

Detecting Linguistic Characteristics of Alzheimer’s Dementia by Interpreting Neural Models

Sweta Karlekar Tong Niu Mohit Bansal
UNC Chapel Hill
{swetakar, tongn, mbansal}@cs.unc.edu

Abstract

Alzheimer’s disease (AD) is an irreversible and progressive brain disease that can be stopped or slowed down with medical treatment. Language changes serve as a sign that a patient’s cognitive functions have been impacted, potentially leading to early diagnosis. In this work, we use NLP techniques to classify and analyze the linguistic characteristics of AD patients using the DementiaBank dataset. We apply three neural models based on CNNs, LSTM-RNNs, and their combination, to distinguish between language samples from AD and control patients. We achieve a new independent benchmark accuracy for the AD classification task. More importantly, we next interpret what these neural models have learned about the linguistic characteristics of AD patients, via analysis based on activation clustering and first-derivative saliency techniques. We then perform novel automatic pattern discovery inside activation clusters, and consolidate AD patients’ distinctive grammar patterns. Additionally, we show that first derivative saliency can not only rediscover previous language patterns of AD patients, but also shed light on the limitations of neural models. Lastly, we also include analysis of gender-separated AD data.

1 Introduction

Alzheimer’s dementia is the most common form of dementia, caused by Alzheimer’s disease (AD). AD cannot be cured or reversed (Glennner, 1990). However, medication can be used to slow or halt degeneration especially when detected at an early stage. Current diagnoses often involve lengthy medical evaluations. One of the early symptoms of AD, cognitive impairment—which can be evidenced by issues with word-finding, impaired reasoning or judgment, and changes in language (McKhann et al., 1984)—is motivating lin-

guists and computer scientists to help quickly diagnose people afflicted by this disease.

This task is challenging because it requires diverse linguistic and world knowledge. For example, the sentence “Well...there’s a mother standing there uh uh washing the dishes and the sink is overflowing...overflowing” is AD-positive. To distinguish this from a control sample, one needs to know that the word “overspill” is not common in American English (Davies, 2009), and the speaker is correcting themselves by saying “overflowing...overflowing”, which hints on signs of confusion and memory loss (Duke et al., 2002). Moreover, different grammar patterns emerge based on the scenario at hand. In addition, the characteristics of AD-affected speech vary between stages of disease progression (König et al., 2015), making it harder for feature-based approaches to adapt.

Motivated by the shortcomings of manual feature-engineering for such a diverse and complex task, we first present three end-to-end neural models to address it. The first two are the widely adopted CNN and LSTM-RNN models, and the third is a stronger joint CNN-LSTM architecture. Our best-performing model requires only minimal feature engineering (namely automatic, commonly-used POS tags that are already present in the dataset) and establishes a new independent benchmark that outperforms previous AD classification scores.

More importantly, we next present interpretation results to explain what AD-relevant linguistic features these neural models are learning, via two visualization techniques: Activation Clustering (Girshick et al., 2014) and First Derivative Saliency (Simonyan et al., 2013), plus our novel approach of automatically discovering grammatical patterns common in different activation clusters. Furthermore, we split our dataset by gen-

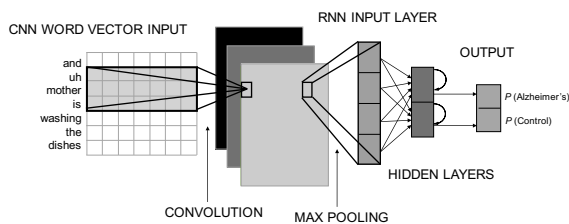


Figure 1: Our CNN-LSTM hybrid neural network.

der and analyze the performance of our model on each subsample of the data to illustrate that the features we find are not gender-specific. These methods not only help rediscover AD linguistic features that have been found by previous works—including short answers, bursts of speech, repeated requests for clarification, and starting with interjections—but also lead to new insights in AD characteristics via our automatic speech pattern extraction method. These findings could potentially help improve the accuracy and speed of medical diagnoses.

