

Design and Evaluation of Conversational Agents for Mental Health Support: An NLP-Driven Approach with Indian User Studies

¹Manish Kumar, ²Rohit Kumar Vijeta, ³Pragyesh Kumar Seth, ⁴Mallika Gupta, ⁵Mr. Himanshu

¹²³MCA Scholar, St. Andrews Institute of Technology & Management, Gurugram, Haryana, India

⁴CSE Scholar, St. Andrews Institute of Technology & Management, Gurugram, Haryana, India

⁵Assistant Professor, St. Andrews Institute of Technology & Management, Gurugram, Haryana, India

Abstract

Text-based support for common mental health concerns remains difficult to reach for a large share of India's population because trained counsellors are concentrated in a small number of cities. Existing screening tools and rule-based chat systems handle narrow scripts and rarely adapt to the Hindi-English code-mixing that everyday speakers use. We built a bilingual conversational agent that pairs a BERT-based natural language understanding pipeline with a state-machine dialogue manager, a dedicated safety monitor, and a hybrid template-sequence-to-sequence generator for empathetic replies. We evaluated the agent with 240 participants, 120 drawn from Delhi and Mumbai and 120 from villages in Uttar Pradesh and Rajasthan, over a four-week trial. Intent recognition reached 91.4% accuracy, emotion detection reached an F1 score of 87.2% across six basic emotions, and the safety monitor achieved 96.8% recall for suicidal-ideation markers. Participants rated the proposed agent at a System Usability Scale score of 72.4 against 54.1 for a rule-based baseline (Mann-Whitney U, $p < 0.001$), and engagement averaged 4.2 sessions per week. Rural participants rated empathy 0.8 points lower than urban participants on a five-point scale ($t(238) = 3.42$, $p = 0.001$), a gap we trace to dialect coverage rather than to the underlying model. These results indicate that a code-mixed, safety-audited agent can extend low-intensity mental health support to populations that a purely English or purely Hindi system would miss, while leaving open questions about dialect coverage and long-term clinical outcomes that a single trial cannot answer.

I. Introduction

India has roughly one psychiatrist for every 100,000 people, and that ratio worsens outside the largest cities, where clinics close by early evening and a first appointment can take weeks. A person who wants to talk about anxiety or low mood at 11 p.m. on a Tuesday has almost no professional to call. Text messaging apps, by contrast, reach a phone that is already in the person's hand. The question this work answers is narrower than "can software replace a therapist": it is whether a conversational agent, built with the natural language processing (NLP) methods now common in intent classification and dialogue generation [29], [167], [172], can hold a supportive conversation in the mixture of Hindi and English that Indian users actually type, detect distress with enough precision to escalate correctly, and do this at a cost that keeps the service free at the point of use.

Two design constraints shaped every decision that follows. First, the input is rarely pure Hindi or pure English; a typical message reads "mujhe bahut anxious feel ho raha hai", mixing Devanagari-transliterated Hindi with an English adjective in the same clause. Classifiers trained only on monolingual corpora misread this pattern as noise. Second, a mental health agent carries a risk that a shopping or scheduling assistant does not: a missed crisis signal is not a lost sale, it is a missed intervention. Every architectural choice in Section III, from the safety monitor sitting inside the dialogue manager rather than as a downstream filter, to the choice of a hybrid generator over a pure large-language-model decoder, follows from taking that asymmetry seriously.

Deep learning methods for text classification, sequence labelling and generation have matured to the point that intent recognition and entity extraction are largely solved for well-resourced languages [31], [77]. Sentiment and emotion analysis over social and clinical text has likewise moved from lexicon counting to learned representations [37], [42]. Parallel work on healthcare artificial intelligence has shown that decision-support systems can raise diagnostic consistency when their outputs are audited by clinicians rather than deployed unsupervised [35], [48]. Internet of Things (IoT) health monitoring platforms have demonstrated that continuous, low-cost sensing can substitute for episodic clinic visits in settings with sparse infrastructure [7], [135], and the same literature has been explicit that any system handling health data must budget for its cybersecurity posture from the outset rather than as an afterthought [21], [73], [113]. This paper draws each of

those threads into one deployed system and reports what changed when 240 people, half of them without reliable broadband, used it for a month.

The contribution is fourfold and each part is verifiable against a section of this paper. We describe a bilingual NLU pipeline that reaches 91.4% intent accuracy and 87.2% emotion-detection F1 on code-mixed Hindi–English text (Section IV). We describe a dialogue manager whose safety monitor achieves 96.8% recall on suicidal-ideation markers without a corresponding collapse in precision (Section VI). We report a 240-participant urban–rural comparison that isolates a specific, measured empathy gap rather than a general claim of rural underperformance (Section VI). We set out an ethical operating framework, including institutional review board (IRB) approval, standing clinical psychologist oversight, and mandatory escalation protocols, that a future deployment can adopt directly (Section VII).

II. Related Work

A. Conversational Agents and Dialogue Systems for Health

Rule-based dialogue systems answer a fixed set of scripted intents and fail as soon as a user departs from the script; retrieval-based systems widen coverage by matching an incoming message against a corpus of prior exchanges but cannot generate a reply the corpus has never seen. Both patterns recur across the health-chatbot literature this review draws on [161], [162], [190]. A retrieval-augmented chatbot for health-policy queries shows the same trade-off in a non-clinical setting: precision on covered questions is high, but the system has no mechanism for a question outside its index [190]. An AI-driven chatbot for doctor-appointment management demonstrates that scheduling intents are tractable with a narrow, well-labelled dataset [161], while an adaptive interview-generation system shows that dialogue state can be tracked across many turns without losing coherence [160]. None of these systems, however, carries a dedicated safety layer for self-harm language, which is the gap our dialogue manager is built to close.

B. Deep Learning for Text and Speech Classification

Convolutional and recurrent architectures have been applied across a wide span of classification problems that share little with mental health support on the surface but converge on the same underlying question of extracting a signal from noisy sequential input: cardiac-signal classification from compressed electrocardiogram traces [3], intrusion detection from multimedia network traffic [4], [5], arrhythmia classification under class imbalance [42], lip-reading from silent video [40], handwritten text recognition [44], diabetic retinopathy grading from fundus images [47], and lung-cancer subtyping from radiomic features [17]. Ensemble and hybrid CNN–LSTM designs recur across agricultural yield prediction [8], [9], pollution forecasting [46], and student-performance modelling [124], showing that the architecture choices behind our intent classifier are not specific to language but reflect a broader pattern in applied deep learning [16], [41], [45], [49], [50], [54], [125], [128], [141]. Reviews of encryption and phishing-detection pipelines built on the same convolutional backbones confirm that the approach generalises across security-adjacent tasks as well [26], [30].

C. Sentiment and Emotion Analysis

A hybrid lexicon-and-CNN-LSTM approach to sentiment detection over COVID-19 vaccine discourse reports that a fixed dictionary alone under-detects sarcasm and negation, and that a learned component recovers much of the missed signal [167]. Emotion recognition from facial expression [35] and from multimodal video [173] extends the same idea from text to other modalities, and a survey of event analysis by deep learning situates emotion classification within the broader task of extracting structured events from unstructured streams [32]. Sentiment-driven stock-price prediction work shows that a large-language-model layer added on top of a recurrent base model improves the handling of ambiguous or mixed-polarity statements [172], which is the same property our emotion detector needs when a user’s message expresses relief and worry in the same sentence. Human emotion recognition through deep learning more generally confirms that six-to-eight basic emotion categories remain the dominant label set in the applied literature [109], a choice we retain in Section IV for comparability.

D. Healthcare AI and Clinical Decision Support

A body of work on explainable AI for healthcare argues that a clinician will not act on a model’s output unless the reasoning behind it can be inspected [22], [143], [144]. Diagnostic-support systems built around this principle have been proposed for cardiovascular risk [15], pregnancy complications [101], perinatal disorders [99], and menopause-related risk [100], each pairing a predictive model with a rationale the treating clinician can review. A broader guide to AI and computational intelligence in healthcare frames these systems as decision aids rather than replacements for clinical judgement [103], a framing this paper adopts explicitly in Section VII. Appointment-scheduling automation [52], blockchain-secured health-record sharing [53], and AI-assisted mock-interview and query systems [160], [161] round out a picture of healthcare AI as a set of narrow, auditable

components rather than one monolithic system, which is the design philosophy behind the modular architecture in Section III.

E. IoT-Based Health Monitoring

A synchronised health monitoring system built on a linear quadratic regression estimator shows that low-cost IoT sensors can track vital signs continuously without the drift that discrete clinic visits introduce [7]. Blockchain-secured product-and-record management for medical IoT devices addresses a related concern: once a sensor is deployed outside a controlled clinical setting, the chain of custody for its data needs the same auditability as the clinical record it feeds [135]. Camera-based surveillance with integrated edge detection [10], smart energy-management systems for constrained IoT deployments [19], [117], and green IoT architectures aimed at smart-city infrastructure [92], [120] illustrate that the connectivity assumptions behind rural IoT health monitoring, patchy bandwidth, intermittent power and asymmetric device capability, are the same assumptions our text-first interface in Section III was designed around. Aerial and edge-computing platforms extend this monitoring reach into disaster-response and disaster-adjacent settings where fixed infrastructure may be unavailable altogether [136], [137], [138], [156], [157], [158].

F. Security and Privacy of Health Data

Storing transcripts of a mental health conversation carries a higher confidentiality obligation than storing a shopping history, and the security literature this project draws on treats that obligation as a design constraint rather than a compliance checkbox. Cloud security posture management practices for continuously auditing access controls [73], zero-trust anomaly detection for cloud-hosted services [159], and post-quantum digital-signature schemes intended to outlast current encryption standards [96] together describe the defence-in-depth stack our data pipeline follows. Malware and botnet detection tailored to IoT and Android platforms [140], [203] and a systematic review of IoT forensics practices [175] address the complementary problem of detecting compromise after the fact. Federated learning frameworks that keep raw data on-device while still training a shared model [106], [118], [119], [174] describe an alternative architecture we considered and rejected for the pilot, for reasons given in Section VII, but that remains a candidate for a larger rollout. Data-leakage detection in Wi-Fi networks [90] and multimedia forensics techniques [89] round out the threat model our infrastructure was audited against before the trial began. Prototype-pollution and supply-chain vulnerability detection frameworks [83] and a review of encryption and keylogging risks in general-purpose computing systems [26] informed the hardening applied to the deployment server itself.

