

Detecting Cognitive Impairments by Agreeing on Interpretations of Linguistic Features

Zining Zhu^{1,2}, Jekaterina Novikova¹, Frank Rudzicz^{3,5,2,4,1}

¹ Winterlight Labs ² University of Toronto

³ Li Ka Shing Knowledge Institute, St Michael's Hospital, Toronto ON, Canada

⁴ Vector Institute for Artificial Intelligence ⁵ Surgical Safety Technologies

{zining, jekaterina}@winterlightlabs.com, frank@cs.toronto.edu

Abstract

Linguistic features have shown promising applications for detecting various cognitive impairments. To improve detection accuracies, increasing the amount of data or the number of linguistic features have been two applicable approaches. However, acquiring additional clinical data can be expensive, and hand-crafting features is burdensome. In this paper, we take a third approach, proposing Consensus Networks (CNs), a framework to classify after reaching agreements between modalities. We divide linguistic features into non-overlapping subsets according to their modalities, and let neural networks learn low-dimensional representations that agree with each other. These representations are passed into a classifier network. All neural networks are optimized iteratively.

In this paper, we also present two methods that improve the performance of CNs. We then present ablation studies to illustrate the effectiveness of modality division. To understand further what happens in CNs, we visualize the representations during training. Overall, using all of the 413 linguistic features, our models significantly outperform traditional classifiers, which are used by the state-of-the-art papers.

1 Introduction

Alzheimer's disease (AD) and its usual precursor, mild cognitive impairment (MCI), are prevalent neurodegenerative conditions that inhibit cognitive abilities. Cognitive impairments are traditionally diagnosed only with standard clinical tests like MoCA (Nasreddine et al., 2005) and the Rey-Auditory Verbal learning Test (Rey, 1941), but hiring clinicians to administer these tests and analyze their results is costly. Fortunately, many cognitive impairments can be observable in daily life, because they impact one's language abilities. For example, cognitively impaired people tend to

use more pronouns instead of nouns, and pause more often between sentences in narrative speech (Roark et al., 2011).

This insight makes automatic detection possible. Machine learning classifiers can detect cognitive impairments given descriptive linguistic features. In recent work, linguistic features including pronoun-noun-ratios, pauses, and so on, are used to train classifiers to detect cognitive diseases in various tasks. For example, Fraser et al. (2015) achieved up to 82% accuracy on DementiaBank¹, the largest publicly available dataset on detecting cognitive impairments from speech, and Weissenbacher et al. (2016) achieved up to 86% accuracy on a corpus of 500 subjects. Yancheva et al. (2015) estimated Mini-Mental State Estimation scores (MMSEs), describing the cognitive status and characterizing the extent of cognitive impairment.

To improve the accuracy of automated assessment using engineered linguistic features, there are usually two approaches: incorporating more clinical data or calculating more features. Taking the first approach, Noorian et al. (2017) incorporated normative data from Talk2Me² and the Wisconsin Longitudinal Study (Herd et al., 2014) in addition to DementiaBank, which increased AD:control accuracy up to 93%, and moderateAD:mildAD:control three-way classification accuracy to 70%. Taking the second approach, Yancheva and Rudzicz (2016) used 12 features derived from vector space models and reached a .80 F-score on DementiaBank. Santos et al. (2017) calculated features depicting characteristics of co-occurrence graphs of narrative transcripts (e.g., the degree of each vertex in the graph). Their classifiers reached 65% accuracy on DementiaBank

¹<https://talkbank.org/DementiaBank>

²<https://www.cs.toronto.edu/talk2me/>

(MCI versus a subset of Control).

There are limitations in either of the two approaches. On one hand, acquiring additional clinical data can be expensive (Berndt and Cockburn, 2013). Moreover, the additional data should be similar enough to existing training data to be helpful. On the other hand, crafting new features requires creativity and collaboration with subject matter experts, and the implementation can be time consuming. Neither of these approaches is satisfactory.

These limitations motivate us to take a third, novel approach. Instead of using new data or computing new features, we use the existing linguistic features.

If the speaker is cognitively impaired, and their language ability is affected, features from each of the acoustic, syntactic, and semantic modalities should reflect such change (Szatloczki et al., 2015; Moro et al., 2015; Fraser et al., 2015). We therefore need to distill the common information revealed by features from multiple, mainly distinct, modalities.

To utilize information common across different modalities, Becker and Hinton (1992) and de Sa (1994) let classifiers look at each modality and supervise each other. These examples illustrated the effectiveness of multi-view learning in utilizing common information among different observations, but their algorithms fail to train useful classifiers for cognitive impairments in our datasets. Without explicit supervision, self-supervised models almost always converge to a state producing the same predictions for all people, giving trivial classifiers.

Instead of aligning the predictions from modalities, we let the representations of the modalities agree. Generative adversarial networks (GANs) provide an approach. In GANs, a “discriminator” network is trained to tell whether a vector is drawn from the real world or produced synthetically by a “generator” neural network, while the generator is trained to synthesize images as close to real data as possible. We borrow this setting, and encourage the neural networks interpreting different modalities to produce representations of modalities as similar to each other as possible. This leads to our classifier framework, consensus networks (CNs).

