EXPLICITY HEALTH INTELLIGENCE LAB

AI Assisted Healthcare Navigation

Improving Access, Understanding, and Continuity of Care in India

Author	Explicity Health Intelligence Lab	**Prepared for**	Explicity Technologies
Product Perspective	MedBuddy by Explicity Technologies	**Date**	October 2026
Report Classification	Evidence-informed secondary research	**Evidence Scope**	National (India)

Research note: *This report synthesises selected published sources and develops recommendations for Indian healthcare settings. It does not present original patient data or claim that an AI navigation system has been independently evaluated. The MedBuddy implementation framework describes proposed directions; it does not establish the availability, certification, integration status, or effectiveness of particular product features.*

Abstract

Healthcare navigation is the process of helping people find, understand, and complete the steps required to obtain appropriate care. It includes identifying suitable services, preparing for appointments, understanding administrative requirements, completing referrals, obtaining records, and maintaining follow-up. These tasks become difficult when services are fragmented, information is unclear, and patients must coordinate several institutions while dealing with illness, financial pressure, or unfamiliar procedures.

Artificial intelligence could support navigation by interpreting ordinary-language questions, retrieving verified service information, explaining processes, translating communication, and assisting staff with coordination. Its usefulness depends on the quality of underlying information and the availability of real services. A fluent answer does not establish that a provider is suitable, an appointment is available, or a patient has received appropriate care.

This report examines AI assisted healthcare navigation in India through the combined perspectives of access, digital inclusion, service coordination, patient autonomy, and accountability. It distinguishes administrative assistance from activities that influence clinical decisions. It proposes a model in which AI supports defined tasks, current information is verified through appropriate sources, and human teams remain responsible for clinical judgement and unresolved access barriers.

For Explicity Technologies, the report identifies a research opportunity centred on measurable navigation problems. For MedBuddy by Explicity Technologies, it proposes a phased framework connecting provider discovery, clinic workflows, patient information, referrals, and follow-up. Evaluation should focus on completed care, patient understanding, equitable participation, and total burden rather than conversation volume alone.

Keywords: healthcare navigation, artificial intelligence, India, patient access, care coordination, digital inclusion, referrals, Explicity Technologies, MedBuddy.

Executive Summary

Healthcare systems often expect patients to perform substantial coordination work. People must decide where to seek help, determine whether a facility provides the required service, understand costs, arrange appointments, gather records, and interpret the next step. When several providers are involved, each institution may complete its own task while no one takes responsibility for the overall journey.

This burden matters because healthcare is usually sought under conditions of uncertainty. Patients may be anxious, unwell, unfamiliar with terminology, or concerned about expenses. They may rely on relatives, reception staff, community workers, and informal contacts to make sense of the system. The quality of this assistance can influence whether care is delayed, duplicated, or abandoned.

Patient navigation addresses these barriers through practical support. Published guidance describes navigation as a way to help people move through complex services, particularly where care is fragmented or disadvantaged populations face additional obstacles. However, evidence varies by setting and intervention. Findings from human navigation programmes cannot automatically be transferred to an AI product.

AI assisted navigation should therefore be understood as a proposed support model requiring evaluation. Its most credible early applications involve bounded tasks: explaining clinic procedures, retrieving verified information, helping patients prepare documents, identifying administrative barriers, and supporting follow-up. These functions can be useful without assigning the system independent responsibility for diagnosis or treatment.

The report advances seven conclusions.

First, the unit of success should be a completed navigation task linked to appropriate care. A conversation, recommendation, or booking attempt is an intermediate activity.

Second, the system needs reliable operational information. Provider directories, prices, schedules, service availability, and referral requirements change. AI cannot make outdated information current merely by presenting it confidently.

Third, navigation and clinical decision-making overlap at important points. Matching a patient to a specialty, assessing urgency, interpreting a report, or deciding whether a visit can be postponed may influence clinical outcomes. These functions require a different level of oversight from routine administrative assistance.

Fourth, human support should remain accessible. AI may simplify common questions, but unresolved financial, social, clinical, and logistical barriers often require an accountable person.

Fifth, digital inclusion must be built into delivery. Patients may need assistance, familiar languages, accessible interfaces, or non-digital channels. Independent application use should not become a condition for receiving help.

Sixth, product incentives should be transparent. Sponsored listings, commercial referral arrangements, and platform revenue should not be presented as objective evidence of clinical suitability.

Seventh, research claims should match the evidence. An improvement in response speed does not establish reduced treatment delay. A technically functioning recommendation does not establish patient benefit.

For Explicity Technologies, the recommended direction is to study specific navigation failures and evaluate interventions within real clinical workflows. MedBuddy by Explicity Technologies could provide a setting for testing provider information, appointment coordination, patient summaries, referral support, and follow-up assistance, subject to the capabilities actually developed.

The proposed implementation approach begins with dependable service information and staff-assisted navigation. AI can then be introduced for clearly defined tasks, followed by evaluation of errors, workload, patient understanding, and distribution of benefit. Expansion should follow demonstrated performance.

The intended result is a healthcare journey that patients can understand and complete with less avoidable difficulty, while clinical responsibility remains clear and meaningful access is preserved across different circumstances.

1. Introduction: The Navigation Burden in Healthcare

Healthcare access is often discussed in terms of doctors, hospitals, insurance, or technology. These resources matter, but patients also need to know how to use them. A service can exist without being easy to find, understand, or reach. A patient can have an entitlement without knowing which facility accepts it. A referral can be clinically appropriate without providing enough information to complete the next step.

Navigation is the work that connects these elements. It includes understanding a service, preparing for an encounter, moving between institutions, and maintaining continuity. When this work is poorly supported, patients and families become responsible for resolving gaps that originate within the healthcare system.

The burden is particularly visible when a patient must consult several providers. One clinic may request an investigation, another may require a referral document, and a hospital may use a different registration process. Information may be scattered across paper files, messages, and digital accounts. Repeated explanation becomes part of every encounter.

AI offers a possible way to make some of this work easier. A conversational interface can allow people to describe a practical difficulty without knowing the exact menu or administrative term. It can help staff locate information and produce clearer instructions. However, the value depends on whether the answer is accurate and connected to a usable service.

The central issue is therefore institutional coordination. An AI interface cannot independently resolve unavailable appointments, unaffordable treatment, missing transport, or absent specialists. It can identify and explain these barriers, but an operational response remains necessary.

This report examines how AI might support that response in India. It considers patient experience, provider workflows, data quality, accessibility, financial burden, and oversight. It also develops a proposed direction for Explicity Technologies and MedBuddy by Explicity Technologies.

The purpose is to move from broad enthusiasm about conversational AI to a practical account of what navigation assistance should do, what it should not claim, and how its contribution can be measured.

2. Defining Healthcare Navigation

Healthcare navigation is structured assistance that helps a person progress through the services needed for an identified care objective. The objective may be attending a consultation, completing a clinician-issued referral, obtaining an investigation, accessing a record, or returning for follow-up.

Navigation can be provided by several types of worker. Reception teams, nurses, social workers, community health workers, and trained non-clinical navigators may contribute. Their responsibilities differ according to training, experience, and institutional arrangements. A person helping with appointment logistics should not automatically be expected to answer clinical questions.

The European Observatory’s policy brief describes navigators as supporting access and continuity, with roles ranging from identifying barriers to linking patients with providers. It also recognises differences in implementation and the importance of training and sustainable funding. This provides a useful foundation for defining the human functions an AI tool might support.

Navigation should be distinguished from general information provision. A directory tells a patient that a service exists. Navigation helps determine what the patient needs to do to use that service and whether a barrier prevents completion. The difference is practical responsibility for the next step.

