normalized confusion matrix for predictions from the pretrained regression FCN. Predicted continuous scores were rounded to integers. Allowing for deviations of ± 1 (framed diagonal region) yields a relaxed accuracy of 86.9%.

⁶ Results on all patients can be found here: <https://doi.org/10.6084/m9.figshare.8313149.v1>

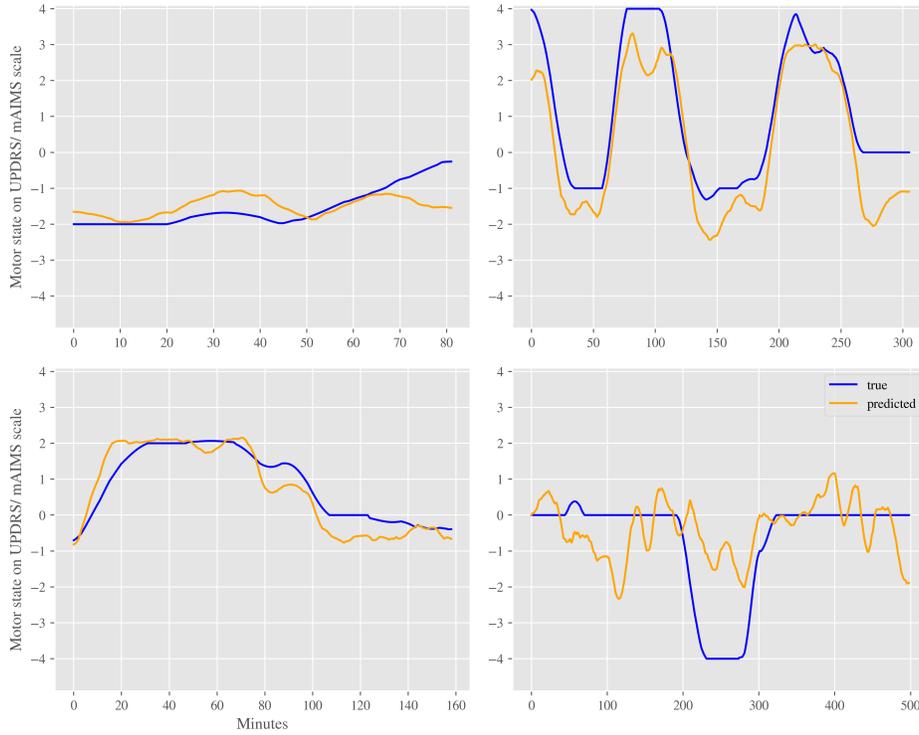


Fig. 5. Comparison of true (blue) and predicted (orange) motor state sequences of four exemplary patients. The label scores are depicted on the y-axis and the minutes on the x-axis. The final model is able to capture the intra-day motor state regime changes of the PwP as shown on the top right plot. Still, the model fails to correctly detect the motor states in some patients e.g. the bottom right one.

6 Conclusion

Different machine learning and deep learning approaches were evaluated on the task to detect motor states of PwP based on wearable sensor data. While the majority of related literature handles the problem as a classification task, the high quality and resolution of the provided data allows evaluation in different problem settings. Framing the problem as a regression task was shown to result in better performance than ordered regression and classification. Evaluation was done using a leave-one-patient-out validation strategy on 28 PwP using a customized performance measure, developed in cooperation with medical experts in the PD domain. The deep learning approaches outperformed the classic machine learning approach. Furthermore, the comparatively simple FCN offered the most promising results. A possible explanation would be that these intricate models call for more available data for successful training. Since high quality labeled

data are scarce and costly in the medical domain, this is not easily achievable. First investigations into transfer learning approaches were successfully employed and showed model improvements for the deep learning approaches.

There exists a plethora of future work to investigate. Computational limitations made it impossible to evaluate all possible models in all problem settings as well as investigate recurrent neural network approaches. The successful usage of a weakly labeled data set for transfer learning suggests further research on the application of semi-supervised learning strategies. This work clearly shows the difficulty in fairly and accurately comparing existing approaches, as available data, problem setting and evaluation criteria differ widely between publications. The introduced performance measure could be a step into the right direction and can hopefully become a reasonable standard for the comparison of such models. In future work, one could directly use this performance measure as a loss function to train deep neural networks instead of using it for evaluation only.

Acknowledgements This work was financially supported by ConnectedLife GmbH and we thank the Schoen Klinik Muenchen Schwabing for the invaluable access to medical expert knowledge and the collection of the data set. This work has been partially supported by the German Federal Ministry of Education and Research (BMBF) under Grant No. 01IS18036A, and by an unrestricted grant from the Deutsche Parkinson Vereinigung (DPV) and the Deutsche Stiftung Neurologie.

