



HHS Public Access

Author manuscript

IEEE J Biomed Health Inform. Author manuscript; available in PMC 2016 September 01.

Published in final edited form as:

IEEE J Biomed Health Inform. 2015 September ; 19(5): 1610–1616. doi:10.1109/JBHI.2015.2429556.

A Robust Deep Model for Improved Classification of AD/MCI Patients

Feng Li¹, Loc Tran¹, Kim-Han Thung³, Shuiwang Ji², Dinggang Shen^{3,4}, and Jiang Li¹

¹Department of Electrical and Computer Engineering, Old Dominion University, Norfolk, VA 23529

²Department of Computer Science, Old Dominion University, Norfolk, VA 23529

³Department of Radiology, University of North Carolina at Chapel Hill, NC 27599

⁴Department of Brain and Cognitive Engineering, Korea University, Seoul, Republic of Korea

Abstract

Accurate classification of Alzheimer's Disease (AD) and its prodromal stage, Mild Cognitive Impairment (MCI), plays a critical role in possibly preventing progression of memory impairment and improving quality of life for AD patients. Among many research tasks, it is of particular interest to identify noninvasive imaging biomarkers for AD diagnosis. In this paper, we present a robust deep learning system to identify different progression stages of AD patients based on MRI and PET scans. We utilized the dropout technique to improve classical deep learning by preventing its weight co-adaptation, which is a typical cause of over-fitting in deep learning. In addition, we incorporated stability selection, an adaptive learning factor, and a multi-task learning strategy into the deep learning framework. We applied the proposed method to the ADNI data set and conducted experiments for AD and MCI conversion diagnosis. Experimental results showed that the dropout technique is very effective in AD diagnosis, improving the classification accuracies by 5.9% on average as compared to the classical deep learning methods.

Index Terms

Alzheimer's Disease; Early Diagnosis; MRI; PET; Deep Learning

I. Introduction

Alzheimer's disease is the sixth-leading cause of death in the United States [1]. AD patients usually undergo progressive stages of cognitive and memory function impairment, including prodromal, MCI and AD. For each of these stages, significant amount of research has been conducted aiming to understanding the underlying pathological mechanisms. In addition, imaging biomarkers have been identified using different imaging modalities such as magnetic resonance imaging (MRI) [2], positron emission tomography (PET) [3], and functional MRI (fMRI) [4]. Imaging biomarkers are a set of indicators computed from

image modalities and can be used for early detection of AD disease. It has been shown that fusing these different modalities may lead to more effective imaging biomarkers [6].

The first successful deep learning framework, auto-encoder, was developed in 2006 [7]. It was subsequently used in other application fields and achieved state-of-the-art performance in speech recognition, image classification and computer vision [8]. Deep learning itself also evolves after 2006. For instance, the multimodal deep learning framework boosted speech classification by learning a shared representation between video and audio modalities [9]. A dropout technique further improved zip code recognition, document classification, and image recognition [10], [11].

In this paper, we developed a robust deep learning framework for AD diagnosis by fusing complementary information from MRI and PET scans. These 3D scans were preprocessed and their features were further extracted. Specifically, we first applied principal component analysis (PCA) to obtain PCs as new features. We then utilized the stability selection technique [13] together with the least absolute shrinkage and selection operator (Lasso) method [14] to select the most effective features. The selected features were subsequently processed by the deep learning structure. Model weights in the deep structure were first initialized by unsupervised training and then fine-tuned by AD patient labels. During the fine-tuning phase, the dropout technique was employed to improve the model's generalization capability. Finally, the learned feature representation was used for AD/MCI classification by a support vector machine (SVM).

In addition to discrete patient labels (AD, MCI or Healthy), there are two additional clinical scores, namely Minimum Mental State Examination (MMSE) and Alzheimer's Disease Assessment Scale-Cognitive subscale (ADAS-Cog) associated with each patient. MMSE is a 30-point questionnaire widely used to measure cognitive impairment [15]. It is used to estimate the severity and progression of cognitive impairment, instead of providing any AD information. ADAS-Cog is the most popular cognitive testing instrument to measure the severity of the most important symptoms of AD, including the disturbances of memory, language, praxis, attention and other cognitive abilities, which have been referred as the core symptoms of AD [16]. The information from these scores is related and identifying the commonality among them may help AD diagnosis. We configured the deep learning structure as a multi-task learning (MTL) framework, and treated the learning of class label, MMSE and ADAS-Cog as related tasks to improve the prediction of main task (class label).

We evaluated the proposed method on the ADNI¹ data set and compared it with a baseline method and a similar deep learning system, where the auto-encoder was used as a feature extractor for AD diagnosis [6]. The baseline method contains feature selection and SVM steps but does not use deep learning. We also evaluated the impact on performance of each of the components in the proposed system. A brief version of this paper was published in [17].

