

AI Virtual Medical Representatives for Scalable and Responsible Physician Engagement: An Evidence-Based Product Analysis of Hidoc UNNATI

Narrative review, governance framework and prospective evaluation protocol for conversational pharmaceutical engagement

Dr Sadhana Gadia; Varun Gadia

Hidoc Journal | Volume 1, Issue 1 | 4 August 2026

Structured Abstract

Background: Pharmaceutical field engagement remains an important channel through which healthcare professionals receive product information, scientific updates, meeting invitations and access to medical-information services. However, conventional detailing is constrained by representative capacity, geography, timing, variable message delivery and the growing preference of many clinicians for concise, on-demand digital interaction. Hidoc Dr describes UNNATI as an artificial-intelligence-enabled virtual medical representative that can automate outbound calls, conduct personalized conversations, collect physician feedback, operate continuously and update customer-relationship-management systems. **Objective:** This article evaluates the potential role of an AI virtual medical representative in a modern pharmaceutical engagement model and defines the controls required for scientific validity, privacy, promotional compliance, pharmacovigilance and human accountability. **Methods:** A narrative review was undertaken using current peer-reviewed literature, official regulatory and professional guidance, and public product information available through 4 August 2026. Evidence was synthesized across pharmaceutical detailing, digital engagement, conversational agents, speech recognition, AI acceptance, data protection, telecom regulation and ethical medicines promotion. **Results:** Conversational automation can plausibly increase coverage, consistency, responsiveness and structured data capture. Evidence from healthcare conversational agents indicates generally acceptable feasibility, while research on pharmaceutical detailing and digital interaction demonstrates both informational value and material risks: promotional bias, information overload, weak source transparency, privacy concerns and erosion of trust. Speech systems also show accent- and language-related performance variation. The defensible design is therefore not an unrestricted generative sales agent. It is a bounded, approved-content system with explicit identity disclosure, recorded consent status, source-linked answers, adverse-event and product-complaint detection, off-label safeguards, human escalation and auditable change control. **Conclusion:** UNNATI should be evaluated as an augmentation layer within an omnichannel medical-commercial operating model. Its success should be measured not only through reach and conversion, but also through scientific accuracy, consent compliance, escalation sensitivity, physician utility, equity and downstream operational quality. Prospective deployment should begin with low-risk use cases and a staged validation protocol before wider promotional autonomy is considered.

Keywords: artificial intelligence; conversational AI; virtual medical representative; pharmaceutical detailing; physician engagement; medical information; pharmacovigilance; digital health; compliance; omnichannel marketing

Disclosure and scope: Varun Gadia is associated with Hidoc Dr. The article is a product-centred, literature-informed narrative review and proposed evaluation framework. It does not report a clinical trial or claim that UNNATI has demonstrated improvements in prescribing, clinical practice or patient outcomes. Publicly described product capabilities are distinguished from evidence derived from independent literature.

1. Introduction

The pharmaceutical representative has historically occupied a complex position between education, promotion, relationship management and market intelligence. In a conventional interaction, a representative identifies an appropriate clinician, requests time, presents an approved product narrative, responds to a limited range of questions, captures follow-up requirements and returns information to a customer-relationship-management system. The encounter can be valuable when it is timely, scientifically accurate and relevant to a clinician's practice. It can also create risk when selective evidence, poorly contextualized claims or incentives distort clinical judgment. Systematic reviews have repeatedly shown that interactions between physicians and pharmaceutical companies can influence attitudes and prescribing, while the quality and completeness of

information supplied during commercial detailing are variable [4-6]. The appropriate question is therefore not whether engagement should occur, but how it can be made more useful, transparent, consistent and accountable.

Digital channels have altered the practical constraints around this interaction. Clinicians now encounter educational and promotional information through email, portals, webinars, social platforms, messaging applications and remote meetings. A 2026 qualitative study in German rheumatology found that digital interactions with the pharmaceutical industry were valued for accessibility and efficiency but were also associated with information overload, uncertainty about roles and a perceived loss of personal connection [8]. These findings are directly relevant to conversational automation. A virtual representative can reduce the friction of scheduling and extend engagement

beyond a physical field force; it can just as easily increase unwanted contact or deliver superficially personalized but repetitive communication. Scale is not intrinsically beneficial. It magnifies both good design and weak governance.

Hidoc Dr publicly describes UNNATI as an AI Virtual Medical Representative designed to automate outbound calls to doctors, conduct personalized human-like conversations, collect interests and feedback, operate continuously, integrate with CRM systems and reduce outreach cost [1]. The company timeline records a June 2025 launch for pharmaceutical brands, and a Hidoc post states that the product received second runner-up recognition at CAHOTTECH 2025 [2,3]. These descriptions establish the intended operating model, not clinical or commercial effectiveness. No controlled outcome data were identified in the public materials reviewed for this manuscript. Accordingly, the article treats UNNATI as a product architecture to be evaluated rather than as a proven intervention.

