

Real-Time Sign Language Recognition

Sign language is the primary communication medium for millions of deaf and hard-of-hearing individuals, yet the communication barrier with non-signers remains near-total. This project implements a real-time sign language recognition system using commodity hardware:

Google's MediaPipe framework extracts 21 three-dimensional hand landmarks per frame from a standard webcam feed, and a hybrid CNN-LSTM classifier maps these landmark sequences to sign classes. Static signs (alphabets, digits) are classified from single-frame landmark vectors by the convolutional branch, while dynamic gestures are modeled as 30-frame sequences by the LSTM branch, capturing the temporal structure of motion-based signs.

A one-click dataset collection tool allows recording custom sign samples, making the vocabulary extensible and giving students a genuine dataset-design story for their viva. Recognized signs are assembled by a sentence builder and vocalized through text-to-speech, completing a two-way communication aid. The system is evaluated on a held-out test set with per-class accuracy and a confusion matrix, and the report documents failure modes honestly — lighting sensitivity, occlusion, and inter-signer variation. Achieving real-time inference at 30 fps on a standard laptop CPU, the project demonstrates applied deep learning from landmark extraction through sequence modeling to deployment. Built-to-order with source code, pre-trained models, dataset tools, full project report, presentation PPT, and viva Q&A; support, this is a high-empathy, high-demo-value final-year AI submission in one of the largest student search niches.