

Detecting Dementia in Mandarin Chinese using Transfer Learning from a Parallel Corpus

Bai Li^{1,2}, Yi-Te Hsu^{1,2,4}, Frank Rudzicz^{1,2,3}

¹ University of Toronto, Toronto, Canada

² Vector Institute, Toronto, Canada

³ Toronto Rehabilitation Institute, Toronto, Canada

⁴ Academia Sinica, Taipei, Taiwan

{bai, eeder, frank}@cs.toronto.edu

Abstract

Machine learning has shown promise for automatic detection of Alzheimer’s disease (AD) through speech; however, efforts are hampered by a scarcity of data, especially in languages other than English. We propose a method to learn a correspondence between independently engineered lexicosyntactic features in two languages, using a large parallel corpus of out-of-domain movie dialogue data. We apply it to dementia detection in Mandarin Chinese, and demonstrate that our method outperforms both unilingual and machine translation-based baselines. This appears to be the first study that transfers feature domains in detecting cognitive decline.

1 Introduction

Alzheimer’s disease (AD) is a neurodegenerative disease affecting 5.7 million people in the US (Association et al., 2018), and is the most common cause of dementia. Although no cure yet exists, early detection of AD is crucial for an effective treatment to delay or prepare for its effects (Dubois et al., 2016). One of the earliest symptoms of AD is speech impairment, including a difficulty in finding words and changes to grammatical structure (Taler and Phillips, 2008). These early signs can be detected by having the patient perform a picture description task, such as the *Cookie Theft* task from the Boston Diagnostic Aphasia Examination (Goodglass and Kaplan, 1983).

Previous models have applied machine learning to automatic detection of AD, for example, Fraser et al. (2016) extracted a wide variety of lexicosyntactic and acoustic features to classify AD and obtained 82% accuracy on the DementiaBank (DB) dataset. However, clinical studies of AD are expensive, so datasets of patient data are often scarce. Noorian et al. (2017) augmented DB with

more a much larger corpus of normative data and improved the classification accuracy to 93% on DB. Similar linguistic differences between healthy and AD speech have been observed in Mandarin Chinese (Lai et al., 2009), but machine learning has not yet been applied to detecting AD in Mandarin.

Daume III (2007) proposed a simple way of combining features in different domains, assuming that the same features are extracted in each domain. In our case, ensuring consistency of features across domains is challenging because of the grammatical differences between Mandarin and English. For example, Mandarin doesn’t have determiners or verb tenses, and has classifiers, which don’t exist in English (Chao, 1965). Another method trains a classifier jointly on multiple domains with different features on each domain, by learning a projection to a common subspace (Duan et al., 2012). However, this method only accepts labelled samples in each domain, and cannot make use of unlabelled, out-of-domain data. Other work from our broader group (Fraser et al., 2019) combined English and French data by extracting features based on conceptual “information units” rather than words, thus limiting the effects of multilingual differences.

In the current work, we train an *unsupervised* model to detect dementia in Mandarin, requiring only the English DB dataset and a large parallel Mandarin-English corpus of normative dialogue. We extract lexicosyntactic features in Mandarin and English using separate pipelines, and use the OpenSubtitles corpus of bilingual parallel movie dialogues to learn a correspondence between the different feature sets. We combine this correspondence model with a classifier trained on DB to predict dementia on Mandarin speech. To evaluate our system, we apply it to a dataset of speech from Mandarin-speakers with dementia, and demon-




Figure 1: Diagram of our model. We train two separate models: the first is trained on OpenSubtitles and learns to map Mandarin features to English features; the second is trained on DementiaBank and predicts dementia given English features. During evaluation, the two models are combined to predict dementia in Mandarin.

strate that our method outperforms several baselines.

