ISSN: 2347-4688, Vol. 14, No.(2) 2026, Pg. 265-278

Current Agriculture Research Journal

www.agriculturejournal.org

Generative AI and Chatbot-Based Advisory Systems in Agricultural Extension: A Comprehensive Review with Insights from Rajasthan, India

SURESH KUMAR SHARMA[1]* and PARUL DHULL[2]

[1]Department of Statistics, Mathematics and Computer Science, SKN College of Agriculture, Sri Karan Narendra Agriculture University, India

[2]Department of Computer Science, College of Dairy Science and Technology, Sri Karan Narendra Agriculture University, India

Abstract

Agricultural extension in most developing countries suffers from a persistent 'last-mile' problem: too few trained personnel, uneven coverage of remote and arid regions, and advice that rarely matches the specific agro-climatic conditions of an individual farm. Generative artificial intelligence (GenAI), and in particular large language model (LLM)-based conversational chatbots, has recently been put forward as a scalable complement to human extension networks. Drawing on a structured narrative synthesis of 30 peer-reviewed, pre-print, and institutional sources identified through systematic database and grey-literature searches, this review synthesises the fast-growing literature on GenAI and chatbot-based advisory in agriculture, tracing its evolution from rule-based expert systems and SMS/IVR advisory to retrieval-augmented, multilingual, voice-first conversational agents. It examines the enabling technologies — large language models, retrieval-augmented generation, and multimodal, multilingual natural language processing — and reviews global and Indian deployments, including Digital Green's Farmer. Chat, IFPRI's Generative AI for Agriculture (GAIA) initiative, and India's newly launched Bharat-VISTAAR platform, inaugurated in Jaipur, Rajasthan, in February 2026. Evidence on farmer trust, adoption, and socio-technical barriers is critically evaluated, with particular attention to the semi-arid, drought- and locust-prone agro-ecology of Rajasthan and to unresolved challenges around low-resource Indic and dialectal language coverage, data sovereignty, and human-oversight requirements. Concretely, the reviewed evidence shows automatic-speech-recognition error rates for Indian languages climbing sharply outside Hindi, multilingual large language models measurably more prone to factual error in low-resource Indic languages than in English, and a 2026 producer survey in which just under half of respondents report weekly use of general-purpose AI tools while only about

Article History

Received: 06 August 2026

Accepted: 26 August 2026

Keywords

Agricultural Extension;
Bharat-VISTAAR;
Chatbot Advisory;
Digital Agriculture;
Generative Artificial Intelligence;
Large Language Models;
Rajasthan;
Smallholder Farmers.

CONTACT Suresh Kumar Sharma suresh.cs@sknau.ac.in Department of Statistics, Mathematics and Computer Science, SKN College of Agriculture, Sri Karan Narendra Agriculture University, India

Doi: http://dx.doi.org/10.12944/CARJ.14.2.2

a quarter say they trust its agricultural recommendations without independently cross-checking them; a large-scale randomised evaluation of non-GenAI digital advisory in Odisha, by contrast, found only modest average yield gains that were markedly larger for farmers facing weather shocks — a reminder that platform scale and demonstrated impact do not automatically track one another. The review identifies persistent gaps — thin longitudinal impact evidence, negligible coverage of regional dialects such as Marwari and Mewari, unresolved liability questions for erroneous advice, and weak integration with local Krishi Vigyan Kendra (KVK) networks — and sets out a research and implementation agenda for context-specific, human-in-the-loop GenAI advisory suited to arid and semi-arid smallholder agriculture.

Abbreviations

GenAI	Generative Artificial Intelligence;
LLM	Large Language Model;
RAG	Retrieval-Augmented Generation;
NLP	Natural Language Processing;
ASR	Automatic Speech Recognition;
ATMA	Agricultural Technology Management Agency;
KVK	Krishi Vigyan Kendra;
DPI	Digital Public Infrastructure;
UPI	Unified Payments Interface;
ICT	Information and Communication Technology;
IVR	Interactive Voice Response;
SMS	Short Message Service;
FAO	Food and Agriculture Organization;
IFPRI	International Food Policy Research Institute;
GAIA	Generative AI for Agriculture;
CGIAR	Consultative Group on International Agricultural Research;
ICAR	Indian Council of Agricultural Research;
PIB	Press Information Bureau.

Introduction

Agriculture remains the primary livelihood for close to half of India's workforce, and in a semi-arid state such as Rajasthan — where roughly four-fifths of cultivated land is rain-fed and farming is repeatedly disrupted by drought, hailstorms, frost, and locust incursions — the timeliness and accuracy of agronomic advice can decide whether a season ends in a viable harvest or a loss.[1] This advice has historically travelled through public extension systems built around the Training-and-Visit model and, more recently, the Agricultural Technology Management Agency (ATMA) framework, supplemented by Krishi Vigyan Kendras (KVKs), Kisan Call Centres, and a growing mix of private and non-governmental advisory services. Independent assessments, however, have long shown that only a small fraction of farmers ever receive advice directly from a public extension worker, a shortfall traced to chronic staff shortages, vacant extension posts, and the sheer geographic dispersion of India's roughly 140 million operational holdings.[2]

Digital technologies have been proposed as a partial remedy for over a decade now, beginning with SMS- and Interactive Voice Response (IVR)-based push advisory and moving through smartphone applications for pest and disease identification. What sets the current moment apart is generative AI: large language models (LLMs) capable of open-ended, conversational, and increasingly multimodal interaction, able in principle to fuse agronomic science, real-time weather and market data, and a farmer's own unstructured query into personalised advice delivered on demand, in a local language, at near-zero marginal cost per interaction.[3,4] Pilot systems such as Digital Green's Farmer.Chat,[5] FBN's Norm,[6] and government platforms including India's Bharat-VISTAAR — launched from Jaipur in February 2026[7] — indicate that this is no longer a purely experimental technology but one already being deployed at national scale.