Consensus networks constitute a framework using adversarial training to utilize common information among modalities for classification. In

this framework, several neural networks (“ePhysicians”) are juxtaposed, each learning the representation of a partition of linguistic features for each transcript. Being trained towards producing agreed representations, we show they are increasingly able to capture common information contained across disparate subsets of linguistic features.

We empirically add two extensions to CN that improve the classification accuracies, called the “noise modality” and “cooperative optimization” respectively, as explained below. To illustrate the effectiveness of the consensus mechanisms, we present two ablation studies. First, we compare neural networks built with consensus (CN) and those without (MLP). On partial or complete modalities, CN outperforms MLP significantly. Second, we compare CNs built with linguistic features divided into random subsets. Division according to their natural modalities train better consensus networks. We also visualize the representations during training procedure, and show that when the representations agree, their distributions appear symmetric.

Overall, taking all 413 linguistic features, our models significantly outperform traditional classifiers (e.g., support vector machines, quadratic discriminant analysis, random forest, Gaussian process), which are used by the state-of-the-art.

2 Related Works

Generative Adversarial Networks The idea of aligning representations by making them indistinguishable is inspired by GAN (Goodfellow et al., 2014), where a generator produces fake images (or other data) that are as similar to real data as possible. However, our model does not have a generator component as GANs do. Instead, we only compress features into representations while trying to align them.

Multi-view Learning Learning from multiple modalities is also referred to as multi-view learning. Becker and Hinton (1992) set up multiple neural networks to look at separate parts of random-dot stereograms of curved surfaces, and urge their prediction to equal each other. The trained neural networks were able to discover depth without prior knowledge about the third dimension. de Sa (1994) divided linguistic features into two modalities, and passed them to two separate neural networks. The two neural networks supervised

each other (i.e., output labels that are used to train the other) during alternative optimization steps to reach a consensus. Their self-supervised system reached $79\pm 2\%$ accuracy using the Peterson-Barney vowel recognition dataset (Peterson and Barney, 1952). Benediktsson et al. (1997) computed multiple views from the same feature sets and classified by taking their majority votes. Pou-Prom and Rudzicz (2018) used canonical correlation analysis (CCA) to classify using multiple aspects. Contrary to that work, our consensus networks take in distinct subsets of features as modalities. Co-training (Blum and Mitchell, 1998) and tri-training (Zhou and Li, 2005) use distinct subsets of features, but they use them to train distinct classifiers, and let the results directly supervise each other. Their approach ‘bootstrapped’ classifications based on a few labeled data, but our method explicitly uses a modality discriminator that enforces alignments between modalities.

Domain Adaptation In domain adaptation and multi-task learning, there have been many attempts to learn indistinguishable embeddings between domains. For example, Ganin et al. (2016) and Joty et al. (2017) applied a gradient reversal layer to let encoders minimize the domain classification loss. Baktashmotlagh et al. (2013) minimized the maximum-mean discrepancy (MMD) loss in a reproductive kernel Hilbert space (RKHS) of the latent representations. Motiian et al. (2017) used semantic similarity loss between latent representations of different class data to encourage alignments between domains. Liu et al. (2017) and Chen and Cardie (2018) used shared and private networks to learn information contained either commonly in domains or domain-specific. Our work is unique. First, there is only one domain in our problem setting. Second, we use iterative optimization to encourage discrepancies between domains. Third, we have two empirical improvements (noise modality and cooperative optimization) that make our Consensus Networks outperform traditional classifiers.

3 Methods

3.1 Dataset

We use DementiaBank, the largest publicly available dataset for detecting cognitive impairments. It includes verbal descriptions (and associated transcriptions) of the Cookie Theft picture description

task from the Boston Diagnostic Aphasia Examination (Becker et al., 1994). The version we have access to contains 240 speech samples labeled Control (from 98 people), 234 with AD (from 148 people), and 43 with MCI (from 19 people)³. All participants were older than 44 years.

3.2 Linguistic Features

The dataset contains narrative speech descriptions and their transcriptions. We preprocess them by extracting 413 linguistic features for each speech sample. These linguistic features are proposed by and identified as the most indicative of cognitive impairments by various previous works, including Roark et al. (2007), Chae and Nenkova (2009), Roark et al. (2011), Fraser et al. (2015), and Hernández-Domínguez et al. (2018). After calculating these features, we use KNN imputation to replace the undefined values (resulting from divide-by-zero, for example), and then normalize the features by their z -scores. The following are brief descriptions of these features, grouped by their natural categories. More detailed descriptions are included in the Appendix.

There are 185 acoustic features (e.g., average pause time), 117 syntactic features (e.g., Yngve statistics (Yngve, 1960) of the parse tree, computed by the LexParser in CoreNLP (Manning et al., 2014)), and 31 semantic features (e.g., cosine similarity between pairs of utterances) Moreover, we use 80 part-of-speech features that relate to both syntax and semantics but are here primarily associated with the latter.