¹Available at <http://www.loni.ucla.edu/ADNI>.

II. Materials and Methods

The proposed system consists of multiple components, including PCA, stability selection, unsupervised feature learning, multi-task deep learning and SVM training, as shown in Fig. 1. We detail each of these components in the following subsections.

A. Data preprocessing

We utilized the public ADNI data set to validate our proposed deep learning framework. The data set consists of MRI, PET, and CSF data from 51 AD patients, 99 MCI patients (43 MCI patients who converted to AD (MCI.C), and 56 MCI patients who did not progress to AD in 18 months (MCI.NC)) as well as 52 healthy normal controls. In addition to the crisp diagnostic result (AD or MCI), this data set contains two additional clinical scores, MMSE and ADAS-Cog, for each patient. A typical procedure of image processing was applied to the 3D MRI and PET images [2], [18], [19] including anterior commissure-posterior commissure correction, skull-stripping, cerebellum removal, and spatially normalization. Finally, we extracted 93 region-of-interest (ROI) based volumetric features from MRI and PET images, respectively, which together with three CSF biomarkers, i.e., $A\beta_{42}$, t -tau, and p -tau, sum up to 189 features for each subject.

B. Principal component analysis

Principal component analysis (PCA) is a linear orthogonal transformation that converts a set of features into linearly uncorrelated variables in which each of the new variable is a linear combination of all original features [5]. The first principal component (PC) is defined as the one that can explain the largest variance in the original data set. The second PC has the second largest variance under the constraint that it is orthogonal to the first component. If correlations exist among features, the number of PCs that can be found is usually less than the number of features in the original data. PCA is optimal for preserving energy and it is often used for dimensionality reduction by just keeping the first few PCs.

Let \mathbf{F} denote a feature data set with a size of $n \times p$, where n is the number of data samples and p is the number of features in the data, and each column in \mathbf{F} is centered. PCA can be achieved by performing the singular value decomposition (SVD) on \mathbf{F} as

$$\mathbf{F} = \mathbf{U} \Sigma \mathbf{V}^T, \quad (1)$$

where \mathbf{U} is an $n \times n$ matrix with orthogonal unit columns (left singular vectors of \mathbf{F}), Σ is an $n \times p$ diagonal matrix consisting of singular values of \mathbf{F} from the largest to least, and \mathbf{V} is an $p \times p$ matrix whose columns are orthogonal unit vectors (right singular vectors of \mathbf{F}).

To achieve dimensionality reduction, the first l columns in \mathbf{V} corresponding to the first l largest singular values of \mathbf{F} can be used as a transformation matrix to be applied on \mathbf{F} ,

$$\mathbf{x} = \mathbf{F} \mathbf{V}_l, \quad (2)$$

where \mathbf{V}_l consists of the first l columns of \mathbf{V} .

Geometrically, PCA analysis rotates data to align its maximum variance direction of the data with the coordinate system as illustrated in Fig. 2. PCA is an effective tool for dimensionality reduction but the preserved PCs may not be useful for classification. The two dimensional artificial data set in Fig. 2 consists of ‘blue’ and ‘red’ classes. After PCA, the whole data set was rotated and its main axis was aligned with the coordinate system. However, even though PC 1 has the largest variance, it does not contain any discriminating information for the two classes. For the purpose of classification, PC 2 is preferred and a feature selection step is necessary. This example shows that feature selection may be applied after PCA to retain discriminating information for classification.

C. Stability selection

In this paper, we first applied PCA to the 189 features and used the resulting PCs as new features. We then applied Lasso [14] to identify the most effective features for AD diagnosis. Lasso tries to minimize the following cost function for feature selection:

$$\min_{\mathbf{s}} \|\mathbf{t} - \mathbf{x}\mathbf{s}\|_2^2 + \lambda \|\mathbf{s}\|_1, \quad (3)$$

where $\mathbf{t} \in \{+1, -1\}^n$ is a class label vector of size $n \times 1$ associated with the feature matrix \mathbf{x} of size $n \times l$, where l is the number of features (PCs) found in PCA, $\mathbf{s} = [s_1, s_2 \dots s_l]^T$ is the weight vector associated with the l features (columns in \mathbf{x}), λ is a regularization parameter, and $\|\cdot\|_2$ and $\|\cdot\|_1$ denote L_2 and L_1 norms, respectively. Because of the L_1 norm constraint on the weight magnitude, the solution minimizing the above cost function is usually sparse, meaning that if a feature is not correlated with the target class label, the feature will have a zero value for its weight. Features having nonzero weights will be selected and otherwise will be excluded.

