

Abstract

Mild Cognitive Impairment (MCI) is a condition that negatively affects some older adults' memory, use of language, and judgement. While there are cognitive assessments to diagnose MCI, screening competes with the many other tasks which clinicians must perform, and is rarely done in a primary care setting. As a result, MCI remains difficult to identify and tends to be underdiagnosed. The research in this thesis was conducted in partnership with the Icahn School of Medicine at Mount Sinai, where researchers assembled a dataset of 800 recordings of primary care visits, paired with MCI diagnoses. Using this dataset, I present a deep learning pipeline for detecting MCI from long form clinical audio recordings. This pipeline uses existing neural network models alongside domain knowledge to correctly isolate the patient's speech in 98% of sampled recordings. Then, it uses short-time Fourier transforms and the Mel scale to convert the resulting audio into a more suitable format for neural network classification. The spectrograms are then split into overlapping regular segments, with each segment inheriting the label of the parent audio. The resulting dataset is used to train convolutional neural networks from scratch, as well as to fine-tune models pretrained on larger labeled datasets. These models convincingly exceed a baseline performance at the segment level, with an AUC ranging from 0.57 to 0.65 based on the model architecture and training regimen. A further exploration of model outputs and patient-level performance is also provided.