Modality Division After representing each sample with a 413-dimensional vector \mathbf{x} consisting of all available linguistic features, we divide the vector into M partitions (‘modalities’) of approximately equal sizes $[\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M]$, according to the groups mentioned above. Unless mentioned otherwise (e.g., in the ablation study shuffling

³The version of DementiaBank dataset we acquired contains a slightly different number of samples from what some previous works used. In Control:AD, Fraser et al. (2015) used 233 Control and 240 AD samples; Yancheva and Rudzicz (2016) had 241 Control and 255 AD samples; Hernández-Domínguez et al. (2018) had 242 Control and 257 AD samples (with 10% control samples excluded from the evaluation). In Control:MCI, Santos et al. (2017) used all 43 transcriptions from MCI and 43 sampled from the Control group. With no clear descriptions of the sampling procedures, the constituents of the Control group might differ from our sample. In this paper, we run our models on the same tasks (i.e., Control:AD) and compare to the results of models used in the literature.

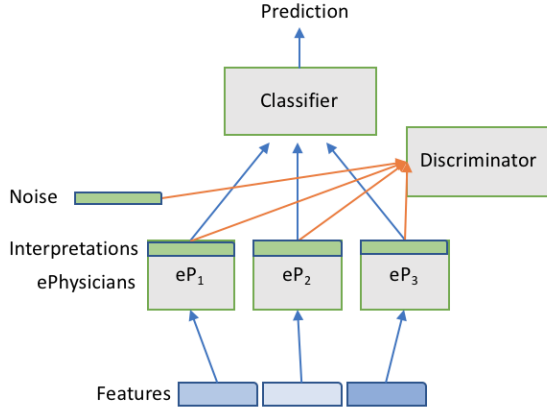


Figure 1: Overview of model structure when features (blue rectangles) are divided into three modalities (non-overlapping subsets). Each subset of features are passed into an “ePhysician” neural network whose outputs (green rectangles) are the representations. They are passed (one by one) into a “Discriminator” neural network and (after combined) into a “Classifier” network, respectively.

modalities), this is our default choice for assigning modalities.

3.3 Model

Figure 1 is an example of our model structure (with $M = 3$ modalities), and this section elaborates the inference procedure, the training algorithm, and our two improvements.

Inference With the extracted linguistic features divided into subsets by their modalities, each speech sample is described by $M = 3$ non-overlapping feature vectors $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$. These feature vectors are then passed into corresponding ePhysician networks, each outputting a vector \mathbf{i}_m , which is a distilled representation of the subject from a modal-specific perspective (e.g., the semantic). We also refer to it as the *interpretation vector* and use them interchangeably. Formally, the m^{th} ePhysician can be written as a function, $f_m(\cdot)$ generating the representation:

$$\mathbf{i}_m = f_m(\mathbf{x}_m)$$

To challenge the similarity of representations from different modalities, we let a discriminator neural network $f_D(\cdot)$ take in each of the M representations and predict the likelihood of the originating modality m .

$$P(m = k | \mathbf{i}) = \frac{e^{f_D(\mathbf{i})_k}}{\sum_k e^{f_D(\mathbf{i})_k}}$$

where $k = 1, \dots, M$.

To attempt a diagnosis, a classifier network $f_C(\cdot)$ takes in the combination of M representations of each speech sample, and outputs a prediction probability for detection result y :

$$P(y = l | \mathbf{x}) = \frac{e^{f_C(\mathbf{i}_{1..M})_l}}{\sum_l e^{f_C(\mathbf{i}_{1..M})_l}}$$

where $l \in \{0, 1\}$ for two-class classification (i.e., 0 for healthy and 1 for dementia). The predicted class corresponds to those with the highest probability:

$$\hat{y} = \arg \max_l P(y = l | \mathbf{x})$$

Optimization The training procedure optimizes the adversarial objective, and the conventional classifier objective:

- The adversarial objective sets up the ePhysicians and the Discriminator to work in an adversarial manners. The ePhysicians try to produce indistinguishable representations, while the discriminator tries to tell them apart.

$$\min_D \max_P \mathcal{L}_D \text{ where}$$

$$\mathcal{L}_D = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{m=1..M} \{-\log P(\hat{m} = m | \mathbf{i}_m)\} \quad (1)$$

- Make the classifier network as accurate as possible. This is done by minimizing the cross entropy classification loss:

$$\min_C \mathcal{L}_C \text{ where}$$

$$\mathcal{L}_C = \mathbb{E}_{\mathbf{x}} \{-\log P(\hat{y} = y | \mathbf{i}_{1..M})\} \quad (2)$$

Overall, $\min_C \mathcal{L}_C$ and $\min_D \max_P \mathcal{L}_D$ set up a complex optimization problem. We use iterative optimization steps, similar to Goodfellow et al. (2014).

There are two tricks that we found to improve the performance of the models. Namely, the *noise modality* and the *cooperative optimization*. We explain them below.

Noise Modality For each participant session, we add a “noise modality representation” \mathbf{i}_0 drawn from a Gaussian distribution with the mean and variance identical to those of other representation vectors.

$$\mathbf{i}_0 \sim \mathcal{N}(\mu_{\mathbf{i}_{1..M}}, \sigma_{\mathbf{i}_{1..M}}^2)$$

This additional representation vector is passed into the discriminator, but not passed into the classifier. The first optimization goal (1) is therefore:

$$\min_D \max_P \mathcal{L}_D \text{ where} \quad (3)$$

$$\mathcal{L}_D = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{m=0..M} \{-\log P(\hat{m} = m | \mathbf{i}_m)\}$$

To some extent, the noise representation vector works like a regularization mechanism to refrain the discriminator from making decisions based on superficial statistics. We show in 4.1 that this addition empirically improves classifier performance.

