

GPT-D: Inducing Dementia-related Linguistic Anomalies by Deliberate Degradation of Artificial Neural Language Models

Changye Li¹, David Knopman², Weizhe Xu³, Trevor Cohen³, and Serguei Pakhomov⁴

¹Institute of Health Informatics, University of Minnesota

²Mayo Clinic, Rochester, MN

³Biomedical and Health Informatics, University of Washington

⁴Pharmaceutical Care and Health Systems, University of Minnesota

^{1,4}{lix3013, pakh0002}@umn.edu

²{knopman}@mayo.edu

³{xuweizhe, cohenta}@uw.edu

Abstract

Deep learning (DL) techniques involving fine-tuning large numbers of model parameters have delivered impressive performance on the task of discriminating between language produced by cognitively healthy individuals, and those with Alzheimer's disease (AD). However, questions remain about their ability to generalize beyond the small reference sets that are publicly available for research. As an alternative to fitting model parameters directly, we propose a novel method by which a Transformer DL model (GPT-2) pre-trained on general English text is paired with an artificially degraded version of itself (GPT-D), to compute the ratio between these two models' *perplexities* on language from cognitively healthy and impaired individuals. This technique approaches state-of-the-art performance on text data from a widely used "Cookie Theft" picture description task, and unlike established alternatives also generalizes well to spontaneous conversations. Furthermore, GPT-D generates text with characteristics known to be associated with AD, demonstrating the induction of dementia-related linguistic anomalies. Our study is a step toward better understanding of the relationships between the inner workings of generative neural language models, the language that they produce, and the deleterious effects of dementia on human speech and language characteristics.

1 Introduction

Alzheimer's disease (AD) dementia affects every aspect of cognition, including language use. Over 50 million people are currently diagnosed with AD dementia, and this number is expected to triple by 2050 (Organization et al., 2017; Patterson, 2018; Prince et al., 2016). Furthermore, over half of the individuals living with dementia are undiagnosed (Lang et al., 2017). While AD has no known cure, timely diagnosis can prevent or alleviate adverse

outcomes ranging from anxiety over unexplained symptoms to family discord and catastrophic events (Stokes et al., 2015; Boise et al., 1999; Bond et al., 2005). However, diagnosis of AD dementia is time-consuming and challenging for patients and physicians alike, and currently relies on patient and caregiver reports, extensive neuropsychological examinations, and invasive imaging and diagnostic procedures (Patterson, 2018). Automated analysis of spoken language can potentially provide accurate, easy-to-use, safe and cost-effective tools for monitoring AD-related cognitive markers. In particular, studies have demonstrated that supervised machine learning methods can learn to differentiate accurately between patients with dementia and healthy controls (Fraser et al., 2016; Orimaye et al., 2017), with particularly strong performance from recent deep learning (DL) models (Balagopalan et al., 2020; Roshanzamir et al., 2021). However, the large number of parameters employed in DL presents a danger of overfitting to the small datasets concerned, and hinders interpretability of model predictions - both critical concerns for clinical artificial intelligence applications (Graham et al., 2020).

As an alternative to fitting model parameters directly, we propose a novel method by which a pre-trained Transformer (Vaswani et al., 2017) model, GPT-2 (Radford et al., 2019) is paired with an artificially degraded version of itself (GPT-D), to compute the ratio of model perplexities on language from cognitively healthy and impaired individuals. We anticipate that semantic information lost with dementia progression may be localized to particular layers of a neural language model, and that one can simulate this information loss by systematically modifying parameters in these layers. Specifically, we hypothesize that impairing certain layers of a DL model can result in linguistic deficits that are also observed in dementia. We further hypothesize

that unlike prior work fitting model parameters to labeled “Cooke Theft” transcripts, this approach will detect task-agnostic linguistic anomalies, permitting evaluation of language from casual conversations. We evaluate these hypotheses by targeting individual layers for induction of dementia-related linguistic anomalies, resulting in a degraded model – GPT-D. We then assess the ability of a paired perplexity approach combining GPT-2 with GPT-D to identify transcripts from participants with dementia. In addition, we assess generalization performance, and consider the extent to which the best-performing degraded model reflects linguistic anomalies known to occur in AD dementia: usage of higher frequency words, and repetitiveness. The contributions of this work can be summarized as follows: a) we develop a novel method for automated detection of dementia-related linguistic anomalies, involving deliberate degradation of a pre-trained Transformer model; b) this method exhibits state-of-the-art (SOTA) within-set performance for models trained on text alone, and is distinguished by its ability to generalize from cognitive tasks to conversational data; c) the degradation process induces linguistic anomalies observed in dementia in language generated by GPT-D¹.

2 Background

Building on a rich body of evidence that machine learning methods can learn to distinguish between language from healthy controls and dementia patients (for a review, see [Lyu \(2018\)](#); [Petti et al. \(2020\)](#)), recent work leveraging pre-trained Transformer models has demonstrated improvements in performance over prior approaches. [Balagopalan et al. \(2020\)](#) fine-tuned the BERT ([Devlin et al., 2019](#)) model on the training set of the AD Recognition through Spontaneous Speech (ADReSS) Challenge ([Luz et al., 2020](#)), which was developed, in part, to address the lack of standardized train/test splits and subset definitions in prior work using DementiaBank ([Becker et al., 1994](#)) (DB). [Balagopalan et al. \(2020\)](#) report an accuracy of 83.3% on the test set, an improvement over machine learning models with expert-defined features. Performance can also be further boosted by introducing more data from the same picture description task ([Guo et al., 2021](#)). These findings suggest a promising direction, as models can be developed without

¹Our code is available at <https://github.com/LinguisticAnomalies/hammer-nets>

extensive feature engineering. However, additional task-specific data are not always available. DL models with millions of parameters are vulnerable to overfitting with small data sets, which may be difficult to detect as they are hard to interpret.

