



A Prospective Feasibility Evaluation of an AI-Assisted Clinical Decision Support Intervention for Diagnosis and Triage in a Semi-Urban Healthcare Setting

Dr. Vishala Bodetti^{1*}, Nrip Nihalani¹, Aditya Patkar², Dr. Kiranjit Singh², Kishan Gor³, Dr. Praful Tamrakar³, Dr. Naryan Prasad Tamrakar⁴, Dr. Vimal Pumposh⁴, Anvit More⁵, Sourabh Lalwani⁵

¹AI Medical Lead, Plus91 Technologies Pvt Ltd, Pune, Maharashtra, India.

¹Director, Plus91 Technologies Pvt Ltd, Pune, Maharashtra, India.

²CEO, Plus91 Technologies Pvt Ltd, Pune, Maharashtra, India.

²Senior Consultant, Jehangir Hospital, Pune, Maharashtra, India.

³VP, AI Solution Architect, Plus91 Technologies Pvt Ltd, Maharashtra, Pune, India.

³Director, Mahadev Nursing Home, Funda, Chhattisgarh, India.

⁴Managing Director, Mahadev Nursing Home, Chhattisgarh, India.

⁴Medical Director, Plus91 Technologies Pvt Ltd, Pune, Maharashtra, India

⁵Junior AI Engineer, Plus91 Technologies Pvt Ltd, Pune, Maharashtra, India.

How to cite this paper:

Dr. Vishala Bodetti^{1*}, Nrip Nihalani², Aditya Patkar³, Dr. Kiranjit Singh⁴, Kishan Gor⁵, Dr. Praful Tamrakar⁶, Dr. Naryan Prasad Tamrakar⁷, Dr. Vimal Pumposh⁸, Anvit More⁹, Sourabh Lalwani¹⁰, "A Prospective Feasibility Evaluation of an AI-Assisted Clinical Decision Support Intervention for Diagnosis and Triage in a Semi-Urban Healthcare Setting", IJIRE-V7I4-07-12.



Copyright © 2026
by author(s) and
Fifth Dimension
Research

Publication. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>

Abstract: **Background:** In semi-urban and rural India, a severe shortfall of specialist physicians leaves junior doctors and community health workers making high-stakes diagnostic decisions without senior clinical support. AI-assisted clinical decision support systems have been proposed to bridge this gap, yet real-world evidence of their diagnostic performance in resource-limited settings remains scarce. **Objective:** To evaluate the diagnostic accuracy of Bhairav AI against senior clinician judgement in active semi-urban practice, and to assess its utility as an AI-assisted support tool for community health workers operating under medical supervision. **Methods:** A one-month prospective feasibility study at Mahadev Nursing Home, Chhattisgarh, India. Participants were three senior consultants, five junior doctors, and ten community health workers. Bhairav AI was deployed on desktop devices for junior doctors and on mobile devices for CHWs at field and satellite clinics under medical supervision. Primary outcomes were diagnostic concordance with senior consultant assessments, consultation duration, and tool uptake. **Results:** Full diagnostic concordance with senior consultants was achieved in 85.0% of 100 cases, with partial concordance in 8.0% - a combined agreement of 93.0% - and discordance in 7.0%. All 13 critical findings were correctly flagged; no critical condition was missed. Mean consultation time fell from 9.53 (SD 0.78) to 6.19 minutes (SD 0.61), a 35.1% reduction ($p < 0.001$). Junior doctor uptake rose from 25% on day one to 90% by the end of the month. CHWs using the mobile version under supervision reported improved confidence in identifying and communicating urgent presentations. **Conclusion:** Bhairav AI demonstrated clinically meaningful diagnostic concordance and significant workflow gains in real-world semi-urban practice. Junior doctors retained full clinical authority, using the tool as a diagnostic aid; CHWs benefited from structured triage support under medical supervision. These findings justify a larger, multi-site randomised study to confirm diagnostic performance, safety, and cost-effectiveness across a broader range of primary care settings in India.

Key Words: Clinical decision support; Artificial intelligence; Diagnostic triage; Semi-urban healthcare; Primary care; Implementation science.