This shift has produced a correspondingly large and fast-moving literature, spread across computer science venues concerned with model architecture and multilingual natural language processing, agricultural economics studies of extension impact, and development-policy analyses of digital public infrastructure. Much of it, though, remains fragmented. Technical papers on retrieval-augmented generation for agricultural question answering rarely engage with the adoption literature, and adoption studies rarely dig into the model limitations that drive farmer distrust in the first place.[8,9] Existing reviews have also tended to treat AI in agriculture as one broad, largely technology-agnostic category dominated by computer-vision applications such as pest and disease detection,[10,11] leaving little room for a focused synthesis of the conversational-advisory sub-domain on its own terms. Concretely, the two most recent broad reviews of AI in agriculture are bibliometric mappings of agri-robotics and of deep learning for crop monitoring, respectively — exercises that chart publication trends and technique clusters across the field as a whole rather than critically synthesising the substantive evidence on any one application.[10,11] Neither engages the technical architecture of conversational-advisory systems, the adoption and trust literature, or the governance questions that this review treats as its core subject matter, and neither offers a state- or region-specific reading of the kind attempted here for Rajasthan. The contribution of the present review is accordingly threefold rather than purely topical: it is the first synthesis, to the authors' knowledge, to read the GenAI-advisory literature specifically through an arid, drought- and locust-prone smallholder agro-ecology; it explicitly cross-references the technical RAG/multilingual literature (Section 4) against the adoption and trust literature (Section 6), which the fragmentation noted above shows rarely happens in the same paper; and it treats Bharat-VISTAAR's Rajasthan launch as an analytic anchor rather than a passing example.

This review is written to close that gap. It deliberately narrows in on generative AI and chatbot-based advisory in agricultural extension — setting aside the already well-reviewed literature on AI in precision agriculture, robotics, and remote sensing — and reads the evidence through the specific agro-ecological and institutional lens of Rajasthan. Three objectives follow: first, to trace the technological and institutional path from static expert systems to conversational GenAI advisory; second, to pull together evidence on real deployments, trust, and adoption from both global and India-specific sources; and third, to name concrete, underexplored research gaps — particularly around low-resource regional dialects, human-in-the-loop safeguards, and longitudinal impact evaluation — that can inform both future research and the design of platforms such as Bharat-VISTAAR as they expand into arid and semi-arid states.

Review Methodology

This review follows a structured narrative synthesis rather than a formal meta-analysis, a choice that reflects how heterogeneous the underlying evidence base still is — a mix of peer-reviewed journal articles, pre-print technical reports, and grey literature from development organisations and government sources that is typical of a technology domain moving faster than the peer-review cycle can keep up with. Literature was located through structured searches, conducted between March and July 2026, of Scopus- and Web of Science-indexed journals (Computers and Electronics in Agriculture, Artificial Intelligence in Agriculture, Agricultural Systems, and Frontiers in Artificial Intelligence among them), the arXiv pre-print repository, Google Scholar, and institutional sources (FAO, IFPRI/CGIAR, the Government of India's Press Information Bureau). Searches combined the terms 'generative AI', 'large language model' or 'LLM', 'chatbot', 'conversational AI', and 'retrieval-augmented generation' or 'RAG' with 'agricultural extension', 'farm advisory', and 'digital advisory', restricted for the India-specific search arm with 'India' or 'Rajasthan' — for example, the Boolean string ("generative AI" OR "large language model*" OR chatbot OR "conversational AI") AND ("agricultural extension" OR "farm advisory" OR "digital advisory") AND (India OR Rajasthan). Reference lists of retrieved articles and of recent reviews of AI in agriculture were also hand-searched for additional eligible sources (backward snowballing).

Records identified through these searches were first screened by title and abstract against four eligibility criteria: sources had to report on (i) the technical architecture of agricultural conversational AI systems; (ii) empirical deployment or evaluation of chatbot or voice advisory tools with farmers or extension agents; (iii) adoption, trust, or socio-technical barrier studies relevant to digital or AI-based advisory; or (iv) ethics, data-governance, or policy analysis specific to AI in agricultural advisory. Sources concerned mainly with computer-vision applications (image-based disease detection, for instance) or autonomous field robotics were excluded at this stage unless they intersected directly with conversational advisory delivery, consistent with this review's narrower scope, as were sources with no substantive agricultural application. Records passing title/abstract screening were then read in full to confirm relevance and to extract the specific claims synthesised below. This process yielded the 30 sources cited in this review, drawn predominantly from work published between 2019 and mid-2026 and weighted toward 2024–2026 given how quickly the underlying LLM technology has been changing; a small number of foundational pre-GenAI digital-extension studies were retained outside this window because they establish baseline evidence on ICT-advisory impact against which GenAI-specific claims can be judged.[2,12,13]

Because the evidence base spans several tiers of scholarly rigor, each source was additionally classified by type — peer-reviewed journal article, arXiv or other pre-print, institutional or government report, industry/survey report, or general-audience press coverage — and this classification is carried through into how findings are reported below: peer-reviewed and pre-registered empirical results (for example, the randomised evaluation of digital extension in Odisha discussed in Section 3.2, or the automatic-speech-recognition benchmarking in Section 4.3) are presented as established findings, while claims drawn from platform-reported metrics, industry surveys, or descriptive government and press sources are explicitly flagged as such rather than treated as independently verified evidence. The same evidence-type classification is carried into Table 1, which tags the predominant evidence base behind each row.

One caveat is worth stating plainly: because several key platforms discussed here — Bharat-VISTAAR most notably — launched only months before this review was written, peer-reviewed, empirical impact evidence for the newest deployments is thin almost by definition, and the sources available for them are necessarily government press releases and descriptive reporting rather than independent evaluation. Where that is the case, it is flagged rather than papered over, and such sources are treated as descriptive rather than as proof of effectiveness; claims about Bharat-VISTAAR's design and rollout in Sections 5.2 and 5.3 are accordingly drawn from official government communications and are explicitly distinguished from the peer-reviewed evidence used elsewhere in this review.