2 Related Works

Language-based Alzheimer’s Detection: Previous works using language to detect AD relied mainly on hand-crafted features from transcripts (Orimaye et al., 2017, 2015), occasionally using acoustic data (König et al., 2015; Rudzicz et al., 2014). The challenge with feature-based approaches is that they rely heavily on the researchers’ linguistics and medical expertise, and are also hard to generalize to other progression stages and disease types, which may correspond to different linguistic features. Hand-picked features may also become outdated as language and culture evolves. Moreover, some features may be too nuanced for humans to detect, especially at early stages of AD. In order to address these issues, Orimaye et al. (2016) adopted a deep neural network language model. However, a neural approach is usually a black-box and it is hard to interpret its reasoning for the final classification decisions. To make our approaches more interpretable while harvesting the benefits of neural approaches, we present three accurate neural models and include multiple visualization techniques to illustrate both their effectiveness and limitations.

CNN-LSTM on NL Classification: CNN and LSTM (Hochreiter and Schmidhuber, 1997) have both been leveraged extensively for extracting features in natural language. Our best performing model CNN-LSTM is closely related to C-LSTM by Zhou et al. (2015), where an LSTM is laid on

top of a CNN model. This model has been shown to perform better at sentiment classification than either of its integral parts alone.

Visualization Techniques for Neural Models:

There have been various visualization techniques proposed for neural networks in both Computer Vision (Krizhevsky et al., 2012; Simonyan et al., 2013; Zeiler and Fergus, 2014; Samek et al., 2017; Mahendran and Vedaldi, 2015) and NLP (Li et al., 2015; Kádár et al., 2017). In this work, we adopt two visualization techniques: Activation Clustering (Girshick et al., 2014) following the politeness interpretation work of Aubakirova and Bansal (2016), which leads to insight on sentence-level patterns, and First Derivative Saliency (Simonyan et al., 2013) following Li et al. (2015) and Aubakirova and Bansal (2016), which provides insight to the importance of each word in deciding the final classification label.

3 Classification Models

CNN: For each sentence, we apply an embedding and a convolutional layer, followed by a max-pooling layer (Collobert et al., 2011). The convolution features are obtained by applying filters of varying window sizes to each window of words. The result is then passed to a softmax layer that outputs probabilities over two classes.

LSTM-RNN: CNNs are not specialized for capturing long-range sequential correlations (Pascanu et al., 2013a). We thus also experimented with an LSTM-RNN model, which consists of an embedding layer followed by an LSTM layer. The final state, containing information from the entire sentence, is fed to a fully-connected layer followed by a softmax layer to obtain the output probabilities.

CNN-LSTM: Observing that both models achieve results comparable to previous best performing approach, and considering that they each have their own complementary strengths, we experimented with a combined architecture, laying an LSTM layer on top of CNN (See Figure 1). This CNN layer is identical to the vanilla CNN before the max-pooling layer, and the LSTM layer is identical to the vanilla LSTM-RNN after the embedding layer. More details are provided in the appendix.

4 Experimental Setup

Dataset This study utilizes Dementia-Bank (Boller and Becker, 2005), the largest publicly available dataset of transcripts and

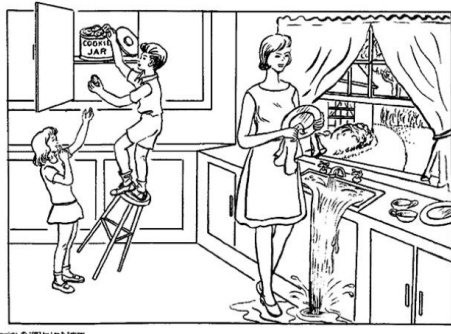


Figure 2: Boston cookie theft description task. Participants were asked to describe all events in the image.

audio recordings of AD (and control) patient interviews.¹ Patients were asked to perform various tasks; for example, in the “Boston Cookie Theft” description task, patients were shown an image and asked to describe what they see (See Figure 2). Other tasks include the ‘Recall Test’ in which patients were asked to recall attributes of a story they had been told previously. Each transcript in DementiaBank comes with automatic morphosyntactic analysis, such as standard part-of-speech tagging, description of tense, and repetition markers.² Note that these features are generic, automatically-extracted linguistic properties and are not AD-specific. We broke each transcript into individual utterances to use as data samples. Note that we also removed utterances that did not have accompanying POS tags. This balancing reduced the amount of data but ensured fair comparison between models with tagged and untagged setups.