1. INTRODUCTION

Picture a junior doctor at a semi-urban clinic in India, seeing her fortieth patient before lunch. She has no specialist to call, no time to look things up properly, and a waiting room that keeps filling. The real pressure is not whether to refer -it is whether she is getting the diagnosis right in the first place. A patient walks in with a fever and fatigue. Is it viral? Typhoid? Early TB? Something she cannot afford to miss? She has five or six minutes, a history, an examination, and her own clinical instincts -and that is all. This is not an unusual situation. Close to 60% of India's population lives in semi-urban or rural areas, yet the specialist workforce has never followed them there [1]. What that leaves behind is a generation of young doctors and community health workers who must make the hardest diagnostic calls alone -where the gap between a right diagnosis and a missed one is filled by guesswork, not guidance.

The numbers behind this are stark. India's Ministry of Health data shows a shortfall of over 80 percent of specialist doctors at rural Community Health Centres -surgeons, physicians, paediatricians, obstetricians -positions that simply sit

vacant [2]. In states like Madhya Pradesh and Gujarat, the gap exceeds 88 percent [3]. What this means in practice is that a junior doctor at a semi-urban clinic is not occasionally without specialist backup; she is almost always without it. This is where a clinical decision support tool changes the equation -an approach increasingly endorsed in global digital health strategy [4]. Not as a replacement for clinical judgement [5], but as something closer to a specialist-guided assistant -a system that takes the symptoms in front of her, runs them against a structured diagnostic framework, and gives her a ranked list of what this is most likely to be and what she should do right now. It does not make the decision. It gives her the kind of prompt a senior colleague might offer if one were standing in the room: “have you considered this, have you ruled that out, this pattern needs urgent attention.” For a junior doctor seeing forty patients alone, that prompt can be the difference between a diagnosis made and one missed.

Bhairav AI was built with exactly that clinic in mind. It runs on low bandwidth, it works for users ranging from junior doctors to community health workers with no clinical background, and it gives real-time diagnostic and triage suggestions based on whatever the clinician enters at the point of care. An expert panel had looked at the underlying diagnostic logic and found it sound. But bench-level validation and field performance are different things, and real-world evidence from comparable low-resource deployments remains limited [6,7], and nobody had actually watched Bhairav AI in use at a busy semi-urban facility to see whether it helped, whether clinicians would use it, and whether anything went wrong.

This paper reports a one-month feasibility study designed to start answering those questions. We wanted to know whether the tool shortened consultations, whether its suggestions lined up with what a senior clinician would have done, and what made staff take to it -or not. We went in expecting at least a 10% reduction in consultation time; what we found was considerably more, alongside findings about adoption and safety that we think are worth sharing even at this early stage.

II. MATERIAL AND METHODS

A. Study Design and Setting: This was a prospective, single-centre feasibility study built around an observational before-and-after comparison, carried out at Mahadev Nursing Home in a semi-urban part of Chhattisgarh, India. Our aim was to probe feasibility, uptake, usability, and diagnostic agreement rather than to test comparative efficacy. The facility draws on a catchment of roughly 85,000 people from nearby villages and townships and runs 40 inpatient beds, a basic laboratory, and an outpatient service. We chose it because it typifies semi-urban private care: a steady flow of patients, modest technical infrastructure, and clinical staff at a range of experience levels.

B. Study Duration and Participants: The study spanned one month. Three senior consultants, averaging eighteen years in practice, acted as the reference standard and later joined the validation work. Five junior doctors, each holding a post-graduate diploma with two to four years of experience, formed the main group using the tool. Ten community health workers, trained in basic assessment under national protocols, made up the field-level group. Mahadev Nursing Home does not maintain a registered Institutional Ethics Committee; therefore, formal ethics committee review was not undertaken. The study involved healthcare professionals, anonymized operational data, and did not influence final clinical decision-making, which remained under the treating clinician's authority

C. Intervention Deployment: Bhairav AI is deployed in two types. At field and satellite clinics, CHWs used the mobile version of Bhairav AI to review common presentations -not independently, but under the oversight of supervising clinicians throughout. A desktop version was used by junior doctors and specialists working on the main facility's computers. A pre-structured entry format with age, presenting symptoms, vital signs and relevant history generates diagnostic suggestions and triage categories and presents them in a ranked list with flagging for urgent referral where appropriate. The suggestions are up to the clinician to ignore, but if he or she does, the system prompts him/her to write down the reason for the decision.