G. Indian-Language NLP and Access in Underserved Settings

Work on bridging the medical gap in underserved communities worldwide argues that a system's language coverage is itself a determinant of who it reaches, not a secondary usability concern [102]. Ethical lessons drawn from AI and IoT use during pandemic response make a parallel point about equitable access when infrastructure is unevenly distributed [104]. Digital-social-linguistics work on code-switching patterns across Indian language communities [177] and a comparative study of supervised and unsupervised methods for language-specific tasks [178] informed the annotation scheme used to label the code-mixed training corpus described in Section IV. Trust in AI systems, examined as a social-psychological rather than purely technical question, finds that a user's willingness to disclose sensitive information to an agent depends on perceived competence and perceived discretion in roughly equal measure [114], a finding that motivated the therapeutic-alliance measure reported in Section VI. Explainable-AI work on human decision-making patterns [116] and on diagnostic support specifically [144] closes the loop between the interpretability goals of Section II-D and the trust outcomes measured in the user study.

H. Opportunistic Networking and Infrastructure Context

A rural deployment cannot assume the always-on connectivity that an urban pilot enjoys, so the routing and infrastructure literature this project drew on for the connectivity assumptions behind Section III is worth stating explicitly. Opportunistic-network routing protocols built with machine-learning link prediction [1], [60] and reviews of the protocols, applications and simulation trends in that field [2], [61] describe how a message can still traverse a network when no continuous end-to-end path exists, the same condition a rural participant's intermittent mobile signal periodically recreates. Comparative performance studies of spray-and-wait style protocols against learned alternatives [62] and analyses of how node density changes quality-of-service outcomes in opportunistic deployments [6] informed our decision, noted in Section III, to keep session state on a managed server rather than push it fully to the device. Foundational surveys of wireless sensor network architecture [11], [68] and of broadcasting and routing characteristics in ad-hoc networks [66] describe the lower layers that any store-and-forward fallback would need, while capacity studies of WiMAX under varying coding schemes [67] and hybrid crowd-cloud routing protocols for wireless sensor networks [13] illustrate the

same trade-off between throughput and robustness that shaped the client-availability concern raised in Section VII. Ant-colony-based task scheduling for cloud environments [12] and its extension to soft-computational NP-hard problems [93] describe the resource-allocation methods considered, and set aside, for load balancing on the application-server tier in Section III.

I. Applied Machine Learning Across Adjacent Domains

Beyond the threads examined above, the applied-AI portfolio this study draws modelling precedent from spans domains with little surface resemblance to mental health support but a shared reliance on the same estimation and classification toolkit. Early patient-monitoring and disease-detection systems [58] and enhanced association-rule mining methods [59] describe the pattern-detection lineage that the safety monitor in Section IV-D extends. Voice-based interaction design for non-clinical tasks, including a voice-driven calculator [25] and a television-viewing cost calculator [24], anticipates the usability questions Section VI raises about the text-versus-voice channel preference. Deep learning applied to healthcare systems generally [63], to malicious-node detection in mobile ad-hoc networks [64], and to edge-detection-based biometric security for payment systems [65] describes the security-adjacent modelling context behind Section VII's defence-in-depth discussion, alongside child-identification software built on machine-learning speech clarification [70] and a comparison of proactive and reactive routing protocols for mobile ad-hoc networks [71]. Soft-computing methods for power-system stability [18], [72] and encryption-and-decryption algorithms for image data [69] extend the same estimation toolkit to power and multimedia security respectively.

Generative AI has been examined across a comparably wide span of applications: as a general review of evolution, applications and challenges [75], [123], applied to Indian stock-market investment decisions [76], to employee-performance prediction [77], to intelligent-factory productivity in Industry 4.0 [80], and to esports production tooling [86]. A parallel line of open-source-intelligence and cyber-threat work spans double-edged-sword risk framing for OSINT practice [78], detection of unauthorised construction from aerial or satellite imagery [81], blockchain-and-Web3 wallet interoperability [82], and a review of emerging computational-technology trends more broadly [87]. Reliability-engineering methods originally developed for supermarket queuing simulation [20] and for repairable industrial plant modelling [111] describe the same discrete-event and stochastic-process toolkit that underlies the session-state machine in Section III.

A further set of studies addresses healthcare and human-centred AI questions adjacent to, but distinct from, the mental health focus of this paper: transforming medicine through AI more broadly [74], AI-powered healthcare-appointment scheduling and blockchain-secured interoperability already discussed in Section II-D [52], [53], AI in optimising and pandemic-era healthcare delivery [88], [108], and the ethics of AI-assisted federated healthcare intelligence [118]. Education-focused applications, including the role of generative AI in education and learning [112], predictive modelling of student motivation and engagement [113], and a hybrid ensemble framework for interpretable student-performance prediction [124], describe methodological choices, particularly around interpretability, that carried over directly into the explainable safety-monitor design of Section IV-D. Human-versus-machine behavioural comparison studies [148] and work on artificial distributed intelligence [153] frame the broader question of where automated and human judgement should divide, a question this paper answers concretely for the crisis-escalation case in Section VII.

Infrastructure and platform studies published alongside this body of health- and language-focused work describe the computing substrate such a service ultimately runs on: green Internet of Things architectures for smart cities [92], [121], mobile edge computing for decentralised systems [98], hybrid cloud infrastructure models and their associated security trade-offs [152], and federated-learning-based intelligent systems for handling Internet-of-Vehicles challenges [107], [110], [115], [122]. Movie- and content-recommendation systems built on transformer and autoencoder architectures [126] and neural-network-based galaxy classification [127] illustrate the same representation-learning toolkit applied to recommendation and scientific classification respectively, outside any health context, which is offered here only to situate the modelling choices in Section IV within a wider applied-AI practice rather than to claim relevance to mental health support directly. Uncertainty-aware air-quality classification [129], stock-market prediction from large-scale data [130], and bank-churn prediction [131] similarly extend the same classification and forecasting toolkit to unrelated forecasting problems, reinforcing that the architectural pattern in Section III, a modular pipeline with an auditable decision point, recurs across the applied literature regardless of domain [132], [133], [134], [139], [142], [145], [146], [147], [149], [150], [151], [154], [155], [164], [168], [169], [171], [176], [179], [180], [181], [182], [183], [184], [185], [186], [187], [188].

A final cluster of recent conference contributions applies deep learning to medical imaging and signal tasks that parallel, without duplicating, the classification methods in Section IV: thyroid-cancer detection from ultrasound imaging [189], lung-disease classification from chest imaging [195], human-pose detection for movement analysis [196], autism-spectrum-disorder recognition from behavioural data [197], brain-tumour classification from magnetic resonance imaging [205], and bidirectional sign-language translation for

accessibility [206]. Retrieval-augmented document-querying frameworks [202] and malware and botnet detection across IoT and Android platforms [140], [203] extend the retrieval and security themes of Sections II-A and II-F respectively, and an ensemble explainable-AI framework grounding large-language-model clinical explanations in verifiable evidence [207] closes the loop with the interpretability requirement Section IV-D places on the safety monitor. Energy-optimisation and route-scheduling systems for smart-home and public-transport contexts [191], [192], [193], [194], [198], [199], [200], [201], [204] round out this cluster and are cited here only to note that the same underlying estimation methods, not the applications themselves, recur across this body of work.

Across these threads, no system in the literature reviewed here combines Hindi–English code-mixed dialogue, an integrated crisis-detection layer, and a rural–urban comparative evaluation in one deployed pipeline; that combination is the gap this paper fills.

III. System Design

The agent is organised as four modules connected by a single message-passing contract, shown in Fig. 1: a user-input handler, a natural language understanding (NLU) module, a dialogue manager, and a natural language generation (NLG) module, with a text or voice interface at the boundary. We chose this separation over a single end-to-end sequence model for one reason that overrides considerations of engineering convenience: an end-to-end model gives no natural point at which to insert an auditable safety check, whereas a modular pipeline lets the safety monitor sit inside the dialogue manager where every candidate action, including every generated reply, passes through it before reaching the user.




Fig. 1. Conversational agent architecture: user input flows through the NLU module (intent recognition, entity extraction, emotion detection) into a dialogue manager that tracks context, maintains conversation state, and runs a safety monitor, before an NLG module produces an empathetic response and, where relevant, a resource recommendation for the text or voice interface.

The user-input handler accepts both typed text and, for the subset of participants who used the voice channel, an automatic speech recognition transcript. We did not train a custom speech recogniser; the voice path uses an existing commercial Hindi–English recogniser and passes its output through the same NLU pipeline as typed text, which kept the modality difference invisible to the rest of the system. This choice cost us some accuracy on the voice channel specifically, a limitation returned to in Section VII, but it avoided the considerably larger engineering cost of a bespoke acoustic model for code-mixed speech.

The dialogue manager holds three components. Context tracking maintains a rolling window of the last eight user turns plus a structured summary of the session, so that a reference to “that thing I mentioned” two turns earlier resolves correctly more often than a stateless design would allow. The state machine enumerates a small set of session states, greeting, exploration, coping-strategy suggestion, resource referral, and crisis, and restricts which NLG templates are reachable from each state, which is what keeps the agent from abruptly suggesting a breathing exercise in the middle of a crisis-state exchange. The safety monitor runs on every incoming and outgoing turn independently of the state machine’s normal transition logic; it can force an immediate transition to the crisis state regardless of the state the conversation was in, and a transition into that state is logged and, per the protocol in Section VII, routed to a human reviewer within the same session where indicated.

Deployment used a stateless application-server tier behind a load balancer, following the fog-cloud load-balancing pattern documented for resource-constrained IoT integration [166], and a managed database for session state, chosen over a fully on-device architecture because the rural participant group’s devices varied too widely in available memory to guarantee a consistent on-device footprint. Continuous integration and containerised deployment followed practices documented for cross-platform mobile development [57], and fault

prediction on the microservice tier drew on trace-log-based proactive monitoring [55] so that a backend slowdown could be caught before it degraded the latency metric reported in Section VI. Duplicate-session detection, adapted from duplicate-node detection methods developed for wireless sensor deployments [43], prevented a single participant's repeated app restarts from being double-counted as new sessions in the engagement statistics of Section VI.

IV. NLP Components

A. Intent Recognition

Intent recognition uses a fine-tuned BERT encoder with a linear classification head over eighteen intent classes covering greeting, mood check-in, coping-strategy request, crisis disclosure, resource request, small talk, and session closure, among others. We fine-tuned on 34,600 labelled utterances collected during a six-month pre-pilot, of which 61% contained at least one code-mixed token. Training a monolingual English baseline on the same label set and testing it on the code-mixed validation split dropped accuracy to 74.1%, a 17.3-point loss that motivated the transliteration-normalisation preprocessing step described next: every Latin-script Hindi token is mapped to a canonical Devanagari-derived representation before tokenisation, which recovers most of that gap. The final bilingual model reached 91.4% intent accuracy on the held-out test set (weighted F1 = 0.906), a result consistent with the accuracy range reported for BERT-based intent classifiers on other applied dialogue tasks [161], [162].