Training Details Our CNN model was a 2-D convolutional neural network. Filter sizes of [3, 4, 5] were used. Our LSTM-RNN had 2 layers. The CNN-LSTM model had filter sizes [3, 4, 5, 6] and 1 LSTM hidden layer. For each model, all hyperparameters were tuned using the dev set.

See appendix for dataset and training details.

5 Results

With untagged data, our CNN, LSTM and CNN-LSTM models achieved an accuracy of 82.8%, 83.7% and 84.9%, respectively. When fed with the given POS-tagged data, our best-performing CNN-LSTM model achieved 91.1% in accuracy, setting a new benchmark accuracy for this task

Model	Details	Accuracy
2D-CNN	Non-Tagged Utterances	82.8
LSTM	Non-Tagged Utterances	83.7
CNN-RNN	Non-Tagged Utterances	84.9
CNN-RNN	POS-tagged Utterances	91.1

Table 1: Accuracy results of models. Note that we downsampled the data to remove utterances that did not have accompanying POS tags, so as to allow fair comparison between the tagged and untagged models.

(see Table 1 for more details).³ Compared to other related works, Orimaye et al. (2015, 2017) used AUC instead of accuracy, and König et al. (2015) did not use DementiaBank data. Rudzicz et al. (2014) achieved an accuracy of 67.0% on the DementiaBank dataset using audio features as well as transcripts. Orimaye et al. (2016) achieved an accuracy of 87.5% but only used 36 transcripts. Their test set is thus different from ours, leading to very little data and high variance. To the best of our knowledge, Orimaye et al. (2016) used the original transcripts, which include POS tags. Hence, most previous works on this topic are not directly comparable to our work, but we aim to establish a new independent, strong neural benchmark and then more importantly focus on visualization and interpretability of neural models.

Based on error analysis of the POS-based CNN/RNN model’s classification result, we found that almost all AD-positive results are classified correctly as AD-positive. However, there is more error in classifying non-AD samples, which could be due to the fact that DementiaBank includes patients with probable and possible AD, each exhibiting various degrees of symptoms. Patients who are AD-positive may still have partially unaffected speech (similar to non-AD control patients’ speech). However, because all utterances from AD-positive interviews are tagged as AD-positive, these seemingly unaffected utterances are still tagged AD-positive. To further understand the errors of our model, 10% of the wrongly classified non-AD examples were randomly selected and analyzed. Of this smaller sample, 36.3% were short utterances such as “okay”, “alright”, “oh my”, etc. These forms of speech are utterances that are present in both classes, but more commonly found in AD-positive cases. The remaining 63.7% were examples of speech that could be classified either way without surrounding context, such as “she’s

¹<http://dementia.talkbank.org>

²<http://talkbank.org/manuals/CHAT.docx>

³We also tried using a bidirectional RNN, which gave 84.7%, 86.2% and 91.1% accuracies for our LSTM-RNN, CNN-RNN, and CNN-RNN-tagged models, respectively.

Non-AD Clusters - Cookie Task					
POS	Freq	POS	Freq	POS	Freq
<i>n</i>	0.15	<i>n</i>	0.13	<i>n</i>	0.15
<i>det</i>	0.13	<i>det</i>	0.13	<i>det</i>	0.13
<i>presp</i>	0.07	<i>part</i>	0.09	<i>part</i>	0.10
<i>part</i>	0.05	<i>presp</i>	0.09	<i>presp</i>	0.10

Table 2: Top POS tags and frequencies for three non-AD clusters for the Cookie task.

drying dishes” and *“stool’s tipping over”*. Hence, future work could incorporate context from surrounding samples in each interview to help distinguish between temporarily unaffected speech patterns in AD-positive patients and the continually unaffected speech of non-AD control patients.

