

Deep neural architectures for prediction in healthcare

Dimitrios Kollias¹ · Athanasios Tagaris¹ · Andreas Stafylopatis¹ ·
Stefanos Kollias²  · Georgios Tagaris³

Received: 16 June 2017 / Accepted: 2 November 2017 / Published online: 13 November 2017
© The Author(s) 2017. This article is an open access publication

Abstract This paper presents a novel class of systems assisting diagnosis and personalised assessment of diseases in healthcare. The targeted systems are end-to-end deep neural architectures that are designed (trained and tested) and subsequently used as whole systems, accepting raw input data and producing the desired outputs. Such architectures are state-of-the-art in image analysis and computer vision, speech recognition and language processing. Their application in healthcare for prediction and diagnosis purposes can produce high accuracy results and can be combined with medical knowledge to improve effectiveness, adaptation and transparency of decision making. The paper focuses on neurodegenerative diseases, particularly Parkinson's, as the development model, by creating a new database and using it for training, evaluating and validating the proposed systems. Experimental results are presented which illustrate the ability of the systems to detect and predict Parkinson's based on medical imaging information.

Keywords Deep learning · Convolutional recurrent neural networks · Prediction · Adaptation · Clustering · Parkinson's · Healthcare

Introduction

Current biomedical signal analysis, including medical imaging, is based on signal processing for feature extraction, segmentation, quantitative and qualitative analysis. Recent advances in Machine Learning and Deep Neural Networks (DNNs) have boosted state-of-the-art performance in all related signal processing tasks. DNNs are the state-of-the-art in machine learning and big data analytics, being used in a large number of applications, ranging from defence and surveillance to human computer interaction and question answering systems [12, 21, 22]. DNNs can also be applied as end-to-end-architectures which are composed of different network types and are trained to analyse signals, images, text and other inputs [12, 19]. However, they lack on-line adaptation capability and transparency in decision making. This makes their use difficult in fields such as healthcare, where personalisation and trust are key issues.

The current paper aims at advancing the state-of-the-art, by developing and using DNNs able to perform effective analysis of complex data for healthcare, with focus on neurodegenerative diseases, in particular Parkinson's [3, 11]. For Parkinson's disease (PD), we have the required medical support and expertise and a new public dataset, which enables us to design an end-to-end neural architecture and platform that can be adaptable to patient-specific data. We describe a novel DNN system evaluated on a rich public Parkinson's dataset, which can serve as a model for many other related fields.

Whilst Parkinson's will provide the test-bed for the proposed end-to-end deep neural system, this system will provide an extensible handle for other neurodegenerative diseases. This aligns directly with the Pathway Analysis across Neurodegenerative Diseases described in [16], as 'there is clinical, genetic and biochemical evidence that similar

✉ Stefanos Kollias
skollias@lincoln.ac.uk

¹ School of Electrical and Computer Engineering, National Technical University Athens, Athens, Greece

² School of Computer Science, University of Lincoln, Lincoln, UK

³ Department of Neurology, Georgios Gennimatas General Hospital, Athens, Greece

molecular pathways are met in different neurodegenerative diseases: Alzheimer’s and dementias, Parkinson’s and related disorders, Huntington’s, motor neuron, prion, spinocerebellar ataxia and spinal muscular atrophy’.

The target of this paper was to design and implement end-to-end deep neural architectures that can assist doctors and clinicians in providing improved and more accurate predictions and assessments, while overcoming existing limitations. Focusing on a specific healthcare problem, we design DNN systems integrating imaging, demographic/epidemiological and clinical data, to support doctors in patient-specific prediction and assessment. To achieve this goal, we present a novel approach, developing a combined supervised and unsupervised learning methodology. First, data-driven supervised training of deep neural networks is performed and, then, clustering of the derived network structures is applied to improve the derived results and allow adaptation and handling of new subject cases.

Section “Generation of the Parkinson’s database” presents the new Parkinson’s database, that we have been developing, providing the necessary datasets for training and testing the developed deep neural network systems. Section “Design of deep neural architectures for healthcare” describes the design of DNN architectures for prediction and diagnosis in healthcare applications. The proposed deep neural systems are based on deep Convolutional (CNN) and Recurrent Neural Networks (RNN), which prove to be able to process all types of available data. A novel methodology for network adaptation when facing new subjects, for personalised assessment, as well as for providing transparency to the network’s performance, is presented in Section “A novel method for deep neural network adaptation and transparency”. An experimental study, illustrating the performance of the generated deep neural architectures, is provided in Section “Experimental

study”. Conclusions and further planned work are given in Section “Conclusions and further work” of the paper.

Generation of the Parkinson’s database

We have been creating a novel public dataset, currently composed of 55 patients with Parkinson’s and 23 subjects with Parkinson-related syndromes, including subjects’ MRI, DaT Scans and clinical data. Our target is that the database soon includes 100 patients’ and 40 non-patients’ data. The database is becoming publicly available as *Parkinson Dataset-v1* [27].

MRI data The rapid evolution of non-invasive medical imaging techniques, over the past decades, has opened new possibilities for the analysis of the brain. The basic imaging technique is Magnetic Resonance Imaging (MRI) which can yield from hundreds to even thousands of images per scan. The assessment of this extremely large set of images per patient can be complicated and time-consuming for doctors. In Parkinson’s Disease, the MRI can show the extent to which the different structures of the brain have been degenerated. Figure 1 shows an example of an MRI. Our main interest regarding Parkinson’s is the lentiform nucleus (green line in Fig. 2) and the capita of the caudate nucleus (red line in Fig. 2). Since we focus on volume estimation, we process the image sequences in batches, each composed of 3–4 consecutive frames.

DaT scan The second brain imaging technique included in the database is Dopamine Transporters (DaT) Scan. This examination is a form of Single-Photon Emission Computer Tomography (SPECT) with Ioflupane Iodide-123 as it is contrast agent. In this examination, we can detect the extent of dopaminergic innervations to the Striatum from the Substan-




Fig. 1 A frame of an axial T1 sequence from a brain MRI (right). Location of the previous slice is placed with regard to a sagittal view of the brain (left)




Fig. 2 An image from an axial T1 sequence. The lentiform nucleus is depicted with a green line, while the capita of the caudate nucleus with a red line

tia Nigra. A series of images is produced in this way, as shown in Fig. 3.