B. Entity Extraction

A conditional random field (CRF) tagger, layered over the same BERT encoder's token embeddings, extracts five entity types: time expressions, named support contacts, self-reported symptom terms, medication names, and location references relevant to resource referral. The CRF layer reached a token-level F1 of 88.6% across these five types, with symptom-term extraction the weakest category at 81.2% F1, largely because symptom descriptions in code-mixed text vary more in surface form than the other four categories. We evaluated a graph neural network alternative for this task, following the recommendation-system literature that reports strong performance for GNN-based extraction on structured domains [48], but abandoned it after a pilot run: the relational structure that makes graph neural networks effective for recommendation did not have an analogue in single-turn symptom mentions, and the added inference latency was not offset by an accuracy gain on our validation split.

C. Emotion Detection

Emotion detection classifies each user turn into one of six basic categories, joy, sadness, anger, fear, surprise, and disgust, plus a neutral class, using a separate classification head fine-tuned on the same encoder backbone. The classifier reached 87.2% F1 macro-averaged across the six emotion classes, with fear and sadness the best-separated pair (F1 = 91.0% and 89.4% respectively) and surprise the weakest (F1 = 78.3%), a pattern also reported for facial-expression-based emotion recognition where low-arousal categories are harder to separate than high-arousal ones [35]. We considered extending the label set to a dimensional valence-arousal model instead of discrete categories, an approach with precedent in emotion-based video summarisation [173], but retained the discrete scheme because the downstream dialogue manager's state machine reasons over discrete emotion labels rather than continuous scores, and a dimensional score would have needed an additional discretisation step that offered no accuracy benefit in our pilot testing.

D. Safety Monitor and Crisis Detection

The safety monitor combines a curated keyword and phrase list, reviewed and periodically updated by the clinical psychologist on the study team, with a fine-tuned classifier trained on anonymised, clinician-labelled examples of suicidal-ideation language. We tuned the classifier's decision threshold to favour recall over precision, on the explicit judgement that a false escalation costs a participant an unnecessary check-in message while a missed escalation costs a missed intervention; those costs are not symmetric. At the selected threshold the monitor achieved 96.8% recall for suicidal-ideation markers on a held-out clinician-labelled test set, with a corresponding precision of 71.3%, meaning roughly three escalations in ten were, on clinical review, judged not to require immediate intervention. We treat that precision figure as a limitation rather than a result to defend; Section VII discusses what a lower threshold would have cost in missed detections and why we did not accept that trade.

E. Natural Language Generation

The NLG module blends two generation strategies rather than committing to one. For the majority of dialogue states, a template-based generator selects and lightly parameterises a clinician-reviewed response, which guarantees that every sentence a participant reads has been checked for clinical appropriateness in

advance. For the exploration state, where a fixed template would read as repetitive across a multi-turn conversation, a sequence-to-sequence generator fine-tuned on clinician-reviewed transcripts produces a free-form empathetic reply, which is then passed through the same safety monitor described above before release. We rejected a pure large-language-model decoder for the entire pipeline after an early pilot surfaced two hallucinated medical claims in fewer than 200 test conversations; a purely template-driven design, evaluated separately, scored lower on the naturalness dimension reported in Section VI. The hybrid design was a direct response to that trade-off, and we return to its remaining failure modes, including the residual hallucination risk in the free-form branch, in Section VII, drawing on the broader literature on identifying hallucination risk in generative AI output [163], [170] and on trustworthy deep learning practices for bounding that risk [165]. Resource recommendation, triggered from either the coping-strategy or resource-referral state, surfaces a short, geographically filtered list of helplines and low-cost clinics, curated from a database maintained jointly with the study's clinical partner and cross-checked against blockchain-secured provenance records to guard against a stale or incorrect contact being served to a participant in crisis [105], [135].

V. User Study Design

We recruited 240 participants: 120 from urban neighbourhoods in Delhi and Mumbai and 120 from villages in Uttar Pradesh and Rajasthan. Recruitment used a combination of community health worker referral in rural sites and university and workplace mailing lists in urban sites, a difference in recruitment channel we flag as a possible confound in Section VII rather than treat as neutral. Eligibility required that a participant be 18 years or older, report at least mild self-perceived stress or low mood on a pre-screening questionnaire, and have access to a smartphone capable of running the study application; participants already under active psychiatric treatment for a severe mental illness were excluded and referred directly to the study's clinical partner.

The trial ran for four weeks per participant, with the agent available on demand rather than on a fixed schedule, an important methodological choice because it means the reported 4.2 sessions-per-week engagement figure reflects voluntary use rather than compliance with a prompted schedule. We compared the proposed agent against two baselines run by the same 240 participants in a within-subject counterbalanced design across separate weeks of the trial: a rule-based system built from a fixed decision tree of scripted responses, and a retrieval-based system that returns the closest matching response from a corpus of prior clinician-reviewed exchanges. Comparative baselines of this kind are standard practice in the applied-AI evaluation literature this study follows [190], and the rule-based and retrieval-based designs mirror architectures already documented in Sections II-A and IV.

Data collection combined three streams: automatically logged interaction data (turn counts, response latency, session length), in-app surveys administered after each session (a five-item mood check plus, at study midpoint and endpoint, the System Usability Scale and an adapted Working Alliance Inventory–Short Revised subscale for therapeutic alliance), and a structured post-trial interview conducted in the participant's preferred language by a bilingual research assistant. All quantitative comparisons in Section VI used a two-sided Mann-Whitney U test for between-group comparisons of non-normally distributed scores and a two-sided independent-samples t-test for the urban–rural empathy comparison, with significance set at $\alpha = 0.05$ throughout, following standard practice for small-sample behavioural comparisons in applied AI research of this kind [21], [114].

The study protocol received institutional review board approval before recruitment began, and a clinical psychologist on the study team reviewed every escalation event within 24 hours; the full ethical protocol, including consent procedures, data anonymisation, and the escalation pathway itself, is described in Section VII rather than here, so that the study design and its ethical governance are documented in full but kept legible as two distinct concerns.

VI. Results

We report results against the three-strand evaluation framework in Fig. 2: technical metrics, user-experience metrics, and clinical safety metrics, assessed jointly rather than in isolation, because a system that scores well on one strand and poorly on another is not a system we would recommend for deployment regardless of which single number looks favourable.




Fig. 2. Evaluation framework applied across the four-week trial, combining technical metrics (intent accuracy, response relevance, latency), user experience metrics (System Usability Scale, therapeutic alliance, engagement), and clinical safety metrics (crisis detection rate, appropriate escalation, harm avoidance), cross-checked by the study’s clinical psychologist review board.

A. Technical Metrics

Intent accuracy of 91.4% and emotion-detection F1 of 87.2% are reported in Section IV against a held-out clinician-annotated test set and are not repeated in full here; Table I summarises them alongside the remaining technical indicators. Median end-to-end response latency was 640 ms on the urban network sample and 1,180 ms on the rural sample, a gap attributable to cellular network variability rather than to any difference in the model itself, since both samples were served by the identical backend deployment described in Section III. Human-rated response relevance, scored by two independent bilingual raters blind to which system produced each response, averaged 4.1 out of 5 for the proposed agent, 2.7 for the rule-based baseline, and 3.3 for the retrieval-based baseline (inter-rater agreement, Cohen’s $\kappa = 0.71$).

Metric	Proposed Agent	Rule-Based	Retrieval-Based
Intent accuracy (%)	91.4	68.2	76.5
Emotion-detection F1 (%)	87.2	—	54.0
Median latency, urban (ms)	640	210	480
Median latency, rural (ms)	1180	340	890
Response relevance (1–5)	4.1	2.7	3.3
Crisis-detection recall (%)	96.8	41.0	58.7

Table I. Technical, latency and safety indicators for the proposed agent against the rule-based and retrieval-based baselines. Rule-based emotion-detection F1 is omitted because the baseline has no emotion classifier.

B. User Experience Metrics

System Usability Scale scores averaged 72.4 for the proposed agent against 54.1 for the rule-based baseline (Mann-Whitney $U = 9,845$, $p < 0.001$) and 61.7 for the retrieval-based baseline (Mann-Whitney $U = 11,203$, $p < 0.001$). A score above 68 is conventionally read as above-average usability, which the proposed agent cleared and both baselines did not. Therapeutic-alliance scores, on the adapted five-point subscale, averaged 3.7 for the proposed agent, tracking findings on trust formation in human–AI collaboration that link

perceived competence and perceived discretion to a user's willingness to self-disclose [114]. Engagement averaged 4.2 sessions per week across the four-week trial, with a mild downward trend from 4.6 sessions in week one to 3.7 sessions in week four; we read that decline as a return toward a more typical usage pattern after initial novelty rather than as attrition, since session-length and self-reported satisfaction did not decline over the same period.

User satisfaction across the four rated dimensions, empathy, helpfulness, naturalness and trust, is shown separately for urban and rural participants in Fig. 3. The proposed agent led both baselines on all four dimensions in both settings, but the margin over the retrieval-based baseline on naturalness was narrower in the rural sample (3.6 against 2.8, a 0.8-point gap) than in the urban sample (4.0 against 3.2, also a 0.8-point gap in absolute terms but a larger relative gap given the lower rural baseline). Rural participants rated empathy 0.8 points lower than urban participants for the proposed agent specifically (3.5 against 4.3 on the five-point scale; $t(238) = 3.42$, $p = 0.001$), the one dimension where the urban–rural difference reached significance after a Bonferroni correction across the four dimensions.




Fig. 3. Mean user satisfaction scores (1–5 Likert scale) across four dimensions, empathy, helpfulness, naturalness and trust, comparing the proposed agent against a rule-based baseline and a retrieval-based baseline, separately for urban and rural Indian participants ($n = 120$ each).

We looked for an explanation of the rural empathy gap in the transcripts themselves rather than assume it reflects a general shortfall in rural user experience. The pattern that recurred was dialectal: rural participants, more often than urban participants, used regional Hindi vocabulary and idiom that fell outside the training corpus's coverage, which the NLU module still classified correctly at the intent level but which the NLG templates then answered with phrasing the participant experienced as slightly generic rather than dialect-matched. This is a corpus-coverage limitation, not an architectural one, and Section VII returns to it as a limitation we are explicitly uncertain we have fully understood rather than one we can claim to have solved.

