

ORIGINAL RESEARCH

Open access

Large Language Model for Generating Patient-Friendly Care Navigation Instructions from Referral Orders, Clinic Requirements, Insurance Rules, Preparation Instructions, and Scheduling Constraints

Fernando Diaz¹, Lucia Morales^{1*}, Diego Perez², Valeria Soto¹, Martin Alvarez²

Abstract

Navigating specialty care often requires patients to understand referral reasons, appointment logistics, preparation rules, insurance requirements, and follow-up expectations. These instructions are frequently distributed across separate documents and portals, creating avoidable confusion for patients and caregivers. No unified system currently converts fragmented referral, clinic, insurance, preparation, and scheduling information into one personalized, plain-language care navigation guide. As a result, patients may miss critical steps before appointments or misunderstand what they need to do. This article proposes a conceptual large language model system for generating patient-friendly care navigation instructions from clinical, administrative, and scheduling data. The objective is to describe how such a system could support clearer, safer, and more accessible patient communication. The proposed pipeline would extract relevant facts from referral orders, clinic requirements, insurance rules, preparation instructions, and scheduling constraints. A retrieval-augmented LLM would then synthesize these facts into a cohesive instruction sheet with traceability back to verified institutional sources. Conceptually, the system would generate a clear, step-by-step appointment guide tailored to the patient's language, health literacy needs, and preferred communication channel. The output would be expected to reduce cognitive burden by consolidating complex healthcare logistics into one practical message. An LLM-based patient navigation instruction system could bridge the communication gap between healthcare operations and patient understanding. Responsible deployment would require strong grounding, validation, accessibility design, and human oversight for high-risk instructions.

Keywords Large language models, generative AI, Retrieval-augmented generation, Patient communication, Patient navigation, Health literacy

*Correspondence:

Lucia Morales

lucia.morales@gmail.com

¹ Department of Healthcare Informatics and Smart Systems, Faculty of Medicine, University of Buenos Aires, Buenos Aires, Argentina

² Department of Clinical AI Engineering, Faculty of Engineering, National University of La Plata, La Plata, Argentina

Introduction

Missed appointments, incomplete preparation, and delayed follow-through often arise not only from patient choice but also from unclear navigation instructions that fail to translate healthcare logistics into actionable steps. Patient-facing discharge and appointment materials can contain dense terminology, fragmented sequencing, and assumptions about literacy or digital access, all of which can increase confusion during transitions of care [1, 2]. Studies of AI-generated

discharge and patient-facing explanations suggest that LLMs could help reframe institutional language into clearer formats, but only when outputs remain accurate, complete, and contextually appropriate [1, 3]. For referral-based care, the communication burden is especially high because patients must often coordinate clinical instructions, administrative requirements, and scheduling tasks without a single coherent guide [4, 5].

Referral navigation commonly draws on multiple disconnected information sources, including referral orders, specialty clinic letters, appointment reminders, insurance authorization notices, preparation sheets, and portal messages. Research on patient-facing AI responses and radiology explanations demonstrates that patients often benefit from language that is simplified, organized, and adapted to their immediate questions, yet many health systems still distribute instructions as separate documents rather than as an integrated care pathway [6-8]. A patient may receive one message about appointment time, another about fasting, another about insurance approval, and another about what to bring, creating a fragmented experience that is difficult to manage. An LLM-based navigation assistant could conceptually consolidate these sources into one patient-friendly guide while preserving the institutional facts from which each instruction was derived [9, 10].

Recent advances in LLMs and generative AI show increasing capacity to transform clinical and administrative text into human-centered narratives, including plain-language discharge summaries, patient education materials, and responses to patient questions [11-13]. These systems are not substitutes for clinicians or administrative staff, but they can serve as communication assistants that restructure source information into more usable language. Studies on plain-language medical information and readable patient education materials indicate that generative models may improve accessibility when carefully constrained and reviewed [14-16]. For care navigation, this capability could be directed toward producing stepwise instructions that explain what the appointment is for, where to go, how to prepare, what documents to bring, and whom to contact if questions remain.

The thesis of this article is that a dedicated LLM system could generate safe, accurate, and patient-centered navigation instructions by combining retrieval-augmented generation, structured fact extraction, health-literacy controls, and workflow integration. Unlike a general chatbot, the proposed system would be anchored to institutional source data such as referral content, clinic protocols, insurance rules, and scheduling records, thereby limiting unsupported generation. It would be expected to produce a practical guide rather than clinical advice, with uncertain or high-risk instructions routed to staff review before patient delivery. This approach aligns with emerging evidence that LLMs can support patient communication, but it also recognizes concerns about factuality, bias, readability, privacy, and over-trust in AI-generated health information.