The analysis has four aims. First, it defines the legitimate use cases for a virtual medical representative. Second, it converts general AI-governance principles into concrete controls for pharmaceutical communication. Third, it identifies operational and human factors that may determine physician acceptance. Fourth, it proposes a prospective evaluation protocol capable of separating reach from meaningful, compliant engagement.

2. Review Approach and Evidence Boundaries

This paper is a narrative review rather than a systematic review or meta-analysis. Searches were conducted across PubMed-indexed literature, publisher repositories, official health-agency websites, Indian government portals, telecommunications regulation, international pharmaceutical-industry guidance and Hidoc Dr public webpages. Priority was given to systematic or scoping reviews, recent empirical studies, official legislation, professional standards and primary product descriptions. The evidence domains included physician-industry interaction; commercial and academic detailing; digital pharmaceutical communication; healthcare conversational agents; voice interfaces and automatic speech recognition; clinician acceptance of AI; pharmacovigilance; data protection; unsolicited commercial communication; and ethical promotion.

The available evidence has important limits. Most conversational-agent trials concern patient support, behavior change, history-taking or disease management rather than pharmaceutical promotion. Evidence about conventional sales representatives cannot be transferred mechanically to an AI voice system because the interaction cost, persistence, personalization and auditability differ. Similarly, studies of speech recognition in clinical documentation demonstrate technical risks but do not measure the performance of UNNATI. This article therefore uses triangulation: it draws design implications from adjacent evidence while clearly avoiding fabricated product metrics.

A product-centred narrative review is useful at an early stage because it can expose assumptions before a large deployment. It cannot establish effectiveness. A future

evaluation must use predefined endpoints, independent annotation, comparator groups and sufficient follow-up. It should also make the tested version of the conversational model, approved content pack and escalation policy identifiable, since an AI system can materially change after a model, prompt, knowledge base or telephony update.

3. From Field Detailing to Governed Conversational Engagement

Commercial detailing is a structured form of influence. Its legitimate function is to communicate accurate information about approved products and relevant disease management while enabling clinicians to request additional scientific or service support. Its commercial function is to improve product awareness, consideration and appropriate use. These aims can coexist only when patient welfare and scientific accuracy remain the controlling constraints. The literature does not support a naive assumption that more representative contact produces better care. Reviews have associated exposure to pharmaceutical promotion with changes in prescribing frequency, cost and quality, although heterogeneity and confounding make causal interpretation difficult [4,5]. A comparative study of sales visits in Canada, France and the United States found that safety information was often incomplete, illustrating why a persuasive interaction cannot be treated as equivalent to balanced medical education [6].

A virtual representative changes the economics of detailing. Human field capacity is scarce: travel, territory size, clinic access, leave, turnover and scheduling constrain the number and timing of encounters. Automated voice systems can initiate thousands of conversations, adapt language and timing, and create structured records. This may be particularly useful for geographically dispersed clinicians, low-frequency prescribers, post-event follow-up or service reminders. The marginal cost of another attempt is low. That same feature creates a distinct risk of over-contact. A poorly governed system can call too often, ignore preference signals, repeat irrelevant claims and convert a relationship problem into an automated one.

The more defensible concept is governed conversational engagement. The system should perform bounded tasks under an approved communication policy: identify itself; confirm the intended healthcare professional; state the purpose of the call; verify contact permission; deliver only version-controlled content; capture a limited set of responses; detect questions that require medical-information or pharmacovigilance handling; and terminate or transfer when it reaches a boundary. Generative flexibility may improve natural language but should not expand the substantive claim set. This distinction is central. Natural conversation is a user-interface property, not permission for unrestricted scientific generation.

Recent mapping of commercial detailing emphasizes that prescribing decisions are shaped by multiple contextual and behavioral variables rather than by message exposure alone [7]. Digital engagement research similarly indicates that convenience must be balanced against trust, relevance and information load [8]. Therefore, the virtual representative should be designed as one component of an omnichannel sequence, not as a universal replacement for human relationships. The optimal interaction may be a two-minute

automated qualification call followed by a medically trained human response, an appointment, a verified article or an opt-out—not a prolonged autonomous sales dialogue.

4. UNNATI Product Concept and Reference Architecture

The public description of UNNATI identifies six intended capabilities: automated outbound calls, personalized conversation, response and interest capture, continuous availability, CRM integration and improved outreach efficiency [1]. Translating those capabilities into a health-sector architecture requires more than connecting a language model to a telephone number. A robust implementation contains at least eight layers: audience and consent management; approved evidence and claims; campaign rules; conversational orchestration; speech and language processing; response classification; escalation services; and auditable analytics.

Audience management determines who may be contacted, for what purpose, through which number and at what frequency. It must reconcile CRM records, specialty and territory criteria, communication preference, prior consent, suppression lists, time-zone rules and local telecommunications requirements. Approved evidence is the substantive boundary: prescribing information, approved label, medical-legal-regulatory reviewed scripts, references, contraindications, safety statements and frequently asked questions. The policy engine determines the next permissible action based on the campaign, physician response and risk category.