It should also be distinguished from clinical triage. Triage evaluates urgency and appropriate clinical disposition. Some navigation conversations may reveal a need for triage, but the functions should not be silently merged. When the system moves from explaining a process to influencing medical urgency, the risk and required evidence change.

AI assisted navigation is consequently best defined as a combination of information retrieval, communication support, workflow assistance, and appropriate human escalation. It may reduce repetitive tasks while helping teams focus on cases that require judgement or intervention.

The proposed model does not require AI to resolve every question. It requires the system to recognise its task, provide support within that task, and connect the patient to an accountable person when the issue exceeds its scope.

3. Research Objectives and Methodology

This report has four objectives: to explain the navigation problem, identify appropriate AI applications, develop an inclusion-focused implementation framework, and propose methods for evaluating benefit.

The analysis addresses five questions. Where do patients lose their way in a care pathway? Which barriers can information and coordination support realistically reduce? Which proposed functions influence clinical decisions? What operational information must remain current? How can institutions determine whether an intervention improves access equitably?

The report uses a targeted narrative review of official publications and selected policy guidance. WHO materials inform patient navigation and AI governance. Indian government and ICMR sources provide national context. Digital accessibility and public infrastructure sources inform proposed design requirements.

This is not a systematic review or meta-analysis. It does not estimate a pooled effect of AI navigation, compare every available product, or establish national prevalence of navigation failure. No original interviews, experiments, or patient-level analysis were conducted.

The report separates descriptive evidence from proposed design. A cited publication may support the importance of patient navigation or identify a risk of generative AI. It does not establish that a specific proposed MedBuddy feature will deliver the same benefit.

This distinction is especially important when human navigation evidence is used. A human navigator may build relationships, contact institutions, arrange assistance, and recognise contextual problems. An AI system that answers questions performs only part of that work unless connected to appropriate staff and workflows.

The implementation model is therefore presented as a testable proposal. Its components should be assessed against local needs, clinical responsibilities, and available resources. Recommendations are not guarantees of effectiveness.

The evidence review also retains source dates and scope. Older guidance can remain relevant to concepts, but current service information requires separate verification. Provider schedules and appointment availability should come from operational sources, not from the report itself.

The intended use of this document is to guide research and development, establish clearer boundaries, and support more precise decisions about pilot design.

4. Evidence Baseline: What Existing Research Supports

Existing navigation evidence provides a foundation, but its relevance must be interpreted carefully. WHO’s breast cancer technical brief describes navigation as a personalised intervention addressing barriers across detection, diagnosis, and treatment. Its context is cancer care, not an unrestricted validation of general-purpose AI assistance.

A broader policy brief identifies potential contributions to access and continuity while noting limitations in the evidence base. The appropriate conclusion is that navigation can be useful in defined settings and requires contextual implementation. It would be inappropriate to infer a universal effect across every condition or institution.

AI guidance provides a different type of evidence. WHO identifies risks including inaccurate or incomplete output, bias, overreliance, and cybersecurity problems. These concerns establish evaluation priorities; they do not show that every AI application has the same risk profile.

ICMR’s 2023 ethical guidelines address the development, deployment, and adoption of AI in biomedical research and healthcare. Their coverage includes ethical principles, governance, review, and informed consent. They provide an Indian reference point for responsible research, rather than product-specific approval.

Evidence area	Supported interpretation	Interpretation to avoid
Human navigation	Can help address defined access barriers	Every chatbot provides equivalent support
Cancer navigation	Relevant to specified cancer pathways	Outcomes automatically generalise to all care
AI governance	Requires task-specific oversight	A disclaimer alone establishes safety
Digital health guidance	Technology depends on functioning services	Information access replaces service capacity
Public infrastructure	Can support discovery and exchange	Integration guarantees completed care

For Explicity Technologies, this evidence baseline suggests a disciplined strategy. Begin with a clear barrier, identify the function needed to address it, and evaluate whether the implemented service achieves that function.

The research opportunity lies in the gap between promising capability and demonstrated benefit. Carefully designed pilots can show which parts of navigation are improved by AI, which still require people, and what resources are necessary to connect the two.

5. Why Navigation Matters in Indian Healthcare

Navigation difficulties can arise wherever services operate through separate processes. A patient may need to move between a local clinic, a diagnostic provider, a hospital, a pharmacy, and a follow-up service. Each transition can introduce uncertainty about documents, costs, timing, and responsibility.

India requires an approach capable of accommodating different languages, geography, provider types, and levels of digital confidence. A process suitable for a large urban hospital may be impractical in a small clinic or remote community. Navigation should therefore be designed around local service conditions.

The report does not assume that every patient encounters the same problems. Some need only a clear appointment message. Others need assistance with transport, financial support, communication, or several institutions. A uniform chatbot flow can miss these differences.

A useful navigation service should identify the actual barrier. “I could not attend” may reflect an unavailable appointment, an unaffordable journey, unclear instructions, or a need for caregiving support. Each explanation suggests a different response.

The same principle applies to providers. A missed referral might result from an incomplete document, lack of capacity at the destination, or a process that patients do not understand. Sending additional reminders will not resolve every cause.

AI can help organise such information, but institutions need a way to act on it. If repeated questions reveal that a service listing is inaccurate, the listing should be corrected. If patients repeatedly fail a referral, the pathway should be reviewed.

For MedBuddy by Explicity Technologies, the opportunity is to connect individual questions with clinic-level learning. A navigation system should not merely answer the same confusing question indefinitely; it should help identify why the confusion occurs.

The ultimate goal is a system that becomes easier to use over time. Navigation support should improve the patient journey and reveal operational problems that deserve institutional correction.

6. Mapping the Complete Patient Journey

A navigation model should cover the journey rather than a single conversation. The relevant stages include recognising a need, finding a service, preparing, attending, receiving instructions, completing the next step, and maintaining follow-up.

At service discovery, the patient needs accurate information about what a provider offers. At preparation, they need to know what documents or prior records are relevant. At attendance, they may need registration assistance and clear directions. After consultation, they need an understandable plan.

Each stage should have a defined completion point. Discovery is complete when the patient understands suitable options within the service’s scope. Booking is complete when an appointment is confirmed, not when a request is submitted. Referral is complete when the required receiving service is reached, not when a message is sent.

This report proposes recording pathway status in a way that is useful to authorised staff. The record should distinguish waiting for information, waiting for provider confirmation, awaiting patient action, and requiring human assistance. These states should not be inferred solely from conversational tone.

Patients should also be able to correct the pathway. They may already have completed a test elsewhere, changed their preferred provider, or no longer need administrative assistance. The system should not continue pushing an outdated

sequence.

A complete journey includes unsuccessful attempts. If a patient abandons registration, cannot afford a visit, or fails to receive a result, the experience should remain visible where appropriate for service improvement. Standard analytics often capture only completed steps.

For MedBuddy, the proposed journey model should begin with a small number of clearly defined pathways. Broad coverage can be added after the operational requirements are understood.

The value of this approach is shared responsibility. It helps staff see what remains unresolved and prevents the patient from becoming the only person who understands how several separate transactions fit together.

7. Navigation, Clinical Advice, and Urgency

The boundary between navigation and clinical advice is sometimes clear and sometimes difficult. Explaining clinic hours is administrative. Interpreting symptoms to decide whether a person should wait, attend a particular specialty, or avoid a visit can influence clinical outcomes.

This report recommends classifying functions by their actual effect rather than by the product label. Calling a system a “navigator” does not remove clinical risk if its recommendations affect urgency or treatment decisions.

