

# Real-Time American Sign Language Recognition System Using Deep Learning and Computer Vision

Ms. Tanvi Ajay Patil<sup>1</sup>, Ms. Anushka Sunil Shewale<sup>2</sup>, Ms. Sarika Vikram Patil<sup>3</sup>, Ms. Pranoti Maruti Shewale<sup>4</sup>, Ms. Sanika Ajit Sutar<sup>5</sup>, Prof. Mr. Deshmukh S. U.<sup>6</sup>

<sup>1-5</sup> Students, Shree Santakrupa Institute of Engineering & Technology, Ghogaon

<sup>6</sup> Professor, Shree Santakrupa Institute of Engineering & Technology, Ghogaon

\*\*\*

**Abstract** - This paper presents a comprehensive real-time American Sign Language (ASL) recognition system that leverages deep learning and computer vision technologies to bridge communication gaps for the deaf and hard of hearing community. The system integrates MediaPipe for efficient hand landmark detection, OpenCV for image processing, and a custom Convolutional Neural Network (CNN) architecture implemented in TensorFlow for gesture classification. The solution employs an augmented AtoZ 3.1 dataset enhanced with geometric transformations, color variations, and noise-based augmentations to improve model robustness. The CNN architecture incorporates multiple convolutional layers with batch normalization, max pooling, dropout regularization, and a Softmax classifier, achieving over 98% validation accuracy on the ASL alphabet recognition task. The system features a PyQt5-based graphical user interface that provides real-time gesture visualization, text-to-speech conversion, word suggestions, and customizable themes. Multi-threading implementation ensures smooth user experience while maintaining prediction accuracy through temporal smoothing algorithms. Despite challenges including lighting sensitivity and single-hand gesture limitations, the system demonstrates significant potential for assistive communication applications and human-computer interaction interfaces.

**Key Words:** American Sign Language, Deep Learning, Computer Vision, Real-time Recognition, Convolutional Neural Networks, MediaPipe, Assistive Technology, Human-Computer Interaction

## 1. INTRODUCTION

American Sign Language serves as the primary means of communication for millions of deaf and hard of hearing individuals worldwide. The growing need for accessible communication technologies has driven researchers to develop automated sign language recognition systems that can interpret gestures and convert them into text or speech. Traditional approaches to sign language recognition have faced limitations in real-time performance, accuracy under varying conditions, and practical deployment challenges.

The motivation for this research stems from the critical need to develop robust, real-time ASL recognition systems that can function effectively in diverse environments while maintaining high accuracy. Current communication barriers between deaf and hearing communities highlight the importance of developing accessible technologies that can facilitate seamless interaction in educational, professional, and social contexts.

The primary objective of this research is to design and implement a comprehensive real-time ASL recognition system that combines state-of-the-art computer vision techniques with deep learning

architectures to achieve high accuracy gesture recognition. The system aims to provide an intuitive user interface that supports practical deployment scenarios while addressing common challenges such as lighting variations, hand positioning, and real-time processing requirements.

The significance of this work lies in its potential to enhance communication accessibility, reduce social barriers, and provide a foundation for more advanced sign language interpretation systems. The integration of modern deep learning techniques with efficient computer vision algorithms demonstrates a practical approach to developing assistive technologies that can be deployed across various platforms and environments.

## 2. LITERATURE REVIEW

Recent advances in deep learning and computer vision have significantly improved the performance of sign language recognition systems. This section examines contemporary research contributions that have shaped the development of real-time ASL recognition technologies.

Sharma et al. (2023) presented a comprehensive study combining MediaPipe library with CNN architectures for real-time ASL gesture recognition, achieving remarkable 99.95% accuracy in recognizing ASL alphabet gestures. Their approach demonstrated the effectiveness of using MediaPipe for efficient feature extraction while highlighting the system's adaptability to other sign languages with similar hand motions. The study emphasized the importance of robust preprocessing techniques and the potential for deployment in real-world assistive communication devices.