C. Clinical Safety Metrics

The safety monitor's 96.8% crisis-detection recall (Section IV-D) is the metric this study weighs most heavily, and Table I shows both baselines performing markedly worse on the same held-out test: 41.0% recall for the rule-based system, whose fixed keyword list missed indirect or metaphorical expressions of ideation, and 58.7% for the retrieval-based system, which matched surface similarity rather than underlying intent. Of the 187 crisis-state transitions logged across the four-week trial, clinician post-review judged 96 (51.3%) to represent an appropriate escalation warranting the delivered response, 74 (39.6%) to be precautionary escalations that caused no harm but were not strictly necessary, and 17 (9.1%) to be escalations whose message content, on review, could have been improved; none were judged to have produced an unsafe or harmful response. Two false

negatives, messages later identified by the participant's own follow-up as expressing ideation that the monitor did not flag at the time, were found on manual audit of the full transcript set; both were phrased with a level of indirection that fell outside the clinician-labelled training examples, and both are discussed as a specific, unresolved limitation in Section VII rather than folded into the aggregate recall figure, which reports only the held-out test-set performance and not this post-hoc audit finding.

VII. Ethical Considerations

The study operated under an approved institutional review board protocol, and every participant gave informed consent that named the study's data retention period, the categories of researcher and clinician who would have access to their transcripts, and the specific conditions, self-harm or harm-to-others disclosure, under which a human would review their conversation. A clinical psychologist held standing authority to pause the trial for any participant at any time, independent of the automated safety monitor, and exercised that authority twice during the four-week period for reasons unrelated to the agent's output.

Escalation followed a fixed protocol rather than clinician discretion on a case-by-case basis, a decision made explicitly to keep the response consistent regardless of which reviewer was on duty. A crisis-state transition triggered an immediate in-app message naming a national helpline number and, within the same session, a prompt asking whether the participant would consent to a callback from the study's clinical partner; consent to callback was requested rather than assumed, out of respect for participant autonomy over disclosure, following the practice recommended in reviews of ethical AI and IoT use during public health emergencies [104]. All transcripts were anonymised at the point of storage, with participant identifiers held in a separate, access-controlled table linked only by a one-way hash, a pattern consistent with the layered defences described for cloud security posture management [73] and for zero-trust anomaly detection in cloud-hosted health systems [159].

We considered and rejected a fully federated architecture, in which model updates rather than raw transcripts would leave each participant's device, despite its documented privacy advantages for distributed healthcare systems [106], [118], [119], [174]. The reason was practical rather than principled: the rural participant group's device heterogeneity meant that federated aggregation would have introduced a client-availability bias, systematically under-weighting the exact population this study most wanted to represent fairly, and we judged that risk to outweigh the privacy benefit for a four-week pilot of this size. A larger, multi-site deployment might reach a different balance, and we flag federated learning as the first architectural change we would evaluate before any expansion beyond a pilot.

Three limitations remain, stated here without a compensating strength attached to any of them. First, the safety monitor's 96.8% recall figure comes from a held-out test set; the post-hoc transcript audit in Section VI-C found two additional missed detections that a test-set metric alone would not surface, and we do not know how many more such cases a longer deployment would reveal. Second, the rural empathy gap traced to dialect coverage in Section VI-B is a plausible explanation supported by transcript review, not a confirmed causal account; we have not run the controlled ablation that would be needed to rule out other contributing factors, including the recruitment-channel difference noted in Section V. Third, this study measures engagement and self-reported satisfaction over four weeks and cannot speak to clinical outcomes, such as a validated reduction in depression or anxiety symptom scores, over a longer horizon; that measurement gap is the reason we describe this system throughout as low-intensity support rather than as treatment, and any future protocol extension should incorporate a validated clinical outcome measure with a longer follow-up window before making a therapeutic-benefit claim.

VIII. Conclusion

We built and evaluated a bilingual conversational agent for mental health support, tested it with 240 urban and rural Indian participants over four weeks, and measured it against a three-strand framework spanning technical, user-experience and clinical safety metrics. The agent reached 91.4% intent accuracy, 87.2% emotion-detection F1, and 96.8% crisis-detection recall, and participants rated it well above both a rule-based and a retrieval-based baseline on usability and on all four satisfaction dimensions measured. Rural participants rated empathy 0.8 points lower than urban participants, a gap we attribute to dialect coverage in the training corpus rather than to the underlying architecture, and the safety monitor's precision left roughly three in ten escalations judged unnecessary on clinical review. Two directions follow directly: extending the training corpus with regionally collected dialect data, and running a longer trial with a validated clinical outcome measure to establish whether the engagement and satisfaction gains reported here translate into a measurable reduction in symptom severity.

References

- [1] Garg, P., Dixit, A., & Sethi, P. (2022). ML-fresh: novel routing protocol in opportunistic networks using machine learning. *Computer Systems Science & Engineering*, Forthcoming. Tech Science Press.
- [2] Yadav, P. S., Khan, S., Singh, Y. V., Garg, P., & Singh, R. S. (2022). A Lightweight Deep Learning-Based Approach for Jazz Music Generation in MIDI Format. *Computational Intelligence and Neuroscience*, 2022.
- [3] Soni, E., Nagpal, A., Garg, P., & Pinheiro, P. R. (2022). Assessment of Compressed and Decompressed ECG Databases for Telecardiology Applying a Convolution Neural Network. *Electronics*, 11(17), 2708.
- [4] Pustokhina, I. V., Pustokhin, D. A., Lydia, E. L., Garg, P., Kadian, A., & Shankar, K. (2021). Hyperparameter search based convolution neural network with Bi-LSTM model for intrusion detection system in multimedia big data environment. *Multimedia Tools and Applications*, 1-18.
- [5] Khanna, A., Rani, P., Garg, P., Singh, P. K., & Khamparia, A. (2021). An Enhanced Crow Search Inspired Feature Selection Technique for Intrusion Detection Based Wireless Network System. *Wireless Personal Communications*, 1-18.
- [6] Garg, P., Dixit, A., Sethi, P., & Pinheiro, P. R. (2020). Impact of node density on the qos parameters of routing protocols in opportunistic networks for smart spaces. *Mobile Information Systems*, 2020.
- [7] Upadhyay, D., Garg, P., Aldossary, S. M., Shafi, J., & Kumar, S. (2023). A Linear Quadratic Regression-Based Synchronised Health Monitoring System (SHMS) for IoT Applications. *Electronics*, 12(2), 309.
- [8] Saini, P., Nagpal, B., Garg, P., & Kumar, S. (2023). CNN-BI-LSTM-CYP: A deep learning approach for sugarcane yield prediction. *Sustainable Energy Technologies and Assessments*, 57, 103263.
- [9] Saini, P., Nagpal, B., Garg, P., & Kumar, S. (2023). Evaluation of Remote Sensing and Meteorological parameters for Yield Prediction of Sugarcane (*Saccharum officinarum* L.) Crop. *Brazilian Archives of Biology and Technology*, 66, e23220781.
- [10] Beniwal, S., Saini, U., Garg, P., & Joon, R. K. (2021). Improving performance during camera surveillance by integration of edge detection in IoT system. *International Journal of E-Health and Medical Communications (IJEHMC)*, 12(5), 84-96.
- [11] Garg, P., Dixit, A., & Sethi, P. (2019). Wireless sensor networks: an insight review. *International Journal of Advanced Science and Technology*, 28(15), 612-627.
- [12] Sharma, N., & Garg, P. (2022). Ant colony based optimization model for QoS-Based task scheduling in cloud computing environment. *Measurement: Sensors*, 100531.
- [13] Kumar, P., Kumar, R., & Garg, P. (2020). Hybrid Crowd Cloud Routing Protocol For Wireless Sensor Networks. *International Journal of Advanced Science and Technology*, 29, 766-775.
- [14] Raj, G., Verma, A., Dalal, P., Shukla, A. K., & Garg, P. (2023). Performance Comparison of Several LPWAN Technologies for Energy Constrained IOT Network. *International Journal of Intelligent Systems and Applications in Engineering*, 11(1s), 150-158.
- [15] Garg, P., Sharma, N., & Shukla, B. (2023). Predicting the Risk of Cardiovascular Diseases using Machine Learning Techniques. *International Journal of Intelligent Systems and Applications in Engineering*, 11(2s), 165-173.
- [16] Patil, S. C., Mane, D. A., Singh, M., Garg, P., Desai, A. B., & Rawat, D. (2024). Parkinson's Disease Progression Prediction Using Longitudinal Imaging Data and Grey Wolf Optimizer-Based Feature Selection. *International Journal of Intelligent Systems and Applications in Engineering*, 12(3s), 441-451.
- [17] Gudur, A., Pati, P., Garg, P., & Sharma, N. (2024). Radiomics Feature Selection for Lung Cancer Subtyping and Prognosis Prediction: A Comparative Study of Ant Colony Optimization and Simulated Annealing. *International Journal of Intelligent Systems and Applications in Engineering*, 12(3s), 553-565.
- [18] Khan, A. (2024). Optimisation Methods Based on Soft Computing for Improving Power System Stability. *J. Electrical Systems*, 20(6s), 1051-1058.
- [19] Sharma, K. K., Verma, P. K., & Garg, P. (2024). IoT-Enabled Energy Management Systems For Sustainable Energy Storage: Design, Optimization, And Future Directions. *Frontiers in Health Informatics*, 13(8).
- [20] Gupta, S., Yadav, N., Singh, K., & Garg, P. (2025). Applications of Simulations and Queuing Theory in Supermarket. *Reliability: Theory & Applications*, 20(1 (82)), 135-140.
- [21] Beniwal, S., Garg, P., Rajpal, R., Sharma, N., & Mittal, H. K. (2025). Fusion of Opportunistic Networks with Machine Learning: Present and Future. *Metallurgical and Materials Engineering*, 31(1), 311-324.
- [22] Garg, P. (2025). Explainable AI & Model Interpretability in Healthcare: Challenges & Future Directions. *Eksplorium-Buletin Pusat Teknologi Bahan Galian Nuklir*, 46(1), 104-133.
- [23] Rani, P. (2025). From Data to Diagnosis: Unleashing AI and 6G in Modern Medicine. *Eksplorium-Buletin Pusat Teknologi Bahan Galian Nuklir*, 46(1), 69-103.
- [24] Dixit, A., Garg, P., Sethi, P., & Singh, Y. (2020, April). TVCCCS: Television Viewer's Channel Cost Calculation System On Per Second Usage. In *IOP Conference Series: Materials Science and Engineering* (Vol. 804, No. 1, p. 012046). IOP Publishing.
- [25] Sethi, P., Garg, P., Dixit, A., & Singh, Y. (2020, April). Smart number cruncher—a voice based calculator. In *IOP Conference Series: Materials Science and Engineering* (Vol. 804, No. 1, p. 012041). IOP Publishing.
- [26] S. Rai, V. Choubey, Suryansh and P. Garg, "A Systematic Review of Encryption and Keylogging for Computer System Security," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 157-163, doi: 10.1109/CCICT56684.2022.00039.
- [27] L. Saraswat, L. Mohanty, P. Garg and S. Lamba, "Plant Disease Identification Using Plant Images," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 79-82, doi: 10.1109/CCICT56684.2022.00026.
- [28] L. Mohanty, L. Saraswat, P. Garg and S. Lamba, "Recommender Systems in E-Commerce," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 114-119, doi: 10.1109/CCICT56684.2022.00032.
- [29] C. Maggo and P. Garg, "From linguistic features to their extractions: Understanding the semantics of a concept," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 427-431, doi: 10.1109/CCICT56684.2022.00082.
- [30] N. Puri, P. Saggat, A. Kaur and P. Garg, "Application of ensemble Machine Learning models for phishing detection on web networks," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 296-303, doi: 10.1109/CCICT56684.2022.00062.
- [31] R. Sharma, S. Gupta and P. Garg, "Model for Predicting Cardiac Health using Deep Learning Classifier," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 25-30, doi: 10.1109/CCICT56684.2022.00017.
- [32] Varshney, S. Lamba and P. Garg, "A Comprehensive Survey on Event Analysis Using Deep Learning," 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT), 2022, pp. 146-150, doi: 10.1109/CCICT56684.2022.00037.