A practical classification has three levels. Routine administrative tasks include locations, documents, and confirmed booking information. Context-sensitive tasks include explaining a clinician-issued referral or organising follow-up logistics. Clinical functions include assessing urgency, interpreting investigations, and recommending treatment.

The higher-risk functions require appropriate clinical governance and evidence. The report does not prescribe a universal triage algorithm or symptom list. Such a process should be developed and evaluated by qualified teams for the intended population and setting.

Potentially urgent concerns should not be drawn into a prolonged administrative conversation. The service needs an appropriate escalation process and should avoid implying that an AI exchange has ruled out a serious problem.

Human support must also be described accurately. If immediate clinical review is unavailable, the interface should not present a handoff as though it has occurred. Patients need truthful information about what assistance is available and when.

For Explicity Technologies, documenting scope is a foundational development task. Each feature should specify its intended function, the decisions it can influence, and the circumstances in which it stops or escalates.

A clear boundary protects both usefulness and accountability. The system can provide substantial navigation support without pretending that ordinary-language fluency is equivalent to clinical assessment.

8. Understanding Intent Without Overinterpreting the Patient

AI can help interpret ordinary-language requests, especially when patients do not know formal service names. A person may ask where to collect a report, how to change an appointment, or what a referral document means. The system should identify the practical task before collecting additional information.

Intent interpretation should remain modest. A brief statement may be ambiguous, and the system should not convert uncertainty into a confident assumption. “I need a heart doctor” could refer to an existing referral, a general concern, or a request made for someone else.

This report recommends asking only clarifying questions that materially affect the task. Excessive questioning can delay completion and collect unnecessary health information. The objective is to resolve the user’s practical need, not to construct a comprehensive medical history during every interaction.

Clarification should also allow correction. Patients may use local terminology or mixed-language phrases. The system should confirm important details in understandable language rather than silently forcing the request into a predefined

category.

The design should distinguish inferred information from confirmed information. A guessed location, specialty, or relationship should not become a permanent record as though the patient supplied it explicitly.

For MedBuddy by Explicity Technologies, intent handling could support clinic reception and service discovery. Evaluation should examine whether the interpretation leads to correct task completion, not merely whether the category appears plausible.

Research should include incomplete statements, misspellings, and ordinary conversational language. Testing only well-formed English questions would provide an unrealistic picture of performance.

The best intent system reduces the knowledge required to use healthcare services while preserving uncertainty where needed. It should help users express their objective and connect them to the appropriate process without inventing facts about their circumstances.

9. Verified Provider and Facility Information

Provider discovery depends on the quality of the directory. AI may make search easier, but it cannot independently verify every listing. Names, locations, contact details, schedules, services, and affiliations should come from identifiable sources.

This report proposes separating different verification claims. Confirming a contact number is not the same as confirming professional registration. Confirming registration is not the same as assessing clinical quality. A listing should communicate the scope of what has been checked.

Service-level information matters because a facility’s name does not establish what is available at a particular time. A hospital may provide a specialty on selected days or require appointments through a separate process. Navigation should avoid presenting general capability as current availability.

Information should include a last-confirmed date where useful. Stale records should be flagged for review, and patients should have a simple way to report discrepancies. The organisation needs a responsible process for investigating those reports.

The report recommends using authoritative registries and direct provider confirmation where appropriate, while recognising that each source has limits. Data sources should be combined carefully rather than treated as interchangeable.

For MedBuddy, a directory should support operational decisions. A patient looking for a service needs to know whether the information is confirmed, whether booking is possible, and what action is required.

The system should not invent a provider because the language model expects one to exist. When no verified result is available, it should communicate that limitation and provide an appropriate next step.

Directory quality is a continuing service function. It requires maintenance, correction, and accountability. Without those functions, conversational search can make unreliable information easier to believe and act upon.

10. Recommendation, Ranking, and Patient Choice

A navigation service may present several providers or facilities. The ordering of those options can influence decisions, so ranking should be transparent and appropriate to the task.

Distance, service availability, accessibility, language, cost information, and patient preference may all matter. No single factor establishes the “best” provider for every person. A system should therefore explain the criteria it uses rather than produce an unexplained universal ranking.

Clinical suitability requires particular care. Where the patient already has a clinician-issued referral, the system can help identify providers offering the specified service. Where the request requires medical judgement, the service should not disguise an unsupported inference as objective matching.

Sponsored placement should be clearly identified. Payment to the platform should not be presented as evidence of better care. Ratings and testimonials also need context, because they can reflect experience, convenience, or popularity rather than clinical outcomes.

This report recommends offering meaningful alternatives. A patient may prefer a closer public facility, a particular language, a familiar clinic, or a provider available outside working hours. The system should support these preferences where reliable information exists.

For Explicity Technologies, ranking rules should be documented and reviewed for unequal effects. A model trained mainly on already visible providers may repeatedly favour them while overlooking smaller or less digitised services.

Evaluation should examine whether users understand why an option was suggested. It should also assess whether the suggestion leads to a completed, appropriate service encounter.

The aim is informed choice. Navigation should reduce the difficulty of comparing relevant options without claiming certainty that the evidence cannot support.

11. Current Availability and the Difference Between Search and Booking

Finding a provider is different from obtaining an appointment. A directory can identify a service, but current availability requires an operational source such as a provider schedule or booking system.

The National Health Authority describes UHI as infrastructure supporting health-service discovery and related transactions. This establishes a relevant ecosystem direction, but it does not mean that every service or provider participates, or that a particular product has completed integration.

This report recommends making booking status explicit. A slot displayed, a request submitted, a provider acknowledgement, and a confirmed appointment are different states. Patients should know which state applies.

Availability should not be generated from general knowledge. If the system cannot access current information, it should state that confirmation is required. A confident estimate can create unnecessary travel and lost time.

Booking actions should also reflect patient choice. The system should present relevant details before creating or changing a commitment. Users need to understand the provider, location, timing, and any known payment terms.

Cancellation and rescheduling require clear procedures. An apparent success in the interface should correspond to a change in the provider’s actual schedule. If confirmation is delayed, the patient should receive accurate status information.

For MedBuddy by Explicity Technologies, transactional integration should be evaluated separately from conversational performance. The assistant may explain a request correctly while the booking operation fails.

Testing should include duplicate requests, interrupted sessions, time changes, and provider rejection. Recovery should be understandable and should avoid creating multiple unintended appointments.

The goal is a dependable transaction. AI can simplify how the patient expresses the request, but the booking system must remain the source of truth for whether the appointment exists.

12. Explaining Administrative Requirements

Administrative uncertainty can prevent care even when a service is available. Patients may not know whether they need a referral, identification, previous reports, payment confirmation, or another document.

AI can support this task by retrieving current requirements and explaining them in familiar language. The information should be specific to the facility and service. A general answer about hospitals may be insufficient for a particular department.

This report recommends distinguishing mandatory requirements from helpful preparation. Patients should not be told that an optional document is essential merely because it is commonly useful. Such errors can lead to cancelled visits or repeated

travel.

Instructions should prioritise what the patient must do next. A long explanation of the entire institutional process may obscure the immediate task. The system should offer additional detail when needed without overwhelming the initial response.

Requirements may differ for new patients, follow-up visits, and specific programmes. The system should confirm the relevant category rather than assume a uniform process.

For MedBuddy, clinic staff should be able to update approved administrative information through an accountable process. The system should retain enough history to explain which version was used when a patient received instructions.

Evaluation should include patients who arrive without expected documents. The question is whether the system provided accurate guidance and whether an alternative process was available.

Administrative support should reduce unnecessary friction. It should not add requirements that the institution does not actually impose or collect sensitive information simply because a form allows it.