Background

Patient navigation and communication in healthcare

Patient navigation depends on clear communication across clinical, administrative, and logistical steps, yet many patients receive instructions that are difficult to interpret or incomplete for their specific appointment context. Conversational agents and AI-assisted communication tools have been explored as ways to improve healthcare access, answer common questions, and support patient understanding, although their safety depends on accurate scope definition and escalation pathways [17, 18]. In referral care, clear instructions are not merely educational; they affect whether the patient arrives at the correct location, completes required preparation, brings required documents, and understands next steps. A patient navigation LLM would therefore need to focus on practical, task-oriented communication rather than broad clinical diagnosis or treatment recommendation [19].

Large language models for patient-facing content

LLMs have shown promise in converting technical clinical language into more patient-centered text, including discharge summaries, pediatric admission summaries, radiology explanations, and responses to patient questions [1, 4, 7, 20]. At

the same time, these models can introduce unsupported statements, omit important caveats, or produce fluent text that appears authoritative despite being incomplete. Studies comparing AI-generated responses with clinician-authored communication show that tone and readability may improve, but factual accuracy and clinical appropriateness require careful evaluation [6, 21]. For navigation instructions, this means the model should not invent policies, preparation rules, costs, or appointment details; it should generate only from verified source information.

Transforming structured clinical data into natural language

Natural language generation from structured and semi-structured clinical data provides a foundation for converting orders, summaries, and institutional facts into readable patient communication. Generative models have been used to summarize electronic health record content, extract key clinical information, and transform technical documents into plain-language explanations, suggesting a pathway for referral-to-instruction generation [22, 23]. However, care navigation differs from traditional clinical summarization because it must combine clinical indications with operational facts such as time, location, required documents, and preparation steps. The gap is therefore not simply producing understandable medical language, but producing an integrated, accurate, action-oriented guide from heterogeneous source systems [2, 5].

Health literacy and readability of patient instructions

Health literacy is central to patient navigation because patients must translate written instructions into concrete actions before, during, and after an appointment. Studies of LLM-generated patient education materials have examined readability, simplification, and accessibility, showing that AI-generated text can potentially reduce jargon but may still require human review for completeness and cultural appropriateness [10, 14-16]. Patient-facing instructions should use short sentences, familiar words, direct sequencing, and explicit action verbs, especially when communicating preparation rules or administrative requirements. A navigation LLM should therefore be designed not only to summarize information but also to adapt it for health literacy, language preference, disability access, and patient context [24, 25].

Multi-source grounding and factual accuracy in instruction generation

Multi-source grounding is essential because navigation instructions depend on facts that may vary by clinic, procedure, insurer, appointment type, and patient status. Retrieval-augmented generation offers a conceptual mechanism for limiting hallucination by requiring the LLM to retrieve policy snippets, preparation protocols, scheduling facts, and insurance summaries before composing the final instruction [22, 26]. For patient communication, grounding should be paired with traceability so that staff can verify where each instruction came from and identify claims that are unsupported or outdated. Current reviews of LLMs in patient care emphasize that safe deployment requires evaluation of accuracy, completeness, bias, privacy, and governance rather than reliance on fluency alone [27, 28].

Instruction Generation Pipeline Overview

High-level workflow

When a referral is placed or an appointment is scheduled, the proposed pipeline would trigger a sequence of data extraction, fact alignment, LLM prompt assembly, instruction generation, optional human review, and delivery to the patient. This workflow builds on prior evidence that LLMs can transform discharge and clinical summaries into patient-friendly language, while also recognizing the need for structured safeguards before release [1, 11, 13]. The system would first identify the patient-facing facts needed for navigation, then organize them into a coherent appointment guide with source traceability. High-risk content, such as medication preparation or ambiguous insurance requirements, would be held for staff verification rather than automatically sent.

Core input domains

The core input domains would include referral details, clinic logistics, insurance requirements, preparation instructions, and appointment-specific scheduling information. Referral details would clarify the specialty, reason for referral, urgency, and any clinical note that affects preparation, while clinic logistics would specify location, check-in process, identification requirements, parking, and telehealth access. Insurance information would describe authorization status and patient financial obligations in plain language, and preparation data would translate protocol requirements into stepwise actions. This multi-domain design reflects the broader movement from isolated patient education toward integrated, patient-facing communication that connects clinical content with practical navigation needs [8, 17, 29].

Design principles

The system should be fact-grounded, health-literate, multilingual, accessible, and aligned with the patient’s preferred communication channel. Every factual claim would be expected to map to a verified source, while readability controls would encourage short sentences, familiar words, and clear sequencing for patients with varying literacy levels [10, 24, 25]. Accessibility would require outputs that function in plain text, screen-reader-compatible HTML, large-font formats, and translated versions when appropriate. These design principles position the LLM as an assistant for communication and care coordination rather than an autonomous clinical decision-maker [19, 27].

Figure 1 presents the proposed grounded LLM workflow for transforming fragmented referral, clinic, insurance, preparation, and scheduling data into safe, patient-friendly care navigation instructions.




