

# Forgotten Words: Benchmarking NeoBERT for Dementia Detection in Low-Resource Conversational Filipino and English Speech

Rez Samantha Z. Floresca<sup>1\*</sup>, Edric Castel C. Hao<sup>2†‡</sup>, Hannah Grachiella Buñales<sup>1</sup>, Chelsea Dominique E. Temprosa<sup>1</sup>, Georgianna Z. Reyes<sup>1</sup>, Kervin Gabriel L. Chua<sup>1</sup>

<sup>1</sup>Ateneo de Manila Senior High School    <sup>2</sup>Analog Devices, Inc.  
rez.samantha.floresca@student.ateneo.edu

## Abstract

Dementia detection from spontaneous speech offers a scalable approach to cognitive screening, yet NLP systems remain predominantly English-centric. This limitation is especially acute in the Philippines, where Filipino–English code-switching is pervasive and no prior work has addressed NLP-based dementia detection. We present the first systematic evaluation of transformer-based dementia detection in Filipino speech and the first assessment of NeoBERT in a clinical NLP setting. To separate language from domain effects, we construct a parallel bilingual dataset of 4,000 DementiaBank-derived transcripts, with Filipino translations produced manually to preserve discourse-level markers of cognitive decline. We evaluate five model families, TF-IDF + LogReg, BERT, NeoBERT, XLM-R, and RoBERTa-Tagalog, under monolingual, zero-shot cross-lingual, and bilingual fine-tuning settings. We find that in-domain performance does not transfer across languages, with English-trained BERT dropping to Macro-F1 = 0.455 on Filipino, and that architectural modernization alone does not improve robustness. Bilingual fine-tuning, however, eliminates cross-lingual degradation across all transformer models, converging to Macro-F1 = 0.969–0.973. These results suggest that multilingual clinical NLP performance is driven primarily by linguistic coverage during training rather than model scale or architecture.

## 1 Introduction

Dementia detection from speech is a promising direction for scalable cognitive screening, as linguistic and discourse-level patterns often reflect early cognitive decline (Peled-Cohen and Reichart,

2025). However, most NLP systems for this task are developed and evaluated primarily on English-language data, limiting their robustness in multilingual and low-resource settings. This limitation is particularly relevant in the Philippines, where everyday speech frequently involves Filipino–English code-switching and informal conversational structure. As a result, it remains unclear whether failures arise from cross-lingual transfer limitations or English-centric biases in model pretraining, motivating our focus on conversational transcripts to isolate discourse-level markers of cognitive decline.

Progress in Filipino NLP, while growing, has been largely confined to classification tasks on news articles and social media posts, including sentiment analysis (Villavicencio et al., 2021), named entity recognition (Miranda, 2023), and hate speech detection (Cruz and Cheng, 2022). Even recent evaluation benchmarks such as Batayan (Montalan et al., 2025) and FilBench (Miranda et al., 2025), while broader in scope, remain anchored to common written-text tasks such as reading comprehension, natural language inference, and generation. Critically, none of this work addresses naturalistic speech, clinical discourse, or cognitive diagnostic tasks. In particular, Tagalog–English code-switching, pervasive in informal and clinical Filipino communication (Herrera et al., 2022), is systematically absent from existing evaluations. As a result, it remains unknown whether transformer-based models can capture the discourse-level degradation characteristic of dementia in Filipino speech.

Additionally, transformer-based encoders have become the dominant paradigm in clinical NLP (Devlin et al., 2019; Pappagari et al., 2020), yet their application to dementia detection remains largely confined to architectural variants that differ only in pretraining data. ClinicalBERT (Alsentzer et al., 2019), AD-BERT (Mao et al.,

\*Lead author.

†Provided research mentorship and guidance.

‡This work does not relate to the author’s position at Analog Devices, Inc.

2023), and mBERT (Pires et al., 2019) all inherit the original BERT design unchanged. This leaves open the question of how architectural modifications affect the extent to which models capture cognitively diagnostic linguistic properties relevant to dementia detection.

In this work, we study dementia detection under controlled bilingual conditions using conversational speech transcripts in English and Filipino. The full experimental setup is described in Section 3.

Our contributions are threefold: (1) We present the first evaluation of NLP-based dementia detection in Filipino speech, using conversational transcripts manually translated from English under controlled domain conditions. (2) We provide the first empirical characterization of NeoBERT’s behavior under clinical and cross-lingual conditions. (3) We analyze cross-lingual transfer from English to Filipino under controlled conditions, benchmarking English-only models against multilingual (XLM-RoBERTa) and language-matched (RoBERTa-Tagalog) baselines to quantify the effect of English-centric pretraining under language shift. Our code is publicly available.<sup>1</sup>

## 2 Related Work

### 2.1 Linguistic Markers of Cognitive Decline

Spontaneous speech in dementia and mild cognitive impairment (MCI) exhibits reduced lexical diversity, increased disfluency, syntactic simplification, and degraded referential coherence, patterns that manifest measurably before clinical diagnosis (Fraser et al., 2016; Segkouli et al., 2025) and have been linked to early tau accumulation in language-relevant cortical regions (Marier et al., 2026). A recent systematic review confirms that these markers are consistently recoverable from text alone, with linguistic-only models reaching an average diagnostic accuracy of 83% across 51 NLP studies of dementia detection (Shankar et al., 2025).

However, existing work is heavily concentrated on monolingual English speech elicited under controlled clinical tasks, most commonly the Cookie Theft picture description task from DementiaBank (Becker et al., 1994). Bilingualism introduces competing forces:

<sup>1</sup><https://github.com/rezsam09/Filipino-English-Dementia-Classification>

lifelong code-switching is associated with delayed dementia onset through enhanced cognitive reserve (Berkes and Bialystok, 2022), while the syntactic mixing and lexical borrowing characteristic of Taglish may alter the observable linguistic signals on which current models rely (Flores, 2019). This makes Filipino a theoretically significant yet entirely absent deployment context in the dementia NLP literature.

