


International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)





SignVerse: Turning Hand Signs into Seamless Conversation with Gen AI

Khushi N. Kumawat, Aastha Nitin Ambavkar, Siddhant Ganesh Mandlik,

Sarvesh Pramod Raghatate, Dr. Mukesh Israni

UG Student, Dept. of I.T., Thadomal Sahani College of Engineering, Mumbai, Maharashtra, India

Associate Professor, Dept. of I.T., Thadomal Sahani College of Engineering, Mumbai, Maharashtra, India

ABSTRACT: Although there has been a lot of progress in the area of connecting people via computers and silencing the gap between hearing people and deaf people who use sign language, there are still many barriers in communication between the two groups. Traditional sign language recognition (SLR) systems provide only a way to translate each letter of a word individually, only translating the letter A for a person would appear as an A rather than as "Hello." The process is very slow and does not provide an accurate understanding of how people communicate naturally. Sign Verse is a new application that will primarily be used as a digital dictionary and also as a smart sentence builder. Sign Verse also uses a standard RGB camera that can display live video of American Sign Language (ASL) and India Sign Language (ISL) at the same time in both languages. Sign Verse uses Google's MediaPipe to extract 3D spatial features from the live video feed and uses Tensorflow/Keras machine learning classification to classify signs from the live video feed in real-time with minimal delays. Signs that were recognized from the live video feed will be sent into a Large Language Model (LLM) called Google Gemini to predict contextually relevant full sentence suggestions. Sign Verse will process video locally (on the user's computer) and will transmit only lightweight text to the cloud for the purpose of maintaining the user's privacy and expediting digital communication for the hearing impaired.

KEYWORDS: Sign Language Recognition (SLR), American Sign Language (ASL), Indian Sign Language (ISL), Real-Time Translation, Generative AI.

I. INTRODUCTION

Over 400 million people around the world have disabling hearing loss, the World Health Organization says. For a lot of them, sign language isn't just useful—it's how they communicate every day. It's a rich language that blends fast hand movements, facial expressions, and body posture in ways spoken language just can't capture. Machine learning has come a long way in understanding images and video, but honestly, turning these layered gestures into smooth, natural speech or writing on the fly is still a huge challenge. It takes a ton of computing power, and technology's not quite there yet.

Most SLR tools out there still act like basic translators—they just convert what they see into individual letters. If someone signs "H-O-W," the system doesn't bother piecing it together as a word. It just throws out "H," then "O," then "W," one after the other. You're stuck doing the heavy lifting, stringing each letter together in your head and figuring out the actual meaning. Honestly, this method just slows everything down. People have to put in a lot more effort, both mentally as they decode the letters, and physically as the whole process turns into a tedious back-and-forth. By the time you've spelled out a word or two, the flow of the conversation is gone. It's nothing like the way people naturally communicate. Real dialogue moves quickly—full sentences fly back and forth, meaning is understood instantly, and there's barely any lag. When SLR tools fall short like this, they don't just cause frustration, they widen the gap between signing and speech, making communication much less natural and way more exhausting. It's clear these systems need to get a lot smarter if they're going to keep up with real human interaction.

Sign Verse is more than just a tool for translating. It moves beyond static, word-for-word conversions and actually helps people create sentences that flow naturally. The trick lies in how it combines two powerful technologies. On the one hand, you've got edge-based Computer Vision—so the device itself can recognize and interpret signed gestures



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

quickly, right where you are, without relying on a constant internet connection. On the other hand, there's cloud-based Natural Language Processing. This means the system isn't just looking at individual signs; it's truly understanding how they fit together in context, especially since American Sign Language (ASL) and Indian Sign Language (ISL) have very different grammar rules.

So as you're signing, Sign Verse watches for cues from both languages. It uses the Gemini AI API in the background to suggest how your ideas could become full, nuanced sentences—almost as if you had another speaker there, helping you find the right expression. Instead of being stuck with basic, robotic translations, users can finally express themselves in a way that feels right to them, no matter which sign language or spoken language they're working with. As a result, Sign Verse doesn't just bridge communication gaps—it opens up new possibilities for real, everyday conversations between people who would have struggled to connect before.