However, some DL models can be distilled into a single interpretable feature: language model (LM) perplexity (PPL). PPL is a measurement of how well a language sample fits a trained LM. Intuitively, a model trained on language from cognitively healthy participants should be “surprised” by language from participants with dementia, and the opposite should also be true. Accordingly, the difference between the *paired perplexities* from “cognitively healthy” and “dementia” language models produces SOTA results on the task of identifying transcripts from participants with dementia ([Fritsch et al., 2019](#); [Cohen and Pakhomov, 2020](#)), effectively condensing neural network parameters to a single diagnostically useful feature. Contemporary deep LMs such as GPT-2 are *already* trained on large amounts of text, that has presumably been authored predominantly by cognitively healthy individuals. The difficulty with leveraging these models within the paired perplexity paradigm arises from the lack of a correspondingly large set of text from participants with dementia. We negotiate this difficulty by deliberately degrading a Transformer model to limit its semantic processing capabilities, obviating the need for large amounts of dementia-specific training data. We show that the resulting models can effectively identify transcripts from participants with dementia, generalize across language samples and tasks, and generate text with linguistic characteristics of this condition.

3 Methods

3.1 Data

We used three publicly available datasets²: DB, ADReSS, and the Carolinas Conversation Collection (CCC) ([Pope and Davis, 2011](#)). Dataset characteristics are provided in Table 1. DB is a publicly available compendium of manually transcribed audio recordings of neuropsychological tests administered to healthy participants and patients with dementia. A detailed description is available in [Becker et al. \(1994\)](#). In brief, the tests include a

²While the data used in this paper are publicly available, we are not able to redistribute any of these data as per Data Use agreement with Dementia Bank and the Carolinas Conversation Collection.

| Dataset | | Dementia | | | Healthy Controls | | |
|---------|-------|----------------|----------------|-----------------------------|------------------|----------------|-----------------------------|
| | | N participants | MMSE Mean (SD) | Transcript length Mean (SD) | N participants | MMSE Mean (SD) | Transcript length Mean (SD) |
| ADReSS | train | 54 | 17.1 (5.5) | 104 (63) | 54 | 29.1 (1.9) | 114 (49) |
| | test | 24 | 19.5 (5.4) | 95 (47) | 24 | 28.8 (1.5) | 120 (72) |
| | all | 78 | 17.8 (5.5) | 101 (58) | 78 | 29 (1.2) | 116 (56) |
| DB | | 169 | 20.2 (4.6) | 959 (534) | 99 | 29.1 (1.1) | 1085 (556) |
| CCC | | 234 | NA | 1213 (943) | 48 | NA | 714 (308) |

Table 1: Basic characteristics of datasets

picture description task from the Boston Diagnostic Aphasia Examination (Goodglass and Kaplan, 1983), a widely-used diagnostic test for language abnormality detection. In this task, the participants are presented with a “Cookie Theft” picture stimulus (see Figure 4 in Appendix), and are asked to describe everything they see occurring in the picture. In other words, DB data are from tasks that were explicitly designed to detect language abnormalities in dementia patients. We restricted the original set of 194 participants with any AD diagnosis only to those that were assessed as having probable AD, resulting in a set of 169 patients and 99 controls. The ADReSS set is a subset of DB, which the controls and dementia participants were matched age and gender, resulting in a balanced dataset consisting of a total of 156 samples (78 with dementia and 78 controls) split into training and testing portions. Unlike the two preceding datasets derived from picture description tasks, CCC is a collection of 646 transcribed recordings of interviews of 48 elderly cognitively normal individuals with non-dementia related chronic conditions, and 234 individuals with a diagnosis of dementia. Interview topics vary considerably, and include discussions of the participant’s health.

Additionally, we used a set of six synthetic “Cookie Theft” picture description narratives created by Bird et al. (2000) to study the impact of semantic dementia on verb and noun use in picture description tasks. The transcripts were created to manipulate lexical frequency (which is also relevant in AD dementia, where words with higher lexical frequency tend to feature prominently (Almor et al., 1999)) by first compiling a composite baseline narrative from samples by healthy subjects, and then removing and/or replacing nouns and verbs in that baseline with words of higher lexical frequency (e.g., “mother” vs. “woman” vs. “she”). Lexical frequency was calculated using the Celex Lexical Database (LDC96L14) and words

were aggregated into groups based on four log frequency bands (0.5 - 1.0, 1.0 - 1.5, 1.5 - 2.0, 2.5 - 3.0: e.g., words in the 0.5 - 1.0 band occur in Celex more than 10 times per million). We used these synthetic data to help with interpretation of the effects resulting from artificially impairing the GPT-2 model.

We performed basic pre-processing of transcripts in each dataset by which we removed speech artifact descriptions and converted non-ASCII characters to plain text. We also excluded portions of transcripts that represented speech that did not belong to the participant.

3.2 Modeling and Evaluation

We evaluated models for *classification* performance using the standard ADReSS train/test splits. We then performed *cross-validation* of GPT-D models to assess the stability of the best-performing configurations across folds. For *generalization* performance, we evaluated how well models trained on one corpus performed on others. We also assessed differences in text *generation* between GPT-2 and GPT-D, by estimating repetitiveness and lexical frequency, as well as through salience-based *visualization*.

3.2.1 Artificial Impairment: Locations

We experimented with impairing the GPT-2 (small) model in two locations as illustrated in Figure 1 with various portions. We found that impairing 50% of values in the corresponding location resulting in generally better performance, among 25%, 50%, 75% and 100% impairment. The *embedding layer* (see (1) in Figure 1) is a $50,257 \times 768$ matrix where each row represents a token in the model’s vocabulary. The embedding layer was impaired by randomly masking 50% of the rows of the embedding matrix. The *self-attention mechanism* (denoted (2) in Figure 1) was impaired by masking the first 50% of columns in the Value matrix of the concatenated Query-Key-Value matrices. We




Figure 1: Impairment locations within the GPT-2 (small) model.

found that masking random columns resulted in worse performance in preliminary experiments.

The self-attention mechanism multiplies vectors representing an input sequence by three identically-sized matrices, namely Query (Q), Key (K) and Value (V) each with dimension (d) of 768×768 . Q generates a representation of the current token which is compared with token representations derived from K, to calculate each token’s influence on the contextual representation of the current one. Multiplying by V generates a semantic representation of each token, which is added to the outgoing representation of the current token in accordance with this influence. The attention weights are calculated by Equation 1, and the parameters of the matrices are updated during the training process.

$$attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (1)$$

The GPT-2 model’s attention mechanism in each of the 12 decoder layers contains 12 attention heads that are represented as vectors of 64 parameters. We impaired 50% of those parameters of V in various combinations of attention heads in each decoder layer by masking them as zeroes. We only did this in V matrices, as their parameters directly determine the content of the vectors that are passed on to the subsequent feed-forward layer, while the Q and K matrices determine how this content is weighted when generating the representations to be propagated as weighted sums of vectors that have been transformed by the Value matrix.

3.2.2 Artificial Impairment: Patterns

We also experimented with three ways of introducing artificial impairment into the attention mechanism in single and multiple decoder layers: individual, cumulative, and combination. The *individual* approach was to simply impair all 12 layers one at a time. The *cumulative* approach consisted of impairing decoder layers sequentially starting with the bottom decoder layer (layer 0) and adding impairment to layers above it one at a time up to

layer 11, resulting in total of 12 combinations of impairments. The *combination* approach consisted of impairing all possible combinations of layers, one combination at a time, resulting in 4096 combinations. The degraded models were subsequently used in combination with the original GPT-2 model to calculate the difference and ratio of PPLs between these two models on each input transcript.

3.3 Interpretation of Neural Model Behavior

Classification Performance: For the paired perplexity approach, we estimated the ratio of model PPLs ($\frac{PPL_{GPT-2}}{PPL_{GPT-D}}$) for each transcript. These PPLs were averaged for participants with multiple transcripts. All validation methods commenced with calculating the area under the receiver-operator characteristic (ROC) curve (AUC). From this, accuracy (ACC) was determined at equal error rate (EER), a threshold where the false acceptance rate and false rejection rate from an ROC curves is equal. We also calculated Pearson correlation between the ratio in perplexities of the GPT-2 and GPT-D models and the MMSE scores where available (CORR). We used the original fixed single split between training and testing data provided by the creators of the dataset to compare our results to those published by others on ADReSS.

Cross-validation Performance: For all datasets (including ADReSS), we performed standard cross-validation by which we split each dataset into disjoint folds and first determined which combination of GPT-D attention layers results in best performance on the training portion of each fold and then tested that combination on the test portion of the fold averaging the AUC, ACC and CORR values (if available) across the folds. We selected 5-fold cross-validation due to the relatively small size of the ADReSS, DB, and CCC datasets. To ensure reproducibility across runs, data folds for cross-validation were extracted using the KFold method from the scikit-learn library (Pedregosa et al., 2011) with shuffling and a fixed random seed.

Generalization Performance: We tested gen-

eralizability of the paired perplexity approach by evaluating its performance across datasets. We first determined the best-performing pattern of impairment based on the highest AUC obtained on each dataset, and then applied the model impaired with that pattern to the remaining datasets.

Baseline Models: We compared our model performance on transcript *classification* with the previous text-only SOTA (Balagopalan et al., 2020), which was obtained with a 12-layer BERT model fine-tuned on the ADReSS training set, and evaluated on the test set. To evaluate the *generalization* performance, we followed this work’s hyperparameter choices and fine-tuned BERT and DistilBERT (Sanh et al., 2019)³, a distilled BERT base model that is compact and more efficient. We fine-tuned these models on the entire ADReSS, DB and CCC datasets separately, then evaluate the three resulting models on every other set.

Language Generation: To prompt the GPT-2 and GPT-D models to generate text we utilized Bird et al.’s synthetic “Cookie Theft” picture description narrative that represents a composite of narratives produced by healthy controls. Table 5 (in Appendix) illustrates the text generated by GPT-2 and GPT-D in response to prompt sentences taken from the synthetic narrative. Both GPT-2 and GPT-D models were induced to generate at least 20 additional tokens with a beam search (Wiseman and Rush, 2016) that keeps the top n hypotheses ($n = 5$ in this case) at each time step and eventually returns the sequence of hypotheses that achieved the highest probability after reaching the end-of-sequence token. Beam search also works well when the length of output is not predictable, which fits the nature of the language tasks represented by the corpora we tested. However, one of the challenges of using beam search for text generation is that it tends to generate repeated words. We added a penalty for generating repetitive unigrams and implemented the top-p algorithm (Welleck et al., 2019) to keep the set of potential words as small as possible while the cumulative probability of this set is greater than the specific probability p ($p = 0.9$ in our case). The penalty was applied equally to GPT-2 and GPT-D to avoid potentially biasing one of these models to produce more repetitions. After the models generated five best predictions for each prompt, we chose the first non-empty pair of

outputs from both the GPT-2 and GPT-D models as the final result.

Lexical frequency and repetitiveness: Previous work (Cohen and Pakhomov, 2020) suggests that neural language models are sensitive to lexical frequency. We investigated whether GPT-D generates content of higher lexical frequency than the GPT-2 model. To compute lexical frequency, we split each generated output into tokens with the help of the NLTK⁴. We did not stem the tokens to avoid increasing lexical frequency by artificially merging different tokens with the same stem. In addition to the stopwords provided by NLTK, we treated tokens with following part-of-speech tags a) *PRP* (personal pronoun), b) *PRP\$* (possessive pronoun), c) *WP\$* (possessive wh-pronoun), and d) *EX* (existential there) as stopwords. We also added the $n't$ token and tokens starting with $'$ to the list of stopwords. Log lexical frequency of each qualified generated token was calculated based on occurrence in the SUBTLEX_{us} corpus (Brysbart and New, 2009). Tokens that do not appear in SUBTLEX_{us}, were removed as out-of-vocabulary (OOV) items. To assess the degree of repetition present in the generated text, we calculated the type-to-token ratio (TTR) as the number of word types divided by the number of word instances.

Saliency Visualization: We used the gradient×input saliency proposed in Denil et al. (2014), as implemented with the `ecco`⁵ Python package for visualization. Saliency is defined as $\|\nabla_{x_i} f_c(\mathbf{x}_{1:n}) \cdot x_i\|_2$, which is the L2 normalized back-propagated gradient with respect to a) the dot product of the embedding vector of all previous input tokens ($\mathbf{x}_{1:n}$), and b) the model output of token x_i ($f_c(\mathbf{x}_{1:n})$), where c is the predicted token at time-step i . A previous study (Serrano and Smith, 2019) found that raw attention weights were not interpretable for any intermediate representation of a language model. Instead, Bastings and Filippova (2020) argued that saliency is the preferred method for interpretability as it takes the entire input into account and reveals the relevance of each input token to the next predicted token in the sequence.

To make the visualizations comparable for the two models, we repeatedly prompted both models with the same input until both models generated the same token as the prediction. It is worth

³Available on Huggingface <https://huggingface.co/transformers/index.html>

⁴<https://www.nltk.org/>

⁵<https://github.com/jalammar/ecco>




Figure 2: Effects of artificial impairment on model perplexity in synthetic picture description narratives. Higher values on the x axis indicate transcripts simulating more advanced disease.

| Dataset | Combination Impairment Pattern | | | Cumulative Impairment Pattern | | |
|---------|--------------------------------|-------------|------------------|-------------------------------|-------------|------------------|
| | AUC (SD) | ACC (SD) | r with MMSE (SD) | AUC (SD) | ACC (SD) | r with MMSE (SD) |
| ADReSS | 0.80 (0.06) | 0.71 (0.07) | -0.52 (0.08) | 0.79 (0.02) | 0.68 (0.03) | -0.51 (0.05) |
| DB | 0.81 (0.07) | 0.76 (0.04) | -0.45 (0.06) | 0.83 (0.02) | 0.73 (0.02) | -0.41 (0.14) |
| CCC | 0.77 (0.04) | 0.71 (0.04) | - | 0.72 (0.04) | 0.64 (0.09) | - |

Table 2: Five-fold cross-validation results of all possible combination of impairment pattern (masking 50% of value matrix as zeroes) for cumulative and combination methods on ADReSS, DB, and CCC.

noting that `ecc0` for visualization supports limited text generation arguments compared to the `transformers` package, which we used for language generation task. Consequently, we only used the top-p algorithm currently supported by `ecc0` for our visualizations.

4 Results

Impairment Location: The contrast in the effects of artificial impairment on the embedding and attention layers (locations 1 and 2 in Figure 1, respectively) is illustrated in Figure 2. Impairing embeddings results in a distribution of perplexity values over the range of impairment in the synthetic narratives very similar to that of the GPT-2 model. Impairing attention, however, results in a sharp decrease in PPL on the more perturbed narratives (those narratives simulating more impairment), which yields a monotonically increasing step-like function over $\frac{PPL_{GPT-2}}{PPL_{GPT-D}}$ that lends itself well to thresholding for categorization. These results were confirmed by testing on available corpora the discriminative ability of the paired perplexity approach by artificially impairing only the embedding layer, which resulted in near-random AUCs (close to 0.5 - data not shown). Consequently, in subsequent results we will show attention-based models only.

Classification Performance: For comparison with previous work using the ADReSS dataset, the

best training set performance was obtained by impairing 50% of each attention head in layers 0-5, 6, and 8-9. This pattern achieved an AUC of 0.88 (ACC = 0.75, CORR = -0.55) on the test split. The cumulative impairment method performed slightly better. Impairing 50% of each attention head in the first 9 layers resulted in best performance on the training set, and AUC of 0.89 (ACC = 0.85, CORR = -0.64) on the test split. We note that this accuracy exceeds the average result reported by Balagopalan et al. (2020), and approaches the performance of their best run.

Cross Validation: The results of within-set cross-validation are summarized in Table 2. Both *combination* and *cumulative* methods had small standard deviations (~ 0.1) with over or near 0.7 mean AUC on all sets. Estimates from the paired perplexity approach for both methods were negatively correlated with MMSE on the ADReSS (-0.52, -0.51) and DB (-0.45, -0.41) sets, respectively. The best performance obtained with the *individual* approach resulted in AUC of 0.66 (ACC: 0.64) with impairment of layer 8 on the DB dataset; AUC of 0.70 (ACC: 0.66) with impairment of layer 8 on the ADReSS dataset; and AUC of 0.71 (ACC: 0.63) with impairment in layer 7 on CCC.

Generalization: The results of generalization evaluation are shown in Table 3. Both cumulative and combination methods yielded similar performance on CCC, where both AUC and ACC were

| Training method | Testing dataset | | |
|---------------------------------------|-----------------|-----------|-----------|
| | ADReSS | DB | CCC |
| (Best pattern:AUC) | AUC/ACC | AUC/ACC | AUC/ACC |
| Cumulative Impairment Pattern | | | |
| ADReSS (0-8:0.80) | – | – | 0.77/0.72 |
| DB (0-4:0.82) | – | – | 0.69/0.68 |
| CCC (0-2:0.72) | 0.70/0.63 | 0.74/0.63 | – |
| Combination Impairment Pattern | | | |
| ADReSS (0-6,8:0.80) | – | – | 0.76/0.71 |
| DB (0-6,8:0.80) | – | – | 0.76/0.71 |
| CCC (1-3,5,7,9-11:0.79) | 0.69/0.61 | 0.72/0.67 | – |
| Fine-tuned BERT | | | |
| ADReSS | – | – | 0.64/0.63 |
| DB | – | – | 0.67/0.6 |
| CCC | 0.71/0.66 | 0.7/0.65 | – |
| Fine-tuned DistilBERT | | | |
| ADReSS | – | – | 0.67/0.57 |
| DB | – | – | 0.67/0.6 |
| CCC | 0.65/0.62 | 0.47/0.45 | – |

Table 3: Generalizability of GPT-2/GPT-D approach compared to fine-tuning on BERT and DistilBERT. All evaluation metrics are calculated at EER rate. The best-performing pattern and its performance are indicated with parentheses and separated by colon.

close to or exceeded 0.7. In contrast, fine-tuning BERT and DistilBERT resulted in near-random classification performance on the corresponding validation dataset. While fine-tuning BERT on conversational discourse samples in CCC and applying it to the picture descriptions in ADReSS and DB generalized well as compared to the paired perplexity approach, it did not generalize in the opposite direction when BERT was fine-tuned on ADReSS and DB picture descriptions and applied to conversations in CCC.

Language Generation: Table 4 reports mean lexical frequency estimates for words contained in the text generated by GPT-2 and GPT-D models. The GPT-D model was induced by using the best-performing patterns of impaired layers deter-

mined from cumulative and combination methods for pattern selection on the available datasets. Both GPT-2 and GPT-D generate ~ 1 OOV token on average for each prompt. In general, the resulting GPT-D model generated text consisting of words with higher lexical frequency than words in the text generated by the GPT-2 model across all datasets and methods, even though some of the differences failed to reach statistical significance. All GPT-D models also generated more repetitions, evident as lower TTRs .

| Dataset (Pattern) | LF | | TTR | |
|--------------------|-------|--------|-------|-------|
| | GPT-2 | GPT-D | GPT-2 | GPT-D |
| Cumulative | | | | |
| ADReSS (0-8) | 9.48 | 9.82* | 72% | 50% |
| DB (0-4) | 9.49 | 9.83* | 72% | 49% |
| CCC (0-2) | 9.48 | 9.54 | 72% | 51% |
| Combination | | | | |
| ADReSS/DB (0-6,8) | 9.5 | 9.41 | 72% | 55% |
| CCC (1-3,5,7,9-11) | 9.45 | 9.92** | 73% | 64% |

Table 4: Mean log lexical frequency (LF) and type-to-token ratio (TTR) of the generated text. The best-performing pattern is indicated with parentheses. * indicates p-value < 0.05 and ** indicates p-value < 0.01. P-values were obtained with two-sided Welch’s t-test

Salience Visualization: Figure 3 shows the magnitude of the contribution for each token in the prompt used to initiate text generation to the model’s prediction of the same token ‘the’. The weight of the contribution of each token is shown as a percentage that can be interpreted as the amount of contribution the model derives from it. We observe in Figure 3 that impairing GPT-2’s attention heads leads to the redistribution of the model’s contribution to the words in the prompt when making the prediction of the next word. For the GPT-2 model, tokens ‘boy’, ‘climbed’,

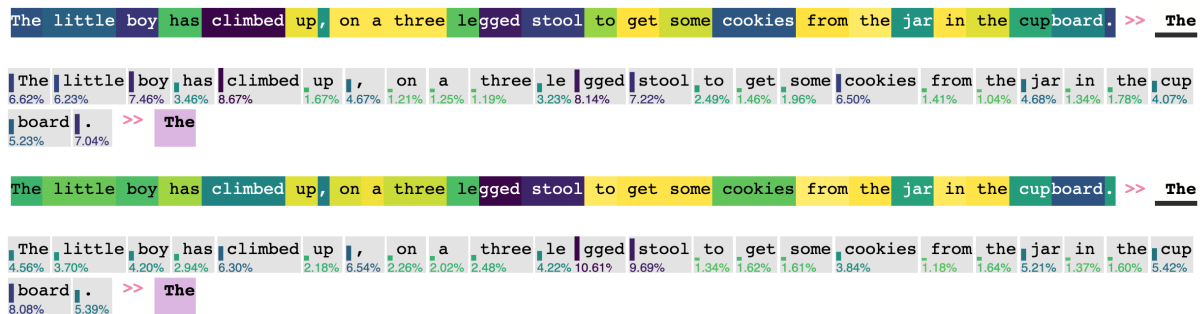


Figure 3: An informal illustration of differences in contributions of input tokens to generating the word “The”, for GPT-2 (top) and GPT-D (bottom) models. Percentages and color represent the degree of contribution.

and 'cookies' contributed more when predicting 'the'. However, for the GPT-D model those word tokens did not clearly stand out as substantially contributing to the prediction in either of these examples. Furthermore, tokens corresponding to function words (e.g., 'on', 'a' and 'from') contributed little to the predictions generated by the GPT-2 model; however, these tokens contributed more for predictions generated by GPT-D model. As evident in the examples in Figure 3, the salience of the words in the prompt is much more diffuse when the GPT-D model is making the prediction - i.e. the model is uncertain with respect to what it should consider as important. In contrast, for the GPT-2 model the key elements of the "Cookie Theft" scenario - 'cookie', 'three-legged stool', 'boy' - stand out as highly salient. These observations, although informal and qualitative, indicate that the impairment of the self-attention mechanism in GPT-2 results in a "behavior" resembling that observed in all stages of AD dementia as a result of impaired selective attention that in turn reduces one's ability to encode new information in episodic memory (see Perry et al. (2000) for a comprehensive review).

5 Discussion

Our key findings are as follows. First, we show that the paired perplexity approach using the ratio between the GPT-2 and GPT-D model perplexities approaches SOTA performance on ADReSS, leveraging GPT-2's extensive pre-training *without* requiring a comparably large data set from dementia patients. Second, this approach generalizes from "Cookie Theft" picture description data to casual conversation, in contrast to BERT/DistilBERT fine-tuning. Finally, artificial impairment of GPT-2's self-attention induces linguistic anomalies observed in dementia.

The best-performing cumulative pattern for the ADReSS training set resulted in accuracy of 0.85 in the test set, exceeding the best BERT results reported on this test set (\bar{x} ACC = 0.833 (Balagopalan et al., 2020)). However, our approach contrasts with approaches that train or fine-tune language models using a specific dataset, and test on held-out components of the same set. While our approach does require some labeled data through which to determine the best-performing layers to impair, our results demonstrate generalization to other datasets and populations as well as a different type of dis-

course - spontaneous conversations. GPT-D is reliably less perplexed by dementia-related linguistic anomalies across all of these sets than GPT-2. This facilitates broader application of the paired perplexity approach than was previously possible, and suggests our approach is more sensitive to task-agnostic dementia-related linguistic anomalies than BERT/DistilBERT fine-tuning.

In contrast to impairing embeddings or individual attention layers, the maximum discriminating effect was achieved by impairing multiple attention layers (either combinatorially or cumulatively), which is consistent with prior observations that Transformer layers encode different syntactic and semantic linguistic features in multiple lower and middle layers (Jo and Myaeng, 2020; Jawahar et al., 2019; Lin et al., 2019). Thus, impairing a single layer may not be enough to achieve the full effect. Since both syntactic and semantic context is encoded in the Transformer decoder layers we expected to find different patterns of artificial impairment to be most effective in vastly different types of discourse represented by the DB and CCC datasets; however, we were surprised to find that only impairing the self-attention layers had the desired effect on the results in contrast to impairing embeddings or feed-forward network components.

The results presented in Table 4 also align with previously published findings that both neural networks trained on language produced by participants with dementia, and the lexical-retrieval processes of patients affected by this condition are sensitive to lexical frequency effects (Cohen and Pakhomov, 2020; Pekkala et al., 2013). Our results suggest that impairing the self-attention mechanism in a Transformer artificial neural network may induce similar sensitivity to lexical frequency. By impairing the attention heads in a GPT-2, we observe significant differences in lexical frequency and TTR characteristics of the text generated by the GPT-2 and GPT-D, with the change in TTR ratio indicating that GPT-D has a greater tendency to produce repeated words when generating text, just as participants with dementia are more prone to repeat words in picture description tasks (Hier et al., 1985).

In other previous work on the DB and the ADReSS datasets, the authors attempted to predict individual MMSE scores in addition to discriminating between cases and controls (Yancheva et al., 2015; Luz et al., 2020). We could not perform a comparable analysis in the current study on account

of focusing on using the paired perplexity measure as a single threshold to distinguish between cases and controls. While predicting MMSE is not the main focus of our study, we did find negative correlations between the paired perplexity measures and the MMSE scores, providing additional evidence that artificially impairing the attention mechanism of the GPT-2 model simulates cognitive effects of dementia detectable in language.

Our findings are also consistent with previous work indicating that Transformer models are able to predict neural responses during language comprehension and generalize well across various datasets and brain imaging modalities (Schrimpf et al., 2021). Thus, our work is another step in the direction of achieving better understanding of the relationship between the inner workings of generative artificial neural language models and the cognitive processes underlying human language production. Impairing how contextual information is stored in the self-attention mechanism *in silico* creates similar deficits to what is observed in dementia. The next important step is perhaps to investigate how contextual information encoding is impaired *in vivo* in AD dementia.

The encouraging results on the CCC dataset point to the possibility of developing a tool for analysing patients' daily spontaneous conversations in a task-agnostic fashion. Generalizable across tasks and domains and easy-to-interpret language-based instruments for detecting anomalies potentially consistent with dementia can be most useful in clinical situations where the patient or family member raise a concern about unexplained changes in cognition. A simple to administer (or self-administer) language-based instrument for objective confirmatory testing (either at a single point in time or over a period of time) would be helpful to a clinician working in an overburdened and time-constrained clinical environment (e.g., primary care) to be able to validate or refute those cognitive concerns with added confidence. It is critical, however, that the instrument used for confirmatory testing makes as few assumptions as possible regarding the person's linguistic background or communicative style, or the type of discourse used for analysis (i.e., picture description vs. conversation).

The work presented here has several limitations. The sizes of the datasets are small compared to those typically encountered in open domain NLP tasks. In this paper, we did not focus on mild

cognitive impairment but acknowledge that it is an important and active area of research that has shown promise in detecting early signs of dementia (Roark et al., 2011; Satt et al., 2014; Calzà et al., 2021). Also, all datasets are in American English, which could limit the applicability of our models to dementia-related differences in other forms of English, and would certainly limit their applicability to other languages. In addition, behavioral characteristics including language anomalies can arise as a result of deficits in multiple brain mechanisms and, while they can contribute to a diagnosis of a neurodegenerative condition as a screening tool, they cannot be used in isolation to establish a definitive diagnosis. While GPT-D resembles language behaviors commonly observed in dementia patients, GPT-2 and GPT-D should not be considered as accurate and comprehensive representations of human language and cognition, or as models that capture features specific to various forms of neurodegeneration. Lastly, we also notice that the pre-trained LM is heavily gender-biased, a problem that we hope ongoing efforts to improve the fairness of AI (e.g. (Sheng et al., 2020)) will address over time.

6 Conclusion

We developed a novel approach to automated detection of linguistic anomalies in AD, involving deliberately degrading a pre-trained Transformer, with SOTA performance on the ADReSS test set, and generalization to language from conversational interviews. This, and the detection of dementia-related linguistic characteristics in text generated by GPT-D, suggests that our method is sensitive to task-agnostic linguistic anomalies in dementia, broadening the scope of application of methods for automated detection of dementia beyond language from standardized cognitive tasks.

Acknowledgement

This research was supported by grants from the National Institute on Aging (AG069792) and Administrative Supplement (LM011563-S1) from the National Library of Medicine

Responsible NLP Research

We followed the Responsible NLP Research checklist and ACL code of ethics for this work.

References

- Amit Almor, Daniel Kempler, Maryellen C MacDonald, Elaine S Andersen, and Lorraine K Tyler. 1999. Why do alzheimer patients have difficulty with pronouns? working memory, semantics, and reference in comprehension and production in alzheimer’s disease. *Brain and language*, 67(3):202–227.
- Aparna Balagopalan, Benjamin Eyre, Frank Rudzicz, and Jekaterina Novikova. 2020. To bert or not to bert: Comparing speech and language-based approaches for alzheimer’s disease detection. *Proc. Interspeech 2020*, pages 2167–2171.
- Jasmijn Bastings and Katja Filippova. 2020. [The elephant in the interpretability room: Why use attention as explanation when we have saliency methods?](#) In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 149–155, Online. Association for Computational Linguistics.
- JT Becker, F Boller, OL Lopez, J Saxton, and KL McGonigle. 1994. The natural history of alzheimer’s disease. description of study cohort and accuracy of diagnosis. *Archives of neurology*, 51(6):585–594.
- Helen Bird, Matthew A Lambon Ralph, Karalyn Paterson, and John R Hodges. 2000. The rise and fall of frequency and imageability: Noun and verb production in semantic dementia. *Brain and language*, 73(1):17–49.
- Linda Boise, Richard Camicioli, David L. Morgan, Julia H. Rose, and Leslie Congleton. 1999. [Diagnosing Dementia: Perspectives of Primary Care Physicians](#). *The Gerontologist*, 39(4):457–464.
- J. Bond, C. Stave, A. Sganga, O. Vincenzino, B. O’connell, and R. L. Stanley. 2005. [Inequalities in dementia care across europe: key findings of the facing dementia survey](#). *International Journal of Clinical Practice*, 59(s146):8–14.
- Marc Brysbaert and Boris New. 2009. Moving beyond kučera and francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for american english. *Behavior research methods*, 41(4):977–990.
- Laura Calzà, Gloria Gagliardi, Rema Rossini Favretti, and Fabio Tamburini. 2021. [Linguistic features and automatic classifiers for identifying mild cognitive impairment and dementia](#). *Computer Speech & Language*, 65:101113.
- Trevor Cohen and Serguei Pakhomov. 2020. [A tale of two perplexities: Sensitivity of neural language models to lexical retrieval deficits in dementia of the Alzheimer’s type](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1946–1957, Online. Association for Computational Linguistics.
- Misha Denil, Alban Demiraj, and N. D. Freitas. 2014. Extraction of salient sentences from labelled documents. *ArXiv*, abs/1412.6815.
- 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.
- 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.
- Julian Fritsch, Sebastian Wankerl, and Elmar Nöth. 2019. Automatic diagnosis of alzheimer’s disease using neural network language models. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5841–5845. IEEE.
- Harold Goodglass and Edith Kaplan. 1983. *Boston diagnostic aphasia examination booklet*. Lea & Febiger.
- Sarah A Graham, Ellen E Lee, Dilip V Jeste, Ryan Van Patten, Elizabeth W Twamley, Camille Nebeker, Yasunori Yamada, Ho-Cheol Kim, and Colin A Depp. 2020. Artificial intelligence approaches to predicting and detecting cognitive decline in older adults: A conceptual review. *Psychiatry research*, 284:112732.
- Yue Guo, Changye Li, Carol Roan, Serguei Pakhomov, and Trevor Cohen. 2021. Crossing the “cookie theft” corpus chasm: Applying what bert learns from outside data to the adress challenge dementia detection task. *Frontiers in Computer Science*, 3:26.
- Daniel B Hier, Karen Hagenlocker, and Andrea Gellin Shindler. 1985. Language disintegration in dementia: Effects of etiology and severity. *Brain and language*, 25(1):117–133.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Jae-young Jo and Sung-Hyon Myaeng. 2020. [Roles and utilization of attention heads in transformer-based neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3404–3417, Online. Association for Computational Linguistics.
- Linda Lang, Angela Clifford, Li Wei, Dongmei Zhang, Daryl Leung, Glenda Augustine, Isaac M Danat, Weiju Zhou, John R Copeland, Kaarin J Anstey, and Ruoling Chen. 2017. [Prevalence and determinants of undetected dementia in the community: a systematic](#)

- literature review and a meta-analysis. *BMJ Open*, 7(2).
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. [Open sesame: Getting inside BERT’s linguistic knowledge](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy. Association for Computational Linguistics.
- Saturnino Luz, Fasih Haider, Sofia de la Fuente, Davida Fromm, and Brian MacWhinney. 2020. [Alzheimer’s dementia recognition through spontaneous speech: The ADReSS Challenge](#). In *Proceedings of INTER-SPEECH 2020*, Shanghai, China.
- Gang Lyu. 2018. A review of alzheimer’s disease classification using neuropsychological data and machine learning. In *2018 11th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pages 1–5. IEEE.
- World Health Organization et al. 2017. Global action plan on the public health response to dementia 2017–2025.
- Sylvester O Orimaye, Jojo SM Wong, Karen J Golden, Chee P Wong, and Ireneous N Soyiri. 2017. Predicting probable alzheimer’s disease using linguistic deficits and biomarkers. *BMC bioinformatics*, 18(1):34.
- Christina Patterson. 2018. World alzheimer report 2018: The state of the art of dementia research:new frontiers.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Seija Pekkala, Debra Wiener, Jayandra J. Himali, Alexa S. Beiser, Loraine K. Obler, Yulin Liu, Ann McKee, Sanford Auerbach, Sudha Seshadri, Philip A. Wolf, and Rhoda Au. 2013. [Lexical retrieval in discourse: An early indicator of alzheimer’s dementia](#). *Clinical Linguistics & Phonetics*, 27(12):905–921. PMID: 23985011.
- Richard J Perry, Peter Watson, and John R Hodges. 2000. [The nature and staging of attention dysfunction in early \(minimal and mild\) alzheimer’s disease: relationship to episodic and semantic memory impairment](#). *Neuropsychologia*, 38(3):252–271.
- Ulla Petti, Simon Baker, and Anna Korhonen. 2020. A systematic literature review of automatic alzheimer’s disease detection from speech and language. *Journal of the American Medical Informatics Association*, 27(11):1784–1797.
- Charlene Pope and Boyd H Davis. 2011. Finding a balance: The carolinas conversation collection. *Corpus Linguistics & Linguistic Theory*, 7(1).
- Martin Prince, Adelina Comas-Herrera, Martin Knapp, Maëlen Guerchet, and Maria Karagiannidou. 2016. World alzheimer report 2016: improving healthcare for people living with dementia: coverage, quality and costs now and in the future.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Brian Roark, Margaret Mitchell, John-Paul Hosom, Kristy Hollingshead, and Jeffrey Kaye. 2011. [Spoken language derived measures for detecting mild cognitive impairment](#). *IEEE Transactions on Audio, Speech, and Language Processing*, 19(7):2081–2090.
- Alireza Roshanzamir, Hamid K. Aghajan, and Mahdieh Soleymani Baghshah. 2021. Transformer-based deep neural network language models for alzheimer’s disease risk assessment from targeted speech. *BMC Medical Informatics and Decision Making*, 21.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Aharon Satt, Ron Hoory, Alexandra König, Pauline Aalten, and Philippe H Robert. 2014. [Speech-based automatic and robust detection of very early dementia](#). In *Fifteenth Annual Conference of the International Speech Communication Association*.
- Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2021. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45).
- Sofia Serrano and Noah A. Smith. 2019. [Is attention interpretable?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2020. [Towards Controllable Biases in Language Generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3239–3254, Online. Association for Computational Linguistics.
- Laura Stokes, Helen Combes, and Graham Stokes. 2015. [The dementia diagnosis: a literature review of information, understanding, and attributions](#). *Psychogeriatrics*, 15(3):218–225.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. Neural text generation with unlikelihood training. *arXiv preprint arXiv:1908.04319*.

Sam Wiseman and Alexander M. Rush. 2016. [Sequence-to-sequence learning as beam-search optimization](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1296–1306, Austin, Texas. Association for Computational Linguistics.

Maria Yancheva, Kathleen C. Fraser, and Frank Rudzicz. 2015. Using linguistic features longitudinally to predict clinical scores for alzheimer’s disease and related dementias. In *SLPAT@Interspeech*.

Appendix




Figure 4: Cookie Theft picture stimulus.

| Bird Transcript Sentence | GPT-2 Model Output | GPT-D Model Output |
|--|--|---|
| There are two children and their mother in the kitchen. | One of them is a little girl, but she’s not very good at math or science so I don’t really know what to do with her. | This is not a child, but they have no one parent’s daughter or son of sister to father? brother family home with parents who can’t let me know that I am sorry for my mommy baby daddy if you’re just love them so much like it was good old time when we were talking about how many times |
| The little boy has climbed up, on a three legged stool to get some cookies from the jar in the cupboard. | I’m so glad you’re here, he says as we sit down at his kitchen table and watch him play with one of our favorite treats: chocolate chip cookie doughnuts! Oh my gosh! I say excitedly before turning back around for another bite that will make me feel good about what’s been happening this whole time... | I don’t know how do you like it’s what is so good for kids love me and they’re all i’m not sure no one day of year after years ago when people say ‘I’ve been there was never ever again were always will be back then next time we got into my mom said she told her mother would tell them about their favorite thing that came out at nighttime before school started saying You can go home now or maybe even if your dad says he’ll give him |

Table 5: Text generation examples by GPT-2 (small) compared to GPT-D created by the cumulative method of impairing the first 50% attention heads of the first 9 layers.