It is well known that the solution of L_1 norm based optimizations are sensitive to the choice of λ , and it is difficult to determine how many features should be kept in the model. A recent breakthrough sheds a light on selecting the right amount of regularization for stability selection [13]. The idea is to repeat the feature selection procedure multiple times based on bootstrapped data sets and compute the probability of the features to be selected. The final selected features are those having probabilities above a predefined threshold t_h . It has been shown experimentally and theoretically that the feature selection results vary little for sensible choices in a range of the cut-off value for t_h [13]. We incorporated the stability selection concept into the AD patient diagnosis in this paper. In particular, we repeated the Lasso procedure 50 times and each time with a different value for the parameter λ (We used the SLEP toolbox for Lasso²). A probability, p_i , for the i th feature was computed by counting the frequency of the feature being selected in the 50 experiments. The i th feature was selected if p_i is larger than a pre-defined threshold t_h .

D. Multi-task deep learning with dropout

In contrast to traditional three-layer neural network (shallow structure), deep learning is based on a deep architecture consisting of many layers of hidden neurons for modeling. A

²Available at <http://www.public.asu.edu/jye02/Software/SLEP/index.htm>

shallow architecture would involve many duplications of effort to express things and such a fat architecture has been shown to suffer from the problem of over-fitting, which leads to a poor generalization capability. Instead, deep architecture could more gracefully reuse previous computations and discover complicated relations of input [20].

To train a deep architecture, the standard Backpropagation (BP) algorithm did not work well with randomly initialized weights because the error feedback becomes progressively noisier as it goes back to lower levels (closer to inputs), making the low-level weight updates less effective. Even though experiments have shown that if top layers have enough units, the deep structure can still bring down training errors small enough, it cannot generalize well to new data [21]. This is because the top layers can be effectively trained by gradient based algorithms but low-levels cannot. The randomly initialized low-level layers behave like random feature detectors so good representations for original data were not achieved leading to degraded generalization capability [21]. In 2006, a breakthrough in deep learning has made deep architecture training possible by utilizing the restricted Boltzmann machine (RBM) to initialize multiple hidden layers one layer at a time in an unsupervised manner [7]. With the unsupervised learning, deep learning tries to understand data first, i.e., to obtain a task specific representation from data so that a better classification can be achieved. It has experimentally proven that the unsupervised learning step plays a critical role in the success of deep learning [8]. The proposed deep model shown in Fig. 3 consists of several components that will be described below.

1) Pre-training with RBM—Each layer in the proposed deep model is an RBM and the deep model used in this paper consists of a stack of RBMs. RBM is an energy-based model in which a scalar energy is associated with each configuration of the variables in the model, and a probability distribution function (PDF) through the energy function is defined. The purpose of learning is to modify the energy function so that a desirable PDF can be achieved, i.e., to have low energy. A basic RBM model having a visible (input) layer and a hidden (output) layer is shown in Fig. 4. The visible layer of the bottom RBM contains real-valued units (receiving data) and all other RBM layers have binary units. Let $v \in R^M$ represent input data (visible units) and $h \in \{0, 1\}^N$ denote binary hidden units for the bottom RBM, we used Gaussian-Bernoulli RBMs to train it [21], [22]. All other RBMs were trained by utilizing Bernoulli-Bernoulli distribution. Variables v and h have a joint probability distribution defined as

$$p(v, h) = \frac{1}{Z} \exp^{-E(v, h)}, \quad (4)$$

where $E(v, h)$ is an energy function and Z is a normalization constant. For real-valued visible layer RBMs, $E(v, h)$ is defined as

$$E(v, h) = \frac{1}{2\sigma^2} \sum_i v_i^2 - \frac{1}{\sigma^2} \left(\sum_i c_i v_i + \sum_j b_j h_j + \sum_{i,j} v_i w_{ij} h_j \right), \quad (5)$$

where c_i and b_j are biases of the i th and j th units in the visible and hidden layers, respectively. w_{ij} is the weight connecting v_i and h_j , and σ^2 is the variance of v . The conditional probability distributions are

$$P(h_j=1|v)=\text{sigmoid}\left(\frac{1}{\sigma^2}\left(\sum_i w_{ij} v_i+b_j\right)\right), \quad (6)$$

$$P(v_i|h)=\mathcal{N}\left(\sum_j w_{ij} h_j+c_i,\sigma^2\right). \quad (7)$$

If both visible and hidden layers are binary, the energy function and conditional probability distributions are defined as

$$E(v, h)=-\left(\sum_i c_i v_i+\sum_j b_j h_j+\sum_{ij} v_i w_{ij} h_j\right), \quad (8)$$

$$P(h_j=1|v)=\text{sigmoid}\left(\sum_i w_{ij} v_i+b_j\right), \quad (9)$$

$$P(v_i=1|h)=\text{sigmoid}\left(\sum_j w_{ij} h_j+c_i\right). \quad (10)$$

Model parameters w , b and c are updated using contrastive divergence [23]. For RBM having a real-valued visible layer, the formulas for updating those parameters during each iteration are