II. LITERATURE REVIEW

Sign language recognition really splits into two groups: hardware-based and vision-based approaches.

Starting with hardware—people first tried sensor gloves packed with flex sensors and accelerometers to capture every finger and hand movement. Sure, they were accurate. But nobody actually wants to wear a giant, wired glove just to talk to someone. They're awkward, expensive, and honestly, most people would rather not bother.

Vision-based methods changed the game once Convolutional Neural Networks like ResNet and VGG16 showed up. With just a camera, these systems could recognize things like the ASL alphabet pretty well. No more gadgets or gloves—just video. That made everything way more comfortable. The problem? Processing all those video frames fast enough sucked up a ton of computing power. The lag ruined the flow, making conversations slow and unnatural. More recently, generative AI models—especially Large Language Models—kicked up the excitement again. LLMs are amazing at catching context and reading between the lines, which is exactly what sign language interpretation needs. But turning all that fancy AI power into something real-time? That's still pretty fresh.

That's where Sign Verse steps in. It uses MediaPipe to skip the messy, heavy video analysis and goes straight for the good stuff—skeletal data pulled from the camera. The system recognizes the sign from this lighter data, then hands it to a generative AI, basically asking, "Okay, what should happen next?" The result: faster, smoother, more natural conversations, and no need for weird gadgets. Sign Verse ditches the gloves and the lag, making real-time sign language chat actually feel real. Signers finally get to communicate just as naturally as everyone else.

III. SYSTEM ARCHITECTURE

The Sign Verse architecture is designed to achieve low latency and high accuracy. It eschews conventional, compute-intensive image processing for low-overhead skeletal tracking and cloud-based natural language processing. System architecture is organized into three interrelated processing pipelines: Data Capture and Preprocessing, Spatial Feature Extraction, and the Deep Learning Classification Model.

A. Hardware interface and Data Preprocessing

The underlying architecture is built on real-time visual acquisition. The system harnesses the OpenCV (cv2) computer vision library in order to interface directly with the user's local web camera.

In order to facilitate a smooth user experience and to obviate any blurring of rapid hand movements, the camera captures a continuous video stream locked at 30 FPS. Once captured, the raw frames are immediately preprocessed by OpenCV. The default BGR (Blue, Green, Red) color matrix generated by OpenCV is transposed into an RGB color space, which is an absolute precondition for the downstream MediaPipe modules to correctly identify human skin tones and hand topologies.

B. Spatial Feature Extraction(MediaPipe)

Traditional computer vision models work by feeding raw image pixels through massive Convolutional Neural Networks (CNNs) which generates catastrophic lag on typical desktop hardware, so in order to truly achieve real-time



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

performance, Sign Verse takes advantage of Google's MediaPipe Hands framework to completely skip pixel-level processing.

Rather than glancing at the entire image, MediaPipe is a super-advanced skeletal tracker. For each and every frame of video, the framework will localize 21 structural landmarks (joints and fingertips) per hand. Each individual landmark is captured with three metrics:

An 'x' coordinate, the horizontal distance of that point, normalized to the width of the video frame.

A 'y' coordinate, the vertical distance of that point, normalized to the height of the video frame.

A 'z' coordinate, the relative depth of that specific joint compared to the position of the wrist.

Since both American Sign Language (ASL) and Indian Sign Language (ISL) often require the use of both hands, the system tracks up to two hands, resulting in 42 total 3D landmarks. This data is then flattened into a one-dimensional array before it is fed to our AI model, providing a crisp, 126-coordinate feature vector for every frame of video.

C. Deep Learning Classification Model (TensorFlow/Keras)

After spatial feature extraction, the 126 data array is then input into a highly optimised, pre-trained deep neural network (sign_language_model.h5) built with the Keras and TensorFlow packages.