The speech layer converts generated text into audio and physician speech into text or intent. This is not a neutral utility. Medical names, doses, abbreviations, accents, code-switching and noisy clinic environments create error. Systematic evidence on AI speech recognition shows variable accuracy across settings, and recent studies demonstrate accent-related disparities even in clinical language [13-15]. The architecture should therefore preserve the audio where lawful and consented, display or store the interpreted transcript, assign confidence scores and avoid acting on low-confidence content without clarification.

The escalation layer links commercial engagement to the functions that carry legal or patient-safety responsibility. A request for an approved brochure may be fulfilled automatically. A question about a patient, an unapproved use, pregnancy exposure, a suspected adverse reaction, a medication error, lack of efficacy or a product-quality complaint must be classified and routed according to policy. The CRM record should not merely state that the call was completed; it should capture version, consent status, purpose, content modules delivered, questions, transfers, suppression requests and closure status. Figure 1 summarizes this reference architecture.




Figure 1. Reference architecture for a governed AI virtual medical representative. The model separates approved evidence, conversational delivery, structured data capture and accountable escalation.

Table 1. Core capabilities, principal risks and minimum controls.

Capability, value, risk and minimum control
Automated outbound calling: extends coverage and timing flexibility. Control excessive contact through consent proof, preference checks, frequency caps and suppression.
Natural-language conversation: reduces interaction friction. Prevent unsupported claims through an approved knowledge base, response boundaries and confidence rules.
CRM integration: creates structured follow-up data. Prevent opaque profiling through purpose limitation, field minimization and access controls.
Continuous operation: supports physician convenience. Use local time windows and doctor-selected scheduling.
Multilingual speech: improves access. Validate by language and escalate uncertain medication names, doses and intents.
Automated analytics: enables optimization. Balance volume with safety, consent, scientific quality and physician usefulness.

5. High-Value Use Cases and Their Boundaries

A virtual representative is most defensible where the task is repetitive, time-sensitive, easily bounded and operationally measurable. The following examples are hypothetical implementation patterns, not claims about completed UNNATI deployments.

Use case 1: post-webinar follow-up. After a clinician attends a webinar, the system may confirm whether the session was useful, offer the approved slide summary, capture a request for a certificate, identify interest in a future topic and schedule a human discussion. The interaction builds on a known engagement and can be classified as service or follow-up communication when consent and local rules support that purpose. It should not improvise additional treatment recommendations beyond the approved material.

Use case 2: invitation and attendance management. The system may invite eligible healthcare professionals to a

scientific meeting, verify specialty, explain the agenda, register consent for reminders and record attendance intent. This can reduce manual coordination, but the invitation must preserve the separation between legitimate scientific exchange and inducement. Hospitality, sponsorship and event support remain subject to applicable codes, including India's UCPMP and international industry standards [18-20].

Use case 3: approved product-information reinforcement. A short call can determine whether a physician wishes to receive prescribing information, an administration guide or a link to an approved clinical paper. The system can ask one or two knowledge-check questions and identify a need for human follow-up. It should not infer prescribing intent from sensitive patient data, exaggerate comparative efficacy or omit material risk information.

Use case 4: medical-information triage. The agent can classify the topic and collect contact details for a response by an appropriately trained medical-information professional. A bounded agent may answer straightforward questions whose exact responses have been approved and version controlled. Questions involving off-label use, complex evidence interpretation or patient-specific decisions should be transferred rather than answered generatively.

Use case 5: pharmacovigilance intake trigger. During any conversation, the agent should recognize potential adverse events and product complaints even when they are mentioned casually. The system can capture the minimum reportable elements, inform the physician that the matter will be routed, and transmit the record immediately. It must not delay safety reporting while attempting to complete a promotional script.

Use case 6: territory prioritization and human handoff. Structured signals—topic interest, request type, preferred time and consented channel—can help representatives or medical teams prioritize follow-up. Such scoring should be used to improve service allocation, not to create opaque pressure tactics or infer sensitive characteristics. A clinician should be able to correct preferences and opt out.

Use case 7: multilingual access. The agent may offer a choice of language and preserve approved scientific terminology across versions. This can extend access, particularly where clinicians prefer discussing operational details in a local language while reading evidence in English. Translation must be medically validated; literal fluency is insufficient when a small error in dose, route, contraindication or comparator reverses meaning.

Table 2. Illustrative deployment cases and non-negotiable boundaries.

Use case, automatable task and mandatory boundary
Post-webinar follow-up: feedback, certificate requests, topic interest and appointment needs; no unapproved clinical expansion.
Meeting invitation: agenda, registration and reminders; maintain event and hospitality code compliance.
Approved information: send label, guide or approved publication; preserve balance and source traceability.
Medical-information triage: classify and route; require human response for complex, off-label or patient-specific requests.
Safety intake: capture minimum elements; route immediately to pharmacovigilance and pause promotion.
CRM qualification: capture preferred channel, time and topic; prohibit sensitive inference or concealed scoring.