2 Datasets

We use the following datasets:

- **DementiaBank** (Boller and Becker, 2005): a corpus of *Cookie Theft* picture descriptions, containing 241 narrations from healthy controls and 310 from patients with dementia. Each narration is professionally transcribed and labelled with part-of-speech tags. In this work, we only use the narration transcripts, and neither the part-of-speech tags or raw acoustics.
- **Lu Corpus** (MacWhinney et al., 2011): contains 49 patients performing the *Cookie theft* picture description, category fluency, and picture naming tasks in Taiwanese Mandarin. The picture description narrations were human-transcribed; patient diagnoses are unspecified but exhibit various degrees of dementia.
- **OpenSubtitles2016** (Lison and Tiedemann, 2016): a corpus of parallel dialogues extracted from movie subtitles in various languages. We use the Traditional Chinese / English language pair, which contains 3.3 million lines of dialogue.

The Lu Corpus is missing specifics of diagnosis, so we derive a *dementia score* for each patient using the category fluency and picture naming tasks. For each category fluency task, we count the number of unique items named; for the picture naming tasks, we score the number of pictures correctly

named, awarding partial credit if a hint was given. We apply PCA to the scores across all tasks, and assign the first principal component to be the *dementia score* for each patient. This gives a relative ordering of all patients for degree of dementia, which we treat as the ground-truth for evaluating our models.

3 Methodology

3.1 Feature Extraction

We extract a variety of lexicosyntactic features in Mandarin and English, including type-token-ratio, the number of words per sentence, and proportions of various part-of-speech tags¹. A detailed description of the features is provided in the supplementary materials (Section A.1). In total, we extract 143 features in Mandarin and 185 in English. To reduce sparsity, we remove features in both languages that are constant for more than half of the dataset.

Due to the size of the OpenSubtitles corpus, it was computationally unfeasible to run feature extraction on the entire corpus. Therefore, we randomly select 50,000 narrations from the corpus, where each narration consists of between 1 to 50 contiguous lines of dialogue (about the length of a *Cookie Theft* narration).

For English, we train a logistic regression classifier to classify between dementia and healthy controls on DB, using our features as input. Using L1 regularization and 5-fold CV, our model achieves 77% classification accuracy on DB. This is slightly lower than the 82% accuracy reported

¹The feature extraction pipeline is open-source, available at: <https://github.com/SPOClab-ca/COVFEEFEE>. The `lex` and `lex_chinese` pipelines were used for English and Chinese, respectively.

by Fraser et al. (2016), but it does not include any acoustic features as input.

3.2 Feature Transfer

Next, we use the OpenSubtitles corpus to train a model to transform Mandarin feature vectors to English feature vectors. For each target English feature, we train a separate ElasticNet linear regression (Zou and Hastie, 2005), using the Mandarin features of the parallel text as input. We perform a hyperparameter search independently for each target feature, using 3-fold CV to minimize the MSE.

3.3 Regularization

Although the output of the ElasticNet regressions may be given directly to the logistic regression model to predict dementia, this method has two limitations. First, the model considers each target feature separately and cannot take advantage of correlations between target features. Second, it treats all target feature equally, even though some are noisier than others. We introduce two regularization mechanisms to address these drawbacks: reduced rank regression and joint feature selection.

Reduced Rank Regression

Reduced rank regression (RRR) trains a single linear model to predict all the target features: it minimizes the sum of MSE across all target features, with the constraint that the rank of the linear mapping is bounded by some given R (Izenman, 1975). Following recommended procedures (Davies, 1982), we standardize the target features and find the best value of R with cross validation. However, this procedure did not significantly improve results so it was not included in our best model.

Joint Feature Selection

A limitation of the above models is that they are not robust to noisy features. For example, if some English feature is useful for predicting dementia, but cannot be accurately predicted using the Mandarin features, then including this feature might hurt the overall performance. A desirable English feature in our pipeline needs to not only be useful for predicting dementia in English, but also be reconstructable from Mandarin features.

We modify our pipeline as follows. After training the ElasticNet regressions, we sort the target

features by their R^2 (coefficient of determination) measured on the training set, where higher values indicate a better fit. Then, for each K between 1 and the number of features, we select only the top K features and re-train the DB classifier (3.1) to only use those features as input. The result of this experiment is shown in Figure 2.