Since we are using coordinate data and not inputting pixels from an image, the model is based on a Multi-Layer Perceptron (MLP) architecture rather than a CNN. The architecture has an input layer designed specifically for the 126 data points, and a number of densely connected hidden layers, optimised with the Rectified Linear Unit (ReLU) activation function which is particularly well-suited to learning the complex, non-linear spatial relationships between moving joint coordinates. The output layer of the neural network uses a Softmax activation function for a clear, easy-to-read, probability distribution. Instead of returning a rough, difficult-to-read 'score' for a gesture (called a 'logit'), it 'normalises' the values into a probability distribution. By calculating the sum of the exponential values for all possible classes against the exponential value of one class, it gives us the proportional likelihood of our input data representing a particular class.

So a percentage score (from 0 to 100) will be applied to each and every ASL and ISL sign in the database. As the app runs, it keeps track of these numbers in real-time and when one of them passes a high-confidence threshold (generally 75) it locks in the prediction and sends the text output to the Gemini API, to put the signs into a

IV. CONTEXT-AWARE LANGUAGE GENERATION VIA

GEMINI AI

The core innovation of SignVerse is the seamless integration of Google Generative AI (Gemini API) to elevate the system from a simple character-recognition tool to an intelligent, context-aware communication engine.

A. Prompt Engineering and NLP Pipeline

Standard dictionaries map a sign to a static word, creating a one-to-one correspondence that doesn't account for natural conversational flow. Sign Verse takes the recognized sign as an initial seed, and forwards it to the Gemini API with the historical context of the current session.

The resulting suggestion set is computed by computing the highest conditional probability for upcoming words, given the context of the sentence built so far, and the new classified letter or word. In practice, this is achieved by prompting the Gemini model to return a probabilistic array of the highest-likelihood conversational branches.

Example Use Case (Standard SLR vs. Sign Verse): Input: User signs the letter "H".

Standard SLR Output: "H".

Sign Verse Output: We prompt Gemini with the seed "H" and no previous context. Gemini returns an array of initial greeting probabilities: ["Hi", "Hello", "How are you", "Help", "Hand"].

Dynamic Contextual Sensitivity: With Gemini's transformer architecture, self-attention is used to determine the weighting of previous inputs. If the user signs "I need" first and then signs "H", the AI will renegotiate the probability weights based on semantic relevance, filtering out greetings and instead offering ["Help", "Hospital", "Hearing"].



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

B. Latency and API Optimization

Restrictive generation parameters are used to optimize API calls so that LLM querying doesn't slow things down. The temperature is kept fairly low (between 0.2–0.3) to encourage common, highly deterministic phrases instead of creative writing. This speeds up token generation and makes suggestions more relevant for everyday communication.

V. IMPLEMENTATION DETAILS AND USER INTERFACE

The desktop app is built with Python and serves as the main link between the local machine-learning processes and the cloud-based NLP. Tkinter is used to make the Graphical User Interface (GUI) so that it works on a wide range of platforms and is light.

A. Asynchronous Architecture and Sentence Construction

Thread blocking is a big problem when trying to connect real-time computer vision with cloud APIs. If the application waits for the Gemini API to respond, the webcam feed will stop working. To fix this, Sign Verse uses a multi-threaded, asynchronous architecture:

Thread 1 (Vision Loop): Uses OpenCV to continuously capture frames, MediaPipe landmarks to extract them, and TensorFlow to run inference to update the current sign prediction state.

Thread 2 (GUI Mainloop): It draws the Tkinter interface, updates the video canvas, and keeps track of what users do.

Thread 3 (Async NLP Worker): Watches for changes in the predicted sign's state and makes asynchronous API calls to Gemini without stopping the vision or GUI threads.

There is a "Sentence Builder" text area on the interface that changes. The Tkinter GUI updates a suggestion row filled by the API as the CV model predicts signs in real time. Users can use these suggestions with just one click or keyboard shortcut.