6 Analysis and Visualization

We first present gender analysis of our results and then present interpretation of the linguistic cues our CNN-LSTM model identified via two visualization strategies: activation clustering (Girshick et al., 2014) and first-derivative saliency heat maps (Simonyan et al., 2013).

6.1 Gender Differences

Many previous works have debated on the difference in language for male versus female patients with Alzheimer’s (Bayles et al., 1999; McPherson et al., 1999; Buckwalter et al., 1996; Ripich et al., 1995; Hebert et al., 2000; Heun and Kockler, 2002). In agreement with some of these previous works, we found that the sets of the top ten most common POS-tags for both AD-positive men and women are the same, i.e., we did not detect a significant difference in the language complexity or syntax of male and female patients with AD in our dataset. Moreover, our best performing model achieved 86.6% classification accuracy on solely the male data and 86.2% accuracy on solely the female data, demonstrating that it found no statistically significant difference between the AD-positive language of men versus women.⁴

6.2 Activation Clusters

Activation clustering (Girshick et al., 2014) treats the activation values of n neurons per input as coordinates in an n -dimensional space. K-means clustering is then performed to group together inputs that maximally activate similar neurons.

⁴Statistical insignificance calculated using the bootstrap test (Noreen, 1989; Efron and Tibshirani, 1994). We split our dataset based on gender and down-sampled the female data subset such that it had the same data-size and non-AD to AD data-ratio as the male data subset.

AD		Non-AD	
POS	Frequency	POS	Frequency
<i>n</i>	0.20	<i>n</i>	0.15
<i>det</i>	0.14	<i>det</i>	0.13
<i>adj</i>	0.05	<i>presp</i>	0.07
<i>adv</i>	0.04	<i>part</i>	0.05

Table 3: Top POS tags in AD cluster and non-AD cluster for Cookie task.

6.2.1 Rediscovering Existing Strategies

Our activation clusters corroborated previous studies, forming clusters around known linguistic characteristics of Alzheimer’s disease (Watson, 1999; Rudzicz et al., 2014).

Short Answers and Bursts of Speech Clusters formed around short answers, which have been split up by natural pauses in speech. {‘okay’, ‘and’, ‘yes’, ‘oh !’, ‘yes’, ‘fine’}

Repeated Requests for Clarification Another cluster formed around clarification questions and confusion about the task, specifically in the past tense. {‘Did I say facts ?’, ‘Did I get any ?’, ‘Did I say elephant ?’}

Starting with Interjections Many clusters contain utterances that start with interjections such as “oh”, “well”, “so”, and “right”. {‘Well I gotta see it’, ‘Oh I just see a lot of uh...’, ‘So all the words that you can...’}

6.2.2 Automatic Cluster Pattern Analysis

Next, we extend activation clustering to perform novel automatic pattern discovery inside different clusters, as opposed to manually looking for patterns as in Aubakirova and Bansal (2016). Finding the most common POS tags in each cluster allows us to better understand which grammatical structures are favored. No two clusters had exactly the same most-common POS tags, but many clusters shared similar top POS-tags on the same task.

Control Clusters An example unaffected speech cluster for the Cookie task has the following as the most common POS tags: [(‘n’, .15), (‘det’, .13), (‘presp’, .07), (‘part’, .05)].⁵ This pattern follows other clusters of unaffected speech, with nouns, determiners, and participles always found in the most-used POS tags. To illustrate this, two other control clusters found for the Cookie task have very similar top POS tags and frequencies: [(‘n’, .13), (‘det’, .13),

⁵POS tags used in this paper: v=verb, n=noun, pro=pronoun, adv=adverb, det=determiner, aux=auxiliary verb, prep=preposition, co=interjection, part=participle, presp=present participle. The frequencies of each POS tag are scaled based on the total number of tags in each cluster.