Results

From Static Expert Systems to Conversational Advisory: An Evolutionary Overview

Figure 1 summarises the four broad phases through which agricultural information delivery in India has passed, from face-to-face extension to today's generative AI chatbots; each phase is discussed in turn below.

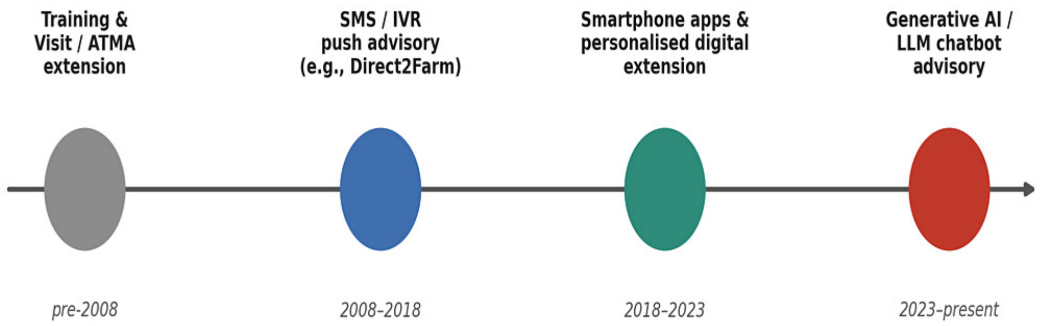

Fig.1: Evolution of agricultural extension and advisory delivery in India, from face-to-face Training-and-Visit/ATMA extension through SMS/IVR push advisory and smartphone-based personalised extension to current-generation generative AI chatbot advisory.

The Limits of Conventional Extension

Public agricultural extension in India has historically run on a one-to-many, top-down transmission model, in which a limited cadre of subject-matter specialists and village-level workers is expected to reach a vastly larger, geographically dispersed farmer population. Nationally representative surveys conducted over the past two decades have repeatedly found that only a small, single-digit percentage of farm households report receiving information from a public extension worker in a given season, with most instead relying on informal networks — other farmers, input dealers, and, increasingly, mobile phones.[2] In Rajasthan specifically, large average landholdings spread across arid and semi-arid terrain, seasonal out-migration, and a KVK network that, while active, cannot realistically conduct frequent individual farm visits across the state's 33-plus districts have all constrained the reach of face-to-face extension.[1]

The ICT-advisory era: SMS, IVR, and mobile apps

The first wave of digital response to this gap, running roughly from 2008 to 2020, centred on relatively low-bandwidth, push-based technologies: bulk SMS advisories, IVR call-in systems, and later smartphone applications. CABI's Direct2Farm programme, piloted from 2014 across six states including Rajasthan, is a representative example — delivering standardised mobile-based agro-advisory content meant to complement, not replace, human extension.[2] Meta-analyses of such services across several countries have generally found modest but positive average effects on yields and practice adoption, though a recent randomised evaluation of a large-scale digital extension programme in Odisha found meaningfully smaller gains at scale than earlier pilot studies had suggested — a useful reminder that impact estimated from a tightly managed pilot does not always survive the jump to a broad, non-targeted rollout.[13] Personalisation — tailoring advice to a specific farmer's crop, soil, and location rather than broadcasting uniform content — emerged from this era as a recurring correlate of stronger outcomes,[12] which helps explain the appeal of the conversational, query-responsive systems that followed.

The Generative AI Turn

The period from roughly 2023 onward marks a genuine break with what came before, driven by the availability of general-purpose LLMs capable of fluent, context-aware natural language generation. Earlier digital advisory was fundamentally a broadcast medium: farmers received pre-authored content matched to broad categories. GenAI chatbots, by contrast, support open-ended, farmer-initiated dialogue — a farmer can describe a specific symptom, ask a follow-up question, or pose a request in colloquial, unstructured language, and receive a synthesised, conversational reply. The literature describes this as a move from static, rule-based expert systems to dynamic foundation models capable of complex reasoning and multimodal analysis,[14] and it is this technological foundation that underlies the case studies examined in Section 5.

Enabling Technologies

Large Language Models and Retrieval-Augmented Generation

Most agricultural GenAI chatbots are built on general-purpose foundation models — the GPT family, LLaMA, and similar transformer architectures — rather than models trained from scratch on agricultural data, since the cost of pre-training from the ground up is prohibitive for most agricultural research or development budgets. To ground these general-purpose models in reliable, domain-specific knowledge and reduce the risk of confidently generated but incorrect responses (hallucinations), the dominant architectural pattern is retrieval-augmented generation (RAG): the model retrieves relevant passages from a curated corpus — agronomic manuals, extension bulletins, crop-specific research — before generating its answer, rather than relying solely on knowledge implicitly baked in during pre-training. Both Digital Green's Farmer.Chat and IFPRI's GAIA initiative adopt RAG architectures over curated CGIAR and CABI content for exactly this reason,[5,15] and work on Bengali-language advisory has shown that cross-lingual RAG pipelines — where a local-language query is translated, matched against an English-language evidence base, and translated back — can deliver source-grounded answers even for languages with limited digital resources.[16] Figure 2 sketches a representative end-to-end pipeline of this kind, including the human-in-the-loop checkpoint for high-stakes advice discussed further in Section 6.2.