A useful assistant makes the process easier to understand while keeping responsibility for the requirements with the institution that establishes them.

13. Cost Information and Financial Navigation

Patients need clear information about the financial consequences of obtaining care. A navigation service can help explain published charges, identify support channels, and prepare questions for a provider. It should not invent a guaranteed total from incomplete information.

Costs can include consultation, investigations, medicines, transport, accommodation, and lost earnings. The relevant amount depends on the care actually required. A listed consultation price is therefore only one part of the episode.

India’s National Health Accounts for 2022–23 reported household out-of-pocket expenditure at 43.4% of total health expenditure. This provides context for financial navigation, but it does not describe the expenses of a particular patient or service.

This report proposes labelling cost information by its source and scope. Published consultation fees, estimated ranges, and confirmed quotations should remain distinguishable. Unknown charges should not be silently treated as zero.

Financial assistance requires similar precision. Explaining an eligibility process is different from determining eligibility, approving a benefit, or guaranteeing a claim. The relevant authority should remain clear.

For MedBuddy, an assistant could help patients identify what information they need from a clinic or assistance programme. It could organise known charges and highlight items requiring confirmation.

Evaluation should assess whether patients understand the limits of the information. A reassuring but incomplete estimate can be more harmful than a clear statement that further confirmation is needed.

The report also recommends tracking whether financial barriers remain after navigation. Better information does not create affordability. Cases requiring assistance should connect to an available human or institutional process.

Financial navigation succeeds when it reduces surprise, helps people understand options, and supports completion of care without misrepresenting benefits or expenses.

14. Language, Voice, and Mixed-Language Communication

A conversational interface should accommodate the language people actually use. In India, this may include regional languages, English, transliteration, and mixed-language speech. A patient should not need formal medical vocabulary to ask where to go or what to bring.

Language support should be evaluated through understanding, not the number of languages listed in product material. A translated answer may be grammatically correct but still unclear, especially when it contains unfamiliar administrative or clinical terms.

This report recommends designing messages around the task. Appointment instructions should clearly identify time, place, and preparation. Referral explanations should preserve the clinician's meaning. Financial information should distinguish confirmed amounts from uncertainty.

Voice can help users who find reading or typing difficult, but it is not automatically inclusive. Speech recognition may struggle with accents, background noise, or unfamiliar names. Audio output may also expose private information in shared spaces.

Important details should be confirmed in an appropriate format. Dates, locations, and provider names may require repetition or a written summary. The patient should be able to correct the system without starting the entire process again.

For Explicity Technologies, language evaluation should include real clinic environments and local expressions. Testing only clean recordings or formal sentences would underestimate practical difficulty.

MedBuddy could support language preferences across the pathway, so that an initial conversation does not end in an inaccessible document. Follow-up messages and visit summaries should be considered as part of the same communication system.

The objective is reliable understanding. Language technology should help patients express needs and complete tasks while making uncertainty visible when the system has not understood an important detail.

15. Accessibility and Assisted Participation

Accessibility should cover the entire navigation process. A readable chat screen does not establish that booking, payment, document access, and account recovery are usable.

W3C's accessibility framework emphasises content being perceivable, operable, understandable, and robust. These principles provide a useful basis for interface design, while healthcare tasks still require testing with intended users.

This report recommends including people with different disabilities, ages, literacy levels, and device access in evaluation. Users should be able to receive assistance without being treated as unsuccessful participants.

Assistance may come from reception staff, a navigator, a caregiver, or a community worker. The role should be clear, and the patient should retain appropriate choice over what information is shared.

Alternative channels should remain practical. If a person cannot use an application, the service should provide a route that leads to equivalent assistance rather than simply directing them back to the same interface.

Time limits and error recovery deserve attention. Patients may need additional time to read, use assistive technology, or consult another person. A process that repeatedly expires can create unnecessary frustration.

For MedBuddy by Explicity Technologies, accessibility should be a documented development requirement. General claims such as "for everyone" should be supported by testing and accompanied by clear information about available assistance.

Evaluation should also examine whether the assisted process creates extra waiting or stigma. An alternative route that is technically available but substantially more difficult may still produce inequity.

The goal is effective participation. AI can provide a flexible interface, but inclusion depends on the broader service arrangement and the ability to complete the necessary care task.

16. Rural Navigation and Local Service Context

Rural navigation requires accurate knowledge of service location, operating arrangements, transport, and referral options. Digital information can help, but its usefulness depends on the local situation.

A patient may need to know whether a service operates on a particular day, whether a sample can be collected nearby, or whether a referral facility requires advance booking. These questions should be answered from verified local information.

Connectivity should influence delivery design. A low-bandwidth message may be more dependable than a complex interface. Assisted access at a local facility may be more useful than expecting every patient to manage remote communication independently.

This report recommends planning navigation together with the physical care pathway. If a remote consultation results in an in-person referral, the patient needs a specific next step. The system should not stop at producing a recommendation.

Local workers can provide context that a model lacks. They may know that a route is disrupted, a clinic is temporarily closed, or a patient needs help arranging transport. The service should support updating this information responsibly.

For Explicity Technologies, rural pilots should include maintenance and support arrangements. A device or connection problem should not leave a clinic dependent on distant technical assistance with no practical fallback.

MedBuddy could support a focused local directory and staff-assisted coordination before attempting broad autonomous navigation. A smaller dependable service may provide greater value than an extensive but unreliable catalogue.

The measure of success should include completed referrals, patient travel burden, and staff experience. A remote answer that is convenient to generate but impractical to act upon does not improve access.

Rural navigation works when digital support fits local capacity and institutions remain responsible for the patient’s next necessary step.

17. Urban Navigation and Fragmented Provider Markets

Urban healthcare can be difficult to navigate despite the concentration of services. Patients may face numerous providers, inconsistent information, unclear costs, and separate administrative processes.

More options do not always mean easier choice. A patient may struggle to distinguish a clinic offering routine care from a facility equipped for a specific referral. Online visibility can influence selection without providing reliable evidence of suitability.

This report recommends presenting options according to defined criteria. Service capability, confirmed availability, travel considerations, accessibility, and patient preference should be communicated transparently.

Work schedules also matter. A digitally confident person may still be unable to attend during standard hours. Navigation should consider timing and the practical consequences of repeated visits.

Fragmented records create another burden. Patients may carry information between small clinics, laboratories, and hospitals. An assistant can help organise documents, but the receiving provider must be able to use them appropriately.

For MedBuddy, urban pilots could examine how patients move across several participating services. The research should identify where information or responsibility is lost, rather than measuring only activity within one clinic.

Commercial ranking requires particular attention in dense markets. Sponsored results should be clearly distinguished, and providers with less digital visibility should not be excluded without explanation.

The report also recommends learning from failed searches. If users repeatedly request a service unavailable in the directory, the organisation should review coverage and communication.

Urban navigation should reduce the burden of comparing and connecting services. Its value lies in helping people complete an appropriate pathway with fewer unnecessary steps, rather than merely presenting a longer list of providers.

18. Referrals and Closed-Loop Coordination

A referral should connect clinical intent with an actual receiving service. In a weak pathway, it becomes a document or instruction that the patient must interpret and fulfil independently.

This report proposes closed-loop referral coordination. The process should distinguish a referral issued, information received, an appointment arranged, the patient attending, and an appropriate response returning to the referring team.

AI can help explain a clinician-issued referral, identify administrative requirements, and organise status information. It should not silently alter the clinical purpose or replace the responsible clinician’s judgement.

The receiving institution needs a role in the workflow. A system cannot confirm referral completion solely because the patient opened a message. Operational confirmation should come from an appropriate source.