User Dictionary Management and Heuristic Caching

To make Sign Verse work better and faster it has a feature called "Save Sentences" that uses a database called SQLite. This database is like a memory that helps the system remember things. When you start signing the system looks at the SQLite database to find phrases that're similar to what you are signing like "I am at the store". The system uses a way of comparing words to find these similar phrases. If it finds a phrase that matches well it does not need to ask the external system for help, which saves time and helps us not use up all the API requests we are allowed to make. This makes Sign Verse work quickly and it is better, for everyone who uses it.

VI. EXPERIMENTAL RESULTS AND PERFORMANCE EVALUATION

Real-Time Latency Pipeline

When we say about accessibility tools the most important thing that we consider is how long does the tool takes to finish from start to end. That is the end-to-end latency. To achieve this we use MediaPipe to extract the coordinates we need by analyzing every pixel in the Convolutional Neural Network. This way our vision pipeline is very fast.

We tested a computer processor similar such as the Intel Core i5 or the AMD Ryzen 5 and measured the time for each step. Here is what we discovered:

- 1] Frame Capture and Preprocessing: takes 3–5 milliseconds
- 2] MediaPipe Landmark Extraction: takes 8–10 milliseconds
- 3] TensorFlow MLP Inference: takes 2–5 milliseconds
- 4] Total Local Vision Pipeline: takes 13–20 milliseconds per frame

MediaPipe and the vision pipeline work very well. The MediaPipe and the vision pipeline can attend at least 30 frames per second. MediaPipe and the vision pipeline beat the 24 frames per second we need for video. The MediaPipe and the vision pipeline are doing a great job and we are glad with the results from MediaPipe and the vision pipeline.

Classification Accuracy and Robustness

The dual-language model was tested on a held out dataset of ASL and ISL spatial coordinates. During training, we used early stopping and connection dropout layers ($p = 0.3$) to help prevent overfitting to particular idiosyncrasies of the training set.



International Journal of Innovative Research in Computer and Communication Engineering (IJRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Our system achieved 94.2 validation accuracy in ideal lighting conditions. The model is able to effectively distinguish the spatial differences between ASL's largely single-handed alphabet and ISL's intricate dual-handed alphabets. The low-light and heavy motion blur situations show some degradation in accuracy, mostly due to MediaPipe dropping landmark confidence scores.

NLP Predictive Efficiency (Keystroke Saving Rate)

The incorporation of Gemini AI greatly enhances the Keystroke Saving Rate (KSR), the mainstream metric in augmentative and alternative communication (AAC) devices.

KSR is the percent reduction in physical inputs to produce a sentence. It is calculated by comparing the total number of individual characters needed to traditionally type a sentence with the reduced number of inputs (signs and clicks) needed to produce a sentence using Sign Verse. Empirical estimates suggest that with 5 contextual suggestions per recognized sign, Sign Verse reduces the number of physical gestures needed to produce a five word sentence by up to 70 relative to traditional letter-by-letter signing.

VII. PRIVACY AND SECURITY PROTOCOLS

Accessibility tools implemented in private, domestic settings need to adhere to strict security for user data. As Sign Verse runs continuously through webcam, it employs Edge-Cloud architecture based on privacy-by-design.

Edge-Centric "Zero-Video-Transmission" Policy

The key to Sign Verse's security is its zero-video-transmission design. All heavy computations and privacy-sensitive computer vision tasks (including raw video capture, biometric hand tracking, spatial landmark mapping, and gesture classification) are performed at the edge (on the user's local device). The system does not have any network ability to transmit raw image tensors, video frames, or audio to the cloud. This ensures that no information about the user's physical appearance, identity, or surrounding background environment can get out to any external servers.

Data Minimization and Obfuscation

To make use of powerful, cloud-based LLMs without exposing private data, Sign Verse uses extreme data minimization. The only data payload allowed to be transmitted from the local machine is obfuscated, serialized text. As an example, if a user signs "Hospital" in a video, the local machine classifies that video and only transmits the text string "Hospital" and the text context of the current session to the Gemini API.