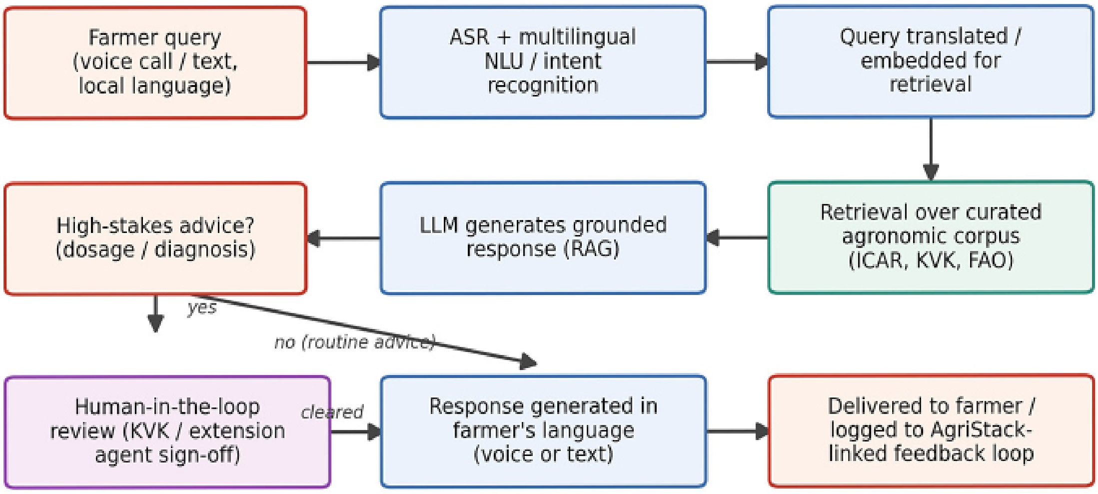

Fig. 2: Representative retrieval-augmented generation (RAG) pipeline for an agricultural advisory chatbot, from farmer query through source-grounded response generation to an optional human-in-the-loop checkpoint for high-stakes advice.

Multimodal and Voice-First Interfaces

A large share of the intended user base for agricultural advisory — older farmers, women, and those with limited formal education in particular — cannot reliably interact with a text-based interface. This has pushed design attention toward voice-first and multimodal systems: Bharat-VISTAAR, for instance, is built around a simple phone call rather than a smartphone app,[7] while several research prototypes accept a photograph of a diseased leaf alongside a voice or text query. Multimodality of this kind adds real technical difficulty, though, since automatic speech recognition (ASR) has to hold up on accented, dialectal, often noisy rural speech before any downstream language understanding can even begin.[17]

Multilingual and Low-Resource Language Challenges

India's linguistic diversity is a distinct and only partially solved problem for agricultural GenAI. LLMs perform comparatively well in high-resource languages such as English and Hindi, but that performance drops markedly for lower-resource regional languages and, more sharply still, for largely unwritten or under-digitised dialects such as Marwari, Mewari, and Dhundhari that are widely spoken across rural Rajasthan yet barely represented in training corpora.[18,19] A 2026 benchmarking study of automatic speech recognition across Indian languages in agricultural settings gives a sense of scale: word error rates for widely used commercial speech-to-text systems, strong on Hindi, climb sharply for lower-resource languages — in one case exceeding the point at which a transcript is practically unusable.[17] A companion strand of evaluation work has found multilingual LLMs measurably more prone to factual error and outright fabrication when queried in low-resource Indic languages than when the identical query is posed in English,[20] a finding with obvious stakes for agronomic advisory, where a fabricated pesticide dosage or planting date is not a trivial mistake. Efforts such as AI4Bharat's IndicNLP resources, IndicBERT, and emerging India-specific foundation models (including Bhashini-linked and Sarvam-based architectures) are narrowing this gap,[19,21] and domain-specific work — the AgriGov multilingual scheme-information corpus[22] and AgriSaathi, which pairs multilingual speech interfaces with acoustic models adapted to rural speech and agricultural vocabulary[18] — represents an early attempt to close it for farm advisory specifically. None of these, however, yet offer validated coverage of Rajasthan's principal regional dialects.

Global and Indian Deployments: An Applied Review

Global Pilots And Platforms

Digital Green's Farmer.Chat, piloted in India and sub-Saharan Africa, is among the most extensively documented GenAI agricultural chatbots.[5] Built on a RAG architecture over structured and unstructured agronomic content — research papers, crop tables, instructional video transcripts — it is designed explicitly to remove dependence on human intermediaries, a design choice aimed at widening reach among women and low-literacy users who have historically been under-served by intermediary-dependent extension.[5] IFPRI's Generative AI for Agriculture (GAIA) initiative, run with CABI, the University of Florida, and Digital Green, used Farmer.Chat deployments in Kenya and India as a testbed for a second, broader phase (2025–2027) aimed at expanding content aggregation, adding real-time data and multimodal crop-health imaging, and building a dedicated GenAI ethics toolkit for advisory contexts[15] — an explicit institutional acknowledgement that technical capability on its own is not enough for responsible deployment. It is worth flagging that Farmer.Chat's reach figures are self-reported by Digital Green rather than independently audited — a caveat that applies to most of the platform-level statistics cited in this section and is revisited in Section 6.2's discussion of the accuracy–trust gap.

In the commercial sector, Farmers Business Network's ‘Norm’ assistant, built on a GPT-3.5 base and grounded in USDA statistical data and proprietary agronomic datasets, illustrates both the promise and the risk calculus of GenAI advisory. Its developers have been blunt that a hallucinated recommendation for a regulated agricultural input is not an acceptable failure mode given the direct financial and safety stakes for a farm operation, and have engineered grounding mechanisms accordingly.[6] A comparable pattern shows up in aquaculture, where systems such as AquaGPT and SmartAqua AI apply the same RAG-based, sensor-integrated advisory approach to fisheries,[23] suggesting the underlying architecture is transferring across agricultural sub-sectors rather than staying crop-specific.

Evidence on the extension-agent side of adoption — as distinct from direct farmer-facing use — is also starting to accumulate. A 2026 study of 240 agricultural extension agents in Benin, applying the Technology Acceptance Model, found that agents' engagement with generative AI tools was associated with measurable changes in advisory performance, while also surfacing familiar constraints around data quality, infrastructure access, and accessibility.[24] This agent-mediated pathway — GenAI as a productivity tool for human extension workers rather than a direct farmer-facing replacement — is an important and comparatively under-examined deployment model relative to the farmer-facing chatbot literature.