$$\Delta W_{ij}^{t+1}=\eta\Delta W_{ij}^t-\varepsilon\left(\left\langle\frac{1}{2}v_i h_j\right\rangle_d-\left\langle\frac{1}{2}v_i h_j\right\rangle_m\right), \quad (11)$$

$$\Delta b_i^{t+1}=\eta\Delta b_i^t-\varepsilon\left(\left\langle\frac{1}{2}v_i\right\rangle_d-\left\langle\frac{1}{2}v_i\right\rangle_m\right), \quad (12)$$

$$\Delta c_j^{t+1}=\eta\Delta c_j^t-\varepsilon\left(\left\langle h_j\right\rangle_d-\left\langle h_j\right\rangle_m\right). \quad (13)$$

where $\langle\cdot\rangle_d$ and $\langle\cdot\rangle_m$ denote the expectation computed over data and model distributions accordingly, t is iteration index, η is momentum and ε is learning rate. For binary RBM, equations (11) and (12) become

$$\Delta W_{ij}^{t+1}=\eta\Delta W_{ij}^t-\varepsilon\left(\left\langle v_i h_j\right\rangle_d-\left\langle v_i h_j\right\rangle_m\right), \quad (14)$$

$$\Delta b_i^{t+1} = \eta \Delta b_i^t - \varepsilon (\langle v_i \rangle_d - \langle v_i \rangle_m). \quad (15)$$

Note that the pre-training of RBM is unsupervised, i.e., class label (classification task) or desired output (regression) is not needed in the training. After the pre-training, we attached the class label on top of the stacked RBMs and utilized an adaptive backpropagation algorithm to fine-tune the weights in the model. All binary layers were also converted to real-valued units by using their continuous activities. Thus the deep learning model turned to be a traditional multilayer perceptron (MLP) but its weights were initialized by RBM.

2) Multi-task learning—In multi-task learning, related tasks are learned simultaneously by extracting and utilizing appropriate shared information across tasks to improve performance. It has received attention in broad areas recently such as machine learning, data mining, computer vision, and bioinformatics [24], [25], [26]. This approach is particularly effective when only limited training data for each task is available. It is worth noting that neural networks can simultaneously model multiple outputs, making deep learning a natural multi-task learning framework if multiple tasks share inputs [7]. The proposed multi-task deep learning framework is shown in Fig. 3, where we treated the predictions of class label, MMSE and ADAS-Cog as three different tasks and modeled them simultaneously. MMSE, and ADAS-Cog were normalized to the range of [0,1] and we used the deep structure as a regression model. The class label was coded by the 1-of- k scheme. To classify an input vector, we checked the corresponding k outputs and assign it to the class having the largest output. One drawback of deep model is over-fitting due to large capacity. This is more prominent if training data is limited. To overcome this limitation, we utilized the dropout technique to improve training.

3) Dropout with adaptive adaptation—Deep learning achieved excellent results in applications where training data size is large. For small sized data sets such as the one in this paper, it is still possible for a deep structure to over-fit the data given the fact that it usually has tens of thousands or even millions of parameters. To improve the generalization capability of the model, the dropout technique tries to prevent weight co-adaptation by randomly dropping out some units in the model during training [10], [11]. We incorporated the dropout technique in the multi-task learning context to improve AD diagnosis as shown in Fig. 3. In the training process, each hidden unit in the model was dropped with a probability of 0.5 when a batch of training cases were present. Previous experiments [10] showed that it is also beneficial if we apply the “dropout” process to the input layer but with a lower probability (i.e., 0.2 in this paper). In the testing procedure, all hidden units and inputs were used to compute model outputs for a testing case with appropriate compensations, i.e., weights between inputs and the first hidden layer were scaled by 0.8 and all other weights were halved.

During the multi-task fine-tuning step, the stochastic gradient descent method with a fixed learning factor is usually utilized as [7],

$$w_{ij} = w_{ij} + \Delta w_{ij} = w_{ij} - \alpha \frac{\partial L}{\partial w_{ij}}, \quad (16)$$

where $\frac{\partial L}{\partial w_{ij}}$ is the gradient of the cost function L and α is a learning factor. Sometimes, the weights update may contain a momentum term [10]. We proposed an adaptive learning factor to speed up the adaptation. The motivation of the adaptive learning is that the learning factor should be large at locations where gradient is small and vice versa. Assume the decrease of L due to the change in w_{ij} is approximated by

$$\Delta L^{ij} = L_{new}^{ij} - L_{old}^{ij} \approx \frac{\partial L}{\partial w_{ij}} \times \Delta w_{ij} = -\alpha \left[\frac{\partial L}{\partial w_{ij}} \right]^2, \quad (17)$$

then L due to all w_{ij} can be computed as

$$\Delta L = -\alpha \sum_i \sum_j \left[\frac{\partial L}{\partial w_{ij}} \right]^2. \quad (18)$$

Suppose we want to decrease L by $\beta\%$, then $L_{new} = (1 - \beta)L_{old}$, and an adaptive learning factor α can be determined as

$$\alpha = \frac{\beta L_{old}}{\sum_i \sum_j \left[\frac{\partial L}{\partial w_{ij}} \right]^2}. \quad (19)$$

We set β as 10% in our experiments in this paper. Once the new feature representation is learned, an SVM classifier [12] was trained using the learned feature representation.