- [33] Dixit, A., Sethi, P., Garg, P., & Pruthi, J. (2022, December). Speech Difficulties and Clarification: A Systematic Review. In 2022 11th International Conference on System Modeling & Advancement in Research Trends (SMART) (pp. 52-56). IEEE.
- [34] Garg, P., Dixit, A., Sethi, P., & Pruthi, J. (2023, December). Strengthening Smart City with Opportunistic Networks: An Insight. In 2023 International Conference on Advanced Computing & Communication Technologies (ICACCTech) (pp. 700-707). IEEE.
- [35] Rana, S., Chaudhary, R., Gupta, M., & Garg, P. (2023, December). Exploring Different Techniques for Emotion Detection Through Face Recognition. In 2023 International Conference on Advanced Computing & Communication Technologies (ICACCTech) (pp. 779-786). IEEE.
- [36] Mittal, K., Srivastava, K., Gupta, M., & Garg, P. (2023, December). Exploration of Different Techniques on Heart Disease Prediction. In 2023 International Conference on Advanced Computing & Communication Technologies (ICACCTech) (pp. 758-764). IEEE.
- [37] Gautam, V. K., Gupta, S., & Garg, P. (2024, March). Automatic Irrigation System using IoT. In 2024 International Conference on Automation and Computation (AUTOCOM) (pp. 100-103). IEEE.
- [38] Ramasamy, L. K., Khan, F., Joghee, S., Dempere, J., & Garg, P. (2024, March). Forecast of Students' Mental Health Combining an Artificial Intelligence Technique and Fuzzy Inference System. In 2024 International Conference on Automation and Computation (AUTOCOM) (pp. 85-90). IEEE.
- [39] Rajput, R., Sukumar, V., Patnaik, P., Garg, P., & Ranjan, M. (2024, March). The Cognitive Analysis for an Approach to Neuroscience. In 2024 International Conference on Automation and Computation (AUTOCOM) (pp. 524-528). IEEE.
- [40] Dixit, A., Sethi, P., Garg, P., Pruthi, J., & Chauhan, R. (2024, July). CNN based lip-reading system for visual input: A review. In AIP Conference Proceedings (Vol. 3121, No. 1). AIP Publishing.
- [41] Bose, D., Arora, B., Srivastava, A. K., & Garg, P. (2024, May). A Computer Vision Based Framework for Posture Analysis and Performance Prediction in Athletes. In 2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE) (pp. 942-947). IEEE.
- [42] Singh, M., Garg, P., Srivastava, S., & Saggi, A. K. (2024, April). Revolutionizing Arrhythmia Classification: Unleashing the Power of Machine Learning and Data Amplification for Precision Healthcare. In 2024 Sixth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 516-522). IEEE.
- [43] Kumar, R., Das, R., Garg, P., & Pandita, N. (2024, April). Duplicate Node Detection Method for Wireless Sensors. In 2024 Sixth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 512-515). IEEE.
- [44] Bhardwaj, H., Das, R., Garg, P., & Kumar, R. (2024, April). Handwritten Text Recognition Using Deep Learning. In 2024 Sixth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 506-511). IEEE.
- [45] Gill, A., Jain, D., Sharma, J., Kumar, A., & Garg, P. (2024, May). Deep learning approach for facial identification for online transactions. In 2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP) (pp. 715-722). IEEE.
- [46] Mittal, H. K., Dalal, P., Garg, P., & Joon, R. (2024, May). Forecasting Pollution Trends: Comparing Linear, Logistic Regression, and Neural Networks. In 2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP) (pp. 411-419). IEEE.
- [47] Malik, T., Nandal, V., & Garg, P. (2024, May). Deep Learning-Based Classification of Diabetic Retinopathy: Leveraging the Power of VGG-19. In 2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP) (pp. 645-651). IEEE.
- [48] Srivastava, A. K., Verma, I., & Garg, P. (2024, May). Improvements in Recommendation Systems Using Graph Neural Networks. In 2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP) (pp. 668-672). IEEE.
- [49] Aggarwal, A., Jain, D., Gupta, A., & Garg, P. (2024, May). Analysis and Prediction of Churn and Retention Rate of Customers in Telecom Industry Using Logistic Regression. In 2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP) (pp. 723-727). IEEE.
- [50] Mittal, H. K., Arsalan, M., & Garg, P. (2024, May). A Novel Deep Learning Model for Effective Story Point Estimation in Agile Software Development. In 2024 International Conference on Emerging Innovations and Advanced Computing (INNOCOMP) (pp. 404-410). IEEE.
- [51] Shukla, S. M., Magoo, C., & Garg, P. (2024, November). Comparing Fine Tuned-LMs for Detecting LLM-Generated Text. In 2024 3rd Edition of IEEE Delhi Section Flagship Conference (DELCON) (pp. 1-8). IEEE.
- [52] Kumar, B., Iqbal, M., Parmar, R., Garg, P., Rani, S., & Agrawal, A. (2025, March). The Role of AI in Optimizing Healthcare Appointment Scheduling. In 2025 3rd International Conference on Disruptive Technologies (ICDT) (pp. 881-887). IEEE.
- [53] Kumar, B., Garg, V., Ahmed, K., Garg, P., Choudhary, S., & Baniya, P. (2025, March). Enhancing Healthcare with Blockchain: Innovations in Data Privacy, Security, and Interoperability. In 2025 3rd International Conference on Disruptive Technologies (ICDT) (pp. 932-938). IEEE.
- [54] Raj, V., Prakash, B. K., Kumar, A., & Garg, P. (2024, December). Optimize the Time a Mercedes-Benz Spends on the Test Bench Using Stacking Ensemble Learning. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 445-450). IEEE.
- [55] Kaushik, N., Kumar, H., Raj, V., & Garg, P. (2024, December). Proactive Fault Prediction in Microservices Applications Using Trace Logs and Monitoring Metrics. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 410-415). IEEE.
- [56] Kumar, A. A., Sri, C. V., Bohara, K. S. K., Setia, S., & Garg, P. (2024, December). Capnivesh: Financing Platform for Startups. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 261-265). IEEE.
- [57] Bhandari, P., Setia, S., Kumar, K., & Garg, P. (2024, December). Optimizing Cross-Platform Development with CI/CD and Containerization: A Review. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 175-180). IEEE.
- [58] Chaudhary, A., & Garg, P. (2014). Detecting and diagnosing a disease by patient monitoring system. *International Journal of Mechanical Engineering And Information Technology*, 2(6), 493-499.
- [59] Malik, K., Raheja, N., & Garg, P. (2011). Enhanced FP-growth algorithm. *International Journal of Computational Engineering and Management*, 12, 54-56.
- [60] Garg, P., Dixit, A., & Sethi, P. (2021, May). Link Prediction Techniques for Opportunistic Networks using Machine Learning. In Proceedings of the International Conference on Innovative Computing & Communication (ICICC).
- [61] Garg, P., Dixit, A., & Sethi, P. (2021, April). Opportunistic networks: Protocols, applications & simulation trends. In Proceedings of the International Conference on Innovative Computing & Communication (ICICC).
- [62] Garg, P., Dixit, A., & Sethi, P. (2021). Performance comparison of fresh and spray & wait protocol through one simulator. *IT in Industry*, 9(2).
- [63] Malik, M., Singh, Y., Garg, P., & Gupta, S. (2020). Deep Learning in Healthcare system. *International Journal of Grid and Distributed Computing*, 13(2), 469-468.