Secure Data in Transit (Encryption Protocols)

The serialized text payload is transmitted to the Google Gemini servers via secure, industry-standard cryptographic protocols. All API requests and responses are routed over HTTPS utilizing Transport Layer Secure (TLS 1.3). This prevents man-in-the-middle (MitM) attacks and ensures that the conversational context being processed by the AI cannot be intercepted by malicious third parties on the network.

Compliance with Privacy Frameworks

By separating the biometric visual processing from the linguistic processing, Sign Verse can closely comply with modern data protection regulations such as the General Data Protection Regulation (GDPR) and Health Insurance Portability and Accountability Act (HIPAA) adjacent data protection frameworks. No Personally Identifiable Visual Information (PII) is ever persisted or processed on third-party data centers, making the application a secure, trustless tool for the deaf and hard-of-hearing population.



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Fig.3. The application's 'About' dialog detailing its core architecture, confirming the use of MediaPipe tracking and fully offline local language models.

A. Concluding Remarks

By fundamentally redefining Sign Language Recognition (SLR) as a dynamic, context-sensing sentence-building process instead of a static, labor-intensive translation exercise, Sign Verse successfully bridges the communication divide for the deaf and hard-of-hearing (DHH) community.

By showcasing a successful blend of localized 3D spatial feature extraction (using MediaPipe and TensorFlow) and cloud-based Natural Language Processing (using the Google Gemini API), we demonstrate that assistive technologies can strike a fine balance between computational lightness for standard consumer hardware, intelligence in predictions, and privacy preservation through a zero-video-transmission approach. In addition, our system's dual-language processing of both American and Indian Sign Languages (ASL and ISL) provides a strong algorithmic foundation that can be scaled globally to fit linguistic demographics.

B. Future Work/; Integration of Non-Manual Markers (NMMS)

While the current version of Sign Verse provides accurate tracking of manually created gestures (hand shapes/positions), natural sign language relies heavily on Non-Manual Markers (NMMs) to encode the grammar and emotional valence of a sentence.

Future versions of Sign Verse will work to mitigate hand-only tracking limitations by integrating real-time facial expression and posture recognition. By integrating the MediaPipe Face Mesh framework to capture a dense 468-point 3D facial geometry, and building upon methodologies taken from the Facial Action Coding System (FACS), the underlying AI will be able to capture subtle but grammatically crucial nuances of facial geometry including:

Eyebrow Positions: Standard statements, "Wh-" Questions (frowned eyebrows), and "Yes/No" Questions (raised eyebrows).

Head Posture: Tracking head tilts and shakes to indicate negations and affirmations, respectively.

Lip Patterns: Recognizing mouth morphemes that are used as adverbs/adjectives in ASL and ISL.

By fusing this facial coordinate array with the manual coordinate array, we will arrive at a far richer multi-modal feature vector, which will drive much more accurate and expressive LLM sentence predictions.

Future Work: Edge-Based Generative AI (Offline Processing)

Currently Sign Verse is using a hybrid Edge-Cloud architecture, where language generation requires an internet connection to query the Gemini API. This is a very secure approach, but it does depend on external networking conditions and it is difficult to use for DHH people in remote or offline environments.

An important part of future development is to make the entire cloud-based NLP component run on the edge. We will implement a local Small Language Model (SLM) directly in the desktop application. Through the use of state of the art model compression techniques, specifically INT4 (4-bit) quantization, we can run extremely capable open-weight models, such as customized variants of Meta's LLaMA or Google's Gemma, on the edge directly on consumer CPUs



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

or GPUs.

This architectural change will have significant benefits:

True Zero-Latency: No network request transit times for instant word prediction.

100 Offline Usability: We ensure that the DHH community can use the software anywhere without needing to connect to the internet.

Ultimate Privacy: We ensure that not only images, but also the resulting textual conversational context never leaves the DHH person's personal machine.

Future Work: Expansion of Sign Language Corpora and Federated Learning

In future iterations of the platform, we will scale the deep learning classification model to handle a broader range of global sign languages (e.g., British Sign Language [BSL] and Auslan). To do so, we must create larger, more demographics-diverse datasets of 3D spatial coordinates to test and train for robustness across variations in signing speed, regional dialects, and anatomical differences.