Patients may encounter barriers between issuance and attendance. Cost, transport, availability, caregiving, and unclear instructions can all matter. The response should address the actual barrier rather than assume that additional reminders are sufficient.

For MedBuddy by Explicity Technologies, a referral pilot should begin with a defined network and service. This allows responsibilities, information requirements, and completion measures to be tested before broader expansion.

The system should also support back-referral or follow-up planning. After a specialist assessment, the patient may need continuing care closer to home. The referring team should receive relevant information through an appropriate process.

Evaluation should measure timely completion and patient understanding, while considering whether the referral was clinically appropriate. A high completion rate is not necessarily beneficial if the pathway encourages unnecessary services.

The purpose is shared responsibility. Closed-loop coordination makes it clear which institution must act next and reduces the amount of uncertainty transferred to patients and families.

19. Follow-Up and Long-Term Care Navigation

Follow-up is a continuing process rather than a sequence of messages. Patients need to understand why a return visit is required, when it should occur, and what to do if they cannot attend.

AI can support reminders, rescheduling, and explanation of an existing plan. It can also help staff identify unresolved administrative tasks. However, it should not independently decide that follow-up is unnecessary or change treatment instructions.

This report recommends distinguishing communication delivery from completion. A message sent, a message opened, and a visit attended are different events. Non-response should not automatically be interpreted as refusal.

The service should ask about communication preferences. Shared phones, work schedules, privacy concerns, and accessibility needs can influence whether a reminder is useful.

Missed follow-up requires proportionate action. Some situations may be resolved through rescheduling, while others need clinical review or practical assistance. The escalation process should reflect the institution’s responsibilities and available capacity.

For MedBuddy, follow-up support could be linked to the clinician-approved plan. The system should preserve who created the instruction and when, so that later changes are not confused with the original recommendation.

Research should assess whether the intervention improves appropriate attendance and continuity. It should also examine burden: repeated notifications can become intrusive, and additional visits can create costs.

Long-term care may involve several conditions and providers. Navigation should help coordinate the combined plan where possible rather than generate separate, conflicting reminders for each programme.

The goal is sustained access to appropriate care. Follow-up assistance should make the plan easier to maintain while preserving clinical judgement and respecting the patient's circumstances.

20. Health Records and Navigation Support

Relevant records can help a patient prepare for a visit and reduce repeated explanation. An assistant may help locate documents, organise them by date, or produce an administrative checklist.

The quality of the underlying information remains crucial. A scanned document may be incomplete, an extracted name may be incorrect, or a historical prescription may no longer represent current treatment. The system should preserve source and date.

This report recommends distinguishing retrieval from interpretation. Finding a report is a navigation task. Explaining its clinical meaning or deciding what action it requires may be a clinical function.

Summaries should be designed for their purpose. A referral summary needs different information from a patient's appointment preparation note. Including every available detail can obscure what matters.

Patients should have a way to report errors. Correction should follow an appropriate review process rather than allow unverified alterations to overwrite clinical records.

For MedBuddy by Explicity Technologies, record assistance should respect existing clinic authority. The system should not extract and store information beyond what is necessary for the agreed navigation task.

Record portability is also important across providers. A patient moving between institutions needs a practical way to present relevant information without depending on an active subscription.

The objective is useful preparation. Records should help patients and clinicians connect care across encounters, while keeping interpretation and decision-making within appropriate clinical boundaries.

21. Privacy, Data Protection, and Health Information Security

Healthcare navigation frequently involves sensitive personal context: symptoms, medical specialties, household dynamics, travel patterns, and financial constraints. Safeguarding this data is both an ethical imperative and a legal obligation under India's Digital Personal Data Protection Act (DPDPA 2023).

Navigation platforms must separate operational metadata from clinical content. An appointment scheduler needs only date, time, facility, and department. It does not require continuous diagnostic history or unredacted medical notes. Systems should enforce purpose limitation and avoid compiling unnecessary longitudinal dossiers during routine navigation interactions.

Consent should be explicit, granular, and easily revocable. A patient consenting to appointment reminders should not be treated as consenting to commercial marketing or unconsented algorithmic training. For vulnerable populations using assisted community devices, consent explanations must be provided in understandable vernacular language without coercive legal jargon.

Data retention schedules should be strict. Transitory navigation logs (such as session audio recordings, address lookups, and temporary OTPs) should be purged automatically after task completion. Audit logs tracking data access must be tamper-evident and accessible to institutional compliance officers.

For MedBuddy by Explicity Technologies, data architecture should embody privacy-by-design. Encryption at rest (AES-256) and in transit (TLS 1.3), role-based access control, and zero-knowledge storage for non-clinical navigation intermediaries ensure that patient trust is systematically preserved.

22. Generative AI Risks: Hallucination, Stale Knowledge, and Safety Guardrails

Generative large language models present unique risks when applied to healthcare navigation. While their conversational fluency makes them accessible, models can generate plausible-sounding but entirely fictitious clinic addresses, non-existent doctor affiliations, expired fees, or erroneous diagnostic timelines.

ICMR's 2023 Ethical Guidelines for AI in Healthcare emphasize validation, algorithmic transparency, and avoidance of harm. In navigation systems, an inaccurate recommendation (such as directing an emergency trauma patient to an outpatient-only dispensary) can produce fatal delays.

To mitigate these risks, generative models should never operate on unrestricted parametric memory for factual healthcare directories. Navigation must rely on Retrieval-Augmented Generation (RAG) strictly grounded in verified institutional databases, real-time registry APIs, and confirmed facility directories.

Furthermore, defensive guardrails must intercept clinical hallucinations. If a model attempts to generate unsolicited diagnostic opinions, medication dosages, or prognosis estimates within an administrative workflow, deterministic safety filters must block the output and route the inquiry to an approved clinical channel.

Explicity Technologies adheres to strict benchmarking: any automated navigation pipeline must maintain zero factual deviation on verified provider metadata and enforce deterministic fallback whenever source confidence drops below verified thresholds.

23. Human-in-the-Loop Architecture and Escalation Pathways

An AI navigation assistant is only as effective as the human support system behind it. Complex healthcare challenges (unmet financial need, insurance denials, rare clinical presentations, severe emotional distress, or sudden symptom deterioration) cannot be resolved through automated conversation alone.

A robust human-in-the-loop (HITL) framework requires deterministic escalation triggers. These triggers should activate based on:

1. Detection of red-flag symptoms (chest pain, dyspnoea, acute neurological deficit, severe trauma, suicidal ideation);
2. Repeated user task failure or expressed confusion;
3. Administrative exceptions (unregistered migrant workers, missing mandatory identity documents);
4. Explicit patient request for a human coordinator.

When an escalation occurs, the transition must be seamless. The AI system must compile a structured handoff brief (containing the identified goal, verified details, and unresolved obstacle) for the receiving human navigator or clinic receptionist, avoiding the frustration of repeated patient interrogation.

Escalation must also be honest. If human staff are offline or unavailable, the system must clearly state operating hours and provide verified emergency contacts rather than creating false expectations of immediate review.

For MedBuddy, human coordinators remain the ultimate anchor of patient care, with AI serving as an administrative amplifier that relieves routine friction while elevating human empathy where it matters most.

24. Frontline Clinic Workflows and Staff-Assisted Navigation Desks

Frontline clinic staff (receptionists, registration clerks, and triage nurses) operate under intense time pressure and high patient volume. In many Indian public facilities and busy charitable trust hospitals, a single receptionist may manage queues exceeding 200 patients per hour.

Poorly integrated technology can exacerbate this bottleneck. If staff must manually cross-reference multiple disjointed portals or input duplicative demographic entries, processing times lengthen and queue tensions escalate.