Cooperative Optimization When optimizing the classifier, we find that allowing gradients to propagate back to the ePhysicians improves the model’s overall performance. During optimization, the ePhysicians need to *cooperate* with the classifier (while adversarial to the discriminator). The second optimization goal (2) is therefore:

$$\min_{C,P} \mathcal{L}_C \text{ where} \quad (4)$$

$$\mathcal{L}_C = \mathbb{E}_{\mathbf{x}} \{-\log P(\hat{y} = y | \mathbf{i}_{1..M})\}$$

Implementation As a note of implementation, all ePhysicians, the classifier, and the discriminator networks are fully connected networks with Leaky ReLU activations (Nair and Hinton, 2010) and batch normalization (Ioffe and Szegedy, 2015). The hidden layer sizes are all 10 for all ePhysician networks, and there are no hidden layers for the discriminator or classifier networks. Although modalities might contain slightly different numbers of input dimensions, we do not scale the ePhysician sizes. This choice comes from the intuition that the ePhysicians should extract into the representation as similar information as possible. We use three Adam optimizers (Kingma and Ba, 2014), each corresponding to the minimization of ePhysician, Discriminator, and the Classifier, and optimize iteratively for no more than 100 steps. The optimization is stopped prior to step 100 if the classification loss \mathcal{L}_C converges (i.e., does not differ from the previous iteration by more than 1×10^{-4}) on training set. The train / validation / test set are divided randomly in 60/20/20 proportions.

Algorithm 1 The overall algorithm

```

1: Initialize the networks
2: for step := 1 to N do    ▷ N is a hyper-param
3:   for minibatch  $\mathbf{x}$  in training data  $\mathcal{X}$  do
4:     for modality  $m := 1$  to  $M$  do
5:        $\mathbf{i}_m = I_m(\mathbf{x}_m)$ 
6:       Sample the noise modality  $\mathbf{i}_0$ 
7:       Calculate  $\mathcal{L}_D$  with  $\mathbf{i}_{0..M}$ 
8:       Concatenate  $\mathbf{i}_{1..M}$  and calculate  $\mathcal{L}_C$ 
9:        $\min_{C,P} \mathcal{L}_C$     ▷ Cooperative optimization
10:       $\max_P \mathcal{L}_D$ 
11:     for  $k:=1$  to  $K$  do    ▷ K is a
                                hyper-param
12:        $\min_D \mathcal{L}_D$ 

```

4 Experiments

We first show the effectiveness of our two improvements to the model. Next, we do two ablation studies on the arrangements of modalities. Then, we evaluate our model against several traditional supervised learning classifiers used by state-of-the-art works. To understand the model further, we also visualize the principal components of the representation vectors throughout several runs.

4.1 Noise Modality Improves Performance

We compare a CN model with a noise modality to one without (with other hyper parameters including hidden dimensions and learning rates identical).

Table 1 shows that in the AD:MCI classification task, the model with an additional noise modality is better than the one without ($p = 0.04$ on 2-tailed t -test with 18 DoF). Here is a possible explanation. Without adding a noise modality, the discriminators may simply look at the superficial statistics, like the mean and variances of the representations. This strategy tends to neglect the detailed aspects encoded in the representation vectors. Adding in the noise modality penalizes this strategy and trains better discriminators by forcing them to *study the details*.

In the following experiments, all models contain the additional noise modality.

4.2 Effectiveness of Cooperative Optimization

The second improvement, cooperative optimization, also significantly improves model perfor-

Model	Micro F1	Macro F1
Noise	.7995 ± .0450	.7998 ± .0449
No noise	.7572 ± .0461	.7577 ± .0456

Table 1: Comparison of models with and without representations in noise modality. The models containing a Gaussian noise modality outperform those without.

mance. We compare Consensus Network classifiers trained with cooperative optimization (i.e., $\min_{C, P_{1..M}} \mathcal{L}_C$) to models with the same hyperparameters but trained non-cooperatively (i.e., $\min_C \mathcal{L}_C$). As shown in Table 2, the cooperative variant produces higher-score classifiers than the non-cooperative one ($p < 0.001$ on a 2-tailed t -test with 18 DoF). With the cooperative optimization setting, the ePhysicians are encouraged towards producing representations both indistinguishable (by the discriminator) and beneficial (for the classifier). Although the representations might agree less with each other, they could contain more *complementary* information, leading to better overall classifier performances.

In other experiments, all of our models use cooperative optimization.

Optimization	Micro F1	Macro F1
Non-coop	.6696 ± .0511	.6743 ± .0493
Cooperative	.7995 ± .0450	.7998 ± .0449

Table 2: Comparison of models optimized in cooperative and non-cooperative manner.

4.3 Effectiveness of Agreement Among Modalities

In this and the next experiment, we illustrate the effectiveness of our models on different configurations of modalities in an ablation study. We show that our models work well because of the effectiveness of the “consensus between modalities” scheme.