Figure 1. Grounded LLM Workflow for Generating Patient-Friendly Care Navigation Instructions from Referral, Clinic, Insurance, Preparation, and Scheduling Data

Data Sources and Information Extraction

Parsing referral orders and clinical indications

Referral orders contain essential patient navigation facts, including the referred specialty, clinical reason, urgency, ordering clinician, and sometimes preparation-relevant notes. An extraction layer could identify these facts and distinguish patient-facing information from internal clinical shorthand, similar to how LLMs and summarization systems have been explored for converting clinical documentation into more understandable forms [4, 22]. For example, a referral reason written in technical language could be translated into a brief explanation of why the visit was scheduled without adding diagnostic certainty beyond the source data. The system should also recognize when clinical indications are sensitive, ambiguous, or unsuitable for direct patient wording without clinician review [21, 28].

Extracting scheduling constraints and location details

Scheduling systems provide practical details that patients need but may not receive in a single coherent message, including appointment date, time, expected duration, building, floor, parking instructions, check-in requirements, and telehealth links. AI-supported patient communication studies suggest that patients benefit when information is presented in direct, organized language rather than scattered across multiple documents or portal notifications [2, 7, 8]. A navigation LLM could retrieve these scheduling facts and place them near the beginning of the instruction sheet so that patients immediately understand where and when to act. The model should not infer missing location or timing details; absent information should be flagged for staff completion or displayed as a prompt to contact the clinic.

Interpreting insurance rules and financial obligations

Insurance rules are often difficult for patients to interpret because authorization status, copay language, coverage restrictions, and required documents may be communicated through administrative terminology. A care navigation LLM could conceptually translate verified insurance information into plain language, such as whether authorization is pending, whether the patient should bring an insurance card, or whether the clinic may contact them if additional information is needed. Because financial obligations and payer rules are highly variable, the system should avoid definitive guarantees unless they are explicitly supported by source data. Evidence on patient-facing LLM communication and ethics emphasizes that AI-generated guidance must manage uncertainty transparently and prevent users from over-trusting unsupported administrative or clinical claims [19, 25, 27].

Retrieving preparation and pre-visit instructions

Preparation instructions may include fasting, medication adjustments, arrival timing, documents to bring, imaging preparation, laboratory requirements, or procedure-specific restrictions. LLM studies in patient education and surgical preparation suggest that generative systems can produce readable preparation guidance, but accuracy and completeness are critical because omissions may affect appointment completion or safety [12, 29]. The proposed system would retrieve the correct protocol based on the order, appointment type, patient factors, and clinic rules before generating the instruction. If the preparation protocol includes medication changes, complex dietary restrictions, or conflicting instructions, the output should be routed for human review before being released to the patient [11, 28].

Table 1 summarizes the input-to-output generative workflow through which institutional source data are converted into a consolidated patient-facing navigation guide.

Table 1. Input-to-Output Generative Workflow for an LLM-Based Patient Care Navigation Instruction System

Workflow component	Operational input	Extraction or processing task	LLM generation function	Patient-facing output	Practical implementation value
Referral order interpretation	Specialty referral order, reason for visit, urgency, referring clinician notes	Identify referred specialty, visit purpose, urgency, and clinically relevant preparation cues	Convert referral rationale into a brief patient-friendly explanation without adding unsupported diagnosis or advice	“You have an appointment with the cardiology clinic because your doctor requested a heart evaluation.”	Helps patients understand why the appointment was scheduled and reduces confusion about referral purpose
Clinic requirement translation	Clinic check-in rules, identification requirements,	Extract required documents, check-in timing, reception	Translate operational requirements into	“Please arrive early, bring your insurance card	Makes clinic logistics actionable and reduces