E. SVM Classifier

Given a set of data pairs $\{\mathbf{r}_i, t_i\}_{i=1}^n$, where $\mathbf{r}_i \in R^M$ is the learned feature representation from subjects, $t_i \in \{+1, -1\}$ is a class label (e.g., AD vs. non-AD) associated with \mathbf{r}_i . An SVM defines a hyperplane

$$f(\mathbf{r}) = \mathbf{k}^T \phi(\mathbf{r}) + e = 0 \quad (20)$$

separating the data points into 2 classes. In equation (20), \mathbf{k} and e are the hyperplane parameters, and $\phi(\mathbf{r})$ is a function mapping the vector \mathbf{r} to a higher dimensional space. The hyperplane (20) is determined using the concept of *Structural Risk Minimization* [12] by solving the following optimization problem,

$$\min_{\mathbf{k}, e, \xi} \left(\frac{1}{2} \mathbf{k}^T \mathbf{k} + C \sum_{i=1}^n \xi_i \right), \quad (21)$$

subject to

$$t_i(\mathbf{k}^T \phi(\mathbf{r}_i) + e) \geq 1 - \xi_i, \xi_i \geq 0, \quad (22)$$

where C is a regularization parameter and ξ_i is a slack variable. After the hyperplane is determined, an AD case is declared if $f(\mathbf{r}_i) > 0$, or otherwise a non-AD case is declared.

III. Results and discussions

A. Experimental setup

1) Ten-fold cross-validation—We consider four classification tasks including AD patients vs Healthy Control subjects (AD vs HC), MCI patients vs HC (MCI vs HC), AD patients vs MCI patients (AD vs MCI) and MCI-converted vs MCI-non converted (MCI.C vs MCI.NC). For each task, we utilized a ten-fold cross-validation (CV) scheme to evaluate the proposed method. In the ten-fold CV, we randomly divided the data set into 10 parts and for one run, we separated one part for testing and applied the proposed framework to the remaining data to train a classification model. This procedure was repeated 10 times so that each part was tested once. Finally, testing accuracies were computed. To obtain a more reliable estimate of the performance, we repeated the ten-fold CV ten times for each task with different random data partitions and computed average accuracy. To compare different classification models, we kept the same data partitions in the ten-fold CV and utilized the paired- t test to evaluate if there is a significant performance difference.

2) Hyperparameter determination—We did preliminary experiments to determine the structure of the deep learning model. It was found that using three hidden layers with hidden units of 100-50-20 worked the best among the candidate structures considered and was thus utilized in our experiments. For the SVM classifier, we tried different kernels and a linear kernel was chosen. We also did a grid search for the “soft margin” parameter in the linear kernel SVM model but it did not improve the classification accuracies. Therefore, in all experiments, we utilized a three hidden-layer model with a structure of 100-50-20 for feature learning and a linear SVM with default soft margin as the classifier.

3) Impact assessment for individual component—There are four components in the proposed framework including PCA, stability selection, dropout and multi-task learning. Inspired by “sensitivity analysis” and “impact assessment” that analyze inputs of or components in a model and identify their impacts on the model objectives by varying the inputs [28]. We incorporated a similar concept to evaluate the impact of each component on model performance by varying the component (presence vs absence). ‘Absence’ means that the component was not included in the model.

4) Methods for comparison—We compared the proposed method with a baseline method and a similar deep learning system proposed in [6]. The baseline method consists of all components in the proposed system except the deep learning step. The work by Suk in [6] is a auto-encoder based deep learning method in which feature representations for MRI, PET and CSF from the same data set were learned separately and combined by a linear SVM classifier. They also combined the learned representations with original features for AD diagnosis.

B. Results

Table I shows the overall performances of the proposed method and the impact of each component in the framework. The proposed method performed the best in diagnosing AD and MCI patients, and discriminating MCI patients from AD patients with accuracies of 91.4%, 77.4% and 70.1%, respectively. It is significantly better than the baseline method that obtained accuracies of 86.4%, 72.1% and 61.5% for the diagnoses. In the MCI conversion diagnosis (MCI.C vs MCI.NC), the PCA component slightly degraded the proposed method (from 58.1% to 57.4%) but it is still significantly better than the baseline method (57.4% vs 50.6%).