India-Specific Developments

India's public digital-extension landscape has moved through several generations: the Kisan Call Centre (a live, human-staffed toll-free advisory service), the mKisan SMS portal, app-based services, and NGO-led programmes such as Direct2Farm, arriving now at a GenAI-native generation. The most significant recent development is Bharat-VISTAAR (Virtually Integrated System to Access Agricultural Resources), announced in the Union Budget 2026–27 and launched in Phase 1 on 17 February 2026 — notably, from Jaipur, Rajasthan, by the Union Agriculture Minister alongside the Rajasthan Chief Minister.[7] Government and independent reporting describe it as an AI-driven, multilingual Digital Public Infrastructure (DPI) for agriculture, delivering personalised crop, weather, pest, market, and government-scheme information through phone calls, a chatbot, and, in a later phase, a mobile app.[7] Phase 1 launched in Hindi and English on 17 February 2026; a follow-up government release the following month indicated that four additional languages were planned within three months of launch and five more within six, though as of this review's writing that expanded rollout has not itself been independently verified.[25] The platform draws on AgriStack and ICAR scientific content as its knowledge base, and government officials at its launch explicitly described it as a Digital Public Infrastructure for agriculture "just like UPI is for payments" — a common, interoperable digital backbone rather than a single closed app.[7]

Complementary India-focused research illustrates both the promise and the specificity this domain demands. A 2026 study on bridging agronomic science and farm-level advisory for rice systems in India describes GenAI as a natural-language interface that can make crop simulation models and remote-sensing diagnostics accessible to farmers, while cautioning explicitly that this is conditional on resolving data-sovereignty questions, closing rural infrastructure gaps, and keeping humans in the advisory loop[14] — cautions just as applicable to a semi-arid, drought-prone state such as Rajasthan as to the irrigated rice systems the study examines. Separately, a 2021 study using primary survey data from Indian smallholders found personalised digital extension — as against generic broadcast advisory — positively associated with farm performance,[12] offering an evidentiary bridge between the pre-GenAI digital-extension literature and the personalisation GenAI now promises.

The Rajasthan Context

Rajasthan presents both an unusually pressing need for, and unusually demanding conditions on, GenAI-based advisory. With roughly four-fifths of cultivated area rain-fed and the state chronically exposed to drought, hailstorm, frost, and — particularly in the western districts of Jaisalmer, Barmer, and Bikaner — locust incursions, timely, location-specific advisory carries outsized economic weight relative to more agro-climatically stable regions.[1] The state government's Rajasthan Agriculture Portal already consolidates services from more than ten departments, spanning crop insurance access, Direct Benefit Transfer integration, and a Soil Health Mission, reflecting an existing institutional push toward digital consolidation that a conversational AI layer such as Bharat-VISTAAR is now being layered onto.[1] That Jaipur was chosen as the national launch site for Bharat-VISTAAR's Phase 1 underscores the state's positioning within this rollout — though it also means Rajasthan is, in effect, an early and largely undocumented test case, since no peer-reviewed, state-specific impact evaluation of the platform yet exists.[7,25]

University- and district-level digital initiatives in Rajasthan, including SKNAU's own portfolio of digital agricultural services developed through its Centre for Information Management and Computer Applications — weather-advisory portals, seed-distribution systems, and examination and extension-education digital infrastructure among them — point to the kind of institutional digital capacity already present at the state agricultural university level that could, in principle, supply verified regional agronomic content for a Rajasthan-tuned RAG knowledge base, directly addressing one of the technical gaps identified in Section 4.3.

Discussion

Adoption, Trust, and Socio-Technical Barriers

Figure 3 summarises three findings from a 2026 producer survey that together capture the shape of the trust problem discussed in this section: substantial weekly use of general-purpose AI tools, but widespread and habitual cross-checking of AI-generated advice, alongside real concern about how farm data might be used downstream.[8]

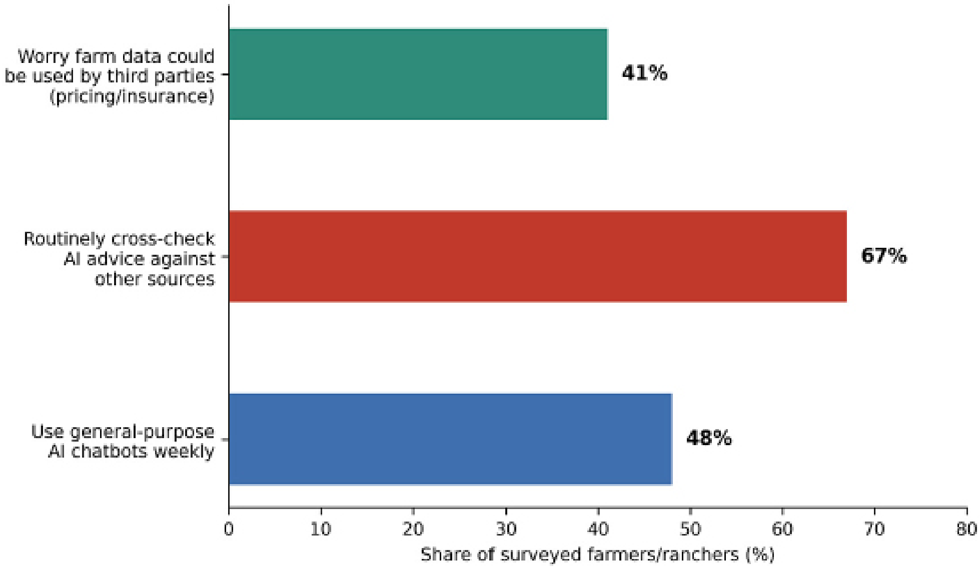

Fig. 3: Farmer/rancher-reported AI usage, trust, and data-privacy concern, drawn from the 2026 MorganMyers/Ag Access producer survey.[8]