4 Experiments

4.1 Baseline Models

We compare our system against two simple baselines:

1. **Unilingual baseline:** using the Mandarin features, we train a linear regression to predict the dementia score. We take the mean across 5 cross-validation folds.
2. **Translate baseline:** The other intuitive way to generate English features from a Mandarin corpus is by using translation. We use Google Translate² to translate each Mandarin transcript to English. Then, we extract features from the translated English text and feed them to the dementia classifier described in section 3.1.

4.2 Evaluation Metric

We evaluate each model by comparing the Spearman’s rank-order correlation ρ (Spearman, 1904) between the ground truth dementia scores and the model’s predictions. This measures the model’s ability to rank the patients from the highest to the lowest severities of dementia, without requiring a threshold value.

4.3 Experimental Results

Model	Spearman ρ
Baselines	
Unilingual	0.385
Google Translate	0.366
Our models	
Feature Transfer	0.319
+ RRR	0.354
+ JFS	0.549

Table 1: Baselines compared with our models, evaluated on the Lu corpus. RRR: Reduced rank regression (3.3), JFS: Joint feature selection (3.3).

²<https://translate.google.com/>

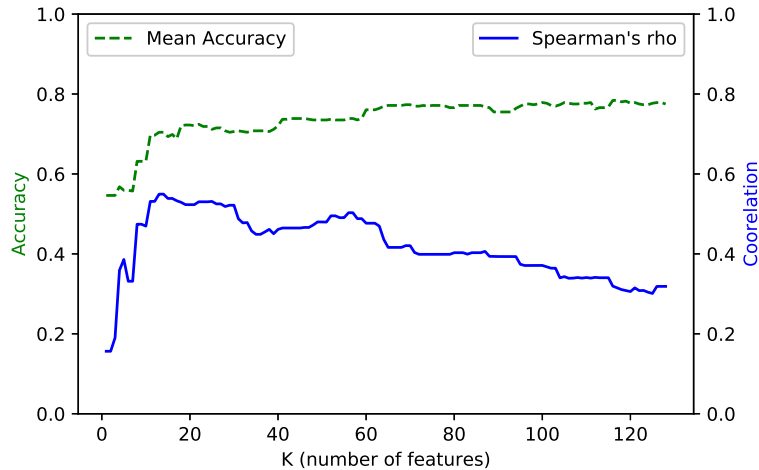


Figure 2: Accuracy of DementiaBank classifier model and Spearman’s ρ on Lu corpus, using only the top K English features ordered by R^2 on the OpenSubtitles corpus. Spearman’s ρ is maximized at $K = 13$, achieving a score of $\rho = 0.549$. DementiaBank accuracy generally increases with more features.

Our best model achieves a Spearman’s ρ of 0.549, beating the translate baseline ($n = 49$, $p = 0.06$). Joint feature selection appears to be crucial, since the model performs worse than the baselines if we use all of the features. This is the case no matter if we predict each target feature independently or all at once with reduced rank regression. RRR does not outperform the baseline model, probably because it fails to account for the noisy target features in the correspondence model and considers each feature equally important. We did not attempt to use joint feature selection and RRR at the same time, because the multiplicative combination of hyperparameters K and R would produce a multiple comparisons problem using the small validation set.

Using joint feature selection, we find that the best score is achieved when we use $K = 13$ target features (Figure 2). With $K < 13$, performance suffers because the DementiaBank classifier is not given enough information to make accurate classifications. With $K > 13$, the accuracy for the DementiaBank classifier improves; however, the overall performance degrades because it is given noisy features with low R^2 coefficients. A list of the top features is given in Table 2 in the supplementary materials.

In our experiments, the correspondence model worked better when absolute counts were used for the Chinese CFG features (e.g., the number of $NP \rightarrow PN$ productions in the narration) rather than ratio features (e.g., the proportion of CFG




Figure 3: Ablation experiment where a various number of OpenSubtitles samples were used for training. The error bars indicate the two standard deviation confidence interval.

productions that were $NP \rightarrow PN$). When ratios were used for source features, the R^2 coefficients for many target features decreased. A possible explanation is that the narrations have varying lengths, and dividing features by the length introduces a nonlinearity that adversely affects our linear models. However, more experimentation is required to examine this hypothesis.