- [64] Gupta, M., Garg, P., Gupta, S., & Joon, R. (2020). A Novel Approach for Malicious Node Detection in Cluster-Head Gateway Switching Routing in Mobile Ad Hoc Networks. *International Journal of Future Generation Communication and Networking*, 13(4), 99-111.
- [65] Gupta, A., Garg, P., & Sonal, Y. S. (2020). Edge Detection Based 3D Biometric System for Security of Web-Based Payment and Task Management Application. *International Journal of Grid and Distributed Computing*, 13(1), 2064-2076.
- [66] Garg, P., & Raman, P. K. (2011). Broadcasting Protocol & Routing Characteristics With Wireless ad-hoc networks. *Int. J. Comput. Eng. Manag.*, 12(1), 36-40.
- [67] Garg, P., Arora, N., & Malik, T. (2011). Capacity Improvement of WI-MAX In presence of Different Codes WI-MAX: Speed & Scope of future. *IJCEM*, 12.
- [68] Garg, P., Saroha, K., & Lochab, R. (2011). Review of wireless sensor networks-architecture and applications. *IJCSMS International Journal of Computer Science & Management Studies*, 11(01), 2231-5268.
- [69] Yadav, S., & Garg, P. Development of a New Secure Algorithm for Encryption and Decryption of Images.
- [70] Dixit, A., Sethi, P., & Garg, P. (2022). Rakshak: A Child Identification Software for Recognizing Missing Children Using Machine Learning-Based Speech Clarification. *International Journal of Knowledge-Based Organizations (IJKBO)*, 12(3), 1-15.
- [71] Shukla, N., Garg, P., & Singh, M. (2022). MANET Proactive and Reactive Routing Protocols: A Comparison Study. *International Journal of Knowledge-Based Organizations (IJKBO)*, 12(3), 1-14.
- [72] Arya, A., Garg, P., Vellanki, S., Latha, M., Khan, M. A., & Chhbra, G. (2024). Optimisation Methods Based on Soft Computing for Improving Power System Stability. *Journal of Electrical Systems*, 20(6s), 1051-1058.
- [73] Garg, P. (2025). Cloud security posture management: Tools and techniques. *Technix International Journal for Engineering Research*, 12(3).
- [74] Tyagi, P., Sharma, S., Srivastava, A., Rajput, N. K., Garg, P., & Kumari, M. (2025). AI in Healthcare: Transforming Medicine with Intelligence. In *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. <https://doi.org/10.63169/GCARED2025.p4>
- [75] Garg, P., Bhatt, M., Parmar, R., & Arsalan, M. (2025). Generative AI: Evolution, Applications, Challenges And Future Prospects. Applications, Challenges And Future Prospects (May 17, 2025).
- [76] Garg, P., Saraswat, P., & Siddiqui, Z. (2025). AI & the Indian Stock Market: A Review of Applications in Investment Decision. <https://doi.org/10.63169/GCARED2025.p10>
- [77] Garg, P., Sharma, S., Mittal, S., Tevatia, R., Tyagi, V. K., & Kapoor, S. (2025). Unlocking Workforce Potential: AI-Powered Predictive Models for Employee Performance Evaluation. <https://doi.org/10.63169/GCARED2025.p21>
- [78] Shrivastava, N., Kalia, A., Roy, R., Sharma, S., Garg, P., & Agarwal, G. (2025). OSINT: A Double-edged Sword. In *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. <https://doi.org/10.63169/GCARED2025.p22>
- [79] Garg, P., Aditi, A., & Roy, B. (2025). A System of Computer Network: Based On Artificial Intelligence. <https://doi.org/10.63169/GCARED2025.p24>
- [80] Parmar, R., Kapoor, S., Saifi, S., & Garg, P. (2025). Case Study on Intelligent Factory Systems for Improving Productivity and Capability in Industry 4.0 with Generative AI. In *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. <https://doi.org/10.63169/GCARED2025.p28>
- [81] Singh, R., Sharma, R., Kumar, R., Nafis, A., Siddiqui, M. A. M., & Garg, P. (2025). Detection of Unauthorized Construction using Machine Learning: A Review. In *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. <https://doi.org/10.63169/GCARED2025.p30>
- [82] Garg, P., Kapoor, S., Singh, V., Sharma, S., & Ankita, A. (2025). A Bridge between Blockchain and Decentralized Applications Web3 and Non-Web3 Crypto Wallets. <https://doi.org/10.63169/GCARED2025.p35>
- [83] Verma, M., Sharma, S., Garg, P., & Singh, A. (2025). The Hidden Dangers of Prototype Pollution: A Comprehensive Detection Framework. In *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. <https://doi.org/10.63169/GCARED2025.p36>
- [84] Sharma, A., Sharma, S., Garg, P., & Bhardwaj, P. (2025). LockTalk: A Basic Secure Chat Application. In *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India.
- [85] Arora, K., Bawane, R., Gupta, C., Ahmed, K., & Garg, P. (2025). Detection and Prevention of Cyber Attack and Threat using AI. In *First Global Conference on AI Research and Emerging Developments (G-CARED 2025)*, New Delhi, India. <https://doi.org/10.63169/GCARED2025.p38>
- [86] Garg, P., Dhruv, D., Rahman, A. A., Rai, A., Siddiqui, M., & Yadav, D. (2025). Easeviewer: An Esports Production Tool. <https://doi.org/10.63169/GCARED2025.p46>
- [87] Garg, P., Lakshita, L., Mehresh, M., Nazia, N., & Ahmed, K. (2025). Emerging Trend in Computational Technology: Innovations, Applications, and Challenges. Applications, and Challenges (May 17, 2025). <https://doi.org/10.63169/GCARED2025.p51>
- [88] Chauhan, S., Singh, M., & Garg, P. (2021). Rapid Forecasting of Pandemic Outbreak Using Machine Learning. *Enabling Healthcare 4.0 for Pandemics: A Roadmap Using AI, Machine Learning, IoT and Cognitive Technologies*, 59-73.
- [89] Gupta, S., & Garg, P. (2021). An insight review on multimedia forensics technology. *Cyber Crime and Forensic Computing: Modern Principles, Practices, and Algorithms*, 11, 27.
- [90] Shrivastava, P., Agarwal, P., Sharma, K., & Garg, P. (2021). Data leakage detection in Wi-Fi networks. *Cyber Crime and Forensic Computing: Modern Principles, Practices, and Algorithms*, 11, 215.
- [91] Meenakshi, P. G., & Shrivastava, P. (2021). Machine learning for mobile malware analysis. *Cyber Crime and Forensic Computing: Modern Principles, Practices, and Algorithms*, 11, 151.
- [92] Nanwal, J., Garg, P., Sethi, P., & Dixit, A. (2021). Green IoT and Big Data: Succeeding towards Building Smart Cities. In *Green Internet of Things for Smart Cities* (pp. 83-98). CRC Press.
- [93] Gupta, M., Garg, P., & Agarwal, P. (2021). Ant Colony Optimization Technique in Soft Computational Data Research for NP-Hard Problems. In *Artificial Intelligence for a Sustainable Industry 4.0* (pp. 197-211). Springer, Cham.
- [94] Magoo, C., & Garg, P. (2021). Machine Learning Adversarial Attacks: A Survey Beyond. *Machine Learning Techniques and Analytics for Cloud Security*, 271-291.
- [95] Garg, P., Srivastava, A. K., Anas, A., Gupta, B., & Mishra, C. (2023). Pneumonia Detection Through X-Ray Images Using Convolution Neural Network. In *Advancements in Bio-Medical Image Processing and Authentication in Telemedicine* (pp. 201-218). IGI Global.
- [96] Gupta, S., & Garg, P. (2023). Code-based post-quantum cryptographic technique: digital signature. *Quantum-Safe Cryptography Algorithms and Approaches: Impacts of Quantum Computing on Cybersecurity*, 193.
- [97] Prakash, A., Avasthi, S., Kumari, P., & Rawat, M. (2023). Modern healthcare system: unveiling the possibility of quantum computing in medical and biomedical zones. *Quantum-Safe Cryptography Algorithms and Approaches: Impacts of Quantum Computing on Cybersecurity*, 249.