The Digital Divide and Differential Adoption

Even where GenAI advisory tools are technically available, adoption is neither automatic nor evenly spread. Cross-country evidence on digital extension more broadly points to infrastructure and network coverage, device and connectivity cost, and low digital literacy as the most persistent barriers to inclusive uptake,[26,27] and 2026 industry survey data shows adoption skewing toward younger, larger-scale, better-resourced operations even in high-income farming contexts[8] — a pattern with obvious implications for smallholder- and dryland-dominated Rajasthan, where landholding size, age profile, and connectivity vary sharply from district to district. Gender is a related and often under-addressed dimension: comparative research on digital extension in Bangladesh and India finds that lower mobile-phone ownership and internet use among rural women, combined with mobility restrictions and limited land ownership, compounds their exclusion from digital advisory generally[27] — a risk that voice-first, phone-call-based designs such as Bharat-VISTAAR may partly ease but do not eliminate.

Trust, Accuracy, and the Hallucination Problem

Trust is a second, distinct barrier from access. Survey evidence from 2026 shows that even among farmers who use general-purpose AI chatbots regularly, a majority routinely cross-check AI-generated guidance against other sources rather than acting on it directly, and producers have self-reported episodes of chatbots hallucinating pesticide rates, misidentifying weeds from photographs, and recommending planting dates that ignore local frost risk in exactly those terms.[8] This caution is not irrational. Unlike a poor restaurant recommendation, an erroneous pesticide dosage or an incorrect sowing window carries direct economic and, in the case of agrochemical misuse, safety consequences.[6] The literature converges on a consistent recommendation here: GenAI advisory should be positioned explicitly as a decision-support tool operating with human oversight and source grounding via RAG (Section 4.1), not as an autonomous decision-maker, particularly for high-stakes recommendations such as agrochemical application rates.[28,29]

The Extension-Agent-Mediated Pathway

As the Benin Technology Acceptance Model study discussed in Section 5.1 suggests, one plausible way to reconcile the trust gap with GenAI's scalability promise is agent-mediated deployment: extension workers, KVK scientists, and input dealers use GenAI tools to augment rather than bypass their own advisory role, keeping human judgement at the point of farmer contact while gaining the speed and breadth of AI-assisted knowledge search.[24] For a state such as Rajasthan, where the KVK network — thin as it is relative to the farmer population — retains substantial farmer trust and local agro-ecological knowledge, this hybrid model may be a more immediately viable adoption pathway than a purely farmer-facing chatbot, at least in an initial deployment phase.

Ethical, Data Governance, and Policy Considerations

Scaling GenAI advisory into public digital infrastructure raises governance questions that go beyond model accuracy. Data ownership and control feature prominently in the literature: farmers surveyed in 2026 industry research worried that farm-level data — yields, input costs, soil conditions — collected through AI advisory interactions could be used by third parties to influence input pricing or insurance terms,[8] echoing a broader ethical-AI concern about the imbalance of power between farmers and the technology providers or aggregators that control both the data and the underlying models.[28,29] For a state-run platform integrated with AgriStack, as Bharat-VISTAAR is, this question carries added weight given the scale of data aggregation involved and the platform's explicit ambition to feed aggregated farmer-interaction data back into national research and policy priorities.[25]

Algorithmic bias and fairness raise a related concern: advisory systems trained mainly on data from higher-resource languages, better-connected regions, or larger and more commercially oriented farms risk producing recommendations that are systematically less accurate or less relevant for exactly the smallholder, low-resource-language, and remote populations digital extension is meant to reach — a risk directly analogous to the multilingual factual-accuracy gap documented in Section 4.3.[20] The ethical-AI literature on agriculture converges on a broadly shared set of principles — transparency about AI involvement in advisory delivery, farmer control over personal and farm data, active bias auditing, and deliberate design for inclusivity of marginalised users[29,30] — that maps closely onto the specific risks identified here for Rajasthan's smallholder- and dryland-dominated farming population. Translating those principles into enforceable governance for a rapidly scaling national platform, though, remains an open policy problem rather than a solved one as of this review.

Critical Synthesis: Opportunities and Limitations

Table 1 draws the review's findings together into a structured assessment of where GenAI chatbot-based advisory currently offers a genuine advantage over earlier digital-extension models, and where real limitations persist, organised around the four cross-cutting themes examined in Sections 4 through 7. A fourth column tags the predominant type of evidence behind each row, following the source-classification scheme set out in the Review Methodology, so that platform-reported and descriptive claims remain visibly distinct from peer-reviewed empirical findings.

Table 1: Synthesis of opportunities and limitations of generative AI chatbot-based agricultural advisory, drawn from the reviewed literature.

Dimension	Opportunity	Persistent limitation	Predominant evidence base
Reach and personalisation	Conversational, on-demand advisory tailored to an individual query, available around the clock without a human intermediary[5,12]	Effective reach still gated by connectivity, device access, and digital literacy; adoption skews toward younger, better-resourced users[8,26]	Platform-reported metrics; one peer-reviewed RCT (Odisha); industry survey
Language and accessibility	Voice-first, multilingual design (e.g., Bharat-VISTAAR's phone -call access) lowers the literacy barrier relative to app- or text-based tools[7]	Performance and factual accuracy degrade sharply for low-resource regional languages and are largely untested for unwritten dialects such as Marwari and Mewari[17,18,20]	Government/institutional descriptive sources; peer-reviewed ASR and LLM benchmarks