Among those components, it is obvious that “dropout” has the most significant impact on the performances. Without “dropout”, deep learning did not significantly improve the baseline method (68.2% vs 67.7% in terms of average acc.). The least important component is “PCA”, i.e., the average acc. slightly dropped from 74.1% to 73.4% without the PCA component. Without “stability selection” and “multi-task learning”, the average accuracy dropped from 74.1% to 72.5% and 72.4%, respectively.

We conducted a paired-*t* test between results by the proposed method and those from classical deep learning (“Dropout”). Table II lists the improvements and *p*-values. The average improvement is 5.9% and the improvements for all the four classification tasks are significant.

The work by Suk [6] on the same data set is also shown in Table II, where “SAEF” corresponds to the method using features learned by a deep auto-encoder and “LLF+SAEF” represents the method that combines original features with the SAEF features for AD diagnosis. The AD vs MCI classification experiment was not conducted in [6]. The proposed method (75.4%) outperformed the SAEF method (with an average accuracy of 70.6%). By combining SAEF with LLF (LLF+SAEF), the average accuracy was increased to 74.2% (Last column in Table II).

C. Discussions

There are usually two ways to increase the generalization capability of a model, adding regularization (L_1 or L_2 norm) on weights or using a committee machine. However, solving the regularization problem is usually challenging especially in the deep learning context. In addition, the committee machine technique requires averaging many separately trained models to compute a prediction for a testing case, which is time consuming for deep learning. The dropout procedure does the both (constraint and committee machine) simultaneously in a very efficient way. 1) Each sub-model in training is a sampled model from all possible ones and all sub-models share weights. The weight sharing property is equivalent to the L_1 or L_2 norm constraint on weights, and 2) the testing procedure is an approximation of averaging all trained sub-models for a testing case but it does not separately store them because they share weights. This is an extremely efficient and smart implementation of a committee machine [10], [11].

The impact evaluation method was inspired by the “sensitivity analysis” and “impact assessment” [28]. We were aiming to identify the impact on performance of each component

in the model by excluding the component from the pipeline. Note that we did not try to decouple the component from the system. This evaluation method may not be a strict sensitivity analysis or impact assessment by means of their definitions, but we can verify each component if it can improve the AD diagnosis when it is included in the proposed system. Our experiments showed that the dropout component has the largest impact on the performance, multi-task learning ranked the second, stability selection the third, and PCA has the least impact on the performance.

In terms of stability selection and computational efficiency, there were usually around 40 features left after the stability selection and it took about 1 hour for a personal computer to conduct a ten-fold CV evaluation for one task. The number of features that were chosen was determined by stability selection, in which the Lasso algorithm ran 50 times with different values of regularization parameter (λ). In each run, Lasso chose different features and a probability of being chosen for each feature was computed in the 50 runs. Finally, a feature was chosen if its probability is larger than 0.5.

It is worth to note that the results by the proposed method in Table I and Table II only used the new representations learned by the deep model. We tried to combine the new representations with the original features but the combination did not improve the performance. In [6], new representations learned from auto-encoder did not perform well unless they were combined with the original features. Our experiment also showed that the deep model without dropout just performed comparably as the baseline method. It seems that traditional deep learning cannot extract information effectively from small data sets unless it is regularized by techniques such as dropout.

In [29], a multi-kernel SVM (MK-SVM) method was applied to the same data set to combine the original LLF features for AD diagnosis, and achieved 93.2% and 76.4% for AD vs HC and MCI vs HC classifications, respectively. The MCI conversion diagnosis and AD vs MCI classification were not conducted. In [6], utilizing the MK-SVM method to combine SAEF features from MRI, PET and CSF boosted the performances to 95.9%, 85.0% and 75.8% for the three tasks (AD vs MCI classification was not performed), respectively. Since the dropout technique improved upon the basic deep learning, we are currently investigating if the MK-SVM method can further boost the performance of the proposed system.

We did not attempt to perform a comprehensive comparison study of the proposed method with others that have been applied to this data set in the literature. Instead, we have evaluated some recently proposed advanced machine learning techniques for AD diagnosis including Lasso, stability selection, multi-task learning, deep learning and dropout. The dropout technique seems to be an effective method of regularization for learning with small data sets. Without dropout, deep learning has no advantage over the baseline method on ANDI data set (68.2% vs 67.7%). Note that dropout is computationally very efficient as compared to either L_1 norm based regularization or committee machine and it can be extended to many models other than the deep model as discussed in this paper.

IV. Conclusion

Our proposed method achieved 91.4%, 77.4%, 70.1% and 57.4% accuracies for AD vs HC, MCI vs HC, AD vs MCI, and MCI.c vs MCI.NC classifications, respectively. The framework consists of multiple components including PCA, stability selection, dropout and multi-task deep learning. We showed that dropout is the most effective one. This is not surprising because the size of ADNI data is relatively small compared to that of the deep structure utilized in this paper. Classical deep learning does not perform well on this small data set, but with the dropout technique, the average accuracy was improved by 5.9% on average. We plan to incorporate MK-SVM [6] into our method for further improving AD diagnosis.