D. Evaluation Metrics and Data Collection: Three primary outcomes were tracked: consultation duration (time from the start of the encounter to a documented disposition decision), diagnostic concordance (whether the tool's primary recommendation matched the senior consultant's independent assessment), and tool uptake (the proportion of eligible consultations in which the tool was actively used). Junior doctors recorded their start and end times for each consultation, cross-checked against system-generated timestamps from the Bhairav AI logs. Diagnostic concordance was assessed by three senior consultants who reviewed cases independently, blinded to each other's assessments. Uptake was captured automatically from system logs and tracked on a day-by-day basis. Secondary data came from a structured 17-item feedback questionnaire administered to all participants at the end of the study, covering ease of use, clinical utility, workflow integration, and safety and trust, each rated on a five-point Likert scale. Instances where a junior doctor chose not to follow the tool's recommendation were logged separately, with the clinician required to document their reasoning.

E. Procedure and Training: All participants -including the three senior consultants -completed two structured training sessions of approximately 45 minutes each, co-facilitated by the Plus91 Technology Training Team and institutional staff. The first session covered tool orientation -system navigation, data entry, and understanding the output format. The second session moved to the bedside, using real patient cases to walk participants through how the tool's suggestions mapped onto actual clinical presentations. Junior doctors were trained to use Bhairav AI as a diagnostic aid alongside their own clinical judgement, with emphasis on when to follow, question, or override a recommendation and how to document their reasoning when they chose not to follow it. Community health workers received focused training on accurate symptom entry and on

interpreting the tool's triage output -understanding what each triage level meant, what action it required, and how to communicate it to patients and families. All CHW assessments during the study were conducted under medical supervision, with no autonomous clinical decisions made without oversight from a supervising clinician.

Consultation length before and after use of the tool was compared with paired t-tests between the same junior doctors, with results presented as proportions and 95% confidence intervals. The use of the tool was plotted and summarized on a day-by-day basis as the percentage of cases using the tool. The ease-of-use ratings were reported as medians (interquartile range). Due to the feasibility design and small sample, we did not conduct a formal power calculation. R version 4.0 (R Foundation for Statistical Computing) was used for analyses, with p-values < 0.05 considered to be statistically significant.

III. RESULTS

A. Workflow Efficiency: Timing data came from 100 consultations handled by the five junior doctors. Before Bhairav AI was introduced, the mean consultation time was 9.53 minutes (SD 0.78; range 8–11 minutes). Once the tool was part of the routine, that fell to 6.19 minutes (SD 0.61; range 5–7 minutes) -a reduction of 3.35 minutes per consultation, or 35.1% ($p < 0.001$). (Figure 1) Across 100 consultations, this translated to 335 minutes of aggregate time saved over the study period. The reduction was consistent across triage categories: routine cases fell from 9.6 to 6.2 minutes (35.3%), urgent cases from 9.6 to 6.3 minutes (34.0%). The single emergency case in the dataset dropped from 8 to 5 minutes, though one case is insufficient to draw conclusions. The narrowing of the standard deviation -from 0.78 at baseline to 0.61 with AI assistance - suggests the tool also reduced variability in consultation length, not just the mean. Time savings were most consistent in straightforward presentations where the tool produced a clear primary recommendation; cases involving multiple comorbidities or ambiguous symptom patterns took longer even with assistance. CHWs reported that field assessments also felt faster with the mobile version, though this was not captured with the same rigour as the junior doctor data.




Figure 1. Mean consultation time before and after Bhairav AI ($n = 100$). Error bars show SD.

B. Diagnostic Performance: Across all 100 cases reviewed independently by the three senior consultants, the tool's primary recommendation was an exact match with the consultant's diagnosis in 85 cases (85.0%). (Figure 2) A further 8 cases (8.0%) showed partial concordance -the tool and the consultant landed in the same clinical category but used different diagnostic labels. Taken together, the senior consultants were in full or partial agreement with Bhairav AI's primary diagnosis in 93 of 100 cases (93.0%), indicating that the tool reached a clinically acceptable diagnosis in the large majority of presentations. Seven cases (7.0%) were fully discordant, involving unrelated diagnoses between the tool and the consultant. These discordant cases each reflected either an atypical presentation outside the tool's training distribution, or a situation where the consultant's reading drew on contextual information -occupational history, family background, a clinical detail in the room - that the structured input fields never captured. Triage concordance followed a similar pattern: 78 cases (78.0%) had an exact triage level match, 19 (19.0%) were within one level, and only 3 (3.0%) were more than one level apart. Of the 13 cases in which the senior consultant identified a critical finding, the tool correctly flagged all 13 for urgent or emergency attention - no critical condition was missed. At the individual doctor level, concordance rates ranged from 80% (JD-001, JD-004) to 95% (JD-003), with mean consultation times across the five doctors ranging from 5.8 to 6.5 minutes. In 65 of 100 cases junior doctors fully followed the tool's recommendation; in a further 29 they partially adopted it; in 6 they overrode it entirely, documenting their clinical reasoning each time. (Figure 3)