The doctor selects the most representative ones (the 8th in the sequence of Fig. 3), and marks the areas corresponding to the head of the caudate nucleus. An automated system then compares these areas with a neutral one (e.g., the cerebellum) and produces the ratios shown at the bottom of Fig. 4. Diagnosis is based on comparison of these ratios with normal ones.

Clinical data These define the patient’s clinical status. We focused on the following scales: UPDRS, the patient’s stage, UDysRS, PDQ-39, FOG, MMSE and two, timed tests [4].

The Unified Parkinson’s Disease Rating Scale (UPDRS) [9] is a metric that examines the patient’s whole clinical per-

formance in 4 parts: motor/non-motor experiences of daily living, motor examination and complications. These contain 13, 13, 18 and 6 elements, respectively, with each ranging from 0 to 4 for a max score of 234.

The patient’s stage [14] represents the evolution of the disease and ranges from asymptomatic (0) to bedridden (5).

The Unified Dyskinesia Rating Scale (UDysRS) [10] was created for evaluating the involuntary movements associated with PD; it has two parts measuring the dyskinesia and dystonia appearing “on” and “off” phases, respectively. The first part has 11 while the second 15 elements, all ranging from 0 (asymptomatic) to 4 (severe symptoms), for a total of 150.


The Parkinson’s Disease Questionnaire consists of 39 questions assessing patient’s functionality and quality of life (PDQ-39) [17]. It can be separated into 8 different categories, while each question represents the frequency of a specific incident, ranging from 0 (never occurring) to 5 (always occurring), for a total of 156.

The “Freezing of Gait” (FOG) [8] is one of the most characteristic PD symptoms. The quantification of this symptom is achieved through the homonymous questionnaire which contains 16 elements for a max rating of 24.

The Mini Mental State Examination (MMSE) [24] is an 11-question questionnaire meant to measure the cognitive impairment associated with PD, with a max rating of 30.

Each of MRI and DaT Scan sets includes sequences/multiple scans. For training, we combine annotated data from both types to create thousands of input data, sufficient to train the proposed systems.

Fig. 3 A sequence of frames from a DaT scan






Fig. 4 DaT scan with expert selection (left). Same image without the markings (right). Ratios, representing the dopamine deficiency, that are used for the diagnosis (bottom)

Design of deep neural architectures for healthcare

Our main goal is to design deep neural architectures and to evaluate their ability to extract correlations in the available datasets, providing a novel platform for assisting doctors in detecting and assessing disease states. Validation is done using the above-described Parkinson's dataset. We also target at endowing our system with adaptation capabilities and to test and validate it when handling new patient cases.

The technologies which we use and extend, in order to develop the novel end-to-end deep neural architecture for diagnosis and prediction are:

Deep convolutional neural networks Deep CNNs are architectures that try to exploit the spatial structure of input information [12]. They have been used with great success in various applications, including image analysis, vision, object and emotion recognition. The most successful CNN was used for classifying millions of images in 1000 classes [21].

Transfer learning Transfer learning [22] is the main approach to avoid learning failure due to overfitting, when training complex CNNs with small amounts of (image) data. In transfer learning, we use networks previously trained with large image datasets (even of generic objects) and fine-tune the whole, or parts of them, using the small training datasets.

Recurrent neural networks RNNs are very powerful for processing sequential data [18]. A very successful model, the

Long Short-Term Memory (LSTM) [25], uses hidden units with gates that explicitly control data flow in terms of both hidden states and inputs. Bidirectional (B-LSTM) models are obtained by combining forward and backward processing of input data. Gated Recurrent Units (GRUs) [2, 12] can be used in place of the BLSTM ones; they have fewer parameters than LSTMs, since they do not include an output gate. Based on our tests with Parkinson's data, GRUs have produced better performances and are used in the experiments of Section "Experimental study".

We propose an end-to-end deep neural architecture including both CNN and RNN components. CNNs derive rich internal representations from input data; B-LSTM/GRU RNNs correlate/analyse time evolution of the inputs, providing the final predictions. The CNN system we consider follows the basic structure of the so called Deep Residual Net (ResNet), which contains 50 layers [13]. This network has won the first places on the tasks of ImageNet detection, ImageNet localization, COCO (object) detection and COCO (object) segmentation.

Following the convolutional and pooling layers we use up to 3 fully connected layers, with the so-called Rectified Linear Units (ReLU) neuron models, i.e., neurons with a linear activation function, for positive input values, and a zeroing function else-where. Other networks such as VGG-16 (e.g. [23]) could also be used, but they have been mainly designed for human face analysis applications. MRI and DaT

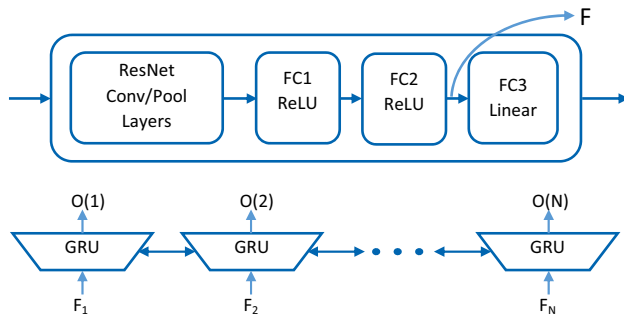


Fig. 5 The CNN part of the CNN-RNN architecture feeds the RNN part which yields the final outputs

Scans are provided at the input of these networks. When epidemiological and clinical data values are to be considered, they will be provided directly to the FC1 layer.