Acknowledgments

This work was partially supported by NIH grants (EB006733, EB008374, EB009634, MH100217, AG041721, AG042599).

References

1. Alzheimer's Association. 2012 Alzheimer's disease facts and figures. *Alzheimer's & Dementia*. 2012; 8(2):131–168.
2. Davatzikos C, Bhatt P, Shaw LM, Batmanghelich KN, Trojanowski JQ. Prediction of MCI to AD conversion, via MRI, CSF biomarkers, and pattern classification. *Neurobiology of Aging*. 2011; 32(12):2322.e19–2322.e27. [PubMed: 20594615]
3. Nordberg A, Rinne JO, Kadir A, Langstrom B. The use of PET in Alzheimer disease. *Nature Reviews Neurology*. 2010; 6(2):78–87. [PubMed: 20139997]
4. Greicius MD, Srivastava G, Reiss AL, Menon V. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: Evidence from functional MRI. *Proceedings of the National Academy of Sciences of the United States of America*. 2004; 101(13):4637–4642. [PubMed: 15070770]
5. Jolliffe IT. *Principal Component Analysis*, Series: Springer Series in Statistics (2). 2002
6. Suk H, Shen D. Deep learning-based feature representation for AD/MCI classification. *MICCAI*. 2013:583–590. [PubMed: 24579188]
7. Hinton GE, Srivastava N, Sutskever I, Salakhutdinov RR. A fast learning algorithm for deep belief nets. *Neural computation*. 2006; 18(7):1527–1554. [PubMed: 16764513]
8. Bengio Y, Courville A, Vincent P. Representation learning: A review and new perspectives. *PAMI*. 2013; 35(8):1798–1828.
9. Ngiam J, Khosla A, Kim M, Nam J, Lee H, Ng A. Multimodal deep learning. *ICML*. 2011:689–696.
10. Hinton GE, Srivastava N, Krizhevsky A, Sutskever I, Salakhutdinov RR. Improving neural networks by preventing co-adaptation of feature detectors. 2012 arXiv 1207.0580.
11. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*. 2014; 15:1929–1958.
12. Cortes C, Vapnik V. Support-vector networks. *Machine Learning*. 1995; 20(3):273.
13. Meinshausen N, Bühlmann P. Stability selection. *J R Statist Soc B*. 2010:417–473.
14. Tibshirani R. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B*. 1996; 58(1):267–288.
15. Pangman VC, Sloan J, Guse L. An Examination of Psychometric Properties of the Mini-Mental State Examination and the Standardized Mini-Mental State Examination: Implications for Clinical Practice. *Applied Nursing Research*. 2000; 13(4):209213.

16. Kolibas E, Korinkova V, Novotny V, Vajdickova K, Hunakova D. ADAS-cog (Alzheimer's Disease Assessment Scale-cognitive subscale)–validation of the Slovak version. *Bratisl Lek Listy*. 2000; 101(11):598–602. [PubMed: 11218956]
17. Li F, Tran L, Thung KH, Ji S, Shen D, Li J. Robust Deep Learning for Improved Classification of AD/MCI Patients. *Machine Learning in Medical Imaging*. :240–247.2014
18. Kabani N, MacDonald D, Holmes C, Evans A. A 3D atlas of the human brain. *NeuroImage*. 1998; 7(4):S717.
19. Hinrichs C, Singh V, Xu G, Johnson SC. Predictive markers for AD in a multi-modality framework: An analysis of MCI progression in the ADNI population. *NeuroImage*. 2011; 55(2): 574–589. [PubMed: 21146621]
20. Erhan D, Bengio Y, Courville A, Manzagol PA, Vincent P, Bengio S. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*. 2010; 11:625–660.
21. Bengio Y. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*. 2009; 2(1):1127.
22. Cho, K.; Ilin, A.; Raiko, T. *Artificial Neural Networks and Machine Learning ICANN 2011*. Springer; 2011. Improved learning of Gaussian-Bernoulli restricted Boltzmann machines; p. 10-17.
23. Hinton G, Osindero S, Teh Y-W. A fast learning algorithm for deep belief nets. *Neural computation*. 2006; 18:1527–1554. [PubMed: 16764513]
24. Heisele B, Serre T, Pontil M, Vetter T, Poggio T. Categorization by learning and combining object parts. *NIPS*. 2001
25. Ji, S.; Ye, J. An accelerated gradient method for trace norm minimization. *Proceedings of the 26th Annual International Conference on Machine Learning*; 2009. p. 457-464.
26. Xue Y, Liao X, Carin L, Krishnapuram B. Multi-task learning for classification with dirichlet process priors. *Journal of Machine Learning Research*. 2007; 8:35–63.
27. Caruana R. Multitask learning: A knowledge-based source of inductive bias. *Machine Learning*. 1997; 28:41–75.
28. Saltelli, A.; Ratto, M.; Andres, T.; Campolongo, F.; Cariboni, J.; Gatelli, D.; Saisana, M.; Tarantola, S. *The Primer*. John Wiley & Sons; 2008. *Global Sensitivity Analysis*.
29. Zhang D, Wang Y, Zhou L, Yuan H, Shen D. Multimodal Classification of Alzheimers Disease and Mild Cognitive Impairment. *Neuroimage*. 2011; 55(3):856–867. [PubMed: 21236349]