Figure 2. Diagnostic concordance between Bhairav AI and the senior consultant (reference standard). Exact match in 85% of 100 cases, partial in 8%, discordant in 7%.




Figure 3. Junior doctor engagement with Bhairav AI (uptake, not diagnostic accuracy). Followed in 65% of cases, partly adopted in 29%, and overridden on clinical judgement in 6%.

C. Safety Profile: No adverse event over the month was traced to a tool recommendation or to any delay caused by using it. In the 6 cases where a junior doctor chose not to follow the tool's recommendation, they documented their reasoning as required. Senior consultants reviewed all 6 and found the clinical reasoning sound in each -the overrides reflected genuine judgement about presentations the tool was not well-suited to handle, not carelessness or distrust. The safety profile across the remaining 94 cases was clean, with no critical condition missed and no patient harm attributable to the intervention.

D. Adoption Trajectory: Uptake started slowly: only about a quarter of eligible consultations used the tool in the first two working days, held back by worries about workflow disruption and doubts over how reliable the algorithm really was. The picture changed after three straight days of senior consultants demonstrating and endorsing it at the bedside. By day ten use had reached 85%, and by the end of the study, roughly 90% of consultations involved at least one query to Bhairav AI. Community health workers came on board faster, with 80% using the mobile version by day three, although uptake in remote satellite clinics was capped by unreliable connectivity -areas with weak coverage managed only 65% throughout.

E. User Feedback: Feedback was collected through a structured 17-item questionnaire at the end of the study. Across four domains, junior doctors rated safety and trust highest (mean 4.4), followed by workflow integration (4.15), clinical utility (4.0), and ease of use (3.88). CHWs rated safety and trust at 4.2, workflow at 3.88, ease of use at 3.76, and clinical utility at 3.7. The single highest-rated item across all 15 participants was "I felt safe using this system for patient care" (overall mean 4.5). The two weakest items were system response time (3.1) and understanding why a particular diagnosis was suggested (3.1) -both flagged consistently across user groups and linked to connectivity gaps and the absence of visible reasoning in the output. On clinical utility, junior doctors rated "the recommendations helped me make better diagnostic decisions" at 4.4, the highest single item in that domain. Several noted in open responses that the differential diagnosis list was the tool's biggest practical value, particularly in fever cases where dengue, malaria, and viral fever needed to be distinguished quickly. One junior doctor noted the tool had changed how he structured consultations mentally -he found himself running through the same checklist even on days he did not use the app. CHWs highlighted the confidence the output gave them when communicating referral decisions to families; one described being able to show the triage result on screen as something that made families take the recommendation more seriously. Connectivity was the most cited frustration: two CHWs reported frequent disruptions that prevented use during home visits in remote areas, and one estimated she could use the tool in only half her cases because of network issues. Four of five junior doctors said they would definitely or probably continue using the system; CHW intent was more varied, with those in better-connected areas expressing stronger willingness to continue.

Domain	Metric	Baseline	With Bhairav AI	p-value
Workflow	Mean consultation time (min)	9.53 (SD 0.78)	6.19 (SD 0.61)	<0.001
	Time reduction per consultation	—	3.35 min (35.1%)	—
	Total time saved (100 cases)	—	335 min	—
Diagnostic	Full concordance (SC vs AI)	N/A	85 / 100 (85.0%)	—
	Partial concordance	N/A	8 / 100 (8.0%)	—
	Discordant	N/A	7 / 100 (7.0%)	—
	Triage concordance (exact match)	N/A	78 / 100 (78.0%)	—
Safety	Critical conditions missed	—	0 / 13 (0%)	—
	AI recommendation overridden	—	6 / 100 (6.0%)	—
	Overrides judged clinically sound	—	6 / 6 (100%)	—
	Adverse events attributed to tool	—	0	—