	arrival policies, building instructions	location, and clinic-specific rules	simple, sequenced instructions	and photo ID, and check in at the third-floor reception desk.”	avoidable front-desk clarification
Insurance and authorization explanation	Eligibility response, prior authorization status, copay notes, payer-specific documentation rules	Identify authorization status, unresolved insurance steps, and patient-facing financial obligations	Explain insurance requirements cautiously using verified source data and uncertainty language where needed	“Your authorization is listed as pending. The clinic may contact you if more information is needed.”	Reduces administrative ambiguity while avoiding unsupported cost or coverage promises
Preparation instruction synthesis	Procedure protocols, fasting requirements, medication-related instructions, pre-visit testing requirements	Match the correct protocol to the appointment type and extract preparation steps	Produce a step-by-step preparation section in plain language	“Do not eat or drink before the visit if your clinic instruction sheet says fasting is required.”	Supports appointment readiness and reduces incomplete preparation risk
Scheduling and location integration	Appointment date, time, duration, clinic address, floor, parking note, telehealth link	Normalize appointment-specific scheduling and access details	Place date, time, location, and access instructions prominently in the guide	“Your appointment is on Monday at 9:00 AM in Building B, Floor 2.”	Consolidates logistical details that are often split across multiple reminders
Patient preference adaptation	Preferred language, communication channel, accessibility needs, caregiver involvement	Retrieve documented communication preferences and accessibility requirements	Adapt tone, language, format, and delivery channel without altering factual content	Portal message, SMS reminder, large-font version, caregiver-facing copy, or screen-reader-compatible HTML	Improves access for patients with language, literacy, disability, or caregiver-support needs
Source-grounded prompt assembly	Extracted facts, retrieved policy snippets, protocol excerpts, scheduling fields	Create a structured generation packet with allowed facts and forbidden assumptions	Instruct the LLM to generate only from verified source content	Draft instruction sheet with internal source links for audit	Reduces hallucination risk and supports staff verification
Structured patient instruction generation	Grounded prompt packet and patient communication rules	Generate a cohesive navigation guide organized by appointment purpose, preparation,	Produce patient-friendly instructions in a clear sequence	Single consolidated care navigation guide	Replaces fragmented messages with one practical, understandable instruction artifact

		logistics, documents, and contact steps			
Delivery and documentation	Patient portal, appointment reminder system, SMS/email platform, printed discharge or referral materials	Store final message, version output, and record delivery channel	Format final approved content for the selected channel	Portal guide, reminder attachment, printed instruction sheet, or caregiver copy	Integrates generated instructions into existing healthcare communication workflows

LLM Architecture and Instruction Customization

Model choice and fine-tuning

The proposed system would use an instruction-tuned LLM adapted to health communication tasks through curated examples of high-quality navigation instructions, health-literacy guidance, and controlled output formats. Prior research on fine-tuned or prompted language models for discharge instructions and patient-facing summaries supports the idea that model behavior can be shaped toward clearer and more structured communication [5, 13]. However, the system should not be described as independently learning clinical policy or administrative rules; it would generate from retrieved institutional facts and predefined templates. Fine-tuning would therefore serve the communication style, not replace source verification or workflow governance.

Retrieval-augmented generation and fact-grounding

Retrieval-augmented generation would be the central safety mechanism because it would require the model to retrieve relevant clinic policies, preparation protocols, referral facts, scheduling details, and insurance summaries before composing the patient-facing output. This approach is consistent with emerging healthcare RAG frameworks that seek to ground generative AI in source documents rather than relying only on model memory [22, 26]. The prompt assembly layer could include structured facts, retrieved text snippets, required wording constraints, and forbidden behaviors such as inventing dates, costs, preparation steps, or clinical advice. The final instruction would be expected to include source-linked assertions for staff audit, even if the patient-facing version uses simple language.

Health literacy and readability control

Health literacy control would guide the model to use plain language, short sentences, direct action verbs, and a clear sequence of steps that patients can follow before the appointment. Research on LLM-generated patient education materials, plain-language evidence summaries, and simplified medical information indicates that readability can improve when models are explicitly prompted and evaluated for patient comprehension [10, 14, 16, 23]. The system could be instructed to avoid unexplained medical jargon, define necessary terms, and separate urgent actions from background information. Although readability formulas can support evaluation, human review remains important because a message can be easy to read while still incomplete, culturally mismatched, or operationally unclear [24, 25].

Multilingual and accessible output generation

Multilingual and accessible output generation would allow the same verified facts to be expressed in the patient's preferred language and format, while preserving meaning across versions. Patient-facing AI systems and conversational agents have been discussed as tools for improving access, but they must account for language preference, disability, digital literacy, and the risk of unequal communication quality across patient groups [17-19]. The model could generate plain-text instructions for SMS, structured HTML for portals, and screen-reader-compatible content for patients using

assistive technologies. Translation and accessibility outputs should still be checked against the same source-grounded facts so that the system does not introduce discrepancies between language versions.

Ensuring Accuracy, Safety, and Health Literacy

Source citation and traceability

Source citation and traceability would be essential because patients may act directly on generated instructions about fasting, arrival time, insurance steps, or required documents. The system should link each factual statement to a verified source, such as a referral order, scheduling record, clinic policy, or preparation protocol, so staff can audit the origin of each instruction [22, 26]. This design would extend prior work on patient-friendly discharge summaries and clinical information extraction by making the generated text not only readable but also accountable to institutional records [1, 13]. Claims that cannot be linked to source data should be highlighted for review rather than included as confident patient-facing guidance.

Hallucination prevention and safety guardrails

Hallucination prevention would require more than careful prompting because fluent LLM output can still contain unsupported or misleading information. A post-generation validator should compare the instruction sheet against extracted source facts, checking whether dates, times, locations, preparation rules, insurance statements, and contact instructions match the retrieved evidence [27, 28]. Contradictions, missing source support, or critical uncertainty should prevent automatic release and route the message to staff for correction. Such safeguards are especially important because studies of patient-facing AI communication show that patients may perceive AI-generated responses as helpful even when factual completeness and appropriateness still require expert oversight [6, 21].