Accuracy and trust	Retrieval-augmented generation grounded in curated agronomic sources can substantially reduce hallucination relative to ungrounded general-purpose chatbots[5,15,16]	Documented cases of hallucinated dosages and misidentified pests persist; farmers routinely and appropriately cross-check AI advice[6,8]	Platform/institutional descriptions; one pre-print; press and industry survey (self-report)
Institutional integration	Potential to augment, rather than replace, KVK and extension -agent workflows, combining AI-assisted knowledge retrieval with retained human judgement[24]	Agent-mediated deployment models remain comparatively under-studied relative to direct farmer-facing chatbots; integration pathways with existing KVK networks are largely undocumented	One peer-reviewed survey (n=240, Benin); no Rajasthan/KVK-specific evidence found
Governance and equity	Explicit ethics toolkits and DPI-style interoperable design (as pursued by GAIA and Bharat-VISTAAR respectively) signal growing institutional awareness of governance needs[15,25]	Data ownership, third-party use of farm data, and algorithmic bias against smallholder and low-resource-language users remain largely unresolved at policy level[8,28,29]	Institutional/government descriptive sources; industry survey; normative ethics literature

Research Gaps and Future Directions

Pulling the preceding sections together, five research gaps stand out as both consequential and currently underexplored, with particular relevance to Rajasthan and comparable semi-arid, linguistically diverse Indian states.

1. Dialectal and low-resource language coverage: existing multilingual agricultural NLP efforts — AgriSaathi, AgriGov, IndicBERT-based systems [18,19,22] — have made progress on major scheduled languages but offer little validated coverage of regional dialects such as Marwari, Mewari, Dhundhari, or Wagdi, which remain the primary spoken register for many rural users even where Hindi is understood as a second language. Building and validating agriculture-domain speech and text corpora for these dialects is a concrete, tractable research priority.
2. Longitudinal, independent impact evaluation: much of the evidence base for GenAI-specific advisory, as distinct from earlier SMS/IVR advisory, remains descriptive or short-horizon. Rigorous, independent evaluation of yield, income, and input-use outcomes attributable specifically to chatbot-based advisory — as opposed to broader digital-extension bundles — is still largely missing,[12,13] and this gap is especially acute for Bharat-VISTAAR given how recently it launched.
3. Human-in-the-loop and liability frameworks: while the literature broadly recommends retaining human oversight for high-stakes recommendations,[28,29] few studies specify concrete operational protocols — for example, which categories of advice (agrochemical dosage, disease diagnosis, credit-scheme eligibility) should require mandatory human sign-off before reaching a farmer, and who bears liability when an AI-generated recommendation causes economic loss.
4. Integration with local knowledge institutions: drawing on state agricultural university and KVK-generated, agro-ecologically localised content — soil type, locally validated crop varieties, region-specific pest calendars — as a RAG knowledge source, rather than relying solely on national or international agronomic corpora, is theoretically attractive but empirically almost entirely untested for arid-zone contexts such as Rajasthan.
5. Gender- and equity-disaggregated adoption research: existing adoption studies rarely disaggregate uptake and outcomes by gender, landholding size, or social category within the

Indian context specifically,[27] limiting the ability of both researchers and platform designers to spot and correct exclusionary design patterns before they become entrenched at national scale.

Conclusion

Generative AI and chatbot-based advisory mark a genuinely new chapter in agricultural extension, not merely an incremental extension of earlier SMS- and IVR-based digital advisory, offering conversational, personalised, on-demand interaction at a scale human extension networks alone cannot match. The evidence reviewed here — spanning Digital Green's Farmer.Chat, IFPRI's GAIA initiative, commercial systems such as FBN's Norm, and India's own Bharat-VISTAAR platform launched from Jaipur, Rajasthan, in February 2026 — shows both the growing technical sophistication of retrieval-augmented, multilingual conversational architectures and the persistence of real, unresolved risks around low-resource language accuracy, farmer trust, data governance, and equitable access. It is worth being explicit about which of these is which: the technical feasibility of RAG-grounded, voice-first advisory is now reasonably well established, and platform-reported reach figures are large, but rigorous, independent evidence that this reach translates into measurable yield, income, or risk-reduction gains specifically attributable to GenAI advisory — as opposed to digital advisory more broadly — remains thin, and broader claims of proven impact for these platforms should be read as promise rather than as demonstrated outcome. For Rajasthan specifically — a state combining high agro-climatic vulnerability with substantial linguistic diversity and an established, if resource-constrained, state agricultural university and KVK network — GenAI advisory holds real promise but needs deliberate, locally grounded adaptation rather than uncritical adoption of platforms designed and validated elsewhere. The research agenda set out above in Research Gaps and Future Directions — dialectal language coverage, rigorous impact evaluation, human-in-the-loop governance, local-institution integration, and equity-disaggregated adoption research — offers a concrete path toward realising that promise responsibly as platforms such as Bharat-VISTAAR continue their national rollout.

Acknowledgement

The authors gratefully acknowledge Sri Karan Narendra Agriculture University (SKNAU), Jobner, and the Centre for Information Management and Computer Applications (CIMCA) for institutional support during the preparation of this review.

Funding Sources

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Conflict of Interest

The authors do not have any conflict of interest.

Data Availability Statement

This statement does not apply to this article, as it is a narrative literature review and does not report new primary datasets.

Ethics Statement

This research did not involve human participants, animal subjects, or any material that requires ethical approval.

Informed Consent Statement

This study did not involve human participants, and therefore, informed consent was not required.

Clinical Trial Registration

This research does not involve any clinical trials.

Permission to Reproduce Material from other Sources

Not Applicable. Figures 1-3 and Table 1 in this review were prepared originally by the authors and do not reproduce previously published material.

Author Contributions

Each author mentioned has significantly and directly contributed intellectually to this review and has given approval for its publication.

- **Suresh Kumar Sharma:** Conceptualization, Literature Review, Writing – Original Draft, Writing – Review & Editing, Supervision.
- **Parul Dhull:** Literature Review, Writing – Review & Editing.