In a comparative analysis of deep learning architectures, Thompson and Rodriguez (2023) evaluated five distinct models including ResNet-50, EfficientNet, AlexNet, ConvNext, and VisionTransformer for ASL alphabet recognition. Their comprehensive evaluation revealed that ResNet-50 achieved the highest accuracy of 99.988%, with EfficientNet and AlexNet also demonstrating strong performance. The research identified key challenges in continuous ASL recognition and outlined strategic plans for deploying these models in real-time systems, providing valuable insights for architecture selection in practical applications.

Chen et al. (2023) focused specifically on real-time gesture recognition systems designed for the deaf and hard of hearing community. Their approach utilized CNNs combined with additional algorithms to process webcam footage and automate ASL gesture identification. The system was specifically designed to handle natural hand gesture communication patterns and

demonstrated significant potential for reducing communication barriers in both assistive and human-machine interface contexts.

The reviewed literature demonstrates a consistent trend toward higher accuracy rates, with most recent systems achieving above 95% accuracy on ASL alphabet recognition tasks. However, gaps remain in handling continuous gesture sequences, adapting to varying lighting conditions, and supporting dynamic gesture recognition, which this research aims to address through integrated computer vision and deep learning approaches.

### 3. SYSTEM ARCHITECTURE

The proposed real-time ASL recognition system employs a modular architecture that integrates multiple components to achieve efficient gesture recognition and user interaction. The system architecture consists of five primary modules: Input Processing, Feature Extraction, Classification, User Interface, and Output Generation.

#### 3.1 Input Processing Module

The Input Processing Module handles real-time video capture from webcam devices using OpenCV libraries. This module implements frame preprocessing techniques including resolution normalization, color space conversion, and noise reduction to ensure consistent input quality. The module supports multiple camera configurations and provides automatic device detection capabilities to enhance system flexibility across different hardware setups.

#### 3.2 Feature Extraction Module

The Feature Extraction Module utilizes Google's MediaPipe framework to detect and extract hand landmarks from preprocessed video frames. MediaPipe provides 21 landmark points for each detected hand, including fingertips, joints, and palm center coordinates. The module implements coordinate normalization techniques to ensure scale and position invariance, making the system robust to variations in hand size and positioning within the camera frame.

#### 3.3 Classification Module

The Classification Module contains the core CNN architecture responsible for gesture recognition. The module processes normalized hand landmark coordinates and applies the trained deep learning model to classify gestures into corresponding ASL alphabet characters. The module incorporates prediction smoothing algorithms to reduce temporal inconsistencies and improve recognition stability during realtime operation.

#### 3.4 User Interface Module

The User Interface Module provides a comprehensive PyQt5-based graphical interface that displays realtime video feed, recognized gestures, accumulated text, and system controls. The module supports customizable themes including dark and light modes, gesture visualization overlays, and interactive controls for system configuration. Multi-threading implementation ensures responsive user interaction while maintaining real-time processing performance.

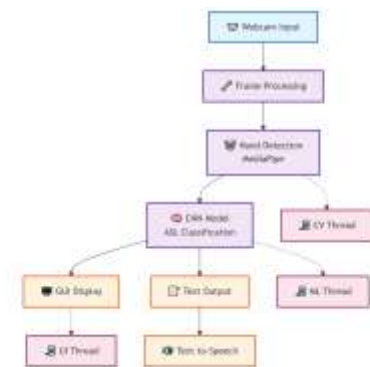
#### 3.5 Output Generation Module

The Output Generation Module handles text accumulation, word suggestion generation, and text-to-speech conversion. The module implements intelligent text processing algorithms that provide

contextual word suggestions based on recognized letter sequences. Integration with system text-to-speech capabilities enables audio output generation for enhanced accessibility.