Human-in-the-loop review for high-risk instructions

Human-in-the-loop review should be reserved for instructions that carry higher clinical, logistical, or financial risk, such as medication adjustments, fasting requirements, anesthesia preparation, prior authorization uncertainty, or complex disability accommodations. Evidence from LLM-generated discharge and preparation materials suggests that generative systems can improve readability, but high-risk patient instructions still require professional verification before release [11, 12, 29]. The review interface should allow staff to see the generated text, source evidence, flagged uncertainties, and suggested corrections in one view. This would position the model as a drafting and consolidation assistant rather than as an autonomous authority over patient preparation or care navigation.

Table 2 outlines the safety, quality-control, human-review, and governance safeguards required for responsible deployment of patient-facing LLM navigation instructions.

Table 2. Safety, Quality Control, Human Review, and Governance Framework for Patient-Facing LLM Navigation Instructions

Governance domain	Safety or quality risk	Control mechanism	Human review trigger	Responsible implementation role	Practical output

Source grounding	The model invents appointment details, preparation steps, costs, or clinic rules	Retrieval-augmented generation with source-linked facts and restricted prompt context	Any generated claim without a linked institutional source	Clinical informatics team, referral operations lead	Source-traceable instruction sheet
Factual consistency	Generated text conflicts with scheduling, insurance, protocol, or referral data	Post-generation validator compares output against extracted facts	Contradiction in date, time, location, preparation rule, insurance status, or contact instruction	Scheduling supervisor, patient navigation team	Corrected or withheld patient message
Preparation safety	Incorrect fasting, medication, laboratory, imaging, or procedure-preparation instructions	High-risk preparation rules are routed through predefined staff review criteria	Medication changes, fasting rules, anesthesia preparation, complex diet restrictions, or unclear protocol match	Nurse, clinic staff, procedure coordinator	Staff-approved preparation section
Health literacy	Output is readable but still confusing, too technical, or poorly sequenced	Readability control, jargon detection, short-sentence prompting, action-step formatting	Complex terminology, long paragraphs, unexplained abbreviations, or unclear patient actions	Patient education specialist, health literacy reviewer	Plain-language navigation guide
Multilingual accuracy	Translation changes the meaning of preparation, location, timing, or insurance instructions	Same source-grounded fact packet used for each language; translation review for high-risk instructions	Non-English output containing preparation, medication, or financial information	Interpreter services, bilingual staff, patient communication team	Language-concordant patient instructions
Accessibility	Instructions are inaccessible to patients using assistive technology or requiring alternative formats	Plain-text structure, screen-reader-compatible HTML, large-font option, caregiver-facing version	Document contains complex formatting, image-only text, or inaccessible portal layout	Patient access team, digital accessibility lead	Accessible communication artifact
Privacy and sensitive content	Generated message includes unnecessary sensitive clinical	Minimum necessary content rule and privacy screen before delivery	Sensitive diagnosis, stigmatizing language, unnecessary financial details, or caregiver disclosure concern	Privacy officer, compliance team	Privacy-preserving patient communication

	or insurance details				
Over-trust and scope creep	Patient treats AI-generated navigation instructions as clinical advice	Explicit assistant role, escalation language, and boundaries between logistics and medical advice	Patient asks clinical, urgent, or symptom-related question	Nurse triage team, clinician, patient navigator	Safe escalation to human care team
Version control	Patient receives outdated instructions after rescheduling or protocol change	Versioned outputs linked to appointment status and current source documents	Appointment rescheduled, location changed, preparation protocol updated, or authorization status changed	Scheduling team, portal operations team	Updated patient-facing instruction message
Equity and monitoring	Instructions work better for some patient groups than others	Review by language, literacy, disability, age, digital access, and communication channel	Recurrent unclear feedback from specific populations or access groups	Health equity lead, patient experience team	Equity-informed refinement plan
Liability and accountability	Unclear responsibility for AI-generated patient instructions	Governance policy defining approval, audit, escalation, and release responsibilities	High-risk instructions, patient complaint, adverse communication event, or staff override	Health system leadership, legal, clinical governance board	Auditable deployment framework
Continuous improvement	Repeated unclear instructions persist across clinics or referral pathways	Patient feedback, staff correction logs, and recurring error review	Frequent clarification calls, repeated patient misunderstanding, or staff edits to same content	Quality improvement team, informatics team	Prioritized improvement backlog

Personalization and Patient Interaction

Learning patient preferences and history

Personalization would allow the system to adapt navigation instructions to stable patient preferences, including language, preferred communication channel, accessibility needs, transportation concerns, and caregiver involvement. Patient-facing AI and chatbot research suggests that conversational systems may improve engagement when they respond to user needs, but they must avoid over-personalization that assumes unsupported facts or reinforces inequities [17-19]. A navigation LLM could use documented preferences to determine whether instructions should be sent through the portal, SMS, email, printed letter, or caregiver-facing format. Personalization should remain transparent, consent-aware, and limited to communication support rather than inference about sensitive patient attributes.

