ADHD Classification from EEG

Subject-Level Generalization on Pediatric EEG

Capstone Project Report

A multi-channel 1D convolutional neural network that classifies children's EEG recordings by study group, using a strict subject-level split so the reported performance reflects genuine generalization to unseen individuals.

Arel · RLcapstone.ai

1. Summary

This project classifies pediatric EEG recordings by study group (ADHD vs. control) using a **multi-channel 1D convolutional neural network** that reads all 19 scalp electrodes simultaneously. With only 121 children but hundreds of thousands of short windows, the decisive choice was how to split the data. Using a strict **subject-level split**, the model attains **91.7% subject-level accuracy** and a **ROC-AUC of 0.965** on 24 previously unseen children.

Unlike the ECG project, this result exhibits no collapse from validation to test, indicating that the ADHD-vs-control signature in this task genuinely transfers across children.

> Not a diagnostic tool. ADHD is a clinical diagnosis made by professionals, never from a brief EEG. This project classifies research-cohort group membership and is strictly educational.

2. Motivation

EEG-based ADHD classification is both a popular topic and a well-known methodological trap: a substantial body of published work mixes windows from the same child across training and test, so the model learns to recognize individuals rather than anything about ADHD. The aim here was to avoid that leakage and report a figure that reflects real generalization.

3. Dataset

The project uses the public **Nasrabadi ADHD/Control EEG** dataset: 61 children with ADHD and 60 controls, ages 7–12, recorded across 19 electrodes (10–20 system) at 128 Hz during a visual attention task. Each recording was segmented into 2-second windows, with **every child kept entirely on one side of the split**.

Split	Children	Windows
Training	83	5,838
Validation	14	948
Test (unseen children)	24	1,614

4. Model and evaluation protocol

The classifier is a **multi-channel 1D CNN** (193k parameters, 775 KB) that reads all 19 channels jointly. Evaluation is performed **per subject, not per window**: each child's window-level predictions are averaged into a single decision for that child, so a child contributing many windows does not dominate the score. The validation set used during training is itself a separate group of children, so every reported figure measures generalization to new individuals.

5. Results

On unseen children	Value

Subject-level accuracy	91.7% (22 of 24)
ROC-AUC	0.965
Window-level accuracy	89.0%

The validation and test figures fall in the same range, indicating that the model captures a genuine group-level signal rather than memorizing individuals. Two caveats maintain a balanced interpretation: 91.7% of 24 subjects is 22 of 24 — a small denominator with a wide confidence interval — and the model classifies **research-group membership**, not a clinical diagnosis.

6. Deployment

The model runs in a browser demonstration that renders all 19 EEG channels like a clinical montage and classifies a recording on-device, with a matching TensorFlow Lite build in the Android application (EEG Explorer). Both reproduce the original model's predictions.

7. Limitations

- Classifies two research groups — it does not diagnose any individual.
- Only 24 test subjects, so 91.7% is 22 of 24; the true accuracy could be meaningfully higher or lower.
- A single cohort and a single task; real-world attention difficulties are far more varied.

8. Future work (v2)

The most valuable follow-ups concern trust rather than headline accuracy: band-limiting the input to identify which **frequency bands** carry the signal, **ablating electrodes** to determine which scalp regions are informative, and running subject-level cross-validation so the headline figure rests on all 121 children rather than a single 24-subject split.