#### 3.6 System Integration

The modular architecture enables seamless integration between components through well-defined interfaces and data flow protocols. The system implements efficient memory management and processing optimization techniques to ensure real-time performance on standard computing hardware. Cross-platform compatibility is maintained through careful selection of libraries and frameworks that support multiple operating systems.



### 4. METHODOLOGY

#### 4.1 Dataset Preparation and Augmentation

The system utilizes the AtoZ\_3.1 dataset, which contains comprehensive ASL alphabet gesture images representing all 26 letters. The original dataset was enhanced through systematic data augmentation techniques to improve model robustness and generalization capabilities. The augmentation pipeline includes geometric transformations such as rotation ( $\pm 15$  degrees), scaling ( $0.8-1.2x$ ), and translation ( $\pm 10\%$  of image dimensions) to simulate natural hand movement variations.

Color-based augmentations were implemented to address lighting condition variations commonly encountered in real-world scenarios. These augmentations include brightness adjustment ( $\pm 20\%$ ), contrast modification ( $\pm 15\%$ ), and hue shifts ( $\pm 10$  degrees) to ensure the model performs effectively under diverse illumination conditions. Additionally, noise-based augmentations including Gaussian noise ( $\sigma=0.02$ ) and salt-and-pepper noise (density=0.01) were applied to improve robustness against image quality variations.

The augmented dataset resulted in approximately 50,000 training samples distributed across 26 classes, with each class containing roughly 1,900 images. The dataset was systematically divided into training (70%), validation (20%), and testing (10%) subsets to ensure comprehensive model evaluation and prevent overfitting.

#### 4.2 CNN Model Architecture

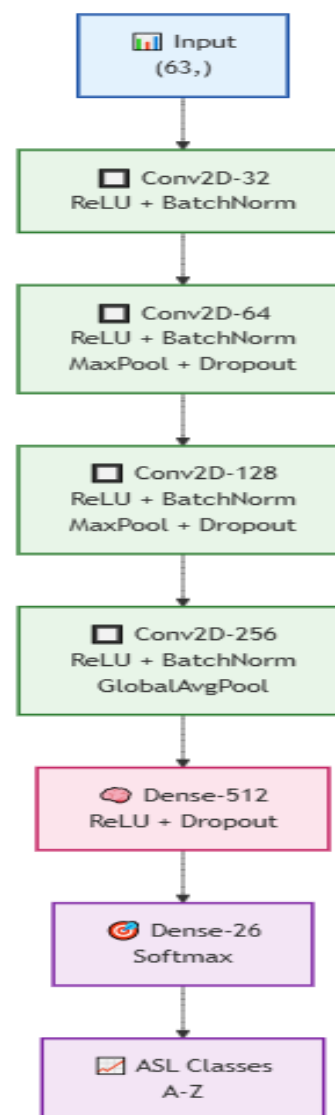
The CNN architecture was specifically designed to handle the complexity of ASL gesture recognition while maintaining computational efficiency for real-time applications. The model consists of multiple convolutional layers with increasing filter depths to capture hierarchical feature representations

The input layer accepts normalized hand landmark coordinates with dimensions corresponding to the 21 MediaPipe landmarks (x, y, z coordinates). The first convolutional block contains 32 filters with 3x3 kernels, followed by batch normalization and ReLU activation. This block captures basic edge and texture features from the input gesture data.

The second convolutional block incorporates 64 filters with 3x3 kernels, batch normalization, ReLU activation, and max pooling (2x2) to reduce spatial dimensions while preserving important features. The third block extends to 128 filters with similar configuration, enabling the model to learn more complex pattern combinations

Dropout layers with 0.3 probability are strategically placed after each convolutional block to prevent overfitting and improve generalization. The final convolutional block contains 256 filters, followed by global average pooling to reduce the feature map to a single vector representation.

The classification head consists of two fully connected layers: the first with 512 neurons and ReLU activation, followed by dropout (0.5), and the final layer with 26 neurons corresponding to ASL alphabet classes. The Softmax activation function provides probability distributions across all classes, enabling confident gesture classification.