Figure 5 shows the CNN–RNN architecture. The CNN part of the neural architecture, using a linear FC3 layer provides continuous clinical data estimation. The CNN feeds the RNN part with the neuron outputs of its second FC layer (F). The RNN accepts $F_1, F_2, F_3, \dots, F_N$ and delivers predicted values $O(1), \dots, O(N)$ through time, at its output. A total of 4 images are given to the architecture as a single input. These include 3 greyscale consecutive frames from an axial T1 MRI and a colour DaT scan.

To implement this architecture, we first perform transfer learning of the weights of the convolutional and pooling parts of, e.g., the ResNet network to it. These parts are then fixed during the training phase, where we only train the fully connected layers of the system. The pre-trained convolutional networks have already learnt to generate rich image representations that have proven adequate for image classification and segmentation. These representations are abstract enough to help with specialized tasks, such as the analysis of MRIs and DaT Scans.

This leaves the fully connected part of the network, which is the only part of the network that we actually train in the CNN case. Many variants of this approach have been designed and tested. We selected to freeze the weights of some of the fully connected layers, particularly those belonging to the first FC layer. We have also considered some additional weights of the network as free parameters, by applying fine tuning (a smaller learning rate value) to the weights of (some of) the convolutional layers of the ResNet network, while using a normal learning rate value for the FC part of it.

We use the TensorFlow Platform as the main tool for generating the software implementation of the presented architecture. TensorFlow is a toolkit which got published by Google, under Apache License 2.0. It is mainly implemented using C++, with a significant bit of Python. Its architecture provides the ability to deploy computation to one or more

CPUs or GPUs in a desktop, server, or mobile device with a single API.

A novel method for deep neural network adaptation and transparency

We aim at providing the deep neural architecture with the ability to adapt to new subject cases, assisting doctors with efficient patient-specific analysis and treatment selection, without forgetting its former knowledge. Our methodology is based on a new network retraining approach which extends the work in [5, 19]. This approach uses clustering [26] of trained system internal representations, in particular, of the neurons' outputs at the last fully connected CNN layer (denoted, in vector form, as F in Fig. 5), or at the last hidden RNN layer (let us denote them, in vector form, as u , and consider them feeding the output units o). We use the centres of these clusters as knowledge extracted from the data-driven supervised training of the DNN architecture.

Whenever a new subject's data are applied to the input of the DNN end-to-end architecture, the latter computes the respective internal representations and provides a prediction at its output. Our approach is next to compute the distances of these representations from the above described cluster centres and use them to validate, or not, the DNN prediction on these new data. If one of these distances is small, compared to some appropriate threshold, then classification of the new data is made in the same category (patient/non-patient) with that of the specific cluster, generally coinciding with the DNN prediction. If all distances are large, then a drift in the DNN modelling procedure is detected. In the case of drift, we need to train again the DNN including the new data. However, we do not perform the usual fine-tuning procedure. We choose to retrain the fully connected CNN layers and/or the RNN hidden and output layers, using, on the one hand, the input (image) data corresponding to the cluster centres (Existing Knowledge) and, on the other hand, the new data.

Following this retraining procedure, we avoid the catastrophic forgetting problem in DNN systems, which occurs when we apply repeated fine-tuning to new data cases. This is so, because we keep both the old knowledge (through the cluster centres' information) and the new information provided by specific subject cases. Following retraining, we update the cluster centres as well, after medical validation of the new data, so as to create personalised system knowledge instances.

In particular, the retraining procedure can be implemented as follows:

Let us first consider that, based on the training of the deep neural architecture for Parkinson's, a specific set, say S_b , including the training input data corresponding to the previously computed cluster centres and the respective annotations

(patient/non-patient), has been created. Let $y(i)$ denote the network output when applied to a new data sample, $i = 1, 2, \dots$, not included in the previous network training data set.

Let w_b include all already computed weights of the fully connected and output layers in a CNN network—and of hidden layers in a CNN–RNN network—before retraining and w_a the new (updated) weight vector which will be obtained through retraining. In particular, let w_b^l and w_a^l , respectively, denote the weights connecting the outputs of the last hidden layer, say u , to the network outputs, y .

A training set S_t is assumed to include the new input (image) data; this will normally include a rather small number of data.

In the proposed retraining procedure, the new network weights, w_a , are computed by minimizing the following error criterion:

$$E_a = E_{t,a} + \eta \cdot E_{f,a} \quad (1)$$

where $E_{t,a}$ denotes the error performed over training set S_t , i.e., over current input information and $E_{f,a}$ is the corresponding error performed over training set S_b , i.e., over previous deep neural network knowledge. Parameter η is a weighting factor accounting for the significance of the current training set compared to the former one. In our approach, we minimize (1) by assuming that a small perturbation of the weights of the fully connected (and/or hidden) layers in the CNN (or CNN–RNN) network is enough to achieve good classification performance in the current conditions. Consequently, we get:

$$w_a = w_b + \Delta w \quad (2)$$

and, similarly,

$$w_a^l = w_b^l + \Delta w^l \quad (3)$$

with Δw and Δw^l being small weight increments. This assumption permits linearization of the nonlinear activation neuron function, using a first-order Taylor series expansion.

It is possible to use the Mean Square Error (MSE) criterion for both quantities in the right-hand side of (1). In this case, we use normal deep learning for CNN and/or RNN networks [12], implemented in the TensorFlow environment. It can be also possible to stress the importance of current data in the minimization of (1). In this case, we replace the first term in the right-hand side of it by the constraint that the actual network outputs $z_a(i)$, after retraining, are equal to the desired ones, i.e.,

$$z_a(i) = d(i), \text{ for all data } i \text{ in } S_t \quad (4)$$

Let us denote the difference of the actual network outputs, after and before retraining, in the case of a CNN network, as follows:

$$\Delta z(i) = z_a(i) - z_b(i) \quad (5)$$

Through linearization and using the fact that the outputs z are weighted averages of the last hidden layer's outputs u , with the w^l weights, it can be shown that

$$z_a(i) = z_b(i) + f_b' \cdot w_b^l \cdot \Delta u^l(i) + \Delta w^l \cdot u_b^l(i) \quad (6)$$

where f' accounts for the derivative of the activation function of the network output neuron(s).