6. Scientific Content Governance

The central control for a pharmaceutical conversational agent is not the general intelligence of the underlying model; it is the integrity of the permitted evidence. Every campaign should begin with a content dossier that identifies the approved indication, target audience, learning or communication objective, key claims, mandatory safety content, references, prohibited statements, expiry or review date and escalation rules. Each deployable response should be traceable to this dossier.

A useful design separates linguistic generation from claim generation. The system may rephrase an approved idea to make the conversation natural, but it should retrieve substantive statements from a controlled knowledge base. Retrieval alone does not eliminate error: the model may combine fragments, lose qualifiers or select an obsolete source. Responses should therefore be constrained by templates or verified response units for high-risk content, with automated source checks and human testing before release.

Medical-legal-regulatory review must cover the entire interaction pattern, not only the opening script. Branching logic, follow-up questions, objection handling, language variants, transfer statements, voicemail, SMS links and post-call summaries can each contain promotional content. Version control should make it possible to reconstruct exactly what policy and evidence were active during a call. When a label or safety communication changes, affected content should be disabled centrally rather than waiting for field retraining.

The WHO ethical criteria for medicinal drug promotion and later industry codes emphasize that promotional information should be accurate, balanced, substantiated and consistent with approved use [19,20,27]. ACCME standards, while directed to accredited continuing education rather than promotion, provide a useful integrity benchmark: patient-care recommendations should be based on current

science and present a fair view of options [28]. UNNATI communications should make the category of interaction explicit. A promotional call should not be disguised as independent continuing education, and an unbranded educational conversation should not quietly transition into product persuasion without disclosure.

A source-linked response model can improve trust. Where practical, the agent should offer to send the approved label, publication or medical-information response and state the version or publication date. It should communicate uncertainty directly. A phrase such as 'I do not have an approved response for that question; I can connect you with medical information' is safer and more credible than a fluent speculative answer.

7. Identity, Consent, Privacy and Telecommunications Compliance

An AI caller should disclose its identity and purpose at the beginning of the interaction. Human-like speech must not be used to create the false belief that the doctor is speaking with a person. Disclosure is important for autonomy, trust and the interpretation of conversational pauses or limitations. The European AI Act establishes transparency duties for certain AI interactions, while WHO guidance emphasizes human autonomy, transparency and accountability [16,17,23]. Even where a local rule does not explicitly require a virtual representative to announce that it is automated, disclosure is a sound baseline.

Consent operates at several levels. The organization needs a lawful basis to store and use physician contact and interaction data. It also needs to respect telecommunications preferences for promotional voice communication. In India, TRAI's commercial-communication framework distinguishes promotional, service and transactional communication, requires traceability, and treats communication outside consent or registered preferences as unsolicited [21]. A healthcare professional's presence in a commercial database is not equivalent to unlimited permission for automated calls. Campaigns should use registered channels, valid consent records, frequency caps and immediate suppression when an opt-out is expressed.

India's Digital Personal Data Protection Act, 2023 requires lawful processing, clear notice, purpose limitation and mechanisms for withdrawal of consent. A November 2025 commencement notification staged the operative provisions, with many substantive obligations scheduled for a later phase; organizations should nevertheless design toward the full standard rather than minimum transitional compliance [22]. The conversation should collect only information necessary for the stated purpose. Recording a broad clinical profile, personal preferences unrelated to engagement or patient-identifiable details creates avoidable risk.

Call recording requires specific attention. Audio is useful for quality review, dispute resolution, transcription improvement and safety detection. It also contains biometric and contextual information and may inadvertently include patient details. The system should state whether the call is recorded, define retention, restrict access, encrypt data, log retrieval and support deletion or correction where applicable. Where recording is unnecessary, structured event logs may provide sufficient auditability with lower privacy risk.

Cross-border deployments require jurisdiction-specific configuration rather than a single global policy. Consent language, recording rules, promotional restrictions, data localization, retention and AI-transparency requirements vary. A campaign should not be launched in a new country merely because the language model supports the local language.

8. Pharmacovigilance, Product Complaints and Medical Information

Any system that speaks with healthcare professionals about medicines is a potential source of safety information. Pharmacovigilance obligations do not disappear because a conversation is automated or commercial. A doctor may mention that a patient developed a rash, stopped treatment, became pregnant, received the wrong dose, experienced lack of effect or found a damaged device. These statements may be incomplete, indirect or embedded within another question. The system must be designed to detect and route them.

The minimum capability is a high-sensitivity trigger library covering adverse events, special situations, medication errors, misuse, overdose, occupational exposure, pregnancy or breastfeeding exposure, lack of efficacy where reportable, and product-quality complaints. Generative semantic classification can improve recall beyond exact keywords, but it requires validation against representative accents, languages and indirect phrasing. The desired trade-off is different from sales qualification: false positives create review work, whereas false negatives may create regulatory and patient-safety risk.

Once a trigger occurs, the promotional flow should pause. The agent can explain that the information may need safety follow-up, collect the minimum elements permitted by local procedure, obtain a preferred contact method and transmit the case to the appropriate function within the required timeline. It should not attempt causality assessment, reassure the reporter that the event is unrelated, or discourage reporting. Quality complaints should be separated from adverse-event reports but linked where both are present.