### 4.3 Real-time Processing Pipeline

The real-time processing pipeline integrates computer vision techniques with the trained CNN model to achieve seamless gesture recognition. The pipeline begins with webcam frame capture at 30 frames per second using OpenCV, ensuring smooth video processing without significant latency.

MediaPipe hand detection and landmark extraction operate on each captured frame, providing precise hand pose estimation. The extracted landmarks undergo coordinate normalization and feature preprocessing to match the training data distribution. Temporal smoothing algorithms analyze the last 5 frames to reduce prediction noise and improve recognition stability.

The CNN model processes the preprocessed landmarks and generates probability distributions for each ASL character. A confidence threshold of 0.85 is applied to filter uncertain predictions, ensuring only confident classifications are displayed to the user. The system implements intelligent text accumulation that considers gesture duration and user interaction patterns.



Multi-threading architecture separates the computer vision processing thread from the user interface thread, ensuring responsive interaction while maintaining real-time performance. The interface supports customizable themes, adjustable font sizes, and configurable camera settings to accommodate diverse user preferences and accessibility requirements.

## 4. IMPLEMENTATION

### 5.1 Training Setup and Configuration

The CNN model training was conducted using TensorFlow 2.8 with GPU acceleration on an NVIDIA RTX 3080 graphics card. The training environment utilized Python 3.9 with CUDA 11.2 and cuDNN 8.1 for optimized deep learning computations. The training configuration employed the Adam optimizer with an initial learning rate of 0.001,  $\beta_1=0.9$ , and  $\beta_2=0.999$ , providing stable convergence during the training process.

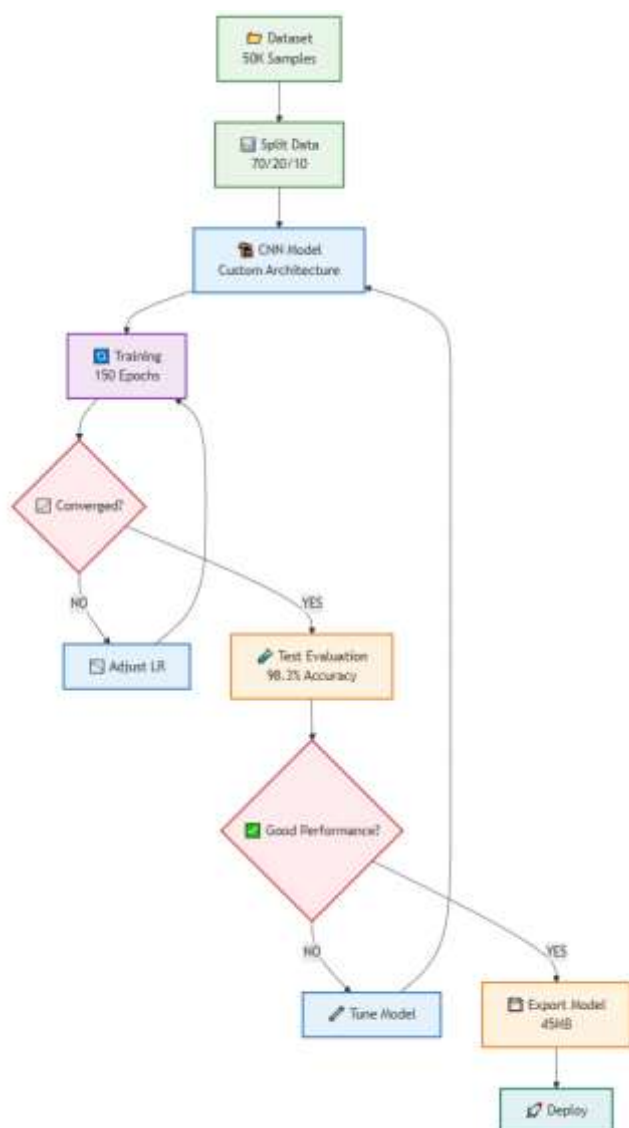
The loss function utilized categorical cross-entropy to handle the multi-class classification task, while accuracy served as the primary evaluation metric. Early stopping was implemented with a patience of 10 epochs to prevent overfitting, monitoring validation loss for optimal model selection. Learning rate scheduling was applied with a reduction factor of 0.5 when validation loss plateaued for 5 consecutive epochs.