4.4 Ablation Study

Next, we investigate how many parallel OpenSubtitles narrations were necessary to learn the correspondence model. We choose various training sample sizes from 10 to 50,000 and, for each training size, we train and evaluate the whole model from end-to-end 10 times with different random

seeds (Figure 3). As expected, the Spearman's ρ increased as more samples were used, but only 1000-2000 samples were required to achieve comparable performance to the full model.

5 Conclusion

We propose a novel method to use a large parallel corpus to learn mappings between engineered features in two languages. Combined with a dementia classifier model for English speech, we constructed a model to predict dementia in Mandarin Chinese. Our method achieves state-of-the-art results for this task and beats baselines based on unilingual models and Google Translate. It is successful despite the stark differences between English and Mandarin, and the fact that the parallel corpus is out-of-domain for the task. Lastly, our method does not require any Mandarin data for training, which is important given the difficulty of acquiring sensitive clinical data.

Future work will investigate the use of automatic speech recognition to reduce the need for manual transcripts, which are impractical in a clinical setting. Also, our model only uses lexicosyntactic features, and ignores acoustic features (e.g., pause duration) which are significant for dementia detection in English. Finally, it remains to apply this method to other languages, such as French (Fraser et al., 2019), for which datasets have recently been collected.

Acknowledgements

We thank Kathleen Fraser and Nicklas Linz for their helpful comments and earlier collaboration which inspired this project.

References

- Alzheimer's Association et al. 2018. 2018 Alzheimer's disease facts and figures. *Alzheimer's & Dementia*, 14(3):367–429.
- Francois Boller and James Becker. 2005. Dementia-bank database guide. *University of Pittsburgh*.
- Yuen Ren Chao. 1965. *A grammar of spoken Chinese*. Univ of California Press.
- Hal Daume III. 2007. Frustratingly easy domain adaptation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 256–263.
- PT Davies. 1982. Procedures for reduced-rank regression. *Applied Statistics*, pages 244–255.
- Lixin Duan, Dong Xu, and Ivor W Tsang. 2012. Learning with augmented features for heterogeneous domain adaptation. In *Proceedings of the 29th International Conference on Machine Learning*, pages 667–674. Omnipress.
- Bruno Dubois, Harald Hampel, Howard H Feldman, Philip Scheltens, Paul Aisen, Sandrine Andrieu, Hovagim Bakardjian, Habib Benali, Lars Bertram, Kaj Blennow, et al. 2016. Preclinical Alzheimer's disease: definition, natural history, and diagnostic criteria. *Alzheimer's & Dementia*, 12(3):292–323.
- Kathleen C. Fraser, Nicklas Linz, Bai Li, Kristina Lundholm Fors, Frank Rudzicz, Alexandra Konig, Jan Alexandersson, Philippe Robert, and Dimitrios Kokkinakis. 2019. Multilingual prediction of Alzheimers disease through domain adaptation and concept-based language modelling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Kathleen C Fraser, Jed A Meltzer, and Frank Rudzicz. 2016. Linguistic features identify Alzheimers disease in narrative speech. *Journal of Alzheimer's Disease*, 49(2):407–422.
- Harold Goodglass and Edith Kaplan. 1983. *Boston diagnostic examination for aphasia*, 2nd edition. Lea and Febiger, Philadelphia, Pennsylvania.
- Alan Julian Izenman. 1975. Reduced-rank regression for the multivariate linear model. *Journal of multivariate analysis*, 5(2):248–264.
- Dan Klein and Christopher D Manning. 2003. Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pages 423–430. Association for Computational Linguistics.
- Yi-hsiu Lai, Hsiu-hua Pai, et al. 2009. To be semantically-impaired or to be syntactically-impaired: Linguistic patterns in Chinese-speaking persons with or without dementia. *Journal of Neurolinguistics*, 22(5):465–475.
- Roger Levy and Christopher Manning. 2003. Is it harder to parse Chinese, or the Chinese treebank? In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pages 439–446. Association for Computational Linguistics.
- Pierre Lison and Jörg Tiedemann. 2016. Opensubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the 10th International Conference on Language Resources and Evaluation*.
- Xiaofei Lu. 2010. Automatic analysis of syntactic complexity in second language writing. *International journal of corpus linguistics*, 15(4):474–496.