Medical-information boundaries are equally important. A product-information request may be fulfilled through an approved document. A scientific question may require a balanced, nonpromotional response prepared by medical information. Off-label questions should be captured reactively and transferred; they should never be introduced by the virtual representative as a means of expanding promotion. The audit trail must distinguish what the doctor asked from what the system volunteered.

The evolving pharmacovigilance function is increasingly data intensive and digitally connected [26]. UNNATI can support this evolution only if safety integration is built into the primary workflow. Retrospective manual review of a small sample of call transcripts is not an adequate substitute for real-time detection, routing and reconciliation.




Figure 2. Risk-based response and escalation pathway. High-risk or ambiguous content is transferred rather than answered through unrestricted generation.

9. Human Factors, Physician Acceptance and Trust

A technically accurate system may still fail if it violates the clinician's workflow or social expectations. Reviews of healthcare AI acceptance consistently identify perceived usefulness, ease of use, trust, transparency, workflow integration, organizational readiness and concern about professional identity as important determinants [24,25]. A physician who is willing to use an AI tool for a thirty-second certificate request may reject the same tool for a complex scientific discussion.

The interaction should therefore be deliberately brief and controllable. The doctor should be able to interrupt, request a human, change language, postpone, obtain a written link or terminate immediately. Long introductory disclosures can reduce usability, but critical identity, sponsor and purpose information cannot be hidden. The script should prioritize the doctor's task rather than the campaign's desire to deliver every message.

Trust is affected by both content and behavior. Repeated calls after an opt-out, failure to recognize a doctor's name, incorrect pronunciation of a medicine or confident answers outside scope will rapidly damage credibility. Conversely, accurate recognition, source-linked information and efficient transfer can make the system useful. Digital interaction research warns that accessibility can coexist with overload and reduced personal connection [8]. Frequency governance and relevance are therefore safety and reputation controls, not merely marketing optimizations.

Human representatives remain important. They interpret context, understand institutional dynamics, support nuanced objection handling and build durable relationships. Medical teams provide nonpromotional scientific exchange. The virtual representative should remove low-value administrative friction and qualify needs so that human time is allocated to conversations requiring judgment. Replacement may be appropriate for selected repetitive tasks; wholesale substitution is neither supported by the evidence nor necessary to achieve scale.

Co-design with physicians is essential. Pilot users should include different specialties, practice settings, age groups, languages, accents and levels of digital familiarity. Interviews should examine not only satisfaction but perceived intrusiveness, clarity of sponsorship, confidence

in accuracy, comfort with recording, preference for escalation and likelihood of future use.

10. Language, Speech Performance and Equity

Voice interaction introduces a chain of possible error: telephony quality, speech recognition, language identification, intent classification, retrieval, response generation and text-to-speech output. A system may sound fluent while misunderstanding the original question. Medical vocabulary is particularly unforgiving because many drug names are acoustically similar and numerical details carry high consequence.

A 2025 systematic review of AI speech recognition for clinical documentation found heterogeneous performance and emphasized the need for evaluation in real clinical environments [13]. A 2026 study demonstrated higher error rates for non-native English clinical speech and showed that post-processing could reduce, but not eliminate, accent-related loss [14]. Research in home healthcare has also identified performance disparities across speaker groups [15]. These findings make an equity audit mandatory for a system intended to engage doctors across India and other multilingual markets.

Validation should be stratified by language, accent, specialty vocabulary, gender where appropriate, device, network condition and background noise. Word error rate alone is insufficient. A transcription may contain many harmless filler-word errors but one dangerous error involving dose or contraindication. The evaluation should therefore include critical concept error rate, medication-name accuracy, numerical accuracy, intent accuracy and escalation sensitivity.

Language choice should not reduce scientific quality. Each approved content pack requires medical translation, back-translation or bilingual expert review, pronunciation dictionaries and test calls. Code-switching should be expected. The safest workflow may use local-language conversation for navigation and an approved written resource for detailed evidence. Doctors should be able to request the original English source.

Equity also concerns who is contacted. Algorithmic prioritization based on prior engagement can create a feedback loop in which already-responsive urban specialists receive more attention while low-data clinicians become invisible. Coverage objectives should therefore include geographic and practice-setting fairness, not only conversion probability.

11. Operational Model and Economics

The commercial appeal of a virtual representative is straightforward: lower marginal outreach cost, extended operating hours, standardized messaging, rapid campaign launch and automatic data capture. However, a credible economic model must include the costs that are often excluded from automation claims: medical review, conversational design, telecom registration, consent management, model hosting, multilingual validation, call recording, cybersecurity, monitoring, pharmacovigilance integration, human escalation, incident response and periodic revalidation.

Efficiency should be measured at the level of completed work, not calls placed. Useful operational endpoints include

verified contact rate, consented conversation rate, task-completion rate, transfer completion, time to medical-information closure, event-registration completion, suppression accuracy, duplicate-contact rate and cost per resolved physician need. A low-cost call that creates a complaint, misses a safety report or requires manual transcript repair is not efficient.