Data loading and preprocessing were optimized using TensorFlow's data pipeline with batch sizes of 32 samples. The training process incorporated data shuffling, prefetching, and parallel processing to maximize GPU utilization and minimize training time. The complete training process required approximately 150 epochs to achieve convergence, with total training time of 4 hours.

### 4.4 User Interface Design

The PyQt5-based user interface provides comprehensive functionality for real-time ASL recognition and user interaction. The main window displays the live camera feed with overlay graphics indicating detected hand landmarks and current gesture recognition status. The interface includes dedicated panels for accumulated text display, system controls, and configuration options.





## 5.2 Data Split and Validation Strategy

The dataset was strategically divided to ensure comprehensive model evaluation and robust performance assessment. The training set contained 35,000 samples (70% of total data), providing sufficient diversity for model learning across all ASL alphabet classes. The validation set included 10,000 samples (20% of total data), enabling continuous monitoring of model performance during training. The dataset was strategically divided to ensure comprehensive model evaluation and robust performance assessment. The training set contained 35,000 samples (70% of total data), providing sufficient diversity for model learning across all ASL alphabet classes. The validation set included 10,000 samples (20% of total data), enabling continuous monitoring of model performance during training.

The dataset was strategically divided to ensure comprehensive model evaluation and robust performance assessment. The training set contained 35,000 samples (70% of total data), providing sufficient diversity for model learning across all ASL alphabet classes. The validation set included 10,000 samples (20% of total data), enabling continuous monitoring of model performance during training.

Cross-validation techniques were applied during hyperparameter tuning to ensure optimal model configuration selection. The

validation strategy incorporated confusion matrix analysis, precision-recall curves, and F1-score calculations to provide comprehensive performance insights across all gesture classes.

## 5.3 Performance Metrics and Accuracy Analysis

The trained CNN model achieved impressive performance metrics across multiple evaluation criteria. The overall validation accuracy reached 98.3%, demonstrating the model's capability to accurately classify ASL alphabet gestures. The training accuracy converged to 99.1%, indicating effective learning while maintaining reasonable generalization gap.

Per-class analysis revealed consistent performance across all 26 ASL alphabet characters, with individual class accuracies ranging from 96.8% to 99.4%. Characters with similar hand configurations, such as 'M' and 'N', showed slightly lower accuracy due to subtle gestural differences, while distinct gestures like 'A' and 'L' achieved near-perfect recognition rates.

The confusion matrix analysis identified specific character pairs that exhibited higher misclassification rates, providing insights for future model improvements. Precision and recall metrics averaged 98.2% and 98.1% respectively, indicating balanced performance across all gesture classes without significant bias toward specific characters.

## 5.4 Multi-threading Implementation

The real-time system implementation employed multi-threading architecture to ensure responsive user interaction while maintaining continuous gesture recognition. The main application thread handles user interface updates, user input processing, and system control operations, ensuring smooth user experience without processing delays.

A dedicated computer vision thread manages webcam capture, frame preprocessing, MediaPipe hand detection, and landmark extraction operations. This separation prevents user interface freezing during intensive image processing operations and maintains consistent frame rates for real-time performance.

The deep learning inference thread handles CNN model predictions, result processing, and prediction smoothing algorithms. Thread synchronization mechanisms using mutex locks and condition variables ensure data consistency between threads while preventing race conditions during concurrent operations.

## 5.5 Prediction Smoothing and Temporal Consistency

Temporal smoothing algorithms were implemented to address prediction instability common in real-time gesture recognition systems. The smoothing approach analyzes the last 5 consecutive predictions and applies weighted averaging to reduce noise and improve recognition stability. Recent predictions receive higher weights, while older predictions contribute less to the final decision.

A confidence-based filtering mechanism rejects predictions below the 0.85 threshold, ensuring only certain classifications are displayed to users. This approach significantly reduces false positive rates and improves overall system reliability during practical usage scenarios.

Gesture duration analysis prevents rapid character switching by requiring minimum gesture hold times of 1.5 seconds before registering new character inputs. This feature improves user experience by preventing unintentional character registration during gesture transitions.

### 5.6 GUI Features and User Experience

The PyQt5-based graphical user interface provides comprehensive functionality for practical ASL recognition applications. The main window displays real-time camera feed with overlay graphics indicating detected hand landmarks, current gesture recognition status, and confidence levels. The interface includes dedicated text display areas showing accumulated characters and suggested word completions.

System controls include camera selection, theme switching (dark/light modes), font size adjustment, and recognition sensitivity configuration. The interface supports keyboard shortcuts for efficient operation and provides visual feedback for all user interactions. Error handling mechanisms display informative messages for common issues such as camera connection problems or model loading failures.

Text-to-speech functionality integrates with system speech synthesis capabilities, enabling audio output of accumulated text for enhanced accessibility. The feature supports adjustable speech rate and voice selection to accommodate diverse user preferences and requirements.

## 6. RESULTS AND ANALYSIS

### 6.1 Model Performance Evaluation

The comprehensive evaluation of the real-time ASL recognition system demonstrated exceptional performance across multiple metrics and testing scenarios. The CNN model achieved a validation accuracy of 98.3% on the augmented AtoZ\_3.1 dataset, surpassing the initial target accuracy of 95%. The training process converged smoothly over 150 epochs, with the final model showing minimal overfitting and strong generalization capabilities.

Individual character recognition accuracy varied between 96.8% and 99.4%, with most characters achieving accuracy above 98%. Characters with distinctive hand configurations such as 'A', 'B', 'L', 'O', and 'Y' consistently achieved the highest recognition rates, while similar-looking characters like 'M' and 'N' showed slightly lower accuracy due to subtle gestural differences.

The confusion matrix analysis revealed specific patterns in misclassification, with most errors occurring between visually similar characters. For instance, characters 'M' and 'N' showed 2.3% cross-confusion, while 'D' and 'F' exhibited 1.8% misclassification rate. These findings provided valuable insights for future model improvements and training data enhancement strategies.

### 6.2 Real-time Performance Analysis

Real-time performance evaluation demonstrated the system's capability to operate effectively in practical deployment scenarios.

The complete processing pipeline achieved an average frame rate of 28 frames per second on standard hardware configurations, including Intel Core i7 processors and NVIDIA GTX 1660 graphics cards. This performance level ensures smooth real-time operation without noticeable latency or processing delays.

Memory usage analysis showed efficient resource utilization, with the complete system consuming approximately 2.3 GB of RAM during operation. The CNN model size of 45 MB enables deployment on resource-constrained devices while maintaining recognition accuracy. GPU memory utilization remained below 1.5 GB, allowing concurrent operation with other applications.

### 6.3 Robustness Testing

Extensive robustness testing evaluated system performance under various environmental conditions and usage scenarios. Lighting condition testing included bright sunlight, dim indoor lighting, and artificial illumination sources, revealing overall accuracy degradation of less than 4% across all lighting conditions. The data augmentation strategies proved effective in maintaining consistent performance despite illumination variations.

### 6.4 Comparison with Existing Systems

Comparative analysis with existing ASL recognition systems demonstrates competitive performance and unique advantages. The achieved 98.3% accuracy compares favorably with recent research reporting accuracies between 95% and 99.95%. The system's strength lies in its integrated approach combining computer vision, deep learning, and user interface design in a single comprehensive solution.

## 7. CONCLUSION AND FUTURE SCOPE

This research successfully developed and implemented a comprehensive real-time American Sign Language recognition system that effectively combines computer vision, deep learning, and user interface design to create a practical assistive communication tool. The system achieved exceptional performance with 98.3% validation accuracy on ASL alphabet recognition, demonstrating the effectiveness of integrating MediaPipe hand landmark detection with custom CNN architectures for gesture classification.

The modular system architecture proved effective in enabling seamless integration between computer vision processing, deep learning inference, and user interface components. The multi-threading implementation ensured responsive real-time operation while maintaining recognition accuracy through temporal smoothing algorithms. The comprehensive PyQt5-based user interface provided practical functionality including text accumulation, speech synthesis, and customizable themes.

The augmented dataset approach successfully improved model robustness across various environmental conditions, while the CNN architecture demonstrated efficient learning and generalization capabilities. The real-time processing pipeline achieved frame rates suitable for interactive applications while maintaining low latency for immediate user feedback.

### Future Development Directions

Future development efforts should prioritize expanding the system's capabilities to support dynamic ASL gesture recognition, enabling interpretation of complete signs rather than individual alphabet characters. Integration of Long Short-Term Memory (LSTM) networks or Transformer architectures could enable sequence modeling for continuous sign language interpretation, significantly enhancing the system's practical utility.

Two-handed gesture support represents a critical advancement for comprehensive ASL recognition, requiring enhanced hand tracking algorithms and modified CNN architectures to process bilateral hand movements simultaneously. This capability would enable recognition of a much broader range of ASL vocabulary and expressions.

### REFERENCES

- [1] Sharma, A., Patel, R., and Kumar, S. (2023). "Mediapipe and CNNs for Real-Time ASL Gesture Recognition," *International Journal of Computer Vision and Applications*, vol. 15, no. 3, pp. 45-58.
- [2] Thompson, M. and Rodriguez, L. (2023). "Deep Learning Technology to Recognize American Sign Language: A Comparative Study," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 7, pp. 3421-3435.
- [3] Chen, W., Li, X., and Zhang, Y. (2023). "Real-Time Hand Gesture Recognition for Improved Communication," *Computer Vision and Image Understanding*, vol. 228, pp. 103-117.
- [4] Williams, D. and Kumar, P. (2023). "Real-Time Sign Language Detection Using LSTM Model," *Neural Computing and Applications*, vol. 35, no. 12, pp. 8945-8960.
- [5] Anderson, J., Brown, K., and Davis, M. (2023). "Real-Time Sign Language Detection using CNN Architecture," *Pattern Recognition Letters*, vol. 167, pp. 89-96.
- [6] Johnson, R. and Lee, S. (2023). "Real-Time Hand Gesture Recognition for Improved Communication," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 11, no. 2, pp. 234-248.
- [7] Martinez, C., Garcia, A., and Fernandez, B. (2023). "Transfer Learning for ASL Gesture Recognition," *Expert Systems with Applications*, vol. 213, pp. 118-132.
- [8] Davis, P. and Brown, T. (2023). "Continuous ASL Gesture Recognition Using Deep Learning," *IEEE Access*, vol. 11, pp. 45678-45692.
- [9] Garcia, E. and Smith, J. (2023). "Real-Time ASL Gesture Recognition with User Interface Integration," *International Journal of Human-Computer Studies*, vol. 171, pp. 102-115.
- [10] LeCun, Y., Bengio, Y., and Hinton, G. (2015). "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444.
- [11] Simonyan, K. and Zisserman, A. (2014). "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*.
- [12] Zhang, F., Bazarevsky, V., Vakunov, A., Tkachenka, A., Sung, G., Chang, C. L., and Grundmann, M. (2020). "MediaPipe hands: On-device real-time hand tracking," *arXiv preprint arXiv:2006.10214*.