References

1. Government of Rajasthan, Department of Agriculture. Rajasthan Agriculture Portal — overview of digital services. Jaipur: Government of Rajasthan.
2. Kansiime MK, Alawy A, Allen C, Subharwal M, Jadhav A, Parr M. Effectiveness of mobile agri-advisory service extension model: evidence from Direct2Farm program in India. *World Dev Perspect.* 2019;13:25-33.
3. Li M, Basit A, Javed A, Song X, Ge Y. Generative artificial intelligence and ChatGPT in agriculture supply chain management: a systematic literature review and future research agenda. *Br Food J.* 2026;128(13):121-141.
4. Ryan M, Isakhanyan G, Tekinerdogan B. An interdisciplinary approach to artificial intelligence in agriculture. *NJAS Impact Agric Life Sci.* 2023;95(1):2168568.
5. Singh N, Wang'ombe J, Okanga N, et al. Farmer.Chat: scaling AI-powered agricultural services for smallholder farmers. arXiv. 2024;arXiv:2409.08916.
6. Macready H. Farmers are tentatively embracing AI. Offrange (formerly Ambrook Research). Published October 6, 2023. Accessed August 2026.
7. Press Information Bureau, Government of India. AI farmer revolution to begin from Jaipur: Union Minister Shri Shivraj Singh Chouhan to launch 'Bharat-VISTAAR'. PIB, Ministry of Agriculture & Farmers Welfare. Published February 17, 2026 (PRID 2228842). Accessed August 2026.
8. MorganMyers, Ag Access. 2026 Farmer Trust in AI Report: nearly half of farmers use AI weekly, but trust hinges on data control and accuracy. MorganMyers; 2026. Accessed August 2026.
9. Joosse T. Building trust in AI farming tools. *CSA News.* 2024;69(2):10-16.
10. Casini S, Ducange P, Marcelloni F, Pollini L. Artificial intelligence in agri-robotics: a systematic review of trends and emerging directions leveraging bibliometric tools. Robotics. 2026;15(1):24.
11. Zhang R, Wu X, Li J, Zhao P, Zhang Q, Wuri L, Zhang D, Zhang Z, Yang L. A bibliometric review of deep learning in crop monitoring: trends, challenges, and future perspectives. Front Artif Intell. 2025;8:1636898.
12. Rajkhowa P, Qaim M. Personalized digital extension services and agricultural performance: evidence from smallholder farmers in India. PLoS One. 2021;16(10): e0259319.
13. Cole S, Goldberg J, Harigaya T, Zhu J. Customised agricultural advice at scale: how digital extension helps Indian farmers grow more and lose less. VoxDev. Published July 2, 2025. https://voxdev.org/topic/agriculture/customised-agricultural-advice-scale-how-digital-extension-helps-indian-farmers. Accessed August 2026.
14. Gakhar S, Koo J, Patnaik GP, Singh RK. Bridging agronomic science and context specific farm-level advisory through generative AI for rice systems in India. Front Artif Intell. 2026;9:1808028.
15. International Food Policy Research Institute. Generative AI for Agriculture (GAIA) — Phase I & II. IFPRI. https://www.ifpri.org/project/generative-ai-for-agriculture-gaia/. Accessed August 2026.
16. Hossain MA, Subhan N, Mahi MR, Nabila JF. Cost-efficient cross-lingual retrieval-augmented generation for low-resource languages: a case study in Bengali agricultural advisory. arXiv. 2026;arXiv:2601.02065.
17. Chandrashekar MS, Singh V, Pedapudi L. Benchmarking automatic speech recognition for Indian languages in agricultural contexts. arXiv. 2026;arXiv:2602.03868.
18. Patne A, Ranjane S, Walke O, Owhal R. AgriSaathi: a comprehensive survey on large language models and multimodal AI for agricultural advisory systems. *Int J Multidiscip Res.* 2025;7(6).
19. Gakhar S, Singh RK. Leveraging large language models (LLMs) for GeoAI-enabled digital agro-advisory. *Front Remote Sens.* 2026;7:1839369.
20. Rohera P, Ginimav C, Sawant G, Joshi R. Better to ask in English? Evaluating

factual accuracy of multilingual LLMs in English and low-resource languages. arXiv. 2025;arXiv:2504.20022.
21. Singh A, Agar A, Kulkarni V, Asati RA, Patwari N. Research gaps in developing fair and inclusive LLMs for India's agricultural sector. *Int Res J Eng Technol.* 2025;12(11).
22. Bilal M, Gopakumar G. AgriGov: a structured multilingual dataset curation for Indian government schemes for farmers. arXiv. 2026;arXiv:2606.08272.
23. Akram W, Ud Din M, Saoud LS, Hussain I. A review of generative AI in aquaculture: foundations, applications, and future directions for smart and sustainable farming. arXiv. 2025;arXiv:2507.11974.
24. Gouroubera MW, Segnon AC, Gouthon M, Moumouni-Moussa I, Zougmoré RB. Generative AI use and advisory performance among agricultural extension agents in Benin. *Front Artif Intell.* 2026;9:1763406.
25. Press Information Bureau, Government of India. Rollout of Bharat-VISTAAR platform. PIB, Ministry of Agriculture & Farmers Welfare. Published March 13, 2026 (PRID 2239788). Accessed August 2026.
26. Food and Agriculture Organization of the United Nations. Guide on digital agricultural extension and advisory services: use of smartphone applications by smallholder farmers. Rome: FAO; 2023.
27. Food and Agriculture Organization of the United Nations. Gender-responsive digital extension and advisory services in Bangladesh and India. Rome: FAO; 2025.
28. Omotayo AO, Adediran SA, Omotoso AB, Olagunju KO, Omotayo OP. Artificial intelligence in agriculture: ethics, impact possibilities, and pathways for policy. *Comput Electron Agric.* 2025;239:110927.
29. Dara R, Hazrati Fard SM, Kaur J. Recommendations for ethical and responsible use of artificial intelligence in digital agriculture. *Front Artif Intell.* 2022;5:884192.
30. Gikunda K. Harnessing artificial intelligence for sustainable agricultural development in Africa: opportunities, challenges, and impact. arXiv. 2024;arXiv:2401.06171.