AI assisted navigation should therefore be designed as a dual-facing workflow tool:

- **Patient-Facing:** Enabling mobile pre-registration, vernacular document preparation, and directional wayfinding before facility arrival.
- **Staff-Facing:** Providing reception desks with an intelligent intake dashboard that pre-validates registration details, surfaces missing documents, auto-populates standardized forms, and suggests appropriate department routing based on clinician referrals.

In assisted clinic desk deployments, trained navigators or community volunteers can operate digital kiosks on behalf of non-literate patients. The navigator verifies the patient's identity token, prints an understandable paper action summary, and directs them to the correct queue.

MedBuddy's clinic workflow module aims to reduce average intake wait times from over 20 minutes to under 4 minutes, transforming chaotic reception lobbies into structured, predictable care pathways.

---

25. Community Health Workers, ASHAs, and Decentralised Navigation

In rural and peri-urban India, Accredited Social Health Activists (ASHAs) and Auxiliary Nurse Midwives (ANMs) are the primary human bridge between households and formal healthcare institutions. They understand household social dynamics, cultural reservations, and local geographic hurdles that no remote digital platform can perceive.

However, community health workers are frequently overburdened with fragmented reporting registers, manual paper logs, and uncoordinated referral tracking. When an ASHA refers a high-risk pregnant woman or a tuberculosis suspect to a Community Health Centre (CHC), she rarely receives structured confirmation of whether the patient completed the visit.

AI assisted navigation can empower community health workers by:

- Providing low-bandwidth mobile tools that translate complex institutional referral guidelines into regional voice checklists;
- Automating status notifications when a referred village patient checks into a district hospital;
- Supplying verified logistical details (transport subsidies, ambulance helpline status, specialist clinic days);
- Supporting local language communication with attending district medical officers.

Rather than replacing community health workers with consumer-facing chatbots, Explicity Technologies champions an **assisted intermediary model**: placing assistive AI into the hands of ASHAs and rural clinic staff to magnify their reach and systemic authority.

---

26. Proposed Implementation Framework for MedBuddy by Explicity Technologies

MedBuddy by Explicity Technologies is conceptualised as an end-to-end healthcare navigation and clinic enablement architecture designed specifically for the Indian healthcare ecosystem. It addresses the systemic fragmentation between patient discovery, facility arrival, clinical consultation, and downstream recovery.

The MedBuddy navigation architecture incorporates four interconnected functional modules:

```

| MEDBUDDY AI HEALTHCARE NAVIGATION ENGINE |
|---|
| |
| 1. Verified Directory & Discovery Bridge |
| - ABDM Health Facility Registry (HFR) & Healthcare Professional |
| Registry (HPR) Synchronisation |
| - Operational Schedule Verification & Transparent Non-Sponsored Rank |
| |
| 2. Dual-Interface Intake & Clinic Flow Controller |
```

- Vernacular Conversational Assistant (Text + Voice in 12 Languages)
- Frontline Staff Assisted Desk Terminal & Smart Token Dispatch
3. Closed-Loop Care Pathway & Referral Coordinator
- Machine-Readable Referral Handoffs (FHIR / ABDM Milestones)
- Automated Provider-to-Provider Notification Loop
4. Patient-Centric Action & Follow-Up Engine
- Plain-Language Vernacular Visit Summaries (Printed + WhatsApp)
- Neutral, Privacy-Preserving Follow-up Reminders

```

Every MedBuddy interaction is bound by strict ethical principles: zero unsolicited clinical diagnosis, deterministic fallback to human staff, complete data sovereignty for patients, and transparent pricing disclosures.

## 27. Universal Health Interface (UHI) and ABDM Integration Architecture

The National Health Authority's (NHA) Ayushman Bharat Digital Mission (ABDM) and Universal Health Interface (UHI) represent India's foundational digital public infrastructure for healthcare. UHI establishes open, interoperable protocols enabling digital health applications to discover providers, schedule teleconsultations, and exchange health records seamlessly.

MedBuddy's navigation framework is architected to align with ABDM milestones and UHI specifications:

- **ABHA Identity Integration:** Facilitating seamless creation and linking of Ayushman Bharat Health Accounts (ABHA) without making ABHA possession a mandatory prerequisite for basic emergency care.
- **HFR and HPR Mapping:** Querying verified facility and professional registries to guarantee that suggested clinics and doctors maintain valid public accreditation.
- **Consent Manager Interoperability:** Integrating with ABDM-compliant Consent Managers to allow patients to share historical diagnostic reports with receiving referral facilities on a time-limited, purpose-bound basis.
- **Open Network for Healthcare Transactions:** Utilizing Beckn-based UHI protocols to submit appointment bookings, receive confirmation webhooks, and coordinate closed-loop referral transfers across heterogeneous hospital management information systems (HMIS).

By anchoring navigation to open public infrastructure rather than closed commercial silos, MedBuddy ensures that patients remain free to choose any empaneled provider without proprietary vendor lock-in.

## 28. Commercial Ethics, Commercial Bias, and Non-Sponsored Transparency

A persistent ethical failure of commercial consumer health portals in India is the covert monetization of search results and patient routing. When platforms accept commissions for directing patients to private specialty hospitals, expensive diagnostic chains, or branded pharmacies, clinical objectivity is severely compromised.

Patients with limited medical literacy cannot easily distinguish an unbiased clinical referral from a sponsored marketing advertisement. This dynamic fuels overtreatment, inflates out-of-pocket expenditure, and erodes public trust in digital health.

Explicity Technologies establishes four non-negotiable ethical commitments for MedBuddy:

1. **Absolute Prohibition of Paid Placement:** No healthcare provider, hospital group, diagnostic laboratory, or pharmaceutical brand can pay to influence search ranking or navigation recommendations.
2. **Transparent Routing Criteria:** Whenever providers are displayed, the ranking criteria (distance, public/private designation, confirmed specialist schedule, accessibility accommodations, published fees) must be explicitly visible to
```

the user.

3. Public Facility Parity: Government primary health centres, community health centres, and district civil hospitals must be indexed with equal prominence alongside private institutions.

4. Zero Commission Referral Model: MedBuddy does not accept per-patient referral fees or transaction cuts from participating healthcare providers.

Transparent navigation protects patients from commercial exploitation and aligns technological innovation with genuine public health objectives.

29. Financial Navigation and Out-of-Pocket Expenditure Reduction

As documented in India's National Health Accounts (NHA) 2022-23, household out-of-pocket expenditure accounts for 43.4% of total health spending. Outpatient consultations, repeated diagnostic scans, and chronic medicine purchases represent the vast majority of this drain and are rarely covered by inpatient hospitalisation schemes like AB-PMJAY.

Financial navigation must move beyond displaying nominal consultation fees. Patients require proactive clarity across the full episode of care:

- **Total Episode Estimation:** Clarifying that a consultation may entail subsequent blood tests, imaging, or prescriptions, helping families budget realistically before travel.
- **Scheme Empanelment Verification:** Confirming whether a hospital accepts specific government welfare schemes (such as AB-PMJAY, state health insurance packages, or employee welfare funds) for the exact clinical specialty required.
- **Generic Medicine & Jan Aushadhi Guidance:** Directing patients to nearby Pradhan Mantri Bhartiya Janaushadhi Kendras where verified generic equivalents of prescribed molecules are available at 50% to 90% lower prices.
- **Public Diagnostic Linkages:** Informing patients when expensive tests (such as MRIs, CT scans, or routine biochemistry) can be performed at subsidized rates in district hospitals or empaneled diagnostic hubs.