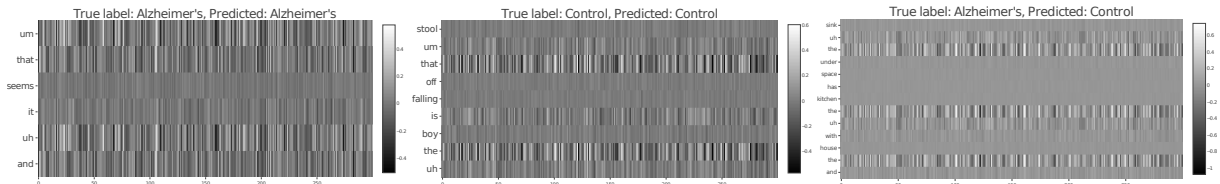


Figure 3: **Left + Middle:** first derivative saliency heat maps for correctly classified Alzheimer’s and control examples. **Right:** first derivative heat saliency maps for incorrectly classified Alzheimer’s example.

(‘*part*’, .09), (‘*presp*’, .09)] and [(‘*n*’, .15), (‘*det*’, .13), (‘*part*’, .10), (‘*presp*’, .10)] (see Table 2 for side-by-side comparison of these clusters). For the same task, non-AD clusters contain the same top four POS tags.

AD Clusters For the Recall task, one cluster of AD patients’ speech shows that the most common POS tags are [(‘*n*’, .15), (‘*co*’, .15), (‘*v*’, .06), (‘*pro*’, .06)]. Across samples from the Recall task, AD-marked clusters contained frequent interjections and verbs. On the other hand, in regards to the Cookie task, the most common POS tags for the AD cluster found are [(‘*n*’, .20), (‘*det*’, .14), (‘*adj*’, .05), (‘*adv*’, .04)], i.e., more adjectives and adverbs. Hence, between different tasks such as Cookie and Recall, the most commonly used POS tags for AD clusters are distinct.

Moreover, qualitative analysis shows dissimilarities in the most common POS tags between the Cookie task’s AD and non-AD cluster(s). Table 3 shows the comparison between representative AD and non-AD clusters for the Cookie task. The AD-positive cluster has only 2 most-used POS tags in common with the non-AD cluster. In fact, none of the 3 non-AD clusters found in Table 2 have adjectives or adverbs in their most-used POS tags list, unlike the AD cluster in Table 3.

6.3 First Derivative Saliency Heat Maps

Saliency heat maps (Simonyan et al., 2013) illustrate which words in an input had the biggest impact on the classification of the whole sentence. This is done by taking the gradient of the final scores w.r.t. the word embeddings of the inputs.

Heat Map Analysis The filler words “uh” and “um” are emphasized in Figure 3 (left), showing that they have a lot of influence on classification. The initial “and” is highlighted as well, corroborating the results of the activation clusters in that starting with a coordinating conjunction is a trait of Alzheimer’s speech. However, in Figure 3 (middle), the “uh” filler word is not highlighted, showing that most control patients do not use filler

words as heavily as Alzheimer’s patients. Instead, words that give structure to a sentence have the biggest impact on classification, such as definite articles and determiners (e.g., “the” and “that”). Figure 3 (middle) shows that the most highlighted words are “the”, “that” and “is”.

Visualizing Limitations Furthermore, by visualizing an incorrectly classified example, we can learn about the limitations of our neural network. Figure 3 (right) illustrates a map of an incorrectly predicted sample. The model misclassified it due to the length of the utterance (activation clustering showed that AD patients tend to have short bursts of speech, see Section 6.2.1) and the heavy use of determiners. However, the repeated “um”s and starting with a coordinating conjunction strongly indicated AD, which confused our model. From Figure 3, the repeated use of filler words (i.e. “uh”) had the second most influence on classification. In future work, with more context data and advanced neural methods, our model’s next steps will be to learn how to better classify samples that strongly exhibit both AD and control features.