The system may also improve field productivity by converting unstructured interactions into standardized signals. CRM integration can identify requested topics, preferred times, language and next action. Yet data abundance can become an operational burden unless fields are limited and ownership is clear. The 2026 rheumatology study described uncertainty about who should monitor and respond to digitally generated information [8]. Every captured intent must therefore have a service-level owner and closure rule.

A hybrid deployment may deliver the best economics. UNNATI can conduct initial outreach, service reminders, qualification and routine fulfillment; field or inside-sales teams can handle relationship-intensive commercial discussions; medical information can answer complex scientific questions; pharmacovigilance can process safety cases. This design treats automation as capacity orchestration rather than head-count arithmetic.

Economic evaluation should include an incremental comparison against the relevant alternative: human call center, field representative, email-only campaign or no follow-up. The analysis should report total program cost, not just telephony or model cost, and should incorporate quality-adjusted outcomes such as verified useful interactions and compliant case closures.

12. Implementation Framework

A staged implementation reduces both technical and regulatory risk. Phase 0 is governance design: appoint accountable owners across commercial, medical, legal, privacy, information security, pharmacovigilance and technology; classify intended use cases; define prohibited functions; and approve a risk taxonomy. Phase 1 is offline validation using synthetic and consented test utterances. The team should challenge the system with accents, interruptions, ambiguous questions, safety statements, off-label prompts, prompt-injection attempts and poor network conditions.

Phase 2 is a closed pilot with a small group of consenting physicians and low-risk tasks such as event registration, certificate fulfillment or appointment scheduling. Calls should be reviewed at a high sampling rate. No autonomous answer should be allowed outside a narrow approved set. Pilot success criteria should include identity disclosure, consent accuracy, task completion, transcript quality, user acceptability and zero unresolved safety triggers.

Phase 3 expands to approved informational interactions and limited CRM integration. Human escalation must be operational before scale increases. Dashboards should display not only commercial outcomes but safety and quality indicators: low-confidence turns, prohibited-answer attempts, opt-outs, complaints, adverse-event triggers, transfer failures and content-version exceptions.

Phase 4 supports multilingual and broader campaign deployment after language-specific validation. Each new therapy area should be treated as a material change because terminology, risk, objections and evidence differ. Continuous monitoring should use random transcript review, automated policy checks, physician feedback, incident investigation and drift detection. A model update should trigger regression testing before production release.

The governance committee should have authority to pause a campaign. Predefined stop criteria may include repeated identity-disclosure failure, a missed serious safety trigger, unauthorized off-label content, abnormal opt-out rates, a data breach, material speech errors or telecom complaints. Commercial urgency should not override these controls.

Training remains necessary even for an automated system. Human teams must understand what the agent can do, how to interpret CRM outputs, when not to rely on transcripts, how to process escalations and how to explain the system to physicians. The organization, not the model vendor, retains accountability for the deployed interaction.

13. Proposed Prospective Evaluation Methodology

A rigorous first study could use a stepped-wedge cluster design or a randomized campaign design across comparable physician segments. Clusters may be defined by territory, specialty or account group. The comparator should reflect current practice: human representative outreach, call-center outreach, email or standard omnichannel follow-up. Randomization should occur before contact attempts, with contamination monitored when the same physician can receive multiple channels.

The primary endpoint should be a composite of useful and compliant task completion rather than raw reach. A completed interaction could require verified healthcare-professional identity, valid communication permission, successful completion of the predefined task, no prohibited content, correct documentation and timely closure of any escalation. Secondary commercial endpoints may include event registration, requested-material delivery, qualified follow-up and cost per completed task. These outcomes should be reported separately from clinical or prescribing outcomes.

Scientific-quality endpoints should be evaluated by independent blinded reviewers using a predefined rubric: factual accuracy, consistency with approved label, balance, source appropriateness, safety completeness, uncertainty handling and absence of unsupported comparative or off-label claims. A stratified sample of calls should include routine, adversarial and edge-case conversations. Inter-rater reliability should be reported.

Safety endpoints should include sensitivity and positive predictive value for adverse-event, special-situation and product-complaint detection; time from first mention to case transfer; completeness of minimum reportable information; reconciliation success; and number of unresolved cases. The study should deliberately include simulated safety utterances in each supported language before real-world launch.

Human-factors endpoints should include physician-rated usefulness, trust, intrusiveness, clarity of AI identity,

perceived scientific credibility, ease of escalation, preference compared with other channels and willingness to receive future calls. Interview data can explain why a technically completed interaction was accepted or rejected.

Equity analysis should compare performance by language, accent, specialty, geography, gender where lawful and relevant, practice setting and network quality. Technical endpoints should include call connection, latency, interruption handling, word and critical-concept error, intent classification, low-confidence rate and transfer reliability.

A sample-size calculation should be based on the expected difference in compliant task completion and adjusted for cluster effects where used. The protocol, analysis plan and versioned system description should be registered or archived before analysis. Results should report failures and complaints, not only successful conversations. No claim about prescribing quality or patient outcome should be made unless those outcomes are prospectively measured with an appropriate design and sufficient follow-up.