Two-way conversational refinement

Two-way conversational refinement would allow patients to ask practical follow-up questions, such as whether they can bring a caregiver, where to park, what to do if they are late, or how to access a telehealth link. Prior work comparing AI and clinician responses to patient questions suggests that LLMs can produce understandable and empathetic language, but escalation rules are needed when questions become clinical, urgent, or outside the retrieved source content [6, 7, 21]. The system should answer only when the relevant policy or scheduling fact is available and should otherwise direct the patient to clinic staff. This design would support patient navigation while preserving boundaries between administrative guidance and clinical decision-making.

Integration with reminder systems and digital front doors

Generated navigation instructions could be attached to appointment reminders, embedded in patient portals, delivered through digital front doors, or sent as SMS and email when appropriate. Conversational agents and AI-supported healthcare communication tools have been studied as mechanisms for improving access and patient engagement, but integration into real workflows is necessary for practical value [18, 19]. A reminder could include the appointment time, location, preparation steps, required documents, and a link to the full instruction guide. The system should ensure that each communication channel preserves privacy, accessibility, and consistency with the source-grounded instruction sheet [17, 25].

Integration Into Care delivery and Patient Portals

Embedding in referral and scheduling workflow

Embedding the system in referral and scheduling workflows would allow instruction generation to occur when the appointment is booked, rescheduled, or updated. The generated guide could be stored in the patient portal, visible to care teams, and versioned so staff can review what was sent and when. Prior studies of LLM-assisted clinical summaries and patient-facing messages show the importance of connecting generated text to actual care processes rather than treating it as a standalone chatbot response [1-3]. If appointment details or preparation rules change, the system should generate an updated message rather than relying on the patient to reconcile conflicting communications.

Patient feedback and continuous improvement

Patient feedback could help identify unclear instructions, missing information, confusing terminology, or recurring navigation problems across clinics. Studies of LLM-generated patient education and plain-language medical information emphasize that readability alone is insufficient; patient understanding, acceptability, and perceived usefulness should also be evaluated [23-25]. Feedback should be used to improve templates, retrieval coverage, staff review rules, and patient communication design, not to allow uncontrolled model self-modification. This approach would support continuous improvement while maintaining governance over the institutional knowledge base and generation workflow.

Evaluation Strategy

Instruction quality metrics: accuracy, completeness, and readability

Evaluation should begin with expert review of generated instructions for accuracy, completeness, readability, and alignment with source documents. Prior studies have assessed AI-generated patient-facing summaries, radiology explanations, discharge materials, and education resources using criteria such as factual correctness, clarity, readability, and appropriateness [4, 8, 10, 16]. For this navigation system, reviewers should examine whether the message includes all required appointment facts, avoids unsupported statements, and expresses preparation steps in plain language. Evaluation should remain conceptual and safety-oriented, focusing on whether the system could be responsibly assessed before operational deployment.

Patient-reported understanding and adherence

Patient-reported evaluation should examine whether patients understand where to go, when to arrive, how to prepare, what to bring, and whom to contact with questions. Studies of patient-facing AI responses and simplified medical communication suggest that perceived clarity and usefulness are important, but they should be interpreted alongside expert review because patient preference does not guarantee factual accuracy [2, 6, 9]. A proposed evaluation could compare standard instructions with LLM-generated navigation guides using surveys or interviews about understanding, confidence, and perceived burden. Such evaluation should avoid claiming benefit before empirical testing and should be designed to include patients with diverse literacy, language, disability, and digital access needs [17, 24].

Operational impact on no-shows and completion rates

Operational evaluation could examine whether the system would be associated with fewer missed appointments, fewer incomplete preparations, improved referral completion, or reduced inbound calls for clarification. However, because this article is conceptual, these outcomes should be framed as future implementation targets rather than reported results. Prior work on patient communication, conversational agents, and LLM-based patient-facing content supports the plausibility of studying operational effects, but real-world impact would depend on workflow design, patient population, and institutional context [18, 19, 27]. Any future pre-post or pragmatic evaluation should account for confounding factors such as staffing changes, reminder policies, portal adoption, transportation barriers, and insurance variation.

Limitations

Data quality and institutional variation

The proposed system would depend on the accuracy, completeness, and timeliness of institutional source data. Clinic policies, preparation protocols, insurance rules, location instructions, and scheduling records may be inconsistent, outdated, or unavailable in machine-readable form, which would reduce the reliability of generated instructions [22, 26]. Institutional variation also means that a model configured for one health system may not transfer safely to another without local validation and governance. These limitations are especially important because LLMs may generate coherent language even when the underlying source data are incomplete or contradictory [27, 28].