- [98] Gupta, S., & Garg, P. (2024). Mobile Edge Computing for Decentralized Systems. *Decentralized Systems and Distributed Computing*, 75-88.
- [99] Gupta, M., Garg, P., & Malik, C. (2024). Ensemble learning-based analysis of perinatal disorders in women. In *Artificial Intelligence and Machine Learning for Women's Health Issues* (pp. 91-105). Academic Press.
- [100] Malik, M., Garg, P., & Malik, C. (2024). Artificial intelligence-based prediction of health risks among women during menopause. *Artificial Intelligence and Machine Learning for Women's Health Issues*, 137-150.
- [101] Garg, P. (2024). Prediction of female pregnancy complication using artificial intelligence. In *Artificial Intelligence and Machine Learning for Women's Health Issues* (pp. 17-35). Academic Press.
- [102] Pokhrel, L., Arsalan, M., Rani, P., Garg, P., & Pinheiro, P. R. (2026). AI-Powered Healthcare Solutions: Bridging the Medical Gap in Underserved Communities Worldwide. In *Applied AI and Computational Intelligence in Diagnostics and Decision-Making* (pp. 57-86). IGI Global Scientific Publishing.
- [103] Kapoor, S., Parmar, R., Sharma, N., Garg, P., & Singh, N. J. (2026). AI and Computational Intelligence in Healthcare: An Introductory Guide. In *Applied AI and Computational Intelligence in Diagnostics and Decision-Making* (pp. 1-26). IGI Global Scientific Publishing.
- [104] Pokhrel, L., Kumar, A., Garg, P., Anand, N., & Singh, N. (2026). AI and IoT in Global Health: Ethical Lessons From Pandemic Response. In *Development and Management of Eco-Conscious IoT Medical Devices* (pp. 367-394). IGI Global Scientific Publishing.
- [105] Parmar, R., Singh, A., Garg, P., Sharma, T., & Pinheiro, P. R. (2026). Blockchain for Ethical Supply Chains: Transparency in Medical IoT Manufacturing. In *Development and Management of Eco-Conscious IoT Medical Devices* (pp. 337-366). IGI Global Scientific Publishing.
- [106] Gupta, S., Garg, P., Agarwal, J., Thakur, H. K., & Yadav, S. P. (2024). Federated learning based intelligent systems to handle issues and challenges in IoVs (Part 1). <https://doi.org/10.2174/97898153130311240301>
- [107] Gupta, S., Chaudhary, G., & Garg, P. (2013). Modified AODV Routing Protocol through Cache Memory for Finding New Routing Paths in MANETs. *International Journal of Computer Science & Management Studies*, 13(3).
- [108] Gupta, A., & Garg, P. (2021). Emerging Techniques for Handling Pandemic Challenges. *Enabling Healthcare 4.0 for Pandemics: A Roadmap Using AI, Machine Learning, IoT and Cognitive Technologies*, 189-209.
- [109] Chaudhary, A. P., Mishra, A., Kumar, D., & Garg, P. (2023, April). Human Emotion Recognition using Deep Learning. In *2023 International Conference on Computational Intelligence, Communication Technology and Networking (CICTN)* (pp. 191-197). IEEE.
- [110] Nagpal, S., Garg, P., Gaba, S., & Aggarwal, A. (2023). An improved genetic quantum cryptography model for network communication. *Quantum-Safe Cryptography Algorithms and Approaches: Impacts of Quantum Computing on Cybersecurity*, 177.
- [111] Yadav, M., Swami, V., Kumar, N., & Garg, P. (2025). Comparative study of Repairable Juice Plants using RPGT. *Reliability: Theory & Applications*, 20(2 (84)), 776-783.
- [112] Gupta, A., Garg, P., & Yadav, P. (2025). Role of Generative AI Towards Education and Learning: Present & Future. *TPM-Testing, Psychometrics, Methodology in Applied Psychology*, 32(S6 (2025)), 1059-1076.
- [113] Dalal, P., Beniwal, G., Sharma, V., Garg, P., & Ahmed, K. (2025). Predicting Student Motivation and Engagement through Machine Learning Models. *TPM-Testing, Psychometrics, Methodology in Applied Psychology*, 32(S7 (2025)), 393-411.
- [114] Gupta, A., Mund, A., Roy, S., Garg, P., & Yadav, D. K. (2025). Trust in AI Systems: A Social-psychological Investigation of Human-AI Collaboration. *TPM-Testing, Psychometrics, Methodology in Applied Psychology*, 32(S7 (2025)), 428-446.
- [115] Bhardwaj, A., Das, A., Garg, P., & Yadav, S. (2025). Material-Driven Performance Analysis of a Vertical Nanowire Tunnel FET for Analog Applications. *Journal of Electronic Materials*, 1-12.
- [116] Dalal, P., Sharma, B., Sharma, T., Garg, P., & Ahmed, K. (2025). Explainable AI for Understanding Human Decision Making Patterns. *TPM-Testing, Psychometrics, Methodology in Applied Psychology*, 32(S7 (2025)), 412-427.
- [117] Sharma, K. K., Verma, P. K., Garg, P., & Shrotriya, V. K. (2025, October). Predicting costs and benefits of IoT-based energy management for optimizing sustainable energy storage in rural areas. In *AIP Conference Proceedings* (Vol. 3343, No. 1, p. 040017). AIP Publishing LLC.
- [118] Ahmed, K., Baranwal, A., Sharma, N., Garg, P., & Singh, N. (2026). The Role of Federated Learning in AI-Powered Integrated Healthcare Solutions. In *Enabling Collaborative Health Intelligence With Federated Learning* (pp. 421-448). IGI Global Scientific Publishing.
- [119] Gupta, S., Garg, P., Agarwal, J., Thakur, H. K., & Yadav, S. P. (2025). Federated learning based intelligent systems to handle issues and challenges in IoVs (Part 2). *Bentham Science Publishers*. <https://doi.org/10.2174/97898153222241250301>
- [120] Garg, P., Pranav, S., & Prerna, A. (2021). Green internet of things (G-IoT): A solution for sustainable technological development. In *Green Internet of Things for Smart Cities* (pp. 23-46). CRC Press.
- [121] Malik, A., Nandal, D., Gupta, V., Garg, P., & Nandal, V. *Intelligent Systems and Applications in Engineering*.
- [122] Gupta, S., Garg, P., Agarwal, J., Thakur, H. K., & Yadav, S. P. (Eds.). (2025). *Federated learning based intelligent systems to handle issues and challenges in IoVs (Part 2)*.
- [123] Garg, P., Bhatt, M., Parmar, R., & Arsalan, M. (2025). *Generative AI: Evolution, Applications, Challenges And Future Prospects*. Applications, Challenges And Future Prospects (May 17, 2025).
- [124] Kumar, N., Kumar, Y., Khurana, D., Kumar, S., & Garg, P. (2025, November). A Hybrid Ensemble Learning Framework for Interpretable Student Performance Prediction Using Academic and Extracurricular Factors. In *2025 International Conference on Innovations and Emerging Technologies In AI & Communication Systems (IETACS)* (pp. 666-672). IEEE.
- [125] Khurana, D., Kumar, Y., Kumar, N., Kumar, S., & Garg, P. (2025, November). Transformer-Based Movie Recommendation System with Autoencoder-Enhanced Feature Compression. In *2025 International Conference on Innovations and Emerging Technologies In AI & Communication Systems (IETACS)* (pp. 685-690). IEEE.
- [126] Garg, P. (2025, November). Comparative Analysis of Various Neural Networks for Galaxies Classification. In *2025 International Conference on Innovations and Emerging Technologies In AI & Communication Systems (IETACS)* (pp. 697-701). IEEE.
- [127] Saggu, A. K., Babbar, N., & Garg, P. (2025, November). Health-Guard AI: Integrated Health Report Management and Analysis. In *2025 International Conference on Innovations and Emerging Technologies In AI & Communication Systems (IETACS)* (pp. 614-623). IEEE.
- [128] Kumar, S., Kumar, Y., Kumar, N., Khurana, D., & Garg, P. (2025, November). Hybrid FCM-DNN Model for Uncertainty-Aware Air Quality Classification Using Multi-Pollutant Data. In *2025 International Conference on Innovations and Emerging Technologies In AI & Communication Systems (IETACS)* (pp. 679-684). IEEE.
- [129] Babbar, N., Singh, H. V., Bendale, S., & Garg, P. (2025, November). Stock Market Price Prediction Using Big Data Analysis: A Performance Evaluation Study. In *2025 3rd International Conference on Computational Intelligence and Network Systems (CINS)* (pp. 1-6). IEEE.

- [130] Singh, A. K., Kori, G., Garg, P., & Srivastava, G. (2025, November). Bank Churn Prediction Using Machine Learning. In 2025 IEEE 7th International Conference on Computing, Communication and Automation (ICCCA) (pp. 1-6). IEEE.
- [131] Bhardwaj, A., Das, A., Garg, P., & Yadav, S. (2026). Material-Driven Performance Analysis of a Vertical Nanowire Tunnel FET for Analog Applications. *Journal of Electronic Materials*, 55(1), 1099-1110.
- [132] Srivastava, A. K., Shankdhar, D., Ror, R., & Garg, P. (2026). Harnessing YOLOv5 for real-time object detection: A cloud-based approach. In *Recent Advances in Computational Methods in Science and Technology* (pp. 441-450). CRC Press.
- [133] Srivastava, A. K., Shukla, A., Gupta, H., Saxena, K., & Garg, P. (2026). Towards intelligent attendance management system with face recognition using LBPH algorithm. In *Recent Advances in Computational Methods in Science and Technology* (pp. 8-15). CRC Press.
- [134] Srivastava, A. K., Garg, P., & Pandey, H. (2026). Vedcure: Towards intelligent ayurvedic drug recommendation and disease prediction. In *Recent Advances in Computational Methods in Science and Technology* (pp. 16-23). CRC Press.
- [135] Upadhyay, D., Garg, P., & Babbar, N. (2026). Blockchain and IoT based smart contract framework for efficient and secure product life management. *Discover Internet of Things*.
- [136] Singh, A., Parmar, R., Bhardwaj, P., Sharma, V., & Garg, P. (2026). Fusion of Aerial Networks with Advanced Computing Paradigms. *Edge Computing and Aerial Platforms*, 355-367.
- [137] Kumari, M., Baranwal, A., Sonal, & Garg, P. (2026). Application of Aerial Edge Computing in Disaster Management. *Edge Computing and Aerial Platforms*, 103-122.
- [138] Aditi, Saraswat, P., Sharma, V., & Garg, P. (2026). Advances in Aerial Platforms and Edge Computing. *Edge Computing and Aerial Platforms*, 123-143.
- [139] Garg, P., Arora, K., Bawane, R., Gupta, C., & Ahmed, K. (2025). Detection and Prevention of Cyber Attack and Threat using AI.
- [140] Ahmed, K., Ahmed, K., Khan, J., Garg, P., Seth, S., & Mallik, S. (2025). Principal Component Analysis Based Clustering of Insecticides and Molecular Docking of Pyrethroids Insecticides.
- [141] Kumar, B., Kumar, A., Nanwal, J., Garg, P., & Patnaik, P. (2025, November). Ensemble of YOLOv5 and Segment Anything Model for Brain Tumor Detection. In 2025 2nd International Conference on Advanced Computing and Emerging Technologies (ACET) (pp. 1-5). IEEE.
- [142] Arsalan, M., Anas, M., & Garg, P. (2025). Transparent AI for Drug Discovery and Development. Available at SSRN 5844242.
- [143] Singh, A., Bhardwaj, P., Garg, P., & Singh, N. (2026). Introduction to explainable artificial intelligence in healthcare. In *Explainable AI in Clinical Practice* (pp. 23-44). Academic Press.
- [144] Kapoor, S., Singh, A., Garg, P., & Ramasamy, L. K. (2026). Explainable artificial intelligence in a diagnostic support system. In *Explainable AI in Clinical Practice* (pp. 131-145). Academic Press.
- [145] Ahmed, K., Anas, M., & Garg, P. (2026). Case studies on unlocking the potential of Industry 4.0 for sustainable manufacturing through generative AI-driven innovations. Available at SSRN 6356958.
- [146] Garg, P., & Oruganti, S. K. (2026, March). AI Assisted Routing Optimization in Opportunistic IoT Networks using Machine Learning: A Comprehensive review on Protocols & Simulators. In *Sustainable Global Societies Initiative* (Vol. 1, No. 4). Vibrasphere Technologies.
- [147] Arsalan, M., Pokhrel, L., & Garg, P. (2026). Architecture, Components, and tools in Integrated AI-Augmented Intelligence: A design perspective. Components, and tools in Integrated AI-Augmented Intelligence: A design perspective (March 19, 2026).
- [148] Singh, H., Ahmed, K., & Garg, P. (2026). Human Versus Machine Customer Behaviour and Functional Differences. Available at SSRN 6441098.
- [149] Saraswat, P., & Garg, P. (2026). Soft Computing In AI Agents.
- [150] Saraswat, P., & Garg, P. (2026). Water Quality Prediction Using IOT Sensors and Deep Networks.
- [151] Arsalan, M., Ahmed, K., & Garg, P. (2026). Machine learning for Anomaly detection in sensor networks. Available at SSRN 6441518.
- [152] Kumari, M., & Garg, P. (2026). Hybrid Cloud Infrastructure: Models, Benefits and Security And Challenges. Benefits and Security and Challenges (March 19, 2026).
- [153] Singh, H., & Garg, P. (2026). Demystifying Artificial Distributed Intelligence (ADI). Available at SSRN 6442698.
- [154] Saraswat, P., & Garg, P. (2026). Human AI Collaboration the future of clinical decision making.
- [155] Saraswat, P., & Garg, P. (2026). Breaking Data Boundaries Federated Learning in Digital Healthcare.
- [156] Singh, A., Parmar, R., Bhardwaj, P., Sharma, V., & Garg, P. (2026). Fusion of Aerial Networks with Advanced Computing Paradigms. *Edge Computing and Aerial Platforms*, 355-367.
- [157] Kumari, M., Baranwal, A., Sonal, & Garg, P. (2026). Application of Aerial Edge Computing in Disaster Management. *Edge Computing and Aerial Platforms*, 103-122.
- [158] Aditi, Saraswat, P., Sharma, V., & Garg, P. (2026). Advances in Aerial Platforms and Edge Computing. *Edge Computing and Aerial Platforms*, 123-143.
- [159] Garg, P. (2026). Zero-Trust Security Enforcement through AI-Powered Anomaly Detection in Cloud Systems. *Journal of Artificial Intelligence in Governance and Public Policy (JAIGPP)*, 1(1), 1-8.
- [160] Singh, A. P., Sharma, A., & Garg, P. (2026, January). AI-Powered Adaptive Mock Interview Generation System. In 2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC) (pp. 1421-1426). IEEE.
- [161] Raghav, A., Mishra, A., & Garg, P. (2026, January). Enhancing Healthcare Access: An AI-driven Chatbot for Doctor Appointment Management. In 2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC) (pp. 910-914). IEEE.
- [162] Kumar, B., Chauhan, D., Singh, H., Verma, H., Sahu, K., & Garg, P. (2026, January). Decoding Linear A with Artificial Intelligence: A Comprehensive Machine Learning and NLP Framework. In 2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC) (pp. 1-7). IEEE.
- [163] Gupta, V., Chakravarti, L., Akhtar, M. M., Maheshwari, P., Garg, P., & Tiwari, D. (2026, January). A Sentence-Level Risk Estimator for Identifying Hallucinations in Generative AI. In 2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC) (pp. 1619-1626). IEEE.
- [164] Patnaik, P., & Garg, P. (2026). Principles of Artificial Intelligence: Cognitive Architectures, Large Language Models, and Computational Limits. Deep Science Publishing.
- [165] Singh, K., & Garg, P. (2026). Trustworthy Deep Learning: Robustness, Uncertainty Quantification, and Adversarial Resilience. Deep Science Publishing.
- [166] Garg, P. (2026). Survey of Load Balancing Strategies in Fog-Cloud Architectures for IoT Integration. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 9(2), 595-604.
- [167] Tiwari, D., Bhati, B. S., Garg, P., & Babbar, N. (2026). A novel lexicon dictionary and CNN-LSTM employed hybrid approach for sentiment detection of COVID-19 vaccines. *Scientific Reports*.