Table 1: Summary of Key Pilot Study Outcomes

IV. DISCUSSION

The central question this study set out to answer was a practical one: can an AI-assisted clinical decision support tool actually help a junior doctor in a semi-urban Indian clinic reach a better diagnosis, faster, without a specialist in the room? On the basis of this one-month feasibility study at Mahadev Nursing Home, the early signals are encouraging. Consultation time fell from 9.53 to 6.19 minutes -a 35.1% reduction -which is no small thing for a facility managing 120 to 150 outpatients a day. The tool's diagnostic recommendations matched the senior consultant's judgement exactly in 85% of cases, with a further 8% showing partial concordance in the same clinical category. These figures are consistent with the 20–40% time savings and broadly comparable diagnostic alignment reported from similar implementation studies in primary care [8,9], though direct comparison across settings and tools remains difficult.

The 85% full concordance rate is a useful starting point, but it deserves careful reading. A senior consultant's assessment is a clinical judgement, not a confirmed diagnosis; true verification requires investigation, imaging, or follow-up over time, none of which fell within this short study. The 7 discordant cases were not random failures -they clustered around atypical presentations and situations where the consultant's reading drew on contextual detail that the structured input fields never captured: an occupational history, a family pattern, something observed in the room. These cases mark the edges of where the algorithm applies rather than flaws in its core logic. What matters more for safety is that all 13 cases with critical findings were correctly flagged for urgent or emergency attention. No critical condition was missed. The 6 cases where junior doctors chose to override the tool's recommendation were reviewed by senior consultants, who judged the clinical reasoning sound in every instance -the overrides reflected genuine judgement, not confusion or distrust.

The adoption trajectory is worth reflecting on. Starting at roughly 25% in the first two days is not a failure -it is what honest deployment looks like. Junior doctors were cautious about an unfamiliar system that interrupted a settled routine [10]. The turning point came not from a manual or a training module, but from senior consultants demonstrating the tool at the bedside with real patients, showing their own thinking alongside the tool's output. That single intervention shifted the group. By day ten, uptake had reached 85%, and by the end of the month roughly 90% of consultations involved the tool. This pattern is well recognised in the implementation literature: in hierarchical settings like healthcare, respected early adopters carry more persuasive weight than any amount of structured instruction [10]. What is notable here is how quickly that credibility transfer happened once it began.

This study adds to a growing body of evidence that the most productive frame for clinical AI is not replacement but collaboration. Junior doctors retained full clinical authority throughout -the tool prompted and ranked, they decided. In the 6 cases where they chose to override, they documented their reasoning, and senior consultants subsequently found that reasoning sound in each case. This is not a minor detail. It means the junior doctors were functioning as genuine clinical gatekeepers, not as passive intermediaries between the algorithm and the patient. The implication is that in resource-limited settings, the value of a decision support tool may depend less on near-perfect algorithmic accuracy and more on having clear human oversight built into how it is used.

Community health workers received less attention in this study than junior doctors, but their experience points to something important. Working under medical supervision at field and satellite clinics, they used the mobile version to review common presentations and interpret triage output -a role quite different from diagnosis but no less consequential. Eight of ten rated the tool positively on safety and trust. The structured output gave them a basis for communicating urgent cases to families, who were reportedly more willing to act when shown a system-generated triage result rather than a worker's personal opinion. That is a meaningful finding in a context where delayed care is often driven not by lack of access but by hesitation. Connectivity remained the hardest constraint: two CHWs reported frequent disruptions that forced them to rely entirely on their own judgement during home visits. Any expansion of this model to field-level use must treat offline

capability as a prerequisite, not a future enhancement.

One practical lesson stands out from this deployment. The technology was the easier part. What made the difference -what turned a slow start into 90% adoption -was senior consultants standing beside the tool and showing junior doctors that its output matched their own thinking. Without that institutional endorsement, the early caution would likely have persisted. Any facility planning a similar rollout should budget as much for change management, clinical credibility-building, and structured training as for the software itself. Tools that are technically sound but organisationally unsupported tend not to survive contact with a busy clinic.