By reducing information asymmetries around healthcare pricing, MedBuddy shields low-income families from unexpected debt and catastrophic financial hardship.

30. Accountability, Audit Logs, and Transparent Verification

When an AI system influences the pathway of a human seeking medical care, accountability must be absolute. If an automated assistant misinforms a patient regarding a surgical preparation guideline (such as fasting before general anesthesia) or fails to convey an essential referral document, the institution must be able to reconstruct exactly what occurred.

MedBuddy enforces an immutable, auditable logging architecture:

- **Version-Controlled Clinical Grounding:** Every generated explanation is cryptographically linked to the exact version of the verified clinical guideline or facility directory entry active at that second.
- **Prompt & Output Traceability:** System interactions are logged with timestamps, source API responses, and active model parameters, enabling retrospective algorithmic review.
- **Patient Verification Checkpoints:** Critical operational instructions (date, location, fasting status, required identification) require positive acknowledgement before the task is marked completed.
- **Incident Reporting Portal:** Patients and healthcare providers have a dedicated, low-friction channel to report incorrect directions, outdated pricing, or unresponsive clinic contacts.

Accountability ensures that technology teams cannot evade responsibility behind probabilistic opacity. Every automated action must remain explainable, auditable, and subject to institutional oversight.

31. Comprehensive Evaluation Framework: Beyond Chatbot Metric Traps

Commercial technology developers frequently evaluate AI assistants through superficial conversational metrics: message counts, session duration, user engagement hours, and aggregate prompt volume. In healthcare, these metrics are misleading and potentially counterproductive.

A long, meandering conversation with an AI bot often indicates patient confusion, ambiguous instructions, or frustrating failure to resolve a simple administrative task. Conversely, a brief two-message interaction that instantly confirms a verified clinic appointment and prints an action summary represents a superior navigation outcome.

This report establishes that AI assisted healthcare navigation must be evaluated through **completed care milestones**:

Evaluation Domain	Metric Description	Data Source
Task Completion Rate (TCR)	Percentage of initiated navigation inquiries that conclude with verified appointment, document, or referral	System Transaction Logs
Referral Closed-Loop Rate	Proportion of clinician referrals where the patient successfully attends receiving facility	Receiving Facility HMIS
Time-to-Care Acceleration	Reduction in days elapsed between initial symptom report and formal specialist evaluation	Longitudinal Patient Cohort
Comprehension Index	Verified patient understanding of visit preparation, location, and post-consultation instructions	Post-Encounter Spot Audits
Frontline Desk Wait Time	Average minutes elapsed during registration and intake at participating clinic desks	Facility Timestamp Telemetry
Financial Surprise Frequency	Percentage of patients encountering unexpected unquoted fees at point of service	Patient Experience Surveys

System success is defined by clinical connectivity and patient relief, not artificial engagement.

32. Equity Metrics and Distributional Impact Assessment

Healthcare innovations frequently suffer from the "inverse care law": they are adopted fastest by affluent, highly literate, urban populations who already enjoy superior access, while deepening the exclusion of rural, low-income, and marginalized communities.

An equity-first evaluation framework must disaggregate performance metrics across vulnerable demographic cohorts:

- **Gender Parity Ratio:** Ensuring that female device users (who face structural barriers in personal smartphone ownership) achieve equal task completion rates through assisted community desks;
- **Linguistic Inclusivity:** Verifying that response accuracy and completion rates in regional Indian languages (Hindi, Bengali, Tamil, Telugu, Marathi, Kannada) match English benchmarks;
- **Rural/Urban Balance:** Monitoring that rural district hospitals and CHCs achieve comparable system uptime and directory verification fidelity as urban tertiary centres;
- **Accessibility Compliance:** Evaluating usability among individuals with visual, auditory, cognitive, or physical impairments according to W3C WCAG 2.1 AA standards.

If a navigation platform improves access for the top socio-economic quintile while failing lower quintiles, it cannot be considered successful. Technological progress must demonstrably narrow, rather than widen, existing societal disparities.

33. Policy Recommendations for Indian Healthcare Institutions and MoHFW

To maximize public benefit while safeguarding patient safety, national health ministries and state health departments should establish coordinated policy frameworks for AI assisted navigation:

1. Establish National Registry Standards: Mandate that all licensed public and private healthcare facilities publish and maintain machine-readable operating schedules, specialty availability, and standardized consultation tariffs through ABDM registries.

2. Formulate Task Boundaries for AI Assistants: Issue statutory guidance distinguishing non-clinical administrative navigation (which may use verified conversational AI) from clinical triage and diagnostic assessment (which requires medical device regulatory review).

3. Fund Assisted Navigation Kiosks in Public Facilities: Allocate National Health Mission (NHM) resources to establish staffed, technology-assisted navigation desks at all District Hospitals and Community Health Centres.

4. Mandate Non-Digital Fallback Protections: Legally prohibit any healthcare institution from denying registration, clinical care, or emergency admission to patients who lack digital identities, smartphones, or online appointments.

5. Enforce DPDPA Compliance in Health Tech: Require commercial health aggregators to undergo third-party data privacy audits ensuring zero monetization of patient symptom searches or navigation histories.

---

34. Operational Recommendations for Hospital and Clinic Administrators

Frontline healthcare administrators must take deliberate operational steps to prepare their facilities for integrated digital navigation:

• **Audit Physical Patient Flow:** Conduct baseline observational studies to identify where arriving patients experience delays, directional confusion, and registration bottlenecks before implementing digital software.

• **Standardize Departmental Taxonomy:** Ensure that physical clinic signage, paper forms, and digital directories use consistent, familiar terminology rather than conflicting medical abbreviations.

• **Train Frontline Staff on Assisted Tools:** Equip receptionists and security personnel with simple digital navigation tablets to assist illiterate or elderly visitors arriving at facility entrances.

• **Establish Closed-Loop Referral Desks:** Designate a specific hospital coordinator responsible for logging incoming referrals and confirming patient attendance back to primary care clinics.

• **Monitor Patient Feedback Loops:** Implement a simple, anonymous mechanism (drop boxes, vernacular voice calls) for patients to report navigation failures and unexpected fees.

---

35. Engineering and Technical Guidelines for Health Tech Innovators

Software engineering teams developing AI health tools must adhere to disciplined architectural practices:

• **Enforce Grounded Retrieval-Augmented Generation:** Restrict LLM generation to deterministic retrieval from verified public registries, institutional schemas, and approved clinical guidance.

• **Implement Deterministic Safety Guardrails:** Deploy regex and embedding-based classification layers that detect medical urgency and bypass generative models instantly to display emergency contact numbers.

• **Design for Intermittent Connectivity:** Build offline-first data caching that allows mobile applications and kiosk terminals to function reliably during cellular network interruptions.

• **Prioritize Audio and Low-Literacy Interfaces:** Incorporate speech-to-text and text-to-speech models calibrated for Indian regional dialects, background ambient noise, and varied audio fidelities.

• **Maintain Open Protocol Compatibility:** Ensure system APIs strictly comply with ABDM, UHI (Beckn), and FHIR standards to enable seamless interoperability across diverse electronic health record platforms.