7 Conclusion

We applied three models to the AD classification task, and our CNN-LSTM model achieves a new benchmark accuracy in classifying AD using neural models. We illustrate with two visualization techniques how these models capture unique linguistic features present in AD patients. We also discussed gender analysis. Potential future work includes using more conversational context and implementing multi-class classification to differentiate among stages of AD. We also plan to apply this generalizable model to other similar neurological diseases, such as Diffuse Lewy Body disease and Huntington’s disease (Heindel et al., 1989).

Acknowledgments

We thank the anonymous reviewers for their helpful comments. This work was supported by a Google Faculty Research Award, a Bloomberg Data Science Research Grant, an IBM Faculty Award, and NVidia GPU awards.

References

- Malika Aubakirova and Mohit Bansal. 2016. Interpreting neural networks to improve politeness comprehension. In *Proceedings of EMNLP*.
- Kathryn A Bayles, T Azuma, RF Cruz, CK Tomoeda, JA Wood, and Jr EB Montgomery. 1999. Gender differences in language of alzheimer disease patients revisited. *Alzheimer disease and associated disorders*, 13(3):138–146.
- Francois Boller and James Becker. 2005. Dementia-bank database guide. *University of Pittsburgh*.
- J Galen Buckwalter, Albert A Rizzo, Richard McCleary, Rodman Shankle, Malcolm Dick, and Victor W Henderson. 1996. Gender comparisons of cognitive performances among vascular dementia, alzheimer disease, and older adults without dementia. *Archives of Neurology*, 53(5):436–439.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Mark Davies. 2009. The 385+ million word corpus of contemporary american english (1990–2008+): Design, architecture, and linguistic insights. *International journal of corpus linguistics*, 14(2):159–190.
- Lisa M Duke, Ben Seltzer, Jennifer E Seltzer, and Jennifer J Vasterling. 2002. Cognitive components of deficit awareness in alzheimer’s disease. *Neuropsychology*, 16(3):359.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*, pages 1019–1027.
- Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587.
- George G Glenner. 1990. Alzheimers disease. In *Biomedical Advances in Aging*, pages 51–62. Springer.
- Alex Graves. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- Liesi E Hebert, Robert S Wilson, David W Gilley, Laurel A Beckett, Paul A Scherr, David A Bennett, and Denis A Evans. 2000. Decline of language among women and men with alzheimer’s disease. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 55(6):P354–P361.
- William C Heindel, David P Salmon, Clifford W Shults, Patricia A Walicke, and Nelson Butters. 1989. Neuropsychological evidence for multiple implicit memory systems: A comparison of alzheimer’s, huntington’s, and parkinson’s disease patients. *Journal of Neuroscience*, 9(2):582–587.
- R Heun and M Kockler. 2002. Gender differences in the cognitive impairment in alzheimer’s disease. *Archives of Women’s Mental Health*, 4(4):129–137.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Akos Kádár, Grzegorz Chrupała, and Afra Alishahi. 2017. Representation of linguistic form and function in recurrent neural networks. *Computational Linguistics*, 43(4):761–780.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Alexandra König, Aharon Satt, Alexander Sorin, Ron Hoory, Orith Toledo-Ronen, Alexandre Dereumaux, Valeria Manera, Frans Verhey, Pauline Aalten, Phillippe H Robert, et al. 2015. Automatic speech analysis for the assessment of patients with predementia and alzheimer’s disease. *Alzheimer’s & Dementia: Diagnosis, Assessment & Disease Monitoring*, 1(1):112–124.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.
- Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2015. Visualizing and understanding neural models in nlp. *arXiv preprint arXiv:1506.01066*.
- Aravindh Mahendran and Andrea Vedaldi. 2015. Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5188–5196.
- Guy McKhann, David Drachman, Marshall Folstein, Robert Katzman, Donald Price, and Emanuel M Stadlan. 1984. Clinical diagnosis of alzheimer’s disease report of the nincds-adrda work group* under the auspices of department of health and human services task force on alzheimer’s disease. *Neurology*, 34(7):939–939.
- Susan McPherson, Carla Back, J Galen Buckwalter, and Jeffrey L Cummings. 1999. Gender-related cognitive deficits in alzheimer’s disease. *International Psychogeriatrics*, 11(2):117–122.
- Sylvester O Orimaye, Jojo SM Wong, Karen J Golden, Chee P Wong, and Ireneous N Soyiri. 2017. Predicting probable alzheimers disease using linguistic deficits and biomarkers. *BMC bioinformatics*, 18(1):34.