Table 3. Proposed domains and example measures for prospective evaluation.

Evaluation domain, measures and interpretation
Useful completion: verified identity, consent, task completion and correct closure; primary operational-value endpoint.
Scientific quality: accuracy, balance, label consistency and unsupported-claim rate; independent blinded review.
Safety: adverse-event sensitivity, routing time, reconciliation and missed triggers; critical threshold metrics.
Human factors: usefulness, trust, intrusiveness and future preference; explains adoption beyond completion.
Equity: performance by language, accent, geography and setting; detects uneven access or error.
Economics: total cost per compliant completed task, including governance and human escalation.

14. Governance Matrix and Key Performance Indicators

A balanced scorecard is necessary because commercial metrics alone reward volume. The first domain is access: eligible physicians attempted, verified contact, consented contact, language availability and geographic coverage. The second is utility: task completion, requested-resource fulfillment, successful appointment or transfer and physician-rated usefulness. The third is scientific integrity: approved-response rate, unsupported-claim rate, label consistency, reference delivery and medical-review findings. The fourth domain is safety: adverse-event detection sensitivity, time to routing, reconciliation completion, product-complaint routing and missed-trigger rate. The fifth

is privacy and telecom compliance: disclosure completion, consent proof, opt-out latency, suppression errors, recording notice and complaint rate. The sixth is technical quality: latency, interruption success, speech error, critical-concept accuracy, uptime and version exceptions. The seventh is human workflow: transfer acceptance, backlog, closure time and field-team adoption.

Targets should be risk weighted. For example, a 95% task-completion rate would not compensate for a missed serious adverse-event trigger or repeated unauthorized promotion. Some indicators are therefore thresholds rather than averages. The system should not proceed to scale until critical controls meet predetermined minimum performance.

Independent review strengthens credibility. Medical and compliance teams should have access to unfiltered audit samples, and safety validation should not be performed solely by the commercial product team. Where feasible, an external clinical or methodological advisor can review the evaluation protocol and claims derived from it.

15. Limitations

This article is based on public product descriptions and adjacent literature. It does not include proprietary UNNATI architecture, training data, call scripts, validation records, deployment volumes or outcome datasets. Consequently, specific technical safeguards described here are recommendations rather than confirmed current features. The evidence on conversational agents largely concerns patient-facing or clinical-support use cases, while evidence on pharmaceutical detailing mostly concerns human representatives. The synthesis is therefore inferential.

Regulatory obligations vary by jurisdiction and can change. The discussion of Indian telecommunications and data-protection requirements is not legal advice. Each campaign requires local legal, privacy and regulatory review. Likewise, an interaction may be classified differently depending on its content, consent history and relationship to an existing service.

Finally, physician engagement is not a proxy for clinical benefit. A system may improve contact and task completion without improving knowledge, prescribing appropriateness or patient outcomes. Those higher-order outcomes require separate studies and should not be implied by marketing analytics.

16. Conclusion

UNNATI represents a credible direction for the automation of selected pharmaceutical-engagement tasks. The strongest case is not an autonomous AI salesperson capable of answering anything. It is a governed virtual representative that expands access to approved information, performs routine follow-up, captures structured preferences and routes questions to the right human function.

The technology can create value through consistency, availability, multilingual reach and CRM integration. It can also amplify the established risks of pharmaceutical promotion and introduce new ones: synthetic persuasion, over-contact, hidden automation, speech error, data misuse and weak accountability. These risks are manageable only when identity disclosure, consent, approved-content boundaries, pharmacovigilance detection, human

escalation, version control and independent audit are embedded in the product.

Deployment should begin with low-risk, high-structure tasks and progress only after prospective validation. Success must be defined by useful, scientifically accurate and compliant task completion—not by call volume alone. Under that model, UNNATI can augment field and medical teams while preserving the human judgment required for complex scientific exchange and patient-safety responsibility.