Using Eq. (4) in (6) we get

$$d(i) - z_b(i) = f_b' \cdot w_b^l \cdot \Delta u^l(i) + \Delta w^l \cdot u_b^l(i) \quad (7)$$

All quantities in Eq. (7) are based on former network values, apart from the updates of the weights Δw^l and of the outputs Δu^l . Thus Eq. (7) relates the targeted weights updates in the network output with the outputs of the last hidden layer.

By continuing linearization of the difference of the u values, towards the previous fully connected layers, we replace the $\Delta u^l(i)$ term with its equivalent in terms of the weights of the former layers. This continues until we reach the last convolutional layer, which we use with no retraining, and therefore Δu is zero.

In this way, similarly to [5] we compute the weight increments Δw by solving a set of linear equations, over all data in S_t :

$$c = A \cdot \Delta w \quad (8)$$

with matrix A being computed in terms of previously trained weights, as was above described, while the elements of vector c are defined as follows:

$$c(i) = d(i) - z_b(i), \text{ for all data } i \text{ in } S_t \quad (9)$$

and $z_b(i)$ denotes the outputs of the originally trained network, when this is applied to the data in S_t .

The size of vector c is smaller than the number of unknown weights Δw , thus many solutions exist for (8). Uniqueness, however, is imposed by an additional requirement which is to select the solution that causes a minimal degradation of the previous network knowledge. This is of great significance in our approach, since this knowledge (and the respective cluster centres) has been, normally, already validated by medical experts and, therefore, should be changed the least possible.

Thus, the retraining problem results in minimization of (1) subject to constraints (3) and the constraint for small weight increments. A variety of methods can be used for this




Fig. 6 MRI scan of a patient without Parkinson's Disease. Axial orientation—T1 sequence




Fig. 7 MRI scan of a patient with Parkinson's Disease. Axial orientation—T1 sequence

minimization. One of them is the gradient projection method, which, starting from a feasible point, moves in a direction which decreases the error criterion and satisfies the above constraints. This is used for CNN network retraining in the TensorFlow environment. Extension in the CNN–RNN case is more complex, also taking into account the time evolution and derivatives of the u values.

In addition to personalized diagnosis and prediction, the proposed approach allows the deep neural architecture to exhibit transparency in its decision making. In particular, for each cluster centre, the respective medical input and desired output data are stored in the database, as representative of all data belonging to this cluster. Whenever, upon presentation of new input data to the DNN, the obtained output vector matches that of a specific cluster centre, then the respective input image and medical data are presented to the clinician/user to illustrate that this similarity has been taken into account by the network in computing its prediction.

Experimental study

The current size of the generated fully annotated database is 78 subjects (over a half of the size to be finally generated), with a ratio of 2:1 between Parkinson's patients and non-Parkinson's patients. At this stage, it consists of MRI and DaT scans, annotated as belonging to subjects with Parkinson's or not.

Dataset generation

We generated a dataset of about 100,000 combinations of color DaT scans with triplets of consecutive MRI gray scale images, covering both patient and non-patient categories. Each input (combination) consists of three MRI images and one RGB DaT scan image. To obtain a balanced dataset, we applied various augmentation techniques, such as over-sampling the latter category, or under-sampling the former [1]. The above were then used as data for designing the end-to-end deep neural architectures.




Fig. 8 DaT scan from a patient without Parkinson's Disease (left). Respective image from a patient with Parkinson's (right)

We used about 70% of this data for training the deep neural architectures. Moreover, we kept the original data (corresponding to the rest 30% of augmented data) of 15 subjects (out of the current 78 in our database) for validation and testing. It should be emphasized that our target has been to test the ability of the networks to learn from a number of patients and generalize their performance to other subjects, who have not been included in the training set. For this reason, the test data consisted of six new subjects, four with Parkinson's (PD patients) and two without (Non-PD patients, denoted NPD), to provide about 1,200 test input samples. The networks had two linear outputs, with targeted values (1,0) and (0,1), respectively, for the two categories.

As a reference, 10 consecutive frames from an axial T1 brain MRI are presented in Fig. 6 for a patient without Parkinson's, and 10 more in Fig. 7 for a patient with the disease.

Figure 8 shows two DaT scans of patients without and with Parkinson's Disease, respectively. The dopamine deficiency can be seen in these images.

Network training

As a first approach, we selected to train the CNN and CNN–RNN deep neural networks from scratch; starting from random initial weights in the convolutional and fully connected (FC) parts of the CNNs, or the convolutional and

Table 1 Performance (on test data) of the trained end-to-end CNN architecture for Parkinson's

CNN architectures: 2 output units (PD/NPD)	Number of fully connected (FC) layers	Number of units in each FC Layer	Accuracy
1	1	1000	0.57
2	1	2622	0.60
3	2	2622–500	0.90
4	2	2622–1000	0.91
5	2	2622–1500	0.94
6	2	2622–2000	0.93

hidden layers of the CNN–RNNs. As a second approach, we adopted transfer learning, i.e., transfer of the weights of the convolutional and pooling layers of a pretrained CNN, to the generated networks. Then, the ‘upper’ FC part of the targeted CNN network, as well as the RNN hidden layers of the CNN–RNN, were designed and trained with the above dataset. For the initialization of these weights, we used the ResNet-50 CNN, which has been pre-trained with millions of general type RGB images for this purpose. A separate system was used for each of the image types in our inputs, i.e., one focusing on the MRI triplets and another focusing on the DaT scan. We concatenated the outputs of these two ResNet substructures at the input of the first FC layer of the CNN network. It is at this layer, that epidemiological data will be concatenated as well, when the whole database will have been generated.

Based on this procedure, we separately trained both a deep CNN network and a deep CNN–RNN network for Parkinson's disease diagnosis.