- Sylvester Olubolu Orimaye, Kah Yee Tai, Jojo Sze-Meng Wong, and Chee Piau Wong. 2015. Learning linguistic biomarkers for predicting mild cognitive impairment using compound skip-grams. *arXiv preprint arXiv:1511.02436*.
- Sylvester Olubolu Orimaye, Jojo Sze-Meng Wong, and Judyanne Sharmini Gilbert Fernandez. 2016. Deep-deep neural network language models for predicting mild cognitive impairment. In *BAI@ IJCAI*, pages 14–20.
- Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2013a. How to construct deep recurrent neural networks. *arXiv preprint arXiv:1312.6026*.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013b. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pages 1310–1318.
- Danielle N Ripich, Stephen A Pettrill, Peter J Whitehouse, and Elaine W Ziol. 1995. Gender differences in language of ad patients a longitudinal study. *Neurology*, 45(2):299–302.
- Frank Rudzicz, Leila Chan Currie, Andrew Danks, Tejas Mehta, and Shunan Zhao. 2014. Automatically identifying trouble-indicating speech behaviors in alzheimer’s disease. In *Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility*, pages 241–242. ACM.
- Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, and Klaus-Robert Müller. 2017. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems*.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research*, 15(1):1929–1958.
- Caroline M Watson. 1999. An analysis of trouble and repair in the natural conversations of people with dementia of the alzheimer’s type. *Aphasiology*, 13(3):195–218.
- Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer.
- Chunting Zhou, Chonglin Sun, Zhiyuan Liu, and Francis Lau. 2015. A c-lstm neural network for text classification. *arXiv preprint arXiv:1511.08630*.

A Supplementary Materials

A.1 Training Details

All models had a vocabulary size of 2396, and used an Adam optimizer (Kingma and Ba, 2014) with a learning rate of $1e^{-4}$. All gradient norms were clipped to 2.0 (Pascanu et al., 2013b; Graves, 2013). For each model, the hyperparameters were tuned using the development set.

CNN We used a 2-D CNN. Filter sizes of [3, 4, 5] were used with 128 filters per filter size. Batch size was set to 128, and a dropout (Srivastava et al., 2014) of 0.80 was applied.

LSTM Our LSTM had 2 layers with 128 hidden units. Batch size was 32, and a dropout of 0.70 was used.

CNN-LSTM Our CNN-LSTM model consisted of an LSTM on top of a CNN. The CNN had 100 filters per filter size of [3, 4, 5, 6]. Embedding dimensions of 300 were used. An LSTM with 300 hidden units was used. Both dropout on the CNN and recurrent dropout (Gal and Ghahramani, 2016) on the LSTM used a dropout rate of 0.65.

A.2 Dataset Details

Each transcript in DementiaBank comes with automatic morphosyntactic analysis, such as standard part-of-speech tagging, description of tense, and markers for repetitions. This automatic tagging is identical to that done on other datasets, such as the CHILDES TalkBank, and is thus not specific to DementiaBank. This dataset features transcripts of 104 different control patients, and 208 different diagnosed dementia patients. There is a total of 1017 Alzheimer’s transcripts and 243 control transcripts. Each of these transcripts were then broken down by sentences and interruptions by the interviewer. We used each utterance by the patient as a data sample. Within the 14362 utterance samples, 11458 come from transcripts of Alzheimer’s-diagnosed interviewees and 2904 from those of control patients. Therefore, a majority-baseline classifier that always guesses *AD-positive* will achieve an accuracy of 79.8% in our dataset. Each utterance has a POS-tagged counterpart in the dataset. A 80/10/10 train/dev/test split was used for each setting.