References

1. Ahlrichs, C., Lawo, M.: Parkinson’s disease motor symptoms in machine learning: A review. *Health Informatics* **2** (2013)
2. Bao, L., Intille, S.S.: Activity recognition from user-annotated acceleration data. *International Conference on Pervasive Computing* pp. 1–17 (2004)
3. Breiman, L.: Random forests. *Machine learning* **45**(1), 5–32 (2001)
4. Casale, P., Pujol, O., Radeva, P.: Human activity recognition from accelerometer data using a wearable device. *Iberian Conference on Pattern Recognition and Image Analysis* pp. 289–296 (2011)
5. Chen, S., Zhang, C., Dong, M., Le, J., Rao, M.: Using ranking-cnn for age estimation. *The IEEE Conference on Computer Vision and Pattern Recognition* (2017)
6. Christ, M., Kempa-Liehr, A.W., Feindt, M.: Distributed and parallel time series feature extraction for industrial big data applications. *arXiv preprint arXiv:1610.07717* (2016)
7. Curtze, C., Nutt, J.G., Carlson-Kuhta, P., Mancini, M., Horak, F.B.: Levodopa is a double edged sword for balance and gait in people with parkinson’s disease. *Movement Disorders* **30**(10), 1361–1370 (2015)
8. Elliott, G., Timmermann, A., Komunjer, I.: Estimation and testing of forecast rationality under flexible loss. *The Review of Economic Studies* **72**(4), 1107–1125 (2005)
9. Eskofier, B.M., Lee, S.I., Daneault, J.F., Golabchi, F.N., Ferreira-Carvalho, G., Vergara-Diaz, G., Sapienza, S., Costante, G., Klucken, J., Kautz, T., et al.: Recent machine learning advancements in sensor-based mobility analysis: deep learning

- for parkinson's disease assessment. *Engineering in Medicine and Biology Society* pp. 655–658 (2016)
10. Fawaz, H.I., Forestier, G., Weber, J., Idoumghar, L., Muller, P.A.: Deep learning for time series classification: a review. *arXiv preprint arXiv:1809.04356* (2018)
 11. Fawaz, H.I., Forestier, G., Weber, J., Idoumghar, L., Muller, P.A.: Transfer learning for time series classification. *arXiv preprint arXiv:1811.01533* (2018)
 12. Fisher, J.M., Hammerla, N.Y., Ploetz, T., Andras, P., Rochester, L., Walker, R.W.: Unsupervised home monitoring of parkinson's disease motor symptoms using body-worn accelerometers. *Parkinsonism & related disorders* **33**, 44–50 (2016)
 13. Frank, E., Hall, M.: A simple approach to ordinal classification. *European Conference on Machine Learning* pp. 145–156 (2001)
 14. Ghika, J., Wiegner, A.W., Fang, J.J., Davies, L., Young, R.R., Growdon, J.H.: Portable system for quantifying motor abnormalities in parkinson's disease. *IEEE Transactions on Biomedical Engineering* **40**(3), 276–283 (1993)
 15. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks pp. 249–256 (2010)
 16. Goetz, C.G., Tilley, B.C., Shaftman, S.R., Stebbins, G.T., Fahn, S., Martinez-Martin, P., Poewe, W., Sampaio, C., Stern, M.B., Dodel, R., et al.: Movement disorder society-sponsored revision of the unified parkinson's disease rating scale (mds-updrs): scale presentation and clinimetric testing results. *Movement Disorders*: **23**(15), 2129–2170 (2008)
 17. Goodfellow, I., Bengio, Y., Courville, A.: *Deep learning* (2016), <http://www.deeplearningbook.org>, book in preparation for MIT Press
 18. Gutierrez, P.A., Perez-Ortiz, M., Sanchez-Monedero, J., Fernandez-Navarro, F., Hervas-Martinez, C.: Ordinal regression methods: survey and experimental study. *IEEE Transactions on Knowledge and Data Engineering* **28**(1), 127–146 (2016)
 19. Hammerla, N.Y., Halloran, S., Ploetz, T.: Deep, convolutional, and recurrent models for human activity recognition using wearables. *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence* (2016)
 20. Hammerla, N.Y., Fisher, J., Andras, P., Rochester, L., Walker, R., Plötz, T.: Pd disease state assessment in naturalistic environments using deep learning. pp. 1742–1748 (2015)
 21. He, H., Garcia, E.A.: Learning from imbalanced data. *IEEE Transactions on Knowledge & Data Engineering* (9), 1263–1284 (2009)
 22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015)
 23. Herbrich, R., Graepel, T., Obermayer, K.: Support vector learning for ordinal regression. *International Conference on Artificial Neural Networks* (1999)
 24. Hssayeni, M.D., Burack, M.A., Ghoraani, B.: Automatic assessment of medication states of patients with parkinson's disease using wearable sensors pp. 6082–6085 (2016)
 25. Hssayeni, M.D., Burack, M.A., Jimenez-Shahed, J., Ghoraani, B., et al.: Wearable-based medication state detection in individuals with parkinson's disease. *arXiv preprint arXiv:1809.06973* (2018)
 26. Keijsers, N.L., Horstink, M.W., Gielen, S.C.: Automatic assessment of levodopa-induced dyskinesias in daily life by neural networks. *Movement Disorders* **18**, 70–80 (2003)
 27. Keijsers, N.L., Horstink, M.W., Gielen, S.C.: Ambulatory motor assessment in parkinson's disease. *Movement Disorders*: **21**, 34–44 (2006)

28. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2014), <http://arxiv.org/abs/1412.6980>
29. Lane, R.D., Glazer, W.M., Hansen, T.E., Berman, W.H., Kramer, S.I.: Assessment of tardive dyskinesia using the abnormal involuntary movement scale. *Journal of Nervous and Mental Disease* (1985)
30. Lonini, L., Dai, A., Shawen, N., Simuni, T., Poon, C., Shimanovich, L., Daeschler, M., Ghaffari, R., Rogers, J.A., Jayaraman, A.: Wearable sensors for parkinson's disease: which data are worth collecting for training symptom detection models. *NPJ digital medicine* (2018)
31. Marsden, C.D., Parkes, J.: "on-off" effects in patients with parkinson's disease on chronic levodopa therapy. *The Lancet* **307**(7954), 292–296 (1976)
32. Microsoft: Microsoft band 2 wearable device (2018), <https://www.microsoft.com/en-us/band>
33. Niu, Z., Zhou, M., Wang, L., Gao, X., Hua, G.: Ordinal regression with multiple output cnn for age estimation. *Proceedings of the IEEE conference on computer vision and pattern recognition* pp. 4920–4928 (2016)
34. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
35. Postuma, R.B., Berg, D., Stern, M., Poewe, W., Olanow, C.W., Oertel, W., Obeso, J., Marek, K., Litvan, I., Lang, A.E., et al.: Mds clinical diagnostic criteria for parkinson's disease. *Movement Disorders* **30**(12), 1591–1601 (2015)
36. Pringsheim, T., Jette, N., Frolkis, A., Steeves, T.D.: The prevalence of parkinson's disease: A systematic review and meta-analysis. *Movement disorders* **29**(13), 1583–1590 (2014)
37. Saeb, S., Lonini, L., Jayaraman, A., Mohr, D.C., Kording, K.P.: The need to approximate the use-case in clinical machine learning. *Gigascience* **6**(5), 1–9 (2017)
38. Sama, A., Perez-Lopez, C., Romagosa, J., Rodriguez-Martin, D., Catala, A., Cabestany, J., Perez-Martinez, D., Rodriguez-Molinero, A.: Dyskinesia and motor state detection in parkinson's disease patients with a single movement sensor. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society* pp. 1194–1197 (2012)
39. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S.E., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. *Conference on Computer Vision and Pattern Recognition* pp. 1–9 (2015)
40. Tsipouras, M.G., Tzallas, A.T., Fotiadis, D.I., Konitsiotis, S.: On automated assessment of levodopa-induced dyskinesia in parkinson's disease. *Engineering in Medicine and Biology Society* pp. 2679–2682 (2011)
41. Um, T.T., Pfister, F.M.J., Pichler, D.C., Endo, S., Lang, M., Hirche, S., Fietzek, U., Kulić, D.: Parkinson's disease assessment from a wrist-worn wearable sensor in free-living conditions: Deep ensemble learning and visualization. *arXiv preprint arXiv:1808.02870* (2018)
42. Wang, Z., Yan, W., Oates, T.: Time series classification from scratch with deep neural networks: A strong baseline. *International Joint Conference on Neural Networks* pp. 1578–1585 (2017)
43. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? *Advances in neural information processing systems* pp. 3320–3328 (2014)
44. Zeng, M., Nguyen, L.T., Yu, B., Mengshoel, O.J., Zhu, J., Wu, P., Zhang, J.: Convolutional neural networks for human activity recognition using mobile sensors. *Mobile Computing, Applications and Services (MobiCASE), 2014 6th International Conference on* pp. 197–205 (2014)