References

- Hidoc Dr. Services: UNNATI - AI Virtual MR. Hidoc Dr website. Accessed 4 August 2026.
- Hidoc Dr. About Us and company timeline. Hidoc Dr website. Accessed 4 August 2026.
- Hidoc Dr. Hidoc - Unnati Wins 2nd Runner-Up at CAHOTTECH 2025. LinkedIn post. 2025. Accessed 4 August 2026.
- Fickweiler F, Fickweiler W, Urbach E. Interactions between physicians and the pharmaceutical industry generally and sales representatives specifically and their association with physicians' attitudes and prescribing habits: a systematic review. *BMJ Open*. 2017;7:e016408. doi:10.1136/bmjopen-2017-016408.
- Spurling GK, Mansfield PR, Montgomery BD, et al. Information from pharmaceutical companies and the quality, quantity, and cost of physicians' prescribing: a systematic review. *PLoS Med*. 2010;7:e1000352. doi:10.1371/journal.pmed.1000352.
- Mintzes B, Lexchin J, Sutherland JM, et al. Pharmaceutical sales representatives and patient safety: a comparative prospective study of information quality in Canada, France and the United States. *J Gen Intern Med*. 2013;28:1368-1375. doi:10.1007/s11606-013-2411-7.
- Bhadraj R, et al. Mapping the Influence of Commercial Detailing on Prescribing Decisions: A Scoping Review. *Cureus*. 2026.
- May S, et al. Digital interactions with the pharmaceutical industry: a qualitative focus group study on the perspectives of rheumatology care providers in Germany. *BMC Rheumatol*. 2026;10:23. doi:10.1186/s41927-026-00623-1.
- Laranjo L, Dunn AG, Tong HL, et al. Conversational agents in healthcare: a systematic review. *J Am Med Inform Assoc*. 2018;25:1248-1258. doi:10.1093/jamia/ocy072.
- Li Y, Liang S, Zhu B, et al. Feasibility and effectiveness of artificial intelligence-driven conversational agents in healthcare interventions: a systematic review of randomized controlled trials. *Int J Nurs Stud*. 2023;143:104494. doi:10.1016/j.ijnurstu.2023.104494.
- Bramanti A, Corallo A, Clemente G, et al. Exploring the role of voice assistants in managing noncommunicable diseases: a systematic review on clinical, behavioral outcomes, quality of life, and user experiences. *Healthcare*. 2025;13:517. doi:10.3390/healthcare13050517.
- Shi H, Huang J, Yang J, Han M. Barriers and facilitators to patient acceptance of artificial intelligence in health care: systematic review. *J Med Internet Res*. 2026;28:e80581. doi:10.2196/80581.
- Ng JJW, Wang E, Zhou X, et al. Evaluating the performance of artificial intelligence-based speech recognition for clinical documentation: a systematic review. *BMC Med Inform Decis Mak*. 2025;25:236. doi:10.1186/s12911-025-03061-0.
- Fatapour Y, Samaan JS, Tatonetti NP; Inclusive AI Research Group. Accent related errors in clinical speech transcription and a LLM-based remedy. *npj Digit Med*. 2026;9:297. doi:10.1038/s41746-026-02490-z.
- Xu Z, Vergez S, Esmaeili E, et al. Voice for All: evaluating the accuracy and equity of automatic speech recognition systems in transcribing patient communications in home healthcare. *Stud Health Technol Inform*. 2025;329:1904-1906. doi:10.3233/SHTI251273.
- World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO; 2021.
- World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. Geneva: WHO; 2025.
- Department of Pharmaceuticals, Government of India. Uniform Code for Pharmaceutical Marketing Practices 2024 and UCPMP portal guidance. 2024-2026.
- International Federation of Pharmaceutical Manufacturers and Associations. IFPMA Code of Practice. Geneva: IFPMA; current edition.
- International Federation of Pharmaceutical Manufacturers and Associations. Note for Guidance on educational and scientific events: 2026 update. Geneva: IFPMA; 2026.
- Telecom Regulatory Authority of India. Telecom Commercial Communications Customer Preference Regulations, 2018, as amended in 2025; Advice to Senders/Telemarketers. New Delhi: TRAI; accessed 4 August 2026.
- Government of India. Digital Personal Data Protection Act, 2023; commencement notifications dated 13 November 2025. India Code.
- European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence, as amended. Official Journal of the European Union.
- Hua D, Petrino N, Young N, Cho JG, Poon SK. Understanding the factors influencing acceptability of AI in medical imaging domains among healthcare professionals: a scoping review. *Artif Intell Med*. 2024;147:102698.
- Barriers to and facilitators of clinician acceptance and use of artificial intelligence in healthcare settings: a scoping review. *BMJ Open*. 2025;15:e092624.
- Stagi L, et al. The evolution of the pharmacovigilance department in pharmaceutical companies: results of a national survey. 2024.
- World Health Organization. Ethical criteria for medicinal drug promotion. Geneva: WHO; 1988.
- Accreditation Council for Continuing Medical Education. Standards for Integrity and Independence in Accredited Continuing Education: Standard 1, Ensure Content is Valid. Accessed 4 August 2026.
- International Consensus Framework for Ethical Collaboration in Health. Responsible use of health data and technology, including artificial intelligence. 2025.
- Dentice S, Moreal C, Chiappinotto S, et al. Professional and patient care outcomes of e-learning in healthcare continuing education: a systematic review. *Nurse Educ Pract*. 2026;94:104866. doi:10.1016/j.nepr.2026.104866.

Author and Publication Statements

Author contributions: Dr Sadhana Gadia contributed clinical perspective, critical review and interpretation. Varun Gadia contributed product framing, literature synthesis, implementation analysis and manuscript development. Both authors approved the final manuscript.

Originality statement: The manuscript was newly written as a synthesis of cited sources and product information. No plagiarism score is guaranteed because similarity platforms use proprietary databases and settings; direct copying was avoided and factual material was paraphrased and attributed.

Funding: No external research funding was reported for this narrative review. **Conflicts of interest:** Varun Gadia is associated with Hidoc Dr; this relationship is disclosed so that product-centred interpretation can be evaluated transparently.