V. CONCLUSION

The gap between a right diagnosis and a missed one is where this tool is designed to operate. In this one-month feasibility study at a semi-urban clinic in Chhattisgarh, Bhairav AI reduced consultation time by 35.1%, achieved full diagnostic concordance with senior consultants in 85% of cases, and missed no critical condition across 100 encounters. Junior doctors used it as a diagnostic prompt rather than a crutch -overriding it six times, on sound clinical grounds each time. Community health workers, working under medical supervision, found structured triage output gave them something they previously lacked: a clear, communicable basis for urgent decisions in the field. None of this settles the question of whether AI-assisted decision support belongs in semi-urban and rural Indian primary care. One site and one month cannot do that. What it does is justify the next step -a larger, multi-site study with a longer window, a broader case-mix, and a proper control arm. If those studies confirm what we have seen here, the implications for the diagnostic capacity of facilities serving the 60% of India's population who live beyond easy reach of specialist care are significant.

Two key limitations of this study should be noted. First, the one-month study window is insufficient to capture seasonal variation in disease patterns, assess long-term adoption behaviour, or determine whether diagnostic performance holds as familiarity with the tool increases over time; future studies should employ longer observation periods to address this. Second, connectivity constraints at remote satellite clinics limited consistent tool use among CHWs during field visits, creating gaps in field-level data that cannot be fully accounted for in the current analysis; offline capability must be treated as a prerequisite in any future deployment at this level of care.

Acknowledgement

The authors wish to express their sincere gratitude to Dr. Prafull Tamarkar and Dr. Narayan Prasad Tamarkar for their invaluable support and guidance throughout the study. Their clinical expertise and institutional backing were instrumental in the successful deployment and evaluation of Bhairav AI at Mahadev Nursing Home. The authors also acknowledge the use of large language models, specifically ChatGPT (OpenAI), Gemini (Google), and Claude (Anthropic), during the preparation of this manuscript. These tools were used solely for grammatical correction, language refinement, and graph design. All study data were collected prospectively from clinical records and user assessments; no data were fabricated, altered, or generated using AI tools. The integrity of the original dataset was maintained throughout.

Declarations

Conflict of Interest: All authors are employed by or affiliated with Plus91 Technologies Pvt Ltd, the developer of the Bhairav AI intervention evaluated in this study. The authors declare this as a competing interest.

References

1. International Institute for Population Sciences (IIPS) and ICF. National Family Health Survey (NFHS-5), 2019–21: India. Mumbai: IIPS; 2021. Available at: <http://rchiips.org/nfhs/>
2. Ministry of Health and Family Welfare, Government of India. Health Dynamics of India (Infrastructure & Human Resources) 2022–23. New Delhi: MoHFW; 2023. Available at: <https://www.mohfw.gov.in/?q=en/reports-23>
3. Ministry of Health and Family Welfare, Government of India. Rural Health Statistics 2021–22. New Delhi: MoHFW; 2022. Available at: <https://ruralindiaonline.org/en/library/resource/rural-health-statistics-2021-22/>
4. World Health Organization. Global Strategy on Digital Health 2020–2024. Geneva: WHO; 2021. Available at: <https://www.who.int/docs/default-source/documents/gsdhdaa2a9f352b0445bafbc79ca799dce4d.pdf>
5. Berner ES, La Lande TJ. Overview of clinical decision support systems. In: Berner ES, editor. *Clinical Decision Support Systems: Theory and Practice*. New York: Springer; 2007. p. 3–22. doi:10.1007/978-0-387-38319-4_1
6. Ndlovu K, Stein N, Gaopelo R, et al. Evaluating the feasibility and acceptance of a mobile clinical decision support system in a resource-limited country: exploratory study. *JMIR Form Res*. 2023;7:e48946. doi:10.2196/48946
7. Xiong S, Lu H, Peoples N, et al. Digital health interventions for non-communicable disease management in primary health care in low- and middle-income countries. *NPJ Digit Med*. 2023;6:12. doi:10.1038/s41746-023-00764-4
8. Moja L, Kwag KH, Lytras T, et al. Effectiveness of computerised decision support systems linked to electronic health records: a systematic review and meta-analysis. *Am J Public Health*. 2014;104(12):e12–22. doi:10.2105/AJPH.2014.302164
9. Kwan JL, Lo L, Ferguson J, et al. Computerised clinical decision support systems and absolute improvements in care: meta-analysis of controlled clinical trials. *BMJ*. 2020;370:m3216. doi:10.1136/bmj.m3216
10. Shabeer Ali Thorakkattil et al. Clinical decision support systems in Indian healthcare settings: benefits, barriers, and future implications. *Healthcare (Basel)*. 2025;13(17):2220. doi:10.3390/healthcare13172220