Experimental evaluation


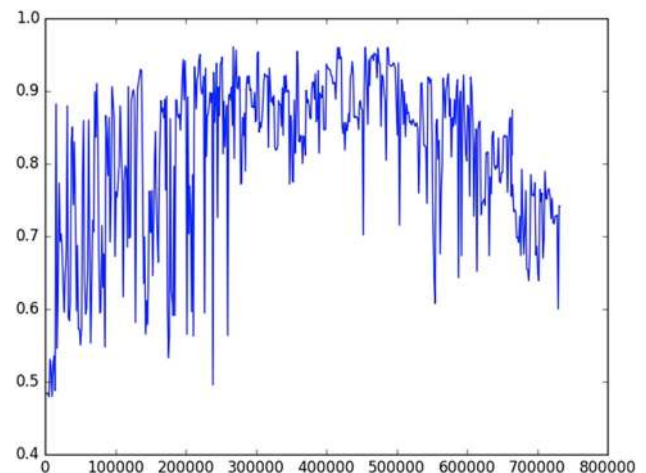
Table 1 summarizes the results obtained through different configurations of the CNN network, i.e., ones with different numbers of hidden layers and hidden units per layer. An accuracy of 96% on training data set was obtained (with network weights selected based on the performance on the validation data set); an accuracy of 94% on testing dataset was obtained in this experiment, as shown in Table 1, which is very satisfactory.

Table 2 summarizes the accuracy obtained by the CNN–RNN (with GRU neuron model) architecture, for different respective structures, with weights selected similarly, based on performance on the validation data set). The addition of the RNN part allows the deep neural architecture to better follow time varying correlations in the MRI sequence of triplets of frames, thus increasing the accuracy of Parkinson's prediction to 98% on the testing data set.

There are some additional metrics obtained in terms of the above results. In the best reported case (line 3 of Table 2),

Table 2 Performance on test data of the trained end-to-end CNN–RNN architecture for Parkinson's

CNN–RNN architectures:	Number of units in the FC layer	Accuracy
1 Fully connected layer		
2 Hidden layers (128 units each)		
2 (linear) output units		
1	500	0.91
2	1000	0.96
3	1500	0.98
4	2000	0.97

**Fig. 9** CNN Performance on validation data, during training epochs**Fig. 10** CNN–RNN Performance on validation data, during training epochs

the MSE value was very low, equal to 0.02. Considering the binary problem examined in this paper (PD/NPD), precision attained was 1.00 and recall was 0.96 (F1 value was 0.98).

Figures 9 and 10 show the accuracy obtained by the end-to-end deep CNN and CNN–RNN architectures, respectively, on the validation/test data set, during training. It can be shown

Table 3 Testing performance of trained CNN–RNN Architecture on each subject

Subject number in the database	Category	Correct classifications (normalized [0, 1])
26	PD	0.90
4	PD	1.00
6	PD	0.985
9	PD	0.956
17	NPD	1.00
21	NPD	1.00

that the best accuracy of the CNN architecture is obtained early in the learning phase, afterwards reaching overfitting conditions. It can also be observed that the Deep CNN–RNN architecture takes longer to derive the best performance than the CNN one.

It should be mentioned that the best performance of the CNN–RNN architecture was 99,97% on the training data and 98% on the test data. The test data set consisted of about 1200 input data (original, i.e., not augmented, MRI triplets and DaT Scans) from six subjects; none of their data had been included in the training data set. About 600 data concerned each one of the PD and NPD categories. The performance on test data was 96% for PD and perfect, i.e., 100%, for NPD patients. In particular, Table 3 shows the percentage of correct classifications for each test subject’s data (combinations of MRIs and DaT scans).

This is an excellent result, which shows the potential of the deep CNN–RNN architecture to provide very accurate predictions of Parkinson’s disease.

We then applied the proposed clustering procedure on the representations (vector of neurons’ outputs) generated at the last hidden layer of the trained CNN and CNN–RNN architectures. The best results were obtained with 5 clusters, 3 corresponding to the Parkinson’s Disease (PD) cases and 2 to the Non-Parkinson’s (NPD) case, as described in the next Section.

Clustering visualization

In order to visually illustrate the distribution of data in categories, Principal Component Analysis (PCA) was performed on the representations obtained through processing of the test data. Focus was put on the derived two main principal components, as shown in Figs. 11a and 12a, for the CNN and CNN–RNN architectures, respectively.

Figure 11a shows the distribution of the representations obtained for PD and NPD subjects, as derived from the CNN architecture. It should be mentioned that the last CNN fully connected layer consisted of 1500 neurons. However, due to the ReLU activation function, only about 30 neurons yielded