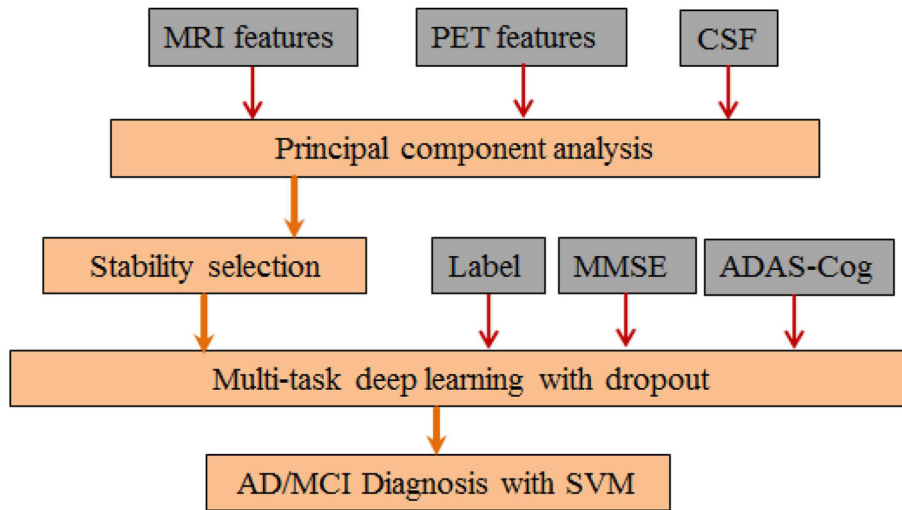


Fig. 1. Diagram of the proposed multi-task deep learning framework.




Fig. 2. Principal component analysis example. PC 1 contains the most energy of the data but does not have any discrimination information for the 'red' and 'blue' classes.




Fig. 3. Multi-task deep learning with dropout. “x” denotes a dropped unit.




Fig. 4.
A basic RBM model.

TABLE I

Performance comparison (in%) of the competing methods. The proposed method consists of four components. “-PCA” stands for “the proposed method without the PCA component” and “-SS” stands for stability selection, “Baseline” denotes the framework without the deep learning component.

Tasks	Proposed	-PCA	-Dropout	-SS	-MultTask	Baseline
AD vs HC	91.4 (1.8)	89.6(1.3)	84.2(3.0)	89.4(1.6)	90.3(1.7)	86.4(2.0)
MCI vs HC	77.4 (1.7)	76.4(1.5)	73.1(3.1)	74.3(1.6)	75.6(1.7)	72.1(3.0)
AD vs MCI	70.1 (2.3)	69.5(2.7)	65.1(3.7)	68.7(2.1)	67.1(2.9)	61.5(2.9)
MCI.C vs MCLNC	57.4(3.6)	58.1 (1.8)	50.2(3.3)	57.7(1.8)	56.7(3.0)	50.6(4.7)
Average	74.1	73.4	68.2	72.5	72.4	67.7

Paired- t test between results of the proposed method vs deep learning without dropout. The methods of “SAEF” and “LLF+SAEF” were proposed by Suk [6]. “SAEF” stands for Stacked Auto-Encoder Features and “LLF” denotes Low Level Features.

TABLE II

Tasks	Proposed	-Dropout	Improvement	p -value	SAEF	LLF+SAEF
AD vs HC	91.4 (1.8)	84.2(3.0)	7.2	$< 10^{-3}$	83.2(2.7)	85.3(3.2)
MCI vs HC	77.4 (1.7)	73.1(3.1)	4.3	0.0034	70.1(2.8)	76.9(2.3)
AD vs MCI	70.1 (2.3)	65.1(3.7)	5.0	0.0017	N/A	N/A
MCLC vs MCLNC	57.4(3.6)	50.2(3.3)	7.2	$< 10^{-3}$	58.4(4.1)	60.3 (2.3)
Overall Average	74.1	68.2	5.9	N/A	N/A	N/A
Average w/o AD vs MCI	75.4	69.2	6.2	N/A	70.6	74.2