Leveraging Artificial Intelligence to Bridge Gaps in Pediatric Oncology Care for Marginalized Spanish-Speaking Communities

Grigorii Khvatskii[1] , **Angelica Garcia Martinez**[1] , **Deng Pan**[1] , **Matthew Belcher**[1] , **Gerónimo Medrano Loera**[2] , **Dayana Pineda Pérez**[3] , **Juan Emmanuel Ferrari Muñoz-Ledo**[4] , **Horacio Márquez-González**[5] , **Nuno Moniz**[1] and **Nitesh V. Chawla**[1]
[1]University of Notre Dame
[2]Universidad Autónoma Metropolitana-Iztapalapa
[3]Secretaría de Educación Pública
[4]Universidad Autónoma Metropolitana-Azcapotzalco
[5]Hospital Infantil de México, Federico Gómez
{gkhvatsk, agarci35, dpan, mbelche2}@nd.edu, sociologogmedrano@gmail.com, dapipe19@gmail.com, j_ferrari_@hotmail.com, horaciohimfg@gmail.com, {nmoniz2, nchawla}@nd.edu,

Abstract

In low-and middle-income countries (LMICs) pediatric cancer patients and their caregivers often suffer from effects of underfunded, fragmented and outdated healthcare systems. One of these effects is a breakdown of communication between hospital staff and caregivers, which is felt stronger among vulnerable populations. Our proposed solution integrates Large Language Models (LLM) and Automatic Speech Recognition (ASR) technologies to enhance communication between caregivers and healthcare providers while integrating community feedback. We combine cutting-edge technology with existing hospital infrastructure to allow for easy deployment and testing. The system will improve access to health, nutrition, and parental care programs, prioritizing caregiver engagement and real-time interaction. Ultimately, our system will pave the way to more equitable access to medical care, and address structural barriers affecting marginalized communities.

1 Introduction

Early childhood, up to 8 years of age [UNICEF, 2025a], is critical for cognitive, social, emotional, and physical development due to the brain's high plasticity and responsiveness during this period. Investing in early childhood development (ECD) is considered highly lucrative [Elango *et al.*, 2015] for ensuring both social and economic progress [World Bank, 2025]. However, international efforts to measure health, nutrition, education, and quality care face significant challenges, particularly in low- and middle-income countries [Nurturing Care, 2025] (LMICs) due to limited data availability. This problem is further exacerbated in less common settings, where limited evidence is the unfortunate reality. For example, in the case of children with cancer, there is little evidence on how treatments affect their ECD in poverty contexts within these regions [Nurturing Care, 2025].

1.1 Deployment Context and Motivation

Our solution will be piloted at the National Institute of Pediatrics, Hospital Infantil de México, Federico Gómez (HIMFG) in Mexico City. HIMFG is the largest pediatric oncology center in Mexico, helping more than 35,000 [HIMFG, 2022] children annually, out of whom more than 6500 are pediatric cancer patients.

Pediatric oncology patients in socially and economically marginalized Spanish-speaking communities face a severe lack of access to healthcare due to fragmented systems and the absence of social health protection programs [Chalfant *et al.*, 2023]. Caregivers of these children often struggle with navigating healthcare processes, understanding medical recommendations, and receiving timely guidance during complications and medical emergencies.

This is compounded by the fact that children with cancer in Mexico and other Latin American countries face significant barriers to timely and quality healthcare. For example, pediatric cancer patients and their caregivers in Mexico suffer from an environment of constrained public investment, absence of health policies, and reliance on external mechanisms within communities to secure dignified healthcare access [Serván-Mori *et al.*, 2023]. The cancellation of state healthcare programs has further widened the gaps in access to pediatric oncology care, leaving children and their caregivers in marginalized communities without the necessary support [Knaul *et al.*, 2023]. These challenges lead to delays in medical care, higher rates of complications, and increased relapses, significantly impacting their ECD or reducing their survival rates.

Additionally, we found that large parts of the healthcare system rely on outdated technology, such as paper files, which compounds the problem of timely triage and coordination of care, leading to a significant social and economic burden on caregivers and families.

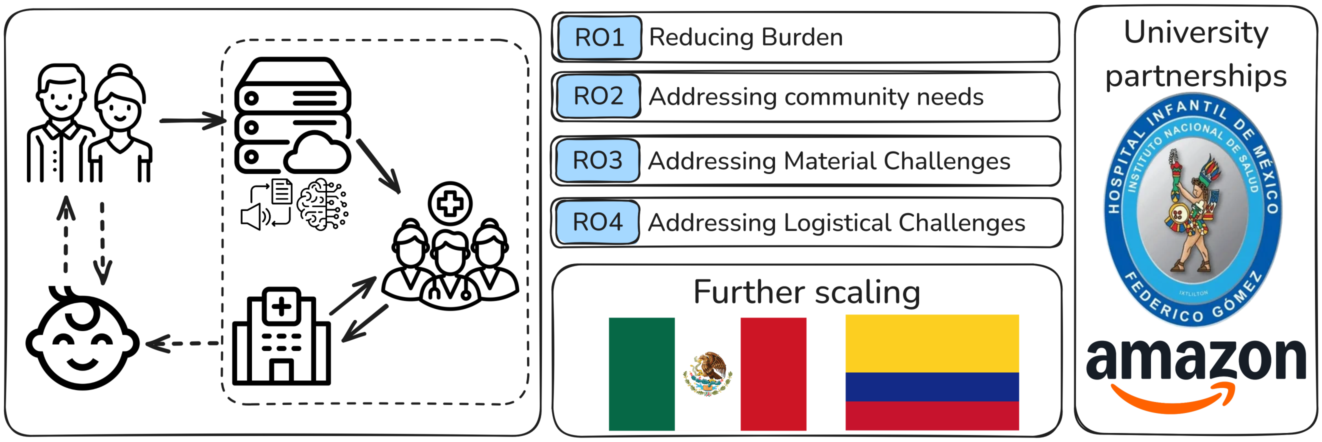

Figure 1: Project overview and vision. A pediatric care patient, discharged from the hospital, communicates with their caregivers, who then communicate the information about the patient's health and SDoH to the hospital via **`SaludASR`**. The medical staff can then use this information to take timely action.

1.2 Vision and Objectives

Automated approaches to enhance information exchange between caregivers, medical professionals. Large Language Models (LLMs) and Automated Speech Recognition (ASR) offer a promising combination of technologies to gather information from caregivers of children with cancer in LMICs. They enable real-time, context-aware communication between all relevant stakeholders, improving engagement with Health Mobile Apps [Qin *et al.*, 2024; Schnur *et al.*, 2024]. These advanced technologies demonstrate significant proficiency in extracting and analyzing medical knowledge, which can potentially improve care delivery and support services for pediatric patients and their families.

Despite the remarkable proficiency of LLM and ASR-based methods for gathering medical information [Yang *et al.*, 2024], balancing the information-gathering capability and privacy remains a significant challenge. Achieving this balance is essential to reduce disparities in accessing and utilizing healthcare services, particularly among vulnerable populations defined by their ethnic, racial, social, or economic circumstances [Dellinger, 2022]. The focus should be on minimizing potential harm or cohesion issues that may arise from using such data, while promoting an ethical and inclusive strategy tailored to the specific needs of the pediatric oncology patients and their caregivers.

1.3 Expected Results and Impact on UN SDGs

By structuring our solution around the WHO/UNICEF Nurturing Care Framework [WHO, 2018] we ensure that the tool not only provides medical guidance but also addresses broader rights and needs, including health access, quality of care, nutrition, and child protection in children under 6 years old with cancer.

Our research will lead to improvement in ECD in pediatric cancer patients, in accordance with the Nurturing Care Framework. This project seeks to benefit not only the children, but their caregivers and medical professionals, in addition to the larger healthcare system. Through the wider deployment of our research products, we seek to make a large positive impact on the lives of tens of thousands pediatric cancer patients in LMICs.

As such, our research directly contributes to UN SDG 3 (Ensure healthