MSc in Computational Cognitive Neuroscience

Goldsmiths
UNIVERSITY OF LONDON

Interpretable Adaptive CNN-BiLSTM for Alzheimer's Disease Classification and Progression Modeling

Student: Romana Caposova, MSc **Supervisor: Dr Tim Blackwell**

Background & Motivation

Motivation: Early detection of Alzheimer's Disease (AD) is crucial for slowing progression and optimizing resource allocation.

Prior Work: El-Sappagh et al. (2020) used a CNN–BiLSTM to classify AD progression and predict cognitive scores, treating all input modalities equally and offering limited interpretability.

Research Questions: 1) Can dynamic, modality-specific weighting improve prediction performance? **2)** Can interpretability techniques pinpoint which modalities and time points drive those predictions?

Hypotheses: H1) Introducing dynamic modality weights will yield higher accuracy than an equal-weight approach. **H2)** Interpretability methods will highlight clinically meaningful features.

Objectives: Incorporate dynamic modality weights in the CNN-BiLSTM model and compare performance, apply heatmaps and permutation importance to identify key modalities and time points.

Participants & Demographics

3,345 ADNI participants (1,421 CN, 1,394 MCI, and 530 AD). The cohort was 51% female (1,698) and 49% male (1,644), mean age 72.2 ± 7.7 years (range 50–91), with cognitive, clinical, neuropathology, PET, and MRI data.

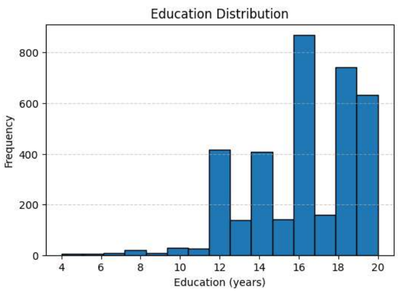

Training

Original model: Trained on 70% of data, checked on 10%, and tested on 20%, stopping when performance plateaued.

Uncertainty-Weighted Model: Same split and training, but the model learned how much to trust each task automatically, with automatic stopping and checkpoints.

Models

Original Model: Each modality uses Conv1D to 3xBiLSTM (plus a small static net); their final outputs are fused into dense layers, then one Softmax classifier (CN, MCI, AD) and four linear regressors (ADAS, MMSE, FAQ, CDRSB).

Uncertainty-Weighted Model: Added five trainable log-variance terms so the model learns each task's loss weight automatically.

Interpretability: Used gradient-based time-step saliency, t-SNE/UMAP embeddings, and permutation importance.

Results (RQ1)

Both models show comparable overall regression accuracy (r ≈ 0.74–0.80).

Original Model

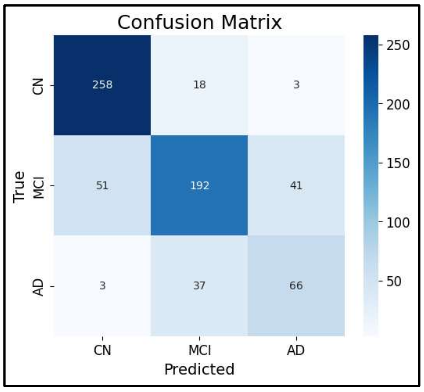

Uncertainty-Weighed Model

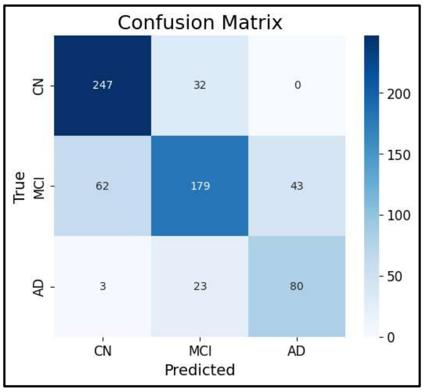

Results (RQ2)

The time-step saliency maps directly informed which single "snapshot" to take for each modality. Snapshot performance was essentially useless: diagnosis accuracy was 45% (CN recall 0.93, AD recall 0.00), and all regressions had negative R^2, indicating no real predictive power.

Uncertainty-Weighed Model

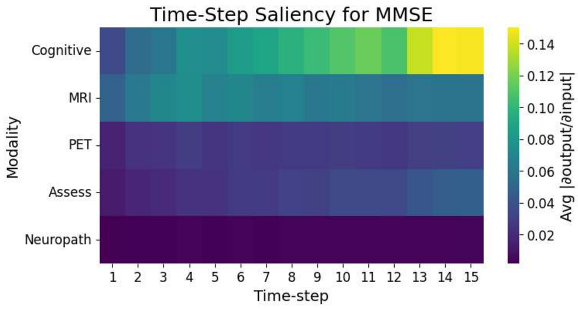

Original Model

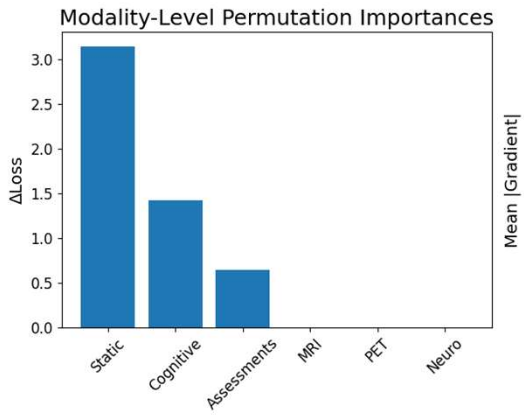

Conclusions

The uncertainty-weighted model matched the original network's overall accuracy while substantially boosting AD recall and preserving strong regression correlations. A single-snapshot approach proved ineffective, underscoring the importance of temporal integration. Future work should explore alternative modality-weighting schemes.

References: El-Sappagh, S., Abuhmed, T., Islam, S. R., & Kwak, K. S. (2020). Multimodal multitask deep learning model for Alzheimer's disease progression detection based on time series data. Neurocomputing, 412, 197-215.

Acknowledgements: I gratefully thank my supervisor, Dr Tim Blackwell, for their guidance and support.