### 2.2 Transformer-Based Models in Clinical NLP

In dementia detection, most systems adapt BERT through domain-specific pretraining: ClinicalBERT (Alsentzer et al., 2019) adapts to clinical notes, AD-BERT (Mao et al., 2023) targets Alzheimer’s progression prediction, and mBERT (Pires et al., 2019) extends to multiple languages. These models differ mainly in pretraining data while sharing the same architecture, leaving open the question of whether architectural changes improve cross-lingual clinical performance.

Representation pooling is another underexplored yet consequential design choice. Most existing systems default to [CLS]-token pooling, a choice optimized for the pretraining objective but not for aggregating task-relevant information across the full sequence. Reimers and Gurevych (2019) demonstrate that mean pooling over final hidden states produces more stable and transferable sentence representations, particularly under domain and language shift, motivating our adoption of mean pooling across all transformer-based models.

### 2.3 Cross-Lingual Dementia Detection

Cross-lingual transfer in dementia detection has been explored almost exclusively between typologically similar European languages. The ADReSS-M challenge (Luz et al., 2024), including work by Tamm et al. (Tamm et al., 2023) and (DeLoach, 2026), established English-to-Greek transfer as a benchmark setting. Greek, however, shares significant Indo-European typological properties with English, limiting the generalizability of these findings to more distant transfer settings such as English-to-Filipino. Studies on dementia detection in low-resource, non-European languages remain rare. Adhikari et al. (2022) applied NLP-based detection to English DementiaBank transcripts translated into Nepali, while Lu and Chen (2025) extended this

line of work to Amis, an indigenous Taiwanese language, finding that augmenting a small dataset of 80 scripts with translated English transcripts improved both robustness and accuracy. Neither study examines cross-lingual transfer nor evaluates multilingual pretraining strategies. No work has examined dementia detection in a Southeast Asian language or in any language spoken by a population where bilingualism and code-switching are the norm rather than the exception. Filipino, spoken by over 100 million people and characterized by pervasive Filipino–English code-switching (Herrera et al., 2022), represents exactly this gap.

## 2.4 Multilingual and Language-Matched Pretraining

Multilingual pretraining has emerged as a common strategy for cross-lingual transfer in low-resource settings. XLM-RoBERTa (Conneau et al., 2020), trained on over 2.5 TB of filtered CommonCrawl data across 100 languages, substantially outperforms mBERT on low-resource cross-lingual benchmarks. However, even strong multilingual models do not fully capture the lexical and syntactic properties of Filipino: RoBERTa-Tagalog (Cruz and Cheng, 2022), pretrained on TLUnified, a large-scale, thematically diverse Filipino corpus, demonstrates consistent gains over XLM-RoBERTa on Filipino classification benchmarks, suggesting that language-matched pretraining captures aspects of Filipino that multilingual models do not adequately model. Whether this advantage transfers to a clinical domain, where training data is scarce and domain shift compounds language shift, remains entirely unknown.

## 2.5 Architectural Modernization of Encoder Models

NeoBERT (Breton et al., 2025) represents the most recent advance in encoder architecture, replacing absolute positional embeddings with Rotary Positional Embeddings (RoPE) (Su et al., 2021), adopting Pre-LayerNorm (Xiong et al., 2020) and RMSNorm (Zhang and Sennrich, 2019) for training stability, introducing SwiGLU in place of GELU activation (Shazeer, 2020), and optimizing depth-to-width ratio for parameter efficiency. Importantly, it remains exclusively English-pretrained on RefinedWeb, a corpus of 600B tokens, 18 times larger than RoBERTa’s training data. Despite outperforming BERT-large,

RoBERTa-large, and ModernBERT on MTEB under identical fine-tuning conditions (Breton et al., 2025), NeoBERT has not been evaluated in any clinical or cross-lingual setting, leaving open the question of whether its architectural advances translate into robustness under the combined domain and language shift characteristic of low-resource clinical NLP.

## 3 Experimental Design

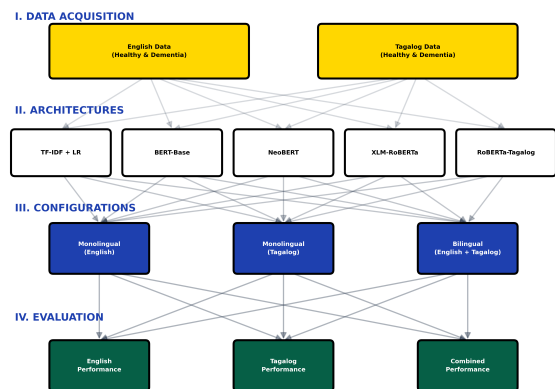


Figure 1: Overview of the experimental pipeline across datasets, model families, training configurations, and evaluation settings.

Performance improvements in NLP are often confounded by simultaneous changes in data, preprocessing, and model architecture (Sculley et al., 2015), a problem amplified in clinical NLP by small dataset sizes, domain sensitivity, and cross-linguistic variability (Fraser et al., 2016; Eyigoz et al., 2020). To mitigate this, all experiments use identical preprocessing pipelines with no external feature engineering beyond model-inherent representations, ensuring that observed performance differences reflect learned representations rather than preprocessing choices. The following subsections describe the dataset construction, preprocessing pipeline, model selection, and evaluation protocol.

### 3.1 Dataset Construction and Preprocessing

We construct a balanced bilingual dataset of 4,000 conversational transcripts for binary dementia classification, comprising 1,000 samples per class per language. All samples originate exclusively from DementiaBank (Becker et al., 1994), ensuring that any observed cross-lingual performance degradation reflects linguistic shift rather than

differences in topic distribution, discourse register, or clinical task structure.

### 3.1.1 English Data Sources

English dementia-positive and healthy control samples are drawn directly from DementiaBank (Becker et al., 1994), a widely used corpus in clinical speech research consisting of elicited speech from structured clinical tasks, most notably the Cookie Theft picture description task. This controlled elicitation procedure yields comparable discourse conditions across participants (Fraser et al., 2016), ensuring that class-level differences reflect genuine cognitive and linguistic variation rather than task-induced confounds.

### 3.1.2 Filipino Data Sources

Filipino samples are constructed through controlled manual translation of the complete set of 2,000 DementiaBank transcripts into Filipino, matching the English subset on clinical content, discourse structure, elicitation task, and class distribution. This ensures that cross-lingual evaluation is not confounded by asymmetric domain exposure. Manual translation is both necessary and principled given the absence of large-scale annotated clinical dementia corpora in Filipino (Mave et al., 2018). Translators are given explicit instructions to retain discourse-level markers of cognitive decline, including repetitions, hesitations, false starts, and syntactic degradation. Machine translation is avoided, since neural systems are known to normalize disfluent speech toward fluent output (Salesky et al., 2019), systematically erasing the features that distinguish dementia from healthy speech. All samples are anonymized prior to training in accordance with standard clinical data handling protocols.

### 3.1.3 Preprocessing Pipeline

All transcripts undergo a standardized preprocessing pipeline applying Unicode normalization, whitespace normalization, and lowercasing. Disfluencies, including filled pauses, repetitions, and hesitation markers, are retained throughout, as these are established linguistic correlates of cognitive impairment (Fraser et al., 2016) and primary classification signals in dementia detection. Stemming, lemmatization, and syntactic parsing are avoided, since such transformations risk normalizing disfluency structures and degrading the diagnostic signal

available to downstream models. All input sequences are truncated to a maximum length of 128 tokens, consistent with the short, task-elicited structure of DementiaBank transcripts and with findings that neurocognitive markers are recoverable from short windows of speech (Fraser et al., 2016).

### 3.1.4 Models and Baselines

We evaluate five models across three pretraining paradigms. TF-IDF with Logistic Regression serves as a lexical baseline to assess whether surface-level token statistics suffice. BERT-base-uncased (Devlin et al., 2019) and NeoBERT (Breton et al., 2025) represent English-only pretraining, differing in architecture and corpus scale to isolate design capacity from language coverage. XLM-RoBERTa (Conneau et al., 2020) provides a 100-language multilingual baseline, while RoBERTa-Tagalog (Cruz and Cheng, 2022) evaluates language-matched pretraining on native Filipino text. These models decompose cross-lingual robustness along two orthogonal axes: architecture and language coverage. Because all models share an identical pipeline, cross-lingual performance variations reflect pretraining language exposure rather than fine-tuning differences.

**Training Configurations:** Models are trained separately for three setups: (1) English-only ( $\mathcal{D}_{EN}$ ), (2) Filipino-only ( $\mathcal{D}_{TL}$ ), and (3) Bilingual ( $\mathcal{D}_{EN} \cup \mathcal{D}_{TL}$ ), enabling controlled analysis of monolingual capacity, cross-lingual transfer, and bilingual transfer.

## 3.2 TF-IDF with Logistic Regression

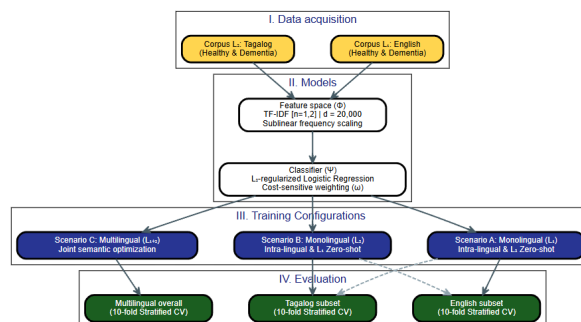


Figure 2: TF-IDF + Logistic Regression baseline training pipeline for evaluating lexical separability under cross-lingual and clinical domain shift conditions

We use TF-IDF with Logistic Regression as a lexical baseline to assess whether dementia-related

signals are captured by surface-level token statistics alone. TF-IDF produces sparse lexical representations by weighting tokens based on frequency statistics while ignoring word order and syntactic structure. Logistic Regression then learns a linear decision boundary over this high-dimensional feature space. This combination is widely used as an interpretable baseline in text classification tasks (Guo et al., 2020; Wang and Manning, 2012; Zhang et al., 2015). Performance of this model provides a reference for lexical separability: strong results suggest that class information is recoverable from surface-level features alone, whereas a weaker performance indicates the need for contextual modeling.

### 3.2.1 TF-IDF Representation

We construct unigram and bigram TF-IDF features ( $n \in 1, 2$ ), capturing both lexical identity and short disfluency-related phrases. Bigrams are sufficient given the short, utterance-level structure of clinical transcripts and have been shown to yield consistent gains in text classification tasks (Wang and Manning, 2012). We apply standard TF-IDF preprocessing settings: (1) sublinear term-frequency scaling ( $1 + \log(tf)$ ) to reduce the influence of repeated tokens common in disfluent speech; (2) a minimum document frequency of 2 (absolute count) to remove rare transcription noise; (3) a maximum document frequency of 0.95 to filter non-informative high-frequency tokens; and (4) a maximum vocabulary size of 20,000, consistent with standard practice in text classification (Zhang et al., 2015). These constraints ensure a stable lexical feature space that is comparable across languages and avoids overfitting in low-resource settings.

### 3.2.2 Logistic Regression Classifier

We use  $\ell_2$ -regularized logistic regression due to its convex optimization properties and strong performance in high-dimensional sparse settings. Optimization is performed using the `liblinear` solver for efficient and deterministic convergence on sparse data, with a maximum of 2,000 iterations to ensure convergence.

## 3.3 Transformer-Based Models

We fine-tune four transformer-based models for binary dementia classification: BERT-base-uncased, NeoBERT, XLM-RoBERTa, and RoBERTa-Tagalog. Together, these models span three pretraining paradigms: English-only,

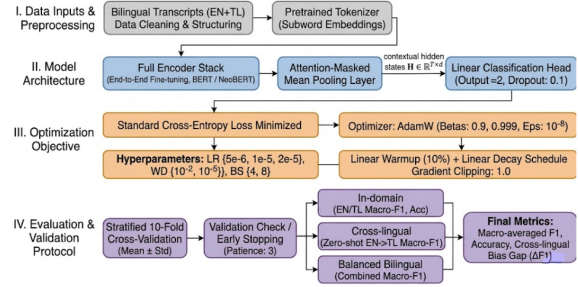


Figure 3: Transformer Models Fine-tuning Pipeline

massively multilingual, and language-matched, enabling a systematic decomposition of cross-lingual robustness along both architectural and pretraining dimensions.

### 3.3.1 Model Architecture and Fine-Tuning

All models follow a standard transformer-based classification pipeline. Given an input transcript  $x$ , tokens are first mapped to subword embeddings using a model-specific pretrained tokenizer, producing input IDs and attention masks. The resulting sequence is passed through the encoder, which is fine-tuned end-to-end without frozen layers. A single sentence representation is obtained via attention-masked mean pooling over the final hidden states:

$$h = \frac{\sum_{i=1}^T m_i H_i}{\sum_{i=1}^T m_i}$$

where  $m_i$  denotes the attention mask, ensuring that padding tokens do not contribute to the sentence representation. Unlike CLS pooling, which is optimized for pretraining objectives, masked mean pooling aggregates distributed token-level information and has been shown to improve robustness under domain shift (Reimers and Gurevych, 2019). The resulting representation is passed to a dropout layer ( $p = 0.1$ ) followed by a linear classification head. The encoder is initialized from pretrained weights, while the classification head is randomly initialized.

### 3.3.2 Optimization Objective

We minimize standard cross-entropy loss:

$$\mathcal{L}(\theta) = -\mathbb{E}_{(x,y)} \log p_{\theta}(y | x)$$

Optimization is performed using AdamW (Loshchilov and Hutter, 2019), which decouples weight decay from gradient updates and improves generalization in transformer fine-tuning settings.

### 3.3.3 Hyperparameter Selection

Hyperparameters are selected via a systematic grid search designed to avoid loss divergence and unstable gradient updates on small clinical datasets. The search space for all BERT models includes: (1) learning rates  $\{5 \times 10^{-6}, 6 \times 10^{-6}, 1 \times 10^{-5}, 2 \times 10^{-5}, 3 \times 10^{-5}\}$  (2) weight decay  $\{10^{-2}, 10^{-5}\}$  (3) batch sizes  $\{4, 8\}$ . These ranges are motivated by the constraints of low-resource clinical fine-tuning. Lower learning rates help prevent catastrophic forgetting of pretrained representations, while varying weight decay controls the balance between task-specific adaptation and robustness to noise introduced by manual transcription and translation. We use a stratified 70–15–15 train/validation/test split. The training set is used for model fitting, the validation set for early stopping, and the test set is used for hyperparameter selection. Small batch sizes (4, 8) are chosen to improve generalization through gradient stochasticity (Masters and Luschi, 2018) while maintaining memory efficiency for fine-tuning NeoBERT. This also increases the number of parameter updates per epoch, which is beneficial in low-data regimes.

Table 1 reports the optimal hyperparameter configurations identified via grid search for each architecture and training regime. Across models, learning rates in the range of  $1 \times 10^{-5}$  to  $3 \times 10^{-5}$  consistently yield stable fine-tuning, with lower rates ( $1\text{--}2 \times 10^{-5}$ ) favored for architecturally complex or multilingual models such as XLM-RoBERTa and NeoBERT, likely reflecting greater sensitivity to catastrophic forgetting. NeoBERT consistently selects a batch size of 4 across all regimes, suggesting that higher gradient stochasticity benefits its fine-tuning under low-resource conditions. Weight decay varies across configurations, with stronger regularization ( $1 \times 10^{-2}$ ) generally preferred in monolingual English and Filipino settings, while smaller values ( $1 \times 10^{-5}$ ) are selected more frequently in bilingual and cross-lingual regimes, consistent with the need for greater representational flexibility under mixed-language supervision.

### 3.4 Training Procedure

Each configuration is trained independently from a fresh pretrained checkpoint to avoid parameter leakage across hyperparameter settings. Training is conducted under the following constraints: (1) a maximum of 10 epochs, which is sufficient for

Architecture	Regime	Batch	LR	WD
BERT	English	8	$2 \times 10^{-5}$	$1 \times 10^{-2}$
BERT	Tagalog	4	$3 \times 10^{-5}$	$1 \times 10^{-5}$
BERT	Bilingual	4	$3 \times 10^{-5}$	$1 \times 10^{-2}$
XLM-RoBERTa	English	8	$1 \times 10^{-5}$	$1 \times 10^{-5}$
XLM-RoBERTa	Filipino	8	$3 \times 10^{-5}$	$1 \times 10^{-2}$
XLM-RoBERTa	Bilingual	8	$2 \times 10^{-5}$	$1 \times 10^{-5}$
RoBERTa-Tagalog	English	8	$3 \times 10^{-5}$	$1 \times 10^{-2}$
RoBERTa-Tagalog	Filipino	4	$3 \times 10^{-5}$	$1 \times 10^{-2}$
RoBERTa-Tagalog	Bilingual	8	$3 \times 10^{-5}$	$1 \times 10^{-5}$
NeoBERT	English	4	$2 \times 10^{-5}$	$1 \times 10^{-5}$
NeoBERT	Tagalog	4	$3 \times 10^{-5}$	$1 \times 10^{-5}$
NeoBERT	Bilingual	4	$3 \times 10^{-5}$	$1 \times 10^{-5}$

Table 1: Optimal hyperparameter configurations for transformers across all training regimes.

convergence in low-resource fine-tuning, and (2) full backpropagation over all model parameters. During hyperparameter search, early stopping with a patience of 3 is applied using validation macro-F1 to prevent overfitting. For final reported results, models are retrained without early stopping to ensure a consistent number of training epochs across experimental conditions.

#### 3.4.1 Cross-Validation Protocol

Clinical NLP datasets are typically small and high-variance, making single train–test splits prone to unstable performance estimates. To evaluate robustness with respect to data variability, we perform stratified 10-fold cross-validation on the data for all models and training configurations. Final cross-validation performance is reported as the mean and standard deviation across folds. Importantly, cross-validation is not used for model selection in the final reported results. Instead, it is used to assess performance stability and sensitivity to data sampling.

#### 3.4.2 Evaluation Protocol and Metrics

For each training configuration, we evaluate under three settings. In-domain evaluation trains and tests within the same language using stratified 10-fold cross-validation. Cross-lingual zero-shot evaluation trains on each fold of one language and evaluates on the full dataset of the other language, which is never seen during training. This ensures a consistent cross-lingual test condition across folds, with variance reflecting sensitivity to training data sampling rather than target-language variability. Bilingual evaluation trains on each fold of the combined corpus ( $DEN \cup DTL$ ) and evaluates on the corresponding held-out test fold. Performance is reported on the full mixed-language test fold as well as separately on English-only

and Filipino-only subsets. Each configuration is trained independently. Cross-lingual generalization is assessed by training on one language and evaluating on the other in a zero-shot setting. The combined setting evaluates models on a mixed-language test set containing both English and Filipino samples. We report Accuracy and Macro-averaged F1-score as the primary evaluation metrics. Macro-F1 is prioritized due to its equal weighting of the Healthy and Dementia classes, making it more appropriate for imbalanced clinical classification tasks where minority-class performance is critical. To quantify robustness under language shift, we define the cross-lingual generalization gap as follows

$$\Delta F1 = |F1_{\text{in-domain}} - F1_{\text{cross-lingual}}| \quad (1)$$

where in-domain refers to evaluation on the same language as training (EN→EN or TL→TL), and cross-lingual refers to zero-shot evaluation on the unseen language (EN→TL or TL→EN). The combined setting is reported separately to assess performance under mixed-language conditions and is excluded from the cross-lingual generalization gap computation. In addition to overall performance, we report class-wise F1 scores and Dementia-class recall to explicitly evaluate clinical sensitivity. These metrics allow us to analyze whether performance degradation is driven by reduced detection of Dementia cases or by general classification instability.

## 4 Results and Discussion

Model	Train	F1 ± std / Acc ± std				
		EN	TL	Comb	Gap	
TF-IDF + LR	EN	0.930 ± 0.013 / 0.930 ± 0.013	0.649 ± 0.008 / 0.684 ± 0.006	0.836 ± 0.005 / 0.839 ± 0.004	0.281	
TF-IDF + LR	TL	0.640 ± 0.005 / 0.677 ± 0.003	0.971 ± 0.015 / 0.971 ± 0.015	0.825 ± 0.003 / 0.830 ± 0.003	0.331	
TF-IDF + LR	EN+TL	0.932 ± 0.014 / 0.932 ± 0.014	0.976 ± 0.010 / 0.976 ± 0.010	0.954 ± 0.011 / 0.954 ± 0.011	0.044	
BERT	EN	0.952 ± 0.014 / 0.952 ± 0.014	0.455 ± 0.012 / 0.561 ± 0.007	0.744 ± 0.008 / 0.755 ± 0.007	0.497	
BERT	TL	0.705 ± 0.033 / 0.728 ± 0.027	0.981 ± 0.010 / 0.981 ± 0.010	0.855 ± 0.014 / 0.858 ± 0.013	0.276	
BERT	EN+TL	0.954 ± 0.009 / 0.954 ± 0.009	0.984 ± 0.009 / 0.984 ± 0.009	0.969 ± 0.007 / 0.969 ± 0.007	0.030	
XLM-RoBERTa	EN	0.948 ± 0.017 / 0.949 ± 0.017	0.936 ± 0.018 / 0.936 ± 0.017	0.942 ± 0.016 / 0.942 ± 0.016	0.013	
XLM-RoBERTa	TL	0.825 ± 0.031 / 0.830 ± 0.029	0.986 ± 0.010 / 0.986 ± 0.010	0.905 ± 0.019 / 0.906 ± 0.019	0.161	
XLM-RoBERTa	EN+TL	0.953 ± 0.010 / 0.954 ± 0.010	0.990 ± 0.007 / 0.990 ± 0.007	0.972 ± 0.006 / 0.972 ± 0.006	0.037	
RoBERTa-Tagalog	EN	0.951 ± 0.014 / 0.951 ± 0.014	0.934 ± 0.005 / 0.934 ± 0.005	0.942 ± 0.015 / 0.943 ± 0.015	0.017	
RoBERTa-Tagalog	TL	0.769 ± 0.045 / 0.781 ± 0.038	0.987 ± 0.006 / 0.987 ± 0.006	0.882 ± 0.018 / 0.884 ± 0.017	0.218	
RoBERTa-Tagalog	EN+TL	0.958 ± 0.010 / 0.958 ± 0.010	0.988 ± 0.007 / 0.988 ± 0.007	0.973 ± 0.006 / 0.973 ± 0.006	0.030	
NeoBERT	EN	0.952 ± 0.013 / 0.952 ± 0.013	0.617 ± 0.109 / 0.667 ± 0.084	0.802 ± 0.045 / 0.808 ± 0.042	0.335	
NeoBERT	TL	0.757 ± 0.060 / 0.772 ± 0.051	0.979 ± 0.011 / 0.979 ± 0.011	0.878 ± 0.031 / 0.880 ± 0.030	0.222	
NeoBERT	EN+TL	0.956 ± 0.015 / 0.956 ± 0.015	0.983 ± 0.009 / 0.983 ± 0.009	0.970 ± 0.007 / 0.970 ± 0.007	0.027	

Table 2: Performance across models and training settings. Cells report  $F1 \pm \text{std} / \text{Acc} \pm \text{std}$ . Cross-lingual gap:  $\Delta F1 = |F1_{\text{in-domain}} - F1_{\text{cross-lingual}}|$

Table 2 reports Macro-F1 and Accuracy across all models and training configurations. Four findings emerge consistently.

### 4.1 Cross-Lingual Failure Persists Even Under Controlled Parallel Data

All models achieve strong in-domain performance but degrade substantially under language transfer. English-trained BERT achieves  $F1 = 0.952$  on English but falls to 0.455 on Filipino ( $\Delta = 0.497$ ), while Filipino-trained BERT achieves  $F1 = 0.981$  on Filipino but drops to 0.705 on English ( $\Delta = 0.276$ ). The asymmetry is informative: transfer into English is consistently easier than transfer into Filipino, even when the model is fine-tuned entirely on Filipino, suggesting that English remains a stronger prior in the representation space due to pretraining exposure. Fine-tuning on Filipino modifies model behavior but does not fully overwrite the English-centric geometry learned during pretraining. These results cannot be explained by domain shift, since the Filipino corpus was constructed through controlled manual translation of the same source transcripts. The degradation, therefore, isolates language transfer rather than demographic or collection confounds. The failure appears representational: models learn dementia-related patterns within a language but fail to align those patterns robustly across languages. The lexical baseline supports this interpretation. TF-IDF + LR achieves  $F1 = 0.649$  on Filipino when trained on English and  $F1 = 0.640$  on English when trained on Filipino, suggesting that some dementia-related surface markers survive translation at the lexical level. The gap relative to transformer-based models indicates that lexical overlap alone is insufficient for stable cross-lingual generalization.

### 4.2 NeoBERT Improves Monolingual Fidelity While Increasing Transfer Instability

Relative to BERT, NeoBERT achieves comparable monolingual performance but does not improve cross-lingual robustness. Despite matching BERT’s in-domain performance ( $F1=0.952$ ), English-trained NeoBERT degrades substantially on Filipino ( $F1=0.617$ ) and exhibits the highest variance across any experiment ( $\sigma = 0.109$ ), suggesting that architectural modernization forms tighter English-side decision boundaries that improve in-domain fidelity while reducing tolerance to linguistic variation. Filipino-trained NeoBERT transfers better into English than BERT ( $F1=0.757$  vs. 0.705) but retains high instability in dementia recall (Table 3). Taken together,

architectural improvements alone do not produce language-invariant clinical representations, and stronger monolingual specialization appears to trade off against transfer stability.

### 4.3 Transfer Behavior of Multilingual and Language-Matched Pretraining

XLM-RoBERTa and RoBERTa-Tagalog exhibit substantially smaller transfer gaps than English-only encoders, but through different mechanisms. English-trained XLM-RoBERTa achieves  $F1 = 0.936$  on Filipino ( $\Delta = 0.013$ ), the smallest gap observed under monolingual training, consistent with multilingual pretraining constructing a shared representation space that allows English fine-tuning to map naturally onto Filipino inputs. However, transfer is asymmetric: XLM-RoBERTa transfers substantially better from English to Filipino than the reverse ( $\Delta = 0.013$  *vs.* 0.161), suggesting that English occupies a disproportionately larger share of the pretraining corpus, producing a representation space calibrated more precisely around English structure. RoBERTa-Tagalog is more surprising. Despite being pretrained exclusively on Filipino, it achieves nearly identical English  $\rightarrow$  Filipino transfer as XLM-RoBERTa ( $F1 = 0.934$ ,  $\Delta = 0.017$ ). We attribute this to the linguistic composition of Filipino itself: conversational Filipino contains extensive English lexical borrowing, code-switching, and Anglicized syntactic structure, meaning a model pretrained deeply on Filipino encounters substantial embedded English structure even without explicit multilingual supervision. The reverse direction supports this interpretation, as RoBERTa-Tagalog degrades more sharply when English becomes the target ( $\Delta = 0.218$ ), suggesting that embedded English structure is sufficient for English-to-Filipino transfer but weaker than explicit multilingual supervision for reconstructing robust English-side representations after Filipino fine-tuning.

### 4.4 Bilingual Training Eliminates Most Cross-Lingual Bias

Bilingual fine-tuning substantially reduces cross-lingual degradation across all architectures, with combined-set Macro-F1 converging to 0.969–0.973 and cross-lingual gaps shrinking to between 0.027 and 0.037. Models that behaved very differently under monolingual transfer become

Model	Train	Healthy			Dementia			Dem. Recall		
		EN	TL	Comb	EN	TL	Comb	EN	TL	Comb
TF-IDF+LR	EN	0.934	0.990	0.807	0.926	0.989	0.853	0.879	0.983	0.968
TF-IDF+LR	TL	0.955	0.971	0.855	0.952	0.970	0.795	0.931	0.952	0.660
TF-IDF+LR	EN+TL	0.972	0.976	0.956	0.971	0.974	0.952	0.957	0.950	0.922
BERT	EN	0.216	0.953	0.690	0.695	0.950	0.798	1.000	0.931	0.966
BERT	TL	0.981	0.786	0.875	0.981	0.624	0.836	0.979	0.455	0.724
BERT	EN+TL	0.984	0.955	0.969	0.984	0.954	0.969	0.982	0.940	0.961
XLM-RoBERTa	EN	0.935	0.950	0.942	0.936	0.947	0.941	0.937	0.920	0.928
XLM-RoBERTa	TL	0.986	0.855	0.914	0.985	0.794	0.896	0.977	0.661	0.815
XLM-RoBERTa	EN+TL	0.990	0.955	0.972	0.989	0.952	0.971	0.983	0.931	0.957
RoBERTa-Tagalog	EN	0.935	0.952	0.943	0.934	0.950	0.942	0.923	0.928	0.925
RoBERTa-Tagalog	TL	0.987	0.821	0.896	0.987	0.717	0.869	0.981	0.563	0.772
RoBERTa-Tagalog	EN+TL	0.988	0.959	0.973	0.987	0.957	0.972	0.986	0.939	0.963
NeoBERT	EN	0.504	0.952	0.774	0.729	0.951	0.830	0.915	0.939	0.927
NeoBERT	TL	0.979	0.816	0.893	0.978	0.699	0.863	0.969	0.544	0.766
NeoBERT	EN+TL	0.983	0.956	0.970	0.983	0.955	0.969	0.985	0.938	0.962

Table 3: Class-level performance across models and training configurations. Cells report F1; Healthy/Dementia are class F1; Dem. Recall is AD recall.

nearly indistinguishable after bilingual supervision. NeoBERT, which exhibited the highest instability under English-only transfer, stabilizes immediately and performs comparably to XLM-RoBERTa and RoBERTa-Tagalog, suggesting that the primary bottleneck is not architectural capacity but representational alignment. The models already possess sufficient capacity to encode dementia-related linguistic structure across both languages; what differs is whether the optimization signal encourages those structures to occupy compatible regions of the embedding space. Multilingual and language-matched models retain a small but consistent advantage, with XLM-RoBERTa and RoBERTa-Tagalog achieving the highest combined-set performance (0.972–0.973), suggesting that pretraining language coverage provides a more favorable initialization for bilingual adaptation. Bilingual TF-IDF + LR substantially reduces the transfer gap ( $\Delta = 0.044$ ) but remains below all transformer models, confirming that surface lexical statistics capture part of the dementia signal but are insufficient to construct stable cross-lingual abstractions.

### 4.5 Clinical Implications Under Language Shift

Table 3 shows that aggregate Macro-F1 partially obscures clinically important failure modes under language transfer. English-trained BERT maintains high dementia recall on Filipino (0.931), but Healthy-class F1 collapses to 0.216, indicating that the model predicts most Filipino samples as dementia-positive rather than transferring meaningful discriminative structure. NeoBERT

exhibits a different pattern: dementia recall on Filipino averages 0.939 but with extremely large variance ( $\sigma = 0.190$ ), reflecting inconsistent decision boundaries across fine-tuning folds rather than stable transfer. In contrast, English-trained XLM-RoBERTa and RoBERTa-Tagalog maintain balanced class F1 scores across languages with substantially lower variance, a property particularly important in screening settings where robustness across patient populations matters more than aggregate accuracy.

Bilingual training resolves most instabilities across all architectures, with all transformer models achieving dementia recall above 0.93 with consistently low variance, indicating substantially more stable decision boundaries than those learned under monolingual training.

#### 4.6 Comparison with Classical Baselines

The lexical baseline shows moderate transfer in both directions: English-trained TF-IDF + LR achieves  $F1 = 0.649$  on Filipino and 0.640 in the reverse direction, suggesting that some dementia-related surface markers are preserved across the translation boundary. Under English-only training, TF-IDF performs comparably to NeoBERT cross-lingually but remains substantially below XLM-RoBERTa and RoBERTa-Tagalog. In the bilingual setting, TF-IDF + LR reaches  $F1 = 0.954$  while transformer models achieve 0.969–0.973, indicating that sparse lexical features are insufficient even under bilingual exposure. These results suggest that transformer gains stem not from stronger in-language modeling alone, but from learning contextual representations that transfer across languages. TF-IDF captures surface-level dementia markers but does not model the contextual interactions through which those markers manifest across languages, whereas multilingual and language-matched models better preserve these dependencies. Overall, TF-IDF retains partial transferability due to the structural persistence of dementia-related cues, but its instability relative to contextual models indicates that robust cross-lingual screening requires representation learning beyond lexical overlap.

## 5 Conclusion

Across all model families, strong in-domain performance does not translate to cross-lingual

robustness, even when domain and clinical content are held constant through parallel dataset construction. This indicates that the primary barrier to multilingual clinical NLP is representational misalignment induced by language-specific pretraining rather than task complexity or domain shift. Architectural modernization alone, as shown by NeoBERT, does not yield consistent cross-lingual gains and may increase sensitivity to language shift through tighter monolingual specialization. Multilingual and language-matched pretraining provide more stable transfer, though through different mechanisms. Bilingual fine-tuning, however, largely eliminates cross-lingual performance gaps across all architectures, suggesting that linguistic coverage during task training is more influential than architectural choice in determining robustness under language shift.

These findings indicate that reliable clinical NLP in low-resource, code-switched settings depends less on model scaling or architectural modification and more on ensuring adequate linguistic coverage during training, with bilingual supervision emerging as the most effective strategy for stable and clinically consistent performance across languages.

## 6 Limitations

Several limitations of this study warrant acknowledgment. First, the Filipino dataset was constructed through controlled manual translation of English DementiaBank transcripts, as no large-scale native clinical dementia corpora exist for Filipino. Although translators preserved structural disfluencies, hesitations, and syntactic fragmentation, the translated text inherently reflects the conversational structure and semantic content of the original English source documents. The dataset’s modest scale of 4,000 samples also restricts the generalizability of evaluation metrics and contributes to cross-validation variance. Acquiring larger, organically produced clinical speech from local Filipino patient cohorts in collaboration with geriatricians remains a priority for validating these findings in local clinical environments.

Second, this study focuses exclusively on text to isolate language-specific discourse patterns, setting aside acoustic and non-verbal features. In real-world clinical settings, vocal features such

as pitch variance, pause duration, and phonation rate provide independent diagnostic markers of neurodegeneration not captured in a text-only framework. Integrating speech encoders such as Wav2Vec 2.0 alongside transformer-based text models represents an important future direction.

Finally, while this evaluation establishes classification baselines across a range of encoder architectures, the mechanisms driving model decisions in multilingual contexts remain opaque. Clinical deployment demands transparency, and applying local feature attribution and interpretability methods is a necessary step before these screening workflows can achieve clinical trust and real-world deployment.

## References

- Surabhi Adhikari, Surendrabikram Thapa, Usman Naseem, Priyanka Singh, Huan Huo, Gnana Bharathy, and Mukesh Prasad. 2022. [Exploiting linguistic information from Nepali transcripts for early detection of Alzheimer’s disease using natural language processing and machine learning techniques](#). *International Journal of Human-Computer Studies*, 160:102761.
- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- James T. Becker, François Boller, 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.
- M. Berkes and Ellen Bialystok. 2022. [Bilingualism as a contributor to cognitive reserve: What it can do and what it cannot do](#). *American Journal of Alzheimer’s Disease & Other Dementias*, 37.
- Lola Le Breton, Quentin Fournier, Mariam El Mezouar, John X. Morris, and Sarath Chandar. 2025. [NeoBERT: A next-generation BERT](#). *Preprint*, arXiv:2502.19587.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jan Christian Blaise Cruz and Charibeth Cheng. 2022. [Improving large-scale language models and resources for Filipino](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6548–6555, Marseille, France. European Language Resources Association.
- Kylar A. DeLoach. 2026. [Dementia detection in low-resource languages: Evaluating translation-assisted transfer learning for multilingual clinical assessment](#). Honors thesis, Mississippi State University.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Elif Eyigoz, Sachin Mathur, Mar Mishkind, Guillermo Cecchi, and Melissa Naylor. 2020. [Linguistic markers predict onset of Alzheimer’s disease](#). *EClinicalMedicine*, 28:100583.
- Eden Regala Flores. 2019. [A study on patterns and functions of tagalog-english code switching in two oral discussions](#). *International Journal of TESOL Studies*.
- Kathleen C. Fraser, Jed A. Meltzer, and Frank Rudzicz. 2016. [Linguistic features identify Alzheimer’s disease in narrative speech](#). *Journal of Alzheimer’s Disease*, 49(2):407–422.
- Yuxia Guo, Xu Dong, Mohammed Ali Al-Garadi, Abeer Sarker, Cecile Paris, and Diego Mollá. 2020. [Benchmarking of transformer-based pre-trained models on social media text classification datasets](#). In *Proceedings of the 18th Annual Workshop of the Australasian Language Technology Association*, pages 104–112.
- Megan Herrera, Ankit Aich, and Natalie Parde. 2022. [TweetTaglish: A dataset for investigating Tagalog-English code-switching](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2592–2601, Marseille, France. European Language Resources Association.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations (ICLR 2019)*, New Orleans, LA, USA. OpenReview.net.
- Chingching Lu and Shih-Wei Chen. 2025. [Improving early detection of dementia in low-resource languages: A connected speech analysis case study in Amis language](#). *Alzheimer’s & Dementia*, 20(Suppl 8):e095345.
- Saturnino Luz, Fasih Haider, Davida Fromm, Ioulietta Lazarou, Ioannis Kompatsiaris, and

- Brian MacWhinney. 2024. [An overview of the ADReSS-M signal processing grand challenge on multilingual Alzheimer’s dementia recognition through spontaneous speech](#). *IEEE Open Journal of Signal Processing*, 5:738–749.
- Chengsheng Mao, Jie Xu, Luke Rasmussen, Yikuan Li, Prakash Adekkanattu, Jennifer Pacheco, Borna Bonakdarpour, Robert Vassar, Li Shen, Guoqian Jiang, Fei Wang, Jyotishman Pathak, and Yuan Luo. 2023. [AD-BERT: Using pre-trained language model to predict the progression from mild cognitive impairment to Alzheimer’s disease](#). *Journal of Biomedical Informatics*, 144:104442.
- Anna Marier, Jaime Fernández Arias, Étienne Aumont, Brandon J. Hall, Arthur C. Macedo, Nesrine Rahmouni, Gleb Bezgin, Paolo Vitali, Pedro Rosa-Neto, and Maxime Montembeault. 2026. [Language deficits across PET-based Braak stages of tau accumulation in Alzheimer’s disease](#). *Alzheimer’s & Dementia*, 22(3):e71286.
- Dominic Masters and Carlo Luschi. 2018. [Revisiting small batch training for deep neural networks](#). *arXiv preprint arXiv:1804.07612*.
- Deepthi Mave, Suraj Maharjan, and Thamar Solorio. 2018. [Language identification and analysis of code-switched social media text](#). In *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*, pages 51–61, Melbourne, Australia. Association for Computational Linguistics.
- Lester James V. Miranda. 2023. [Developing a named entity recognition dataset for Tagalog](#). In *Proceedings of the First Workshop in South East Asian Language Processing*, pages 13–20, Nusa Dua, Bali, Indonesia. Association for Computational Linguistics.
- Lester James Validad Miranda, Elyanah Aco, Conner G. Manuel, Jan Christian Blaise Cruz, and Joseph Marvin Imperial. 2025. [FilBench: Can LLMs understand and generate Filipino?](#) In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2496–2529, Suzhou, China. Association for Computational Linguistics.
- Jann Railey Montalan, Jimson Paulo Layacan, David Demitri Africa, Richell Isaiah S. Flores, Michael T. Lopez II, Theresa Denise Magsajo, Anjanette Cayabyab, and William Chandra Tjhi. 2025. [Batayan: A Filipino NLP benchmark for evaluating large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31239–31273, Vienna, Austria. Association for Computational Linguistics.
- Raghavendra Pappagari, Jaejin Cho, Laureano Moro-Velázquez, and Najim Dehak. 2020. [Using state of the art speaker recognition and natural language processing technologies to detect Alzheimer’s disease and assess its severity](#). In *Proceedings of Interspeech 2020*, pages 2177–2181.
- Lotem Peled-Cohen and Roi Reichart. 2025. [A systematic review of NLP for dementia: Tasks, datasets and opportunities](#). *Transactions of the Association for Computational Linguistics*, 13:1204–1244.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Elizabeth Salesky, Matthias Sperber, and Alex Waibel. 2019. [Fluent translations from disfluent speech in end-to-end speech translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 2786–2792.
- D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo, and Dan Dennison. 2015. [Hidden technical debt in machine learning systems](#). In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 2503–2511. Curran Associates, Inc.
- Sofia Segkoulis, Mara Gkioka, Stylianos Kokkas, Konstantinos Votis, Sergi Valero, Andrea Miguel, Athos Antoniadis, Emily Charalambous, and George Manias. 2025. [Linguistic markers in spontaneous speech: Insights into subjective cognitive decline \(review\)](#). *Healthcare*, 13(22):2888.
- Ravi Shankar, Anjali Bunde, and Amartya Mukhopadhyay. 2025. [A systematic review of natural language processing techniques for early detection of cognitive impairment](#). *Mayo Clinic Proceedings: Digital Health*, 3:100205.
- Noam Shazeer. 2020. [Glu variants improve transformer](#). *arXiv preprint arXiv:2002.05202*.
- Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. [RoFormer: Enhanced transformer with rotary position embedding](#). *Neurocomputing*.
- Bastiaan Tamm, Rik Vandenbergh, and Hugo Van Hamme. 2023. [Cross-lingual transfer learning for Alzheimer’s detection from spontaneous speech](#). In *ICASSP 2023 – 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–2.

- Charlyn Nayve Villavicencio, Julio Jerison E. Macrohon, X. Alphonse Inbaraj, Jyh-Horng Jeng, and Jyh-Cheng Hsieh. 2021. [Twitter sentiment analysis towards COVID-19 vaccines in the Philippines using naïve Bayes](#). volume 12, page 204.
- Sida Wang and Christopher Manning. 2012. [Baselines and bigrams: Simple, good sentiment and topic classification](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 90–94, Jeju Island, Korea. Association for Computational Linguistics.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. 2020. [On layer normalization in the transformer architecture](#). In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*.
- Biao Zhang and Rico Sennrich. 2019. [Root mean square layer normalization](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 649–657. Curran Associates, Inc.