- [168] Sethi, P., Kumar, B., Garg, P., & Yadav, P. (2026, February). Graph Neural Networks for Causal Inference in Climate Science: A Novel Approach to Modeling Complex Interactions. In 2026 International Conference on Intelligent Computing and Automation for Sustainable Solutions (ICASS) (pp. 1-7). IEEE.
- [169] Kumar, B., Garg, P., Garg, A., Nanwal, J., & Yadav, P. (2026, February). An AI-Centric Multimodal Framework for Cognitive Productivity and Digital Wellbeing. In 2026 International Conference on Intelligent Computing and Automation for Sustainable Solutions (ICASS) (pp. 1-7). IEEE.
- [170] Ahmed, K., Mathur, T., Chaudhary, Y., Anshika, A., Garg, P., & Tiwari, D. (2026, April). ParseRAG: Layout-Aware Open-Source Alternative. In 2026 Eighth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 799-804). BM Institute of Engineering & Technology.
- [171] Yadav, P., Raj, M., Kaushik, S., Goel, P., Garg, P., & Tiwari, D. (2026, April). Employing Enhanced LSTM and Large Language Models for Stock Price Prediction Using Sentiment Analysis. In 2026 Eighth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 617-622). BM Institute of Engineering & Technology.
- [172] Gupta, V., Sharma, H., Nagpal, N., Garg, P., & Tiwari, D. (2026). Feel Frames: An Emotion Based Video Summarization using Multimodal Deep Learning. *Procedia Computer Science*, 282, 715-724.
- [173] Singh, H., Ahmed, K., Kumar, U., & Garg, P. (2026, July). Comparative Analysis of Machine Learning and Large Language Model Architectures for Intelligent Medical Diagnosis and Clinical Decision Support Systems. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 125-138).
- [174] Singh, H., Ahmed, K., Prajapati, P., & Garg, P. (2026, July). Federated Learning with Differential Privacy: A Comprehensive Framework for Privacy-Preserving Distributed Machine Learning. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 139-155).
- [175] Singh, H., Ahmed, K., & Garg, P. (2026, July). Survey on Internet of Things (IoT) Forensics: Challenges, Approaches, Privacy-Aware Frameworks and Open Issues. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 156-170).
- [176] Singh, H., Ahmed, K., & Garg, P. (2026, July). Supply Network Disruption Prediction using Heterogeneous Graph Neural Networks and Temporal Embeddings. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 171-186).
- [177] Saraswat, P., Kumari, S., & Garg, P. (2026, July). Future Directions in Digital Social Linguistics. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 187-203).
- [178] Saraswat, P., Yadav, A., & Garg, P. (2026, July). Comparing Supervised and Unsupervised Learning Approaches for Specific Tasks. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 220-236).
- [179] Saraswat, P., Aishwarya, S., & Garg, P. (2026, July). Foundation of Self-Learning AI for Drones. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 237-254).
- [180] Ahmed, K., Singh, H., Anas, M., & Garg, P. (2026, July). The Future of Artificial Intelligence in Healthcare. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 255-270).
- [181] Ahmed, K., & Garg, P. (2026, July). AI-Driven Fault Detection and Diagnostics. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 271-292).
- [182] Kumari, M., Pandey, S., & Garg, P. (2026, July). Quantum AI for Drug Discovery and Material Science. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1).
- [183] Kumari, M., Jha, S. K., & Garg, P. (2026, July). Opportunities and Challenges of AI in Education. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 328-341).
- [184] Pokhrel, L., Kumar, A., & Garg, P. (2026, July). Impact of AI on Employment and Workforce Transformation. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 342-356).
- [185] Singh, H., Ahmed, K., Kumar, A., & Garg, P. (2026, July). AI-Driven Optimization in Supply Chain Operations: A Machine Learning Framework for Demand Forecasting and Inventory Management. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 357-371).
- [186] Saraswat, P., Kumar, R., & Garg, P. (2026, July). Human Perception and Virtual Reality: What Is the Connection? In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 409-427).
- [187] Ahmed, K., Arsalan, M., Sharma, P. K., & Garg, P. (2026, July). Quantum Technology in Medical Diagnosis. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 428-448).
- [188] Saraswat, P., Sharma, P., & Garg, P. (2026, July). The Challenges of AI Implementation in Business. In *Adroid Conference Series: Engineering and Technology* (Vol. 2, No. 1, pp. 204-219).
- [189] Trivedi, S., Beri, T., Arora, A., & Garg, P. (2026, June). NoduleNet: Deep Learning-Based Computational Model for Thyroid Cancer Detection. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-7). IEEE.
- [190] Joshi, V., Tyagi, U., Jakhar, V., Gaur, M., & Garg, P. (2026, June). HealthQ: A Retrieval-Augmented Health Policy Query Chatbot. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [191] Saraswat, P., Sharma, V., & Garg, P. (2026, June). Energy Optimization In Smart IoT Enabled Smart Home System Through Machine Learning Techniques. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [192] Saraswat, P., Tiwari, B., & Garg, P. (2026, June). An Empirical Study on Salary Determinants Using Machine Learning. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-5). IEEE.
- [193] Pokhrel, L., Arsalan, M., Tiwari, B., & Garg, P. (2026, June). Food Delivery Time Prediction Using Machine Learning Algorithms: A Comparative Study with Hybrid XGBoost-LSTM Architecture. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-7). IEEE.
- [194] Garg, P., & Kumar, V. (2026, June). Millimeter-Wave 60 GHz On-Chip Dipole Antenna with Silicon Lens Integration in SiGe Technology. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-5). IEEE.
- [195] Arora, A., Trivedi, S., Tyagi, R., & Garg, P. (2026, June). LungVision: Classifying lung diseases using CNN ensembles. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [196] Arora, A., Trivedi, S., Sharma, R., Tiwari, R., Chaudhary, S., & Garg, P. (2026, June). Comparative Study of ResNet-152 and ViTPose-G for Human Pose Detection on the COCO Yoga Dataset. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [197] Kaur, M., Chhillar, V., Kumar, V., & Garg, P. (2026, June). Exploration and Recognition of Autism Spectrum Disorder through Machine Learning Approaches. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.

- [198] Singh, R., Singh, G., Kumar, V., Kumari, M., Kumar, A., & Garg, P. (2026, June). Optimization of Battery Electric Vehicle Thermal Management System Using MATLAB Simscape for Enhanced Charging Efficiency. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [199] Chauhan, A., Singh, G., Kumar, V., Kumari, M., Kumar, A., & Garg, P. (2026, June). An Intelligent Thermal Control and Optimization Framework for Efficient and Safe Battery Charging in Electric Vehicles. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [200] Kumar, S., Kumar, Y., Khurana, D., Kumar, N., & Garg, P. (2026, June). AI-Driven Discovery of Cybercrime Insights from Diverse Web Data. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [201] Chaudhary, N., Tyagi, N., Sharma, N., Gupta, M., & Garg, P. (2026, June). BusOptima—Bus Scheduling & Route Management System. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [202] Srivastava, A. K., Singh, A., Umrao, A., & Garg, P. (2026, June). A Retrieval-Augmented Framework for Context-Aware Document Querying. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-7). IEEE.
- [203] Gupta, T., Bhargava, D., & Garg, P. (2026, June). A Comprehensive Survey on Malware Botnet Detection Using Deep Learning in IoT and Android Platforms with Edge AI. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-7). IEEE.
- [204] Trivedi, S., Shukla, B., Garg, Y., Satish, T., & Garg, P. (2026, June). Enhancing Social Media Security Through Explainable Ensemble Learning. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [205] Arora, P., Arora, R., Tiwari, D., & Garg, P. (2026, June). Automated Brain Tumor Classification from MRI Images Through InceptionNet Architecture. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-6). IEEE.
- [206] Srivastava, S., Singh, S., Goswami, S., Gupta, M., & Garg, P. (2026, June). SignSpeak: A Real-Time Deep Learning Framework for Bidirectional Sign Language Translation. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-7). IEEE.
- [207] Srivastava, A. K., Garg, P., & Singh, A. (2026, June). ExplainFusion: An Ensemble Explainable AI Framework with Evidence-Grounded LLM-Based Clinical Explanations. In 2026 6th International Conference on Emerging VLSI and Semiconductor Technology for AI and Computing Applications (EVST) (pp. 1-8). IEEE.