Trust and liability concerns

Trust and liability concerns would remain central because patients may rely on generated instructions when deciding how to prepare, whether to attend, or whom to contact. Ethical discussions of generative AI for patient-facing information emphasize the risks of over-trust, privacy exposure, inequitable communication quality, and unclear accountability when AI-generated content affects health-related decisions [19, 25]. The system should therefore disclose its assistant role, maintain human escalation pathways, and preserve a review process for high-risk or uncertain instructions. Legal and governance frameworks for AI-generated patient communication are still developing, so deployment should be cautious, auditable, and institutionally accountable [27, 28].

Conclusion

A large language model for patient-friendly care navigation instructions could transform fragmented referral, clinic, insurance, preparation, and scheduling information into one coherent guide. The system would not replace clinicians, schedulers, or patient navigators, but it could assist them by drafting clear, practical, source-grounded instructions for patients. Its primary role would be communication support: helping patients understand what appointment they have, why it matters, where to go, how to prepare, and what to do next.

The major strength of this approach is consolidation. Instead of expecting patients to reconcile separate referral notes, portal messages, preparation documents, and insurance notices, the system could produce one personalized navigation message tailored to literacy level, language preference, accessibility needs, and communication channel. Traceable fact-grounding would allow staff to verify the source of each instruction, while health-literacy controls would help ensure that the final message remains understandable and actionable. In this way, the model could serve as an assistant-oriented bridge between complex healthcare operations and patient comprehension.

Several challenges would still need to be addressed before such a system could be responsibly deployed. Institutional data may be incomplete or outdated, insurance rules may change quickly, and preparation instructions may vary by procedure, clinic, and patient condition. Real-world validation would be necessary to determine whether generated instructions are accurate, acceptable, accessible, and operationally useful across diverse patient populations. Human oversight would remain essential, especially for medication-related preparation, ambiguous administrative requirements, and any instruction that could affect patient safety.

Future work should focus on pilot implementations in large health systems with diverse referral pathways and patient populations. These pilots should evaluate whether source-grounded LLM navigation instructions improve patient understanding, appointment readiness, and experience without increasing risk or inequity. Responsible implementation would require collaboration among clinicians, schedulers, patient navigators, informaticians, privacy officers, and patient representatives. With careful governance, an LLM-based navigation assistant could become a practical tool for making healthcare instructions clearer, safer, and more patient-centered.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 09 Jan 2025

Revised: 08 Mar 2025

Accepted: 17 Apr 2025

Published online: 20 July 2025

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Zaretsky J, Kim JM, Baskharoun S, Zhao Y, Austrian J, Aphinyanaphongs Y, et al. Generative artificial intelligence to transform inpatient discharge summaries to patient-friendly language and format. *JAMA Netw Open*. 2024;7(3):e240357.
- Huang T, Safranek C, Socrates V, Chartash D, Wright D, Dilip M, et al. Patient-representing population's perceptions of GPT-generated versus standard emergency department discharge instructions: randomized blind survey assessment. *J Med Internet Res*. 2024;26:e60336.
- Liu S, Wright AP, McCoy AB, Huang SS, Jenkins JZ, Peterson JF, et al. Using large language model to guide patients to create efficient and comprehensive clinical care message. *J Am Med Inform Assoc*. 2024;31(8):1665-70.
- Ebby CG, Tse G, Bethel J, Zhao Q, Gerber DM, Kelly MM. Large language models to summarize pediatric admission notes into plain language. *Pediatrics*. 2025;155(6):e2024069515.
- Alqahtani M, Al-Barakati A, Alotaibi F, Al Shibli M, Almousa S. Impact of detailed versus generic instructions on fine-tuned language models for patient discharge instructions generation: comparative statistical analysis. *JMIR Form Res*. 2025;9:e80917.
- Ayers JW, Poliak A, Dredze M, Leas EC, Zhu Z, Kelley JB, et al. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Intern Med*. 2023;183(6):589-96.
- Bernstein IA, Zhang Y, Govil D, Majid I, Chang RT, Sun Y, et al. Comparison of ophthalmologist and large language model chatbot responses to online patient eye care questions. *JAMA Netw Open*. 2023;6(8):e2330320.
- Gordon EB, Towbin AJ, Wingrove P, Shafique U, Haas B, Kitts AB, et al. Enhancing patient communication with ChatGPT in radiology: evaluating the efficacy and readability of answers to common imaging-related questions. *J Am Coll Radiol*. 2024;21(2):353-9.
- Shah YB, Ghosh A, Hochberg AR, Rapoport E, Lallas CD, Shah MS, et al. Comparison of ChatGPT and traditional patient education materials for men's health. *Urol Pract*. 2024;11(1):87-94.
- Moons P, Van Bulck L. Using ChatGPT and Google Bard to improve the readability of written patient information: a proof of concept. *Eur J Cardiovasc Nurs*. 2024;23(2):122-6.
- Rosenberg GS, Magnéli M, Barle N, Kontakis MG, Müller AM, Wittauer M, et al. ChatGPT-4 generates orthopedic discharge documents faster than humans maintaining comparable quality: a pilot study of 6 cases. *Acta Orthop*. 2024;95:152.
- Rust P, Frings J, Meister S, Fehring L. Evaluation of a large language model to simplify discharge summaries and provide cardiological lifestyle recommendations. *Commun Med*. 2025;5(1):208.
- Eisinger F, Holderried F, Mahling M, Stegemann-Philipps C, Herrmann-Werner A, Nazarenius E, et al. What's going on with me and how can I better manage my health? The potential of GPT-4 to transform discharge letters into patient-centered letters to enhance patient safety: prospective, exploratory study. *J Med Internet Res*. 2025;27:e67143.
- Abreu AA, Murimwa GZ, Farah E, Stewart JW, Zhang L, Rodriguez J, et al. Enhancing readability of online patient-facing content: the role of AI chatbots in improving cancer information accessibility. *J Natl Compr Canc Netw*. 2024;22(2D).
- Eid K, Eid A, Wang D, Raiker RS, Chen S, Nguyen J. Optimizing ophthalmology patient education via chatbot-generated materials: readability analysis of AI-generated patient education materials and the American Society of Ophthalmic Plastic and Reconstructive Surgery patient brochures. *Ophthalmic Plast Reconstr Surg*. 2024;40(2):212-6.
- Breneman A, Trager MH, Gordon ER, Samie FH. Readability rescue: large language models may improve readability of patient education materials. *Arch Dermatol Res*. 2024;316(9):669.