In this experiment, we compare our Consensus Network models (i.e., with agreements) with fully-connected neural network classifiers (i.e., without agreements) taking the same partial input features. The networks are all simple multiple layer perceptrons containing the same total number of neurons as the ‘classifier pipeline’ of our models (i.e., ePhysicians plus the classifier)⁴ with batch nor-

⁴For example, for models taking in two modalities, if our

malization between hidden layers. A few observations could be made from Table 3:

1. Some features from particular modalities are more expressive than others. For example, acoustic features could be used for building better classifiers than those in the semantic ($p = .005$ for 2-tailed t -test with 18 DoF) or syntactic modality ($p < .001$ for 2-tailed t -test with 18 DoF). More specifically, the syntactic features themselves do not tell much. We think this is because the syntactic features are largely based on the contents of the speech, and remain similar across speakers. For example, almost none of the speakers asked questions, giving zero values in “occurrences” of corresponding syntactic patterns.
2. Our model is able to utilize multiple modalities better than MLP. For MLP classifiers, combining features from different modalities does not always give better models. The syntactic modality features confuse MLP and “drag down” the accuracy. However, our models built with the consensus framework are able to utilize the less informative features from additional modalities. In all scenarios using two modalities, our models achieve accuracies higher than neural networks trained on *any* of the two individual modalities.
3. Given the same combinations of features, letting neural networks produce representation in agreement *does* improve the accuracy in all four scenarios⁵.

4.4 Dividing Features by Natural Modalities is Preferred

This is the second ablation study towards modality arrangement. We show that dividing features into subsets according to their natural modalities (i.e., the categories in which they are defined) is better than doing so randomly.

In this experiment, we train CNs on features grouped by either their natural modalities, or randomly divided. For natural groupings, we try to

model contain ePhysicians with one layer of 20 hidden neurons, the interpretation vector dimension 10, and classifier 5 neurons, then the benchmarking neural network contains three hidden layers with [20×2, 10×2, 5] neurons.

⁵ $p = 3 \times 10^{-12}$ on syntactic+semantic features, $p = 0.044$ on acoustic + semantic, $p = 0.005$ on acoustic + syntactic, and $p = 0.046$ on all modalities. All 18 DoF one-tailed t -tests.

Models (Modality)	Accuracy
MLP (Acoustic)	.7519 ± .0245
MLP (Syntactic)	.5222 ± .0180
MLP (Semantic)	.6987 ± .0278
MLP (Syntactic + Semantic)	.5819 ± .0216
CN (Syntactic + Semantic)	.7257 ± .0344
MLP (Acoustic + Semantic)	.7002 ± .1128
CN (Acoustic + Semantic)	.7542 ± .0433
MLP (Acoustic + Syntactic)	.6776 ± .0952
CN (Acoustic + Syntactic)	.7574 ± .0361
MLP (All 3 modalities)	.7528 ± .0520
CN (All 3 modalities)	.7995 ± .0450

Table 3: Performance comparison between Consensus Networks and fully-connected neural network classifiers having certain modality information.

let each group contain comparable number of features, resulting in the following settings:

- Two groups, natural: (a) acoustic + semantic, 216 features; (b) syntactic + PoS, 197 features.
- Three groups, natural: (a) acoustic, 185 features; (b) semantic and PoS, 111 features; and (c) syntactic, 117 features. This is the default configuration used in other experiments in this paper.

For random grouping, we divide the features into almost equal-sized 2/3/4 groups randomly. As shown in Table 4. The two natural modality division methods produce higher accuracies than those produced by any of the random modality division methods.

Modality division method	Accuracy
Three groups, random	.7408 ± .0340
Two groups, random	.7623 ± .0164
Four groups, random	.7666 ± .0141
Two groups, natural	.7769 ± .0449
Three groups, natural	.7995 ± .0450

Table 4: Performance comparison between different modality division methods, sorted by accuracy.

4.5 Visualizing the Representations

To further understand what happens inside consensus network models during training, we visualize the representation vectors with PCA. Figure 2 consists of the visualizations drawn from an ar-

bitrary trial in training the model. Each representation vector is shown on the figure as a data point, with its color representing its originating modality (including the noise modality).

Several common themes could be observed:

1. *The clusters are symmetric.* Initially the configurations of representations largely depend on the initializations of the network parameters. Gradually the representations of the same modality tend to form clusters. Optimizing the ePhysicians towards both targets make they compress modalities into representations which are symmetrical in an *aggregate* manner.
2. *The agreements are simple.* The variances explained by the first a few principal components usually increase as the optimizations proceed. When distilling information relevant to detection, the agreement tend to become simple.
3. *The agreements are imperfect.* As shown in Figure 2, the modal representations do not overlap. Also, the discriminator loss is low (usually at 10^{-3} when training is done). This confirms that these representations are still easily distinguishable. This may because the modalities inherently have some *complementary* information, leading to the ePhysicians projecting the modalities differently.
4. *The representations are complex.* Their shapes do not resemble the noise representations (Gaussian) lying at the center of the three petals. This shows that the representations are not simply Gaussian.
5. *The accuracy increases.* The accuracy in validation set generally increases as the training proceeds. Note that the distributions of representation vectors are increasingly similar in shape but remains distinct in spatial allocations. This confirms our conjecture that the information about cognitive impairment resides in complicated details instead of superficial statistics, which neural networks could represent.