Brian MacWhinney, Davida Fromm, Margaret Forbes, and Audrey Holland. 2011. AphasiaBank: Methods for studying discourse. *Aphasiology*, 25(11):1286–1307.

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)*.

Charles Spearman. 1904. The proof and measurement of association between two things. *The American journal of psychology*, 15(1):72–101.

Robyn Speer, Joshua Chin, Andrew Lin, Sara Jewett, and Lance Nathan. 2018. [Luminosinsight/wordfreq: v2.2](#).

Vanessa Taler and Natalie A Phillips. 2008. Language performance in Alzheimer’s disease and mild cognitive impairment: a comparative review. *Journal of clinical and experimental neuropsychology*, 30(5):501–556.

Hui Zou and Trevor Hastie. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320.

A Appendices

A.1 Description of Lexicosyntactic Features

We extract 185 lexicosyntactic features in English and 143 in Mandarin Chinese. We use Stanford CoreNLP to do constituency parsing and part-of-speech tagging (Klein and Manning, 2003; Levy and Manning, 2003). We also use `wordfreq` (Speer et al., 2018) for word frequency statistics in both languages. Our features are similar to the set of features used by Fraser et al. (2016), which the reader can refer to for a more thorough description.

The following features are extracted in English:

- **Narrative length:** Number of words and sentences in narration.
- **Vocabulary richness:** Type-token ratio, moving average type-token ratio (with window sizes of 10, 20, 30, 40, and 50 words), Honoré’s statistic, and Brunét’s index.
- **Frequency metrics:** Mean word frequencies for all words, nouns, and verbs.
- **POS counts:** Counts and ratios of nouns, verbs, inflected verbs, determiners, demonstratives, adjectives, adverbs, function words, interjections, subordinate conjunctions, and

coordinate conjunctions. Also includes some special ratios such as pronoun / noun and noun / verb ratios.

- **Syntactic complexity:** Counts and mean lengths of clauses, T-units, dependent clauses, and coordinate phrases as computed by Lu’s syntactic complexity analyzer (Lu, 2010).
- **Tree statistics:** Max, median, and mean heights of all CFG parse trees in the narration.
- **CFG ratios:** Ratio of CFG production rule count for each of the 100 most common CFG productions from the constituency parse tree.

The following features are extracted in Mandarin Chinese:

- **Narrative length:** Number of sentences, number of characters, and mean sentence length.
- **Frequency metrics:** Type-token ratio, mean and median word frequencies.
- **POS counts:** For each part-of-speech category, the number of it in the utterance and ratio of it divided by the number of tokens. Also includes some special ratios such as pronoun / noun and noun / verb ratios.
- **Tree statistics:** Max, median, and mean heights of all CFG parse trees in the narration.
- **CFG counts:** Number of occurrences for each of the 60 most common CFG production rules from the constituency parse tree.

A.2 Top Joint Features

Table 2 lists the top English features for joint feature selection (most reconstructable from Chinese features), ordered by R^2 coefficients on the Open-Subtitles corpus. The top performing model uses the first 13 features.

#	Feature Name	R^2
1	Number of words	0.894
2	Number of sentences	0.828
3	Brunét's index	0.813
4	Type token ratio	0.668
5	Moving average TTR (50 word window)	0.503
6	Moving average TTR (40 word window)	0.461
7	Moving average TTR (30 word window)	0.411
8	Average word length	0.401
9	Moving average TTR (20 word window)	0.360
10	Moving average TTR (10 word window)	0.328
11	NP → PRP	0.294
12	Number of nouns	0.233
13	Mean length of clause	0.225
14	PP → IN NP	0.224
15	Total length of PP	0.222
16	Complex nominals per clause	0.220
17	Noun ratio	0.213
18	Pronoun ratio	0.208
19	Number of T-units	0.207
20	Number of PP	0.205
21	Number of function words	0.198
22	Subordinate / coordinate clauses	0.193
23	Mean word frequency	0.193
24	Number of pronouns	0.191
25	Average NP length	0.188

Table 2: Top English features for joint feature selection.