DeepScope: A Cascading Biological Plausibility Architecture for Real-Time Deepfake Detection

Rana, Parth (School: Wayzata High School)

Deepfake videos are increasingly convincing, yet current detectors show a frustrating gap between lab performance (over 90% accuracy) and real-world results (50–60% AUC) with latencies too slow for live video. My project, DeepScope, takes a different angle: instead of learning to spot what looks fake, it checks whether a face shows the involuntary biological signals that real human faces always exhibit. I developed 12 Biological Plausibility Metrics (BPM), including micro-tremor analysis in the 8-12 Hz band, facial asymmetry patterns, and muscle group coordination, derived from 10,000 real videos across diverse conditions. I also built Temporal Consistency Fingerprinting (TCF), which uses triplet-loss embeddings to detect manipulation types never seen in training. These feed into a three-stage cascade: a Binary Neural Network pre-filter (1.2M parameters, 7ms) handles ~70% of videos immediately, a Wavelet-Optical Flow stage processes most of the remainder, and a deep ensemble handles the hardest 5%. I evaluated DeepScope on 27,729 videos across FaceForensics++, Celeb-DF-v2, DFDC, and a custom diffusion dataset using five-fold cross-validation. The system achieved 95.3% (±1.2%) within-dataset and 85.7% (±2.1%) cross-dataset accuracy, compared to 65.2% for XceptionNet and 71.3% for FakeCatcher ($p < 0.001$, McNemar's test). Average latency was 14ms, below the 33ms threshold for 30fps video. These results suggest that modeling biological authenticity, rather than memorizing generation-specific artifacts may offer a more reliable path toward practical deepfake detection.

Awards Won:
Third Award of $1,200