4.6 Evaluation Against Benchmarks

With the previous sets of experiments, we have a best working architecture. We now evaluate it against traditional classifiers, which are used by

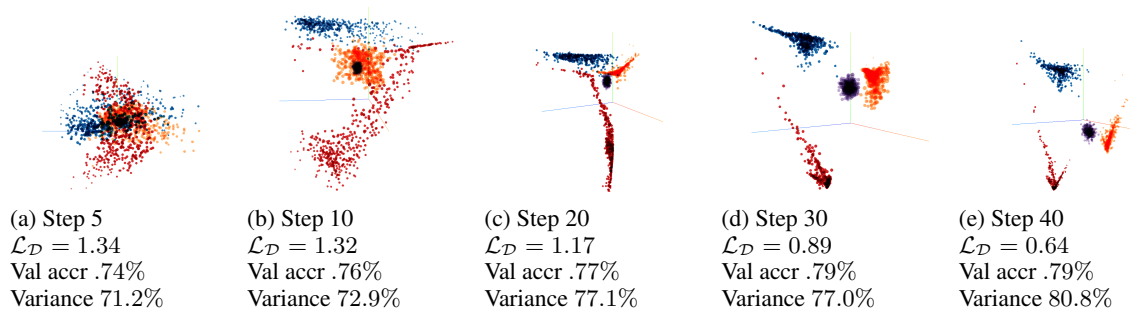


Figure 2: Initially, the representations from the three modalities are mixed. As the training go on, the three modalities gradually form three symmetric “petals”, while the (Gaussian) noise modality stays in the center. These petals do not overlap, as they contain complementary information when combined and passed into the classifier. Instead, their distributions become symmetric. (Best viewed in colours.)

the state-of-the-art papers (Hernández-Domínguez et al., 2018; Santos et al., 2017; Fraser et al., 2015) on our 413 features. Note that the results could be different from what they reported, because the feature sets are different.

We test several traditional supervised learning benchmark algorithms here: support vector machine (SVM), quadratic discriminant analysis (QDA), random forest (RF), and Gaussian process (GP). For completeness, multiple layer perceptrons (MLPs) containing all features as inputs are also mentioned in Table 5. On the binary classification task (healthy control vs. dementia), our model does better than them all.

Classifier	Micro F1	Macro F1
SVM	.4810 ± .0383	.6488 ± .0329
QDA	.5243 ± .0886	.5147 ± .0904
RF	.6184 ± .0400	.6202 ± .0422
GP	.6775 ± .0892	.6873 ± .0819
MLP	.7528 ± .0520	.7561 ± .0444
CN	.7995 ± .0450	.7998 ± .0449

Table 5: Comparison with different traditional classifiers in AD:Control classification task. In particular, our model has higher accuracy than the best traditional classifier, MLP.

5 Conclusion and Future Works

We introduce the consensus network framework, in which neural networks are encouraged to compress various modalities into indistinguishable representations (‘interpretation vectors’). We show that consensus networks, with the noise modality and cooperative optimization, improve upon traditional neural network baselines given the same features. Specifically, with all 413 lin-

guistic features, our models outperform fully-connected neural networks and other traditional classifiers used by state-of-the-art papers.

In the future, the “agreement among modalities” concept may be applied to design objective functions for training classifiers in various tasks, and from other data sets (for example, education and occupation “modalities” for the bank marketing prediction task). Furthermore, the mechanisms that represent linguistic features into symmetric spaces should be analyzed within the context of explainable AI.

References

- Mahsa Baktashmotlagh, Mehrtash T. Harandi, Brian C. Lovell, and Mathieu Salzmann. 2013. Unsupervised Domain Adaptation by Domain Invariant Projection. In *International Conference on Computer Vision (ICCV)*, pages 769–776. IEEE.
- James T Becker, François Boiler, Oscar L Lopez, Judith Saxton, and Karen L McGonigle. 1994. The natural history of Alzheimer’s disease: description of study cohort and accuracy of diagnosis. *Archives of Neurology*, 51(6):585–594.
- Suzanna Becker and Geoffrey E Hinton. 1992. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356):161.
- Jon Atli Benediktsson, Johannes R Sveinsson, Okan K Ersoy, and Philip H Swain. 1997. Parallel consensual neural networks. *IEEE Transactions on Neural Networks*, 8(1):54–64.
- Ernst R Berndt and Iain M Cockburn. 2013. Price indexes for clinical trial research: a feasibility study. Technical report, National Bureau of Economic Research.
- Avrim Blum and Tom Mitchell. 1998. Combining labeled and unlabeled data with co-training. In *Pro-*

- ceedings of the eleventh annual conference on Computational learning theory*, pages 92–100. ACM.
- Jieun Chae and Ani Nenkova. 2009. Predicting the fluency of text with shallow structural features: case studies of machine translation and human-written text. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, pages 139–147. Association for Computational Linguistics.
- Xilun Chen and Claire Cardie. 2018. Multinomial Adversarial Networks for Multi-Domain Text Classification. In *Proc. of NAACL*.
- Michael A Covington and Joe D McFall. 2010. Cutting the Gordian knot: The moving-average type–token ratio (MATTR). In *Journal of quantitative linguistics*, volume 17, pages 94–100. Taylor & Francis.
- Kathleen C Fraser, Jed A Metlzer, and Frank Rudzicz. 2015. Linguistic Features Identify Alzheimer’s Disease in Narrative Speech. *Journal of Alzheimer’s Disease* 49(2016)407-422.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, Victor Lempitsky, Urun Dogan, Marius Kloft, Francesco Orabona, and Tatiana Tommasi. 2016. Domain-Adversarial Training of Neural Networks. *Journal of Machine Learning Research*, 17:1–35.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Proceedings of Advances in neural information processing systems*, pages 2672–2680.
- Curry I Guinn and Anthony Habash. 2012. Language Analysis of Speakers with Dementia of the Alzheimer’s Type. In *AAAI Fall Symposium: Artificial Intelligence for Gerontechnology*, pages 8–13. Menlo Park, CA.
- Pamela Herd, Deborah Carr, and Carol Roan. 2014. Wisconsin longitudinal study (WLS). In *International journal of epidemiology*, volume 43, pages 34–41. Oxford University Press.
- Laura Hernández-Domínguez, Sylvie Ratté, Gerardo Sierra-Martínez, and Andrés Roche-Bergua. 2018. Computer-based evaluation of Alzheimer’s disease and mild cognitive impairment patients during a picture description task. *Alzheimer’s and Dementia: Diagnosis, Assessment and Disease Monitoring*, 10(3):260–268.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 448–456.
- Shafiq Joty, Preslav Nakov, Lluís Màrquez, and Israa Jaradat. 2017. Cross-language Learning with Adversarial Neural Networks: Application to Community Question Answering. In *CoNLL-2017*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Bernd Kortmann and Benedikt Szmrecsanyi. 2004. Global synopsis: morphological and syntactic variation in English. *A handbook of varieties of English*, 2:1142–1202.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2017. Adversarial Multi-task Learning for Text Classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1–10.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Andrea Moro, Valentina Bambini, Marta Bosia, Simona Anselmetti, Roberta Riccaboni, Stefano Cappa, Enrico Smeraldi, and Roberto Cavallaro. 2015. Detecting syntactic and semantic anomalies in schizophrenia. *Neuropsychologia*, 79.
- Saeid Motiian, Marco Piccirilli, Donald A. Adjeroh, and Gianfranco Doretto. 2017. Unified Deep Supervised Domain Adaptation and Generalization. In *International Conference on Computer Vision (ICCV)*.
- Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814.
- Ziad S Nasreddine, Natalie A Phillips, Valérie Bédirian, Simon Charbonneau, Victor Whitehead, Isabelle Collin, Jeffrey L Cummings, and Howard Chertkow. 2005. The Montreal Cognitive Assessment, MoCA: a brief screening tool for mild cognitive impairment. *Journal of the American Geriatrics Society*, 53(4):695–699.
- Zeinab Noorian, Chloé Pou-Prom, and Frank Rudzicz. 2017. On the importance of normative data in speech-based assessment. In *Proceedings of Machine Learning for Health Care Workshop (NIPS MLHC)*.
- Gordon E Peterson and Harold L Barney. 1952. Control methods used in a study of the vowels. In *The Journal of the acoustical society of America*, volume 24, pages 175–184. ASA.
- Chloé Pou-Prom and Frank Rudzicz. 2018. Learning multiview embeddings for assessing dementia. In *Empirical Methods in Natural Language Processing (EMNLP)*.