Chen F, Lan T, Liang J, Zhang R. Toward bridging gaps in patient navigation: a study on the adoption of artificial intelligence technologies. *Med Adv.* 2024;2(3):274-83.

Laranjo L, Dunn AG, Tong HL, Kocaballi AB, Chen J, Bashir R, et al. Conversational agents in healthcare: a systematic review. *J Am Med Inform Assoc.* 2018;25(9):1248-58.

Nadarzynski T, Miles O, Cowie A, Ridge D. Acceptability of artificial intelligence (AI)-led chatbot services in healthcare: a mixed-methods study. *Digit Health.* 2019;5:2055207619871808.

Lyu Q, Tan J, Zapadka ME, Ponnatapura J, Niu C, Myers KJ, et al. Translating radiology reports into plain language using ChatGPT and GPT-4 with prompt learning: promising results, limitations, and potential. *arXiv.* 2023;arXiv:2303.09038.

Liu S, McCoy AB, Wright AP, Carew B, Jenkins JZ, Huang SS, et al. Leveraging large language models for generating responses to patient messages—a subjective analysis. *J Am Med Inform Assoc.* 2024;31(6):1367-79.

Alkhalaf M, Yu P, Yin M, Deng C. Applying generative AI with retrieval augmented generation to summarize and extract key clinical information from electronic health records. *J Biomed Inform.* 2024;156:104662.

Ovelman C, Kugley S, Gartlehner G, Viswanathan M. The use of a large language model to create plain language summaries of evidence reviews in healthcare: a feasibility study. *Cochrane Evid Synth Methods.* 2024;2(2):e12041.

Aydin S, Karabacak M, Vlachos V, Margetis K. Large language models in patient education: a scoping review of applications in medicine. *Front Med.* 2024;11:1477898.

Pal A, Wangmo T, Bharadia T, Ahmed-Richards M, Bhandari MB, Kachhadiya R, et al. Generative AI/LLMs for plain language medical information for patients, caregivers and general public: opportunities, risks and ethics. *Patient Prefer Adherence.* 2025;19:2227-49.

Yang R, Ning Y, Keppo E, Liu M, Hong C, Bitterman DS, et al. Retrieval-augmented generation for generative artificial intelligence in health care. *NPJ Health Syst.* 2025;2(1):2.

Busch F, Hoffmann L, Rueger C, Van Dijk EH, Kader R, Ortiz-Prado E, et al. Current applications and challenges in large language models for patient care: a systematic review. *Commun Med.* 2025;5(1):26.

Williams CY, Miao BY, Kornblith AE, Butte AJ. Evaluating the use of large language models to provide clinical recommendations in the emergency department. *Nat Commun.* 2024;15(1):8236.

Palmarin E, Lando S, Marchet A, Saibene T, Michieletto S, Cagol M, et al. Assessing the accuracy and readability of large language model guidance for patients on breast cancer surgery preparation and recovery. *J Clin Med.* 2025;14(15):5411.