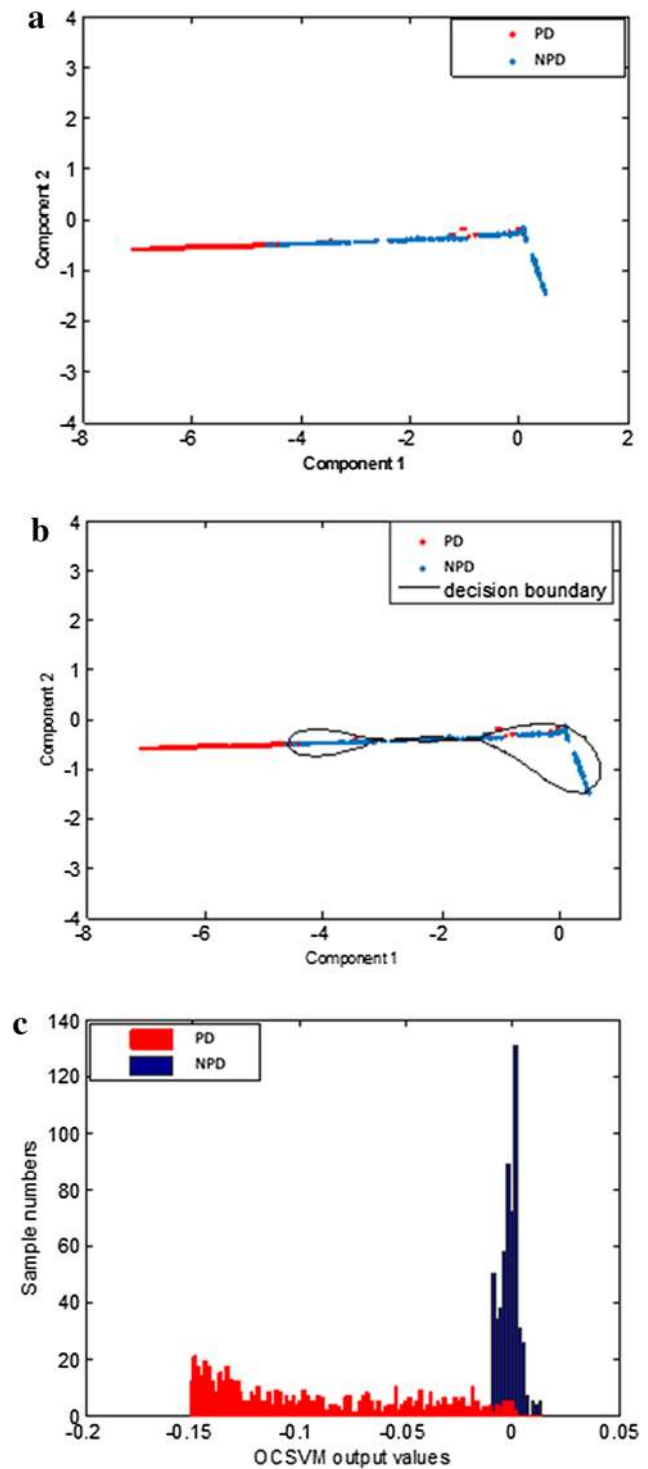


Fig. 11 a The two main principal components of the CNN representation. b Visualization of (three) cluster boundaries for the NPD category provided by an OCSVM approach. c Histogram of the derived OCSVM outputs

non-zero values in this representation. Figure 11b verifies the ability of a one-class support vector machine (OCSVM) [26], to determine clusters corresponding to the NPD class, as shown in Fig. 11b.

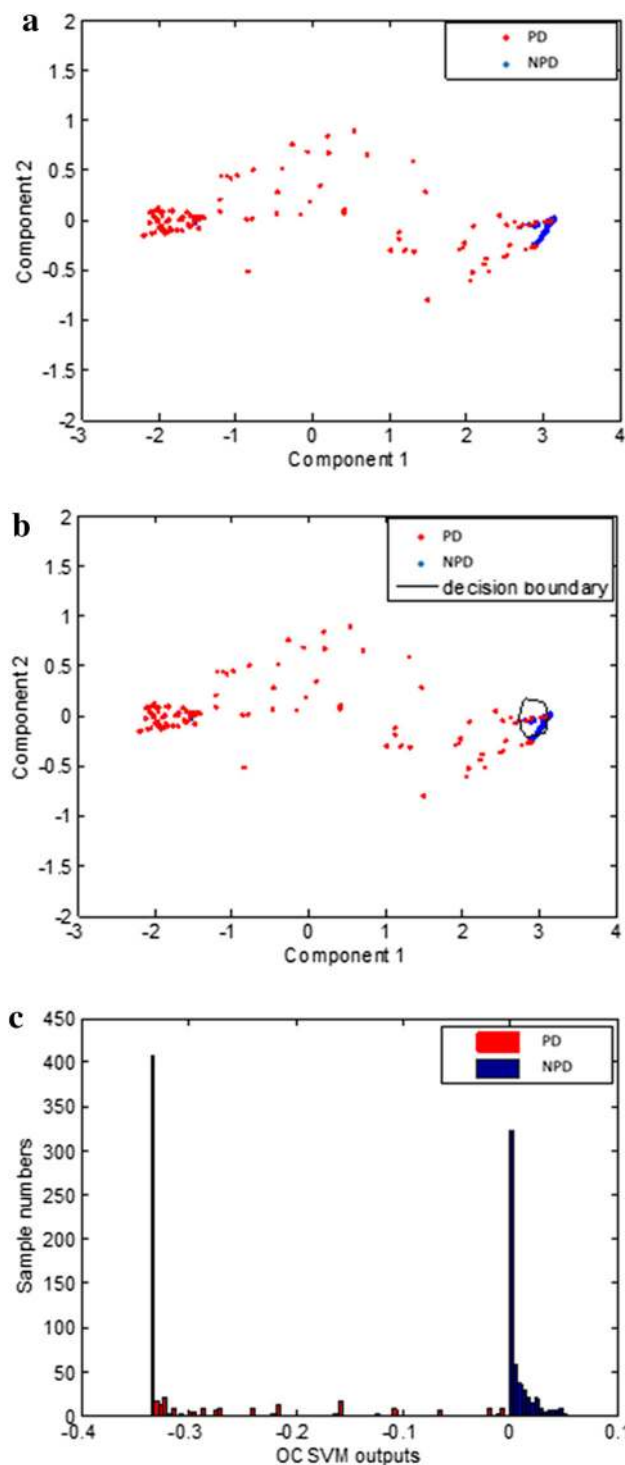


Fig. 12 **a** The two main principal components of the CNN–RNN representation. **b** Visualization of (one) cluster boundary for the NPD category provided by an OCSVM approach. **c** Histogram of the derived OCSVM outputs

It is interesting to mention the variability of the PD cases compared to the NPD ones. This is in accordance with the lower accuracy obtained by the DNN architecture in the PD

Table 4 Cluster precision on the training set

Cluster category	1	2	3	4	5
PD	0	5	18277	1516	18163
NPD	2822	25393	0	0	0




Fig. 13 Projection in 3-D of cluster centres' representations

class, when compared to the NPD case. Figure 11c shows a histogram of the OCSVM values also illustrating this observation.

The respective results obtained for the representations provided by the CNN–RNN architecture are shown in Fig. 12a–c. It should be mentioned that, in this case, the obtained representations consisted of 128 neuron output values, computed through the *tanh* activation function. However, only about 20 of the neurons provided significant non-zero values; the rest yielded very small, practically negligible, values.

By comparing these results with the respective ones in Fig. 11a–c, it is concluded that the CNN–RNN architecture—which has achieved a better performance than CNN—has been able to produce much more compact representations for each category, with well separated clusters.

There were five clusters generated by the proposed approach, three for the PD category and 2 for the NPD one. An indication of the purity (precision) of the clusters in the augmented training data set can be viewed in Table 4. Four clusters have a precision equal to 1.00, with one having a 0.9998 precision.

We computed the cluster centres, as the mean values of all 128-dimensional vector representations included in each cluster. Their projection in 3-D is shown in Fig. 13, showing the significant distance values between them. Moreover, Table 5 shows the corresponding maximum mean squared distance of the representations in each cluster from the corresponding cluster centre.

Table 5 Maximum intra-cluster distance

Cluster	1	2	3	4	5
Distance (MSE)	0.01	0.02	1.565	0.158	0.14

Figure 14a–e illustrate the input images corresponding to the 5 cluster centres that were derived from the CNN–RNN architecture. The clusters have been sorted by the level of degeneration of the basal ganglia (lentiform nucleus, caudate nucleus). The 5 cluster centres roughly represent the 3 stages of DaT loss in PD, as confirmed by medical experts. This provides transparency and is the basis for interpretability of the decision making process implemented and achieved by the proposed deep neural architecture.

Let us now proceed with analysis of new subject data which have not been included in the developed system design phase. Let us consider the test data described in Table 3 for this purpose. Since the six subject cases span different possible scenarios, we will evaluate them in two different steps.

Let us first consider, the 4, 17 and 21 subjects of Table 3 (one from the PD category and two of the NPD category), all data of whom are correctly predicted (100% accuracy) by the CNN–RNN architecture. The internal representations (128-dimensional vectors) generated at the output of the second hidden layer of the RNN were also correctly classified, based on their distances from the centres of the clusters derived from the trained CNN–RNN respective internal representations. All classifications provided by the trained DNN architecture for the data of these three subjects have been, therefore, accepted by our derived end-to-end contextualization approach and formed the finally obtained predictions.

Since the training database has now been increased with three new subject datasets, we can perform an updating of the centres of the clusters to which the new data have been included. Let us assume that a single vector $m[j]$, $j = 1, 128$, is used to update the centre c_i of the i -th cluster composed of N_i members. Then, the new class centre $c_{i,new}$ will be slightly modified, as follows:

$$c_{i,new}[j] = N_i \cdot c_{i,old}[j] / (N + 1) \quad (10)$$

Consequently, an updated, slightly different, system memory is produced, incorporating the new knowledge about the new subjects' data.

Let us now focus on the three remaining cases of Table 3, all referring to PD patients. 10 input combinations, out of 120, 3 out of 204 and 8 out of 184 input combinations, respectively, have been erroneously classified, as NPD cases, both by the CNN–RNN architecture and the cluster-based representation. It should, however, be stressed that in all these cases the distances of the computed representations from the 5 cluster centres have been larger than the respective maximum intra-cluster distances presented in Table 5. This has

been the criterion for considering these cases, as new ones, which require insertion of new cluster centres and retraining of the DNN network with them.

We should mention that, these cases constitute only a 9%, 1.5 and 4.5 of the data obtained by each of these patients, respectively. Thus, we assume that clinician only examines them and provides his/her own diagnosis. Following this validation, two new clusters have been added to the PD existing ones, one of which has been close, but distinct, to the 1st NPD cluster centre and the other close, but distinct, to the 2nd NPD cluster centre.

In addition, we used the adaptation methodology described in Section “A novel method for deep neural network adaptation and transparency” to successfully retrain the DNN architecture so as to accurately classify the new data as well, while keeping the formerly achieved performance. The new dataset in Eqs. (1) and (4) consisted of the above described 21 input data samples. The performance obtained by the network, after weight adaptation, was similar to the one obtained, when retraining the network with all available data in the database.

In all the above experiments, for DNN training, we used the Adam optimizer algorithm, in mini batches, considering the Mean Squared Error (MSE) as cost function.

Hyper-parameter value selection

For the CNN architecture, the hyper-parameter values were selected as follows: a batch size of 30 (15 examples from each category), a constant learning rate of 0.001; 2622 and 1500 hidden units, respectively, in each fully connected layer and dropout after each fully connected layer with a value of 0.5. We also used biases in the fully connected layers.

For the CNN–RNN architectures the hyper-parameters were selected to match the previous ones, apart from the batch size which was 40 (20 examples from each category) and the number of hidden units in the GRU layers, both of which were 128.


The weights of the fully connected layers were initialized from a Truncated Normal distribution with a zero mean and a variance equal to 0.1 and the biases were initialized to 1.

Training was performed on a single GeForce GTX TITAN X GPU and the training time was about 2–3 days.

Conclusions and further work

We have designed novel end-to-end deep neural architectures, composed of CNN and RNN components, appropriately trained with medical imaging data, and have obtained very good performances in diagnosis and prediction of Parkinson's disease. We have been developing and publicizing a new database, which we have used for training and

Fig. 14 **a** The first cluster centre corresponds to a typical frame from a DaT scan of an individual not suffering from PD. **b** The second cluster centre represents an interesting case of an image that seems to be pathological but belongs to a healthy individual. Though the lentiform Nucleus appears to be completely gone, there is no diffusion of the contrast agent in the brain. The latter could be viewed as an indication that the main structures are, in fact, intact. **c** The third cluster represents the early stages (1–2) of the degeneration associated with PD, as both lentiform nuclei appear to be diminishing. **d** The fourth cluster is a typical stage 2 DaT loss. Both lentiform nuclei are completely gone; the only signal is from the caudate, which appear as two almost symmetrical circular areas. **e** The fifth cluster is the most advanced stage of DaT loss, stage 3. Here the basal ganglia appear further degenerated, while there is significant activity in the rest of the brain. This is an indication that these structures have lost their ability to contain the contrast agent and it has diffused throughout the brain



evaluating the performance of the new deep neural architectures.

Moreover, we have proposed a novel unsupervised approach, based on clustering of the trained DNN internal representations, which provides the deep neural architecture

with the ability to adapt to new data cases, without suffering the catastrophic forgetting problem, usually met in DNN fine-tuning adaptation methodologies. This procedure also provides a type of transparency in the decision making process implemented by the deep neural architecture.

In our current research, with the aid of medical experts, we correlate the generated clusters with the medical and clinical data and try to create descriptions relating the DNN decisions with the developed cluster characteristics, as well as using more detailed grading schemes in the data annotation and more categories in the classification task. This is the basis for providing explanations of the network's performance, thus, rendering its use transparent and trustful, while providing more detailed predictions about Parkinson's disease evolution.

A lot of research has been made on neuro-symbolic learning and reasoning, i.e., merging neural networks with knowledge representation, also involving deep neural networks [7, 20] and on extracting rules from trained networks [15]. We will also investigate the use of these methods to provide formal representations of the generated Parkinson's knowledge and/or extract additional rules that may further justify the predictions and assessments of the designed deep neural architectures.

Our future research aims at extending the developments obtained for the Parkinson's case to other degenerative diseases, which are based on similar input medical imaging information. We first target dementias and Alzheimer's, using a recently presented database in [6]. Following the approach proposed in the paper, we will use transfer learning to retrain the DNNs designed for Parkinson's on datasets describing other diseases.

Acknowledgements The work of the NTUA team was financed by the Greek State Scholarships Foundation (IKY) through the "Research Projects for Excellence IKY/Siemens" Programme in the framework of the Hellenic Republic—Siemens Settlement Agreement.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Chawla NV, Japkowicz N, Kotcz A (2004) Special issue on learning from imbalanced data sets. *ACM Sigkdd Explor Newslett* 6(1):1–6
- Cho K, Van Merriënboer B, Bahdanau D, Bengio Y (2014) On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*
- DeMaagd George, Philip Ashok (2015) Parkinson's disease and its management: Part 1: disease entity, risk factors, pathophysiology, clinical presentation, and diagnosis. *Pharm Ther* 40(8):504
- Defer GL, Widner H, Marié RM, Rémy P, Levivier M (1999) Core assessment program for surgical interventional therapies in Parkinson's disease (CAPSIT-PD). *Mov Disord* 14(4):572–584
- Doulamis A, Doulamis N, Kollias S (2000) On-line retrainable NNs: improving the performance of NNs in image analysis problems. *IEEE Trans NNs* 11(1):137–156

- Gao X, Hui R, Tian Z (2017) Classification of CT brain images based on DNNs. *Comput Methods Programs Biomed* 138:49–56
- Garcez AA (2015) Neural-symbolic learning and reasoning: contributions and challenges. *AAAI Spring Symposium*. Stanford University, Stanford
- Giladi N, Shabtai H, Simon ES, Biran S, Tal J, Korczyn AD (2000) Construction of freezing of gait questionnaire for patients with Parkinsonism. *Parkinsonism Relat Disord* 6(3):165–170
- Goetz CG, Tilley BC, Shaftman SR, Stebbins GT, Fahn S, Martinez-Martin P, Dubois B (2008) Movement disorder society-sponsored revision of the Unified Parkinson's Disease Rating Scale (MDS-UPDRS): scale presentation and clinimetric testing results. *Mov Disord* 23(15):2129–2170
- Goetz CG, Nutt JG, Stebbins GT (2008) The unified dyskinesia rating scale: presentation and clinimetric profile. *Mov Disord* 23(16):2398–2403
- Goldman SM, Tanner C (1998) Etiology of Parkinson's disease (1998). In: Jankovic J, Tolosa E (eds) *Parkinson's disease and movement disorders*, 3rd edn. Williams and Wilkins, Baltimore, pp 133–158
- Goodfellow I (2015) Deep learning. *Nature* 521:436–444
- He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 770–778
- Hoehn MM, Yahr MD (1967) Parkinsonism onset, progression, and mortality. *Neurology* 17(5):427–427
- Hu Z (2016) Harnessing deep NNs with logic rules. [arXiv:1603.06318v4](https://arxiv.org/abs/1603.06318v4)
- JPND EU Joint Programme (2017) Neurodegenerative Disease Research. Pathways, <http://jpnd.eu>
- Jenkinson C, Fitzpatrick R, Peto V, Greenhall R, Hyman N (1997) The Parkinson's Disease Questionnaire (PDQ-39): development and validation of a Parkinson's disease summary index score. *Age Ageing* 26(5):353–357
- Kahou SE (2015) Recurrent NNs for emotion recognition in video. *Proc ACM ICM I:2015*
- Kollias D, Tagaris T, Stafylopatis A (2017) On line emotion detection using retrainable deep NNs. In: *IEEE symposium series computational intelligence 2016, IEEE Xplore 13-2-2017*
- Kollias D, Marandianos G, Stafylopatis A (2015) Interweaving deep learning and semantic techniques for HCI. In: *10th International Workshop on Semantics and Adaptation*. Trento, Italy
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*, pp 1097–1105
- Ng HW, Nguyen VD, Vonikakis V, Winkler S (2015 November) Deep learning for emotion recognition on small datasets using transfer learning. In: *Proceedings of the 2015 ACM on international conference on multimodal interaction*, pp 443–449
- Simonyan K et al (2014) CNNs and large-scale image recognition (IR). [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)
- Tombaugh TN, McIntyre NJ (1992) The mini-mental state examination: a comprehensive review. *J Am Geriatr Soc* 40(9):922–935
- Wöllmer M, Kaiser M, Eyben F, Schuller B, Rigoll G (2013) LSTM-modeling of continuous emotions in an audiovisual affect recognition framework. *Image Vis Comput* 31(2):153–163
- Yu M (2013) An on-line one class SVM-based person-specific fall detection system. *IEEE J Biomed Health Inf* 17(6):1002–1014
- https://git.islab.ntua.gr/thanos/ntua_parkinson

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