---

36. Proposed Implementation Roadmap for MedBuddy

A credible innovation roadmap must proceed through disciplined, evidence-based stages rather than premature nationwide deployment. Explicity Technologies proposes a four-phase rollout for MedBuddy:

```

| MEDBUDDY PHASED IMPLEMENTATION ROADMAP |
| --- |
| |
| PHASE 1: Directory Foundation & Deterministic Registry Pilot (Q1-Q2) |
| - Ingest and verify public HFR/HPR provider directories in 2 districts. |
| - Deploy core administrative FAQ engine with deterministic safety rules. |
| - Measure directory accuracy, stale data rates, and user query types. |
| |
| PHASE 2: Assisted Clinic Desk & Vernacular Action Card Pilot (Q3-Q4) |
| - Implement staff-facing intake dashboards in 10 district facilities. |
| - Pilot printed vernacular visit summaries and WhatsApp action cards. |
| - Evaluate reduction in queue waiting times and patient comprehension. |
| |
| PHASE 3: Closed-Loop Referral & UHI Gateway Bridge (Year 2: Q1-Q2) |
| - Integrate Beckn-compliant UHI protocols for cross-provider booking. |
| - Connect primary health centres with district hospital referral desks. |
| - Track referral attendance rates, lost-to-follow-up, and drop-offs. |
| |
| PHASE 4: Multicenter Longitudinal Evaluation & Equity Audit (Year 2) |
| - Conduct independent academic evaluation across 50 urban/rural clinics. |
| - Assess total episode cost changes, out-of-pocket health drain, and |
| long-term chronic illness adherence. |
| - Publish peer-reviewed findings in open-access health policy journals. |

```

Each successive phase is contingent upon independently verified performance benchmarks and bioethical review.

37. Conclusion: Towards Inclusive, Accountable Healthcare Navigation

Healthcare in India is undergoing unprecedented transformation. Massive investments in digital public infrastructure, expanding social insurance coverage, and cutting-edge artificial intelligence offer historic opportunities to democratize clinical excellence.

Yet technology alone cannot heal a fragmented system. A brilliant conversational model cannot cure illness if the patient cannot reach the hospital, afford the medicine, understand the doctor’s instructions, or obtain a follow-up visit. When technology is implemented carelessly, it risks creating a two-tier healthcare reality: effortless convenience for the digitally privileged, and insurmountable administrative hurdles for the vulnerable.

Healthcare navigation must therefore be understood as an act of institutional responsibility and human dignity. By combining verified operational data, grounded artificial intelligence, dedicated frontline clinic workers, and uncompromising bioethical standards, we can build health systems that guide every citizen with empathy, clarity, and competence.

For Explicity Technologies and MedBuddy, this report marks a committed beginning: measuring the hard problems, respecting clinical realities, and designing technologies that ensure no patient is left to wander the labyrinth of healthcare

alone.

Appendix A. Institutional Roles and Navigation Responsibilities

Role / Stakeholder	Primary Navigation Responsibility	Decision Boundaries & Limitations
Patient / Family Caregiver	Express practical care needs, provide accurate history, select preferred provider	Retains full autonomy over clinical choices and consent
Frontline Receptionist	Verify demographic details, manage appointment queues, issue directional tokens	Administrative intake only; no clinical triage
Community Health Worker (ASHA)	Household outreach, assisted digital scheduling, village-level follow-up tracking	Escorts and assists patients; operates under medical officer supervision
Non-Clinical Patient Navigator	Guide patients across complex multispecialty referrals, coordinate financial aid	Resolves administrative/social barriers; does not alter medical prescriptions
MedBuddy AI Navigation Engine	Natural language intent parsing, verified directory lookup, multilingual translation	Strictly non-clinical administrative guidance; deterministic escalation on urgency
Attending Clinician	Clinical diagnosis, prescription formulation, issuing formal referral orders	Retains sole legal and clinical authority over patient care decisions

Appendix B. Navigation Risk Classification Matrix

Risk Tier	Example Interactions	Permitted System Actions	Mandatory Safeguards
Tier 1: Routine Administrative	Clinic location, operating hours, parking, wheelchair access, accepted documents	Fully automated natural language retrieval from verified directory	Regular automated directory audit; user discrepancy reporting button
Tier 2: Transactional Coordination	Booking confirmed appointment slot, rescheduling visit, retrieving verified test report	Execute API transaction via ABDM / UHI gateway with explicit confirmation	Explicit confirmation screen; instant SMS/WhatsApp booking receipt with cancellation link
Tier 3: Clinical Care Plan Context	Explaining doctor-issued referral, pre-procedure fasting guidelines, medication schedules	Retrieve and explain clinician-authored instructions in vernacular plain language	System must cite original clinician prescription; deterministic "Consult Doctor" advisory
Tier 4: Acute Urgency & Emergency	Acute chest pain, breathing difficulty, severe bleeding, sudden paralysis, poisoning	Immediate display of local ambulance helpline (108) and nearest emergency department	Generative conversational response blocked; emergency banner displayed instantly

Appendix C. Proposed MedBuddy Technical and Ethical Guardrails

1. Deterministic Grounding: Generative models are bound by retrieval-augmented generation (RAG) strictly restricted to verified facility records and clinical guidelines.

2. Clinical Disclaimers and Role Clarity: The interface clearly identifies itself as an administrative assistant and explicitly disclaims diagnostic authority.

3. Emergency Override Protocol: Automatic recognition of 40+ acute symptom categories instantly triggers emergency service routing.

4. Data Minimization: No personal health data is extracted or retained beyond what is required for the immediate navigation task.

5. Non-Commercial Neutrality: Commercial ranking, sponsored ads, and provider referral commissions are strictly prohibited.

6. Linguistic Parity: Equal semantic accuracy and response safety guaranteed across all supported Indian languages.

7. Complete Portability: Summaries and referral cards exportable via open standards (PDF, plain text, FHIR) without requiring active app installation.

Appendix D. Core Evaluation Indicators and Benchmarks

Metric Code	Performance Indicator	Pilot Target Benchmark	Measurement Frequency
NAV-01	Directory Information Accuracy	> 98.5% verified active data	Weekly automated verification
NAV-02	Frontline Intake Queue Reduction	> 75% reduction in intake wait time	Real-time facility telemetry
NAV-03	Referral Closed-Loop Attendance	> 70% completed specialist visits	Monthly cross-facility audit
NAV-04	Patient Instruction Comprehension	> 85% verified patient clarity	Randomised post-visit survey
NAV-05	Emergency Red-Flag Interception	100% immediate escalation compliance	Continuous automated audit
NAV-06	Vernacular Task Completion Parity	< 5% variance between languages	Quarterly multi-cohort analysis
NAV-07	Non-Digital Fallback Availability	100% of participating clinics maintain manual desks	Bi-annual on-site inspection

Selected References

1. World Health Organization (2023). *Patient Navigation in Cancer Care: A Technical Brief*. Geneva: WHO.

2. World Health Organization (2024). *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models*. Geneva: WHO.

3. Indian Council of Medical Research (2023). *Ethical Guidelines for Application of Artificial Intelligence in Biomedical Research and Healthcare*. New Delhi: ICMR.

4. Ministry of Health and Family Welfare (2024). *National Health Accounts Estimates for India (2022-23)*. New Delhi: MoHFW.

5. National Health Authority (2024). *Ayushman Bharat Digital Mission: Architecture and Strategy Overview*. New Delhi: NHA.

6. National Health Authority (2025). *Unified Health Interface (UHI): Consultation Paper on Market Architecture and Open Network Protocols*. New Delhi: NHA.

7. European Observatory on Health Systems and Policies (2022). *What is the Evidence on the Impact of Patient Navigation on Healthcare Access and Quality?* Policy Brief 47. Copenhagen: WHO Regional Office for Europe.

8. World Wide Web Consortium (2023). *Web Content Accessibility Guidelines (WCAG) 2.2*. W3C Recommendation.

9. Ministry of Statistics and Programme Implementation (2025). *Comprehensive Modular Survey: Telecom (January-March 2025)*. New Delhi: MoSPI.

10. Ministry of Law and Justice (2023). *The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023)*. New Delhi: Gazette of India.