- A Rey. 1941. L'examen psychologique dans les cas d'encéphalopathie traumatique. (Les problems.). [The psychological examination in cases of traumatic encephalopathy. Problems.]. *Archives de Psychologie*, 28:215–285.
- Brian Roark, Margaret Mitchell, and Kristy Hollingshead. 2007. Syntactic complexity measures for detecting mild cognitive impairment. In *Proceedings of the Workshop on BioNLP 2007: Biological, Translational, and Clinical Language Processing*, pages 1–8. Association for Computational Linguistics.
- Brian Roark, Margaret Mitchell, John-Paul Hosom, Kristy Hollingshead, and Jeffrey Kaye. 2011. Spoken Language Derived Measures for Detecting Mild Cognitive Impairment. In *IEEE transactions on audio, speech, and language processing*. U.S. National Library of Medicine.
- Virginia R de Sa. 1994. Learning classification with unlabeled data. In *Proceedings of Advances in neural information processing systems*, pages 112–119.
- Leandro Santos, Edilson Anselmo Corrêa Júnior, Osvaldo Oliveira Jr, Diego Amancio, Letícia Mansur, and Sandra Aluísio. 2017. Enriching Complex Networks with Word Embeddings for Detecting Mild Cognitive Impairment from Speech Transcripts. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1284–1296, Vancouver, Canada. Association for Computational Linguistics.
- Greta Szatloczki, Ildiko Hoffmann, Veronika Vincze, Janos Kalman, and Magdolna Pakaski. 2015. Speaking in Alzheimer's Disease, is That an Early Sign? Importance of Changes in Language Abilities in Alzheimer's Disease. *Frontiers in Aging Neuroscience*, 7:195.
- Vanessa Taler, Ekaterini Klepousniotou, and Natalie A Phillips. 2009. Comprehension of lexical ambiguity in healthy aging, mild cognitive impairment, and mild Alzheimer's disease. In *Neuropsychologia*, volume 47, pages 1332–1343. Elsevier.
- Davy Weissenbacher, Travis A Johnson, Laura Wojtulewicz, Amylou Dueck, Dona Locke, Richard Caselli, and Graciela Gonzalez. 2016. Automatic prediction of linguistic decline in writings of subjects with degenerative dementia. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1198–1207.
- Maria Yancheva, Kathleen Fraser, and Frank Rudzicz. 2015. Using linguistic features longitudinally to predict clinical scores for Alzheimer's disease and related dementias. In *Proceedings of the 6th Workshop on Speech and Language Processing for Assistive Technologies (SLPAT 2015)*, pages 134–139, Dresden, Germany. Association for Computational Linguistics.
- Maria Yancheva and Frank Rudzicz. 2016. Vector-space topic models for detecting Alzheimer's disease. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 2337–2346.
- Victor H Yngve. 1960. A model and an hypothesis for language structure. In *Proceedings of the American philosophical society*, volume 104, pages 444–466. JSTOR.
- Shunan Zhao, Frank Rudzicz, Leonardo G. Carvalho, Cesar Marquez-Chin, and Steven Livingstone. 2014. Automatic detection of expressed emotion in Parkinson's Disease. *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, (May):4813–4817.
- Luke Zhou, Kathleen C. Fraser, and Frank Rudzicz. 2016. Speech recognition in Alzheimers disease and in its assessment. In *Interspeech 2016*, pages 1948–1952.
- Zhi-Hua Zhou and Ming Li. 2005. Tri-training: Exploiting unlabeled data using three classifiers. *IEEE Transactions on knowledge and Data Engineering*, 17(11):1529–1541.

A Appendices

Linguistic Features

Acoustic

- The fluency of speech. We quantify it with phonation rate, duration of pauses, and number of filled pauses (e.g., "um") of various lengths.
- Following the convention of speech processing literatures (Zhou et al., 2016; Yancheva et al., 2015; Zhao et al., 2014), we compute Mel-scaled cepstral coefficients (MFCCs) containing the amount of energy in 12 different frequency intervals for each time frame of 40 milliseconds, as well as their first- and second-order derivatives. We calculate the mean, variance, kurtosis, and skewness of the MFCCs and include them as features.

Semantic and Lexical

- Lexical norms, including age-of-acquisition, familiarity, imageability, and frequency (Taler et al., 2009). These are averaged over the entire transcript and specific PoS categories, respectively.
- Lexical richness, including moving-average type-token ratio over different window sizes

(Covington and McFall, 2010), Brunet’s index, and Honoré’s statistics (Guinn and Habash, 2012).

- Cosine similarity statistics (minimum, maximum, average, etc.) between pairs of utterances (represented as sparse vectors based on lemmatized words)
- Average word length, counts of total words, not-in-dictionary words, and fillers. The dictionary we use contains around 98,000 entries, including common words, plural forms of countable nouns, possessive forms of subjective nouns, different tenses of verbs, etc.

- The ratio of occurrences of several PoS tags, including noun-pronoun ratio.
- Number of occurrences of words in each of the five categories: subordinate (e.g: “because”, “since”, etc.), demonstratives (e.g: “this”, “that”), function (e.g: words with PoS tag CC, DT, and IN), light verbs (e.g: “be”, “have”), and inflected verbs (words with PoS tag VBD, VBG, VBN, and VBZ), borrowing the categorization method in Kortmann and Szmrecsanyi (2004)

Syntactic

- Composition of languages. We describe it by several features, including the average proportion of context-free grammar (CFG) phrase types⁶, the rates of these phrase types⁷, and the average phrase type length⁸ (Chae and Nenkova, 2009)
- The syntactic complexity of languages. We characterize it by the average heights of the context-free grammar (CFG) parse trees, across all utterances in each transcript. Each tree comes from an utterance parsed by a context free grammar parser (`LexParser` implemented in Stanford CoreNLP (Manning et al., 2014)). In addition, we compute the Yngve scores statistics of CFG parse trees (Yngve, 1960; Roark et al., 2007), where Yngve score is the degree of left-branching of each node in a parsed tree.
- Syntactic components. We describe them by including the number of occurrences of a set of 104 context-free production rules (e.g., $S \rightarrow VP$) in the CFG parse trees.

Part-of-speech

- The number of occurrences of part-of-speech (PoS) tags from Penn-treebank⁹.

⁶number of words in these types of phrases, divided by the total number of words in the transcript

⁷number of occurrences in a transcript, divided by the total number of words in the transcript

⁸number of words belonging to this phrase type in a transcript, divided by the occurrences of this phrase type in a transcript

⁹Using <https://spacy.io>