To scale this at the required level ethically and securely, we propose to use a privacy-preserving, opt-in crowdsourcing mechanism driven by Federated Learning. Instead of uploading raw data from the user to a central server, the system will locally train lightweight model updates based on a user's signing style. These mathematical weight updates will be the only data transmitted to the cloud to improve the global model. This will enable us to grow a corpus of natural, real-life variation while maintaining the edge-centric privacy protocols of this architecture.

Final Conclusion

In the grand scheme of things, the vision of Sign Verse goes way beyond regional translation. The platform aims to create a perpetually-expanding, multi-modal, and multi-lingual architecture that becomes the de facto standard for real-time digital accessibility. As the collaboration between edge computing (for real-time spatial feature extraction) and Generative AI (for contextual language modeling) continues to refine itself, Sign Verse and similar tools will become instrumental in breaking down systemic communication barriers. By deflecting the technological burden from the user to the algorithm, Sign Verse guarantees that the deaf and hard-of-hearing community has the opportunity to participate in this global conversation with maximum fluidity and parity.

REFERENCES

1. World Health Organization, "Deafness and hearing loss," Fact Sheet, Mar. 2023. [Online].
2. C. Lugaresi et al., "MediaPipe: A Framework for Building Perception Pipelines," arXiv preprint arXiv:1906.08172, 2019.
3. M. Abadi et al., "TensorFlow: A System for Large-Scale Machine Learning," in 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), Savannah, GA, USA, 2016, pp. 265–283.
4. G. Bradski, "The OpenCV Library," Dr. Dobbs's Journal of Software Tools, vol. 25, no. 11, pp. 120–123, 2000.
5. Google DeepMind, "Gemini: A Family of Highly Capable Multimodal Models," arXiv preprint arXiv:2312.11805, 2023.
6. A. Vaswani et al., "Attention Is All You Need," in Advances in Neural Information Processing Systems, vol. 30, Curran Associates, Inc., 2017.
7. N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A Simple Way to Prevent Neural Networks from Overfitting," The Journal of Machine Learning Research, vol. 15, no. 1, pp. 1929–1958, 2014.
8. V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," Soviet Physics Doklady, vol. 10, no. 8, pp. 707–710, 1966.
9. H. R. Vaezi Joze and O. Koller, "MS-ASL: A Large-Scale Data Set and Benchmark for Understanding American Sign Language," The British Machine Vision Conference (BMVC), 2019.
10. A. Sridhar, R. Ganesan, P. Kumar, and M. Khapra, "INCLUDE: A Large Scale Dataset for Indian Sign Language Recognition," in Proceedings of the 28th ACM International Conference on Multimedia, 2020, pp. 1366–1375.
11. I. S. MacKenzie and R. W. Soukoreff, "Text Entry for Mobile Computing: Models and Methods, Theory and Practice," Human-Computer Interaction, vol. 17, no. 2, pp. 147–198, 2002.
12. P. Ekman and W. V. Friesen, "Facial Action Coding System: A Technique for the Measurement of Facial Movement," Consulting Psychologists Press, Palo Alto, 1978.



International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

14. W. Z. Khan, E. Ahmed, S. Hakak, I. Yaqoob, and A. Ahmed, "Edge computing: A survey," Future Generation Computer Systems, vol. 97, pp. 219–235, 2019.
15. pp. 219–235, 2019.
16. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentral- ized Data," in Artificial Intelligence and Statistics, PMLR, 2017, pp. 1273–1282.
17. Khushi N. Kumawat, "Sign Verse: Real-Time Bilingual Sign Language Translation and Context-Aware Sentence Construction," Project Report
18. / GitHub Repository, 2026.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



SJIF Scientific Journal Impact Factor



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

 9940 572 462  6381 907 438  ijircce@gmail.com



www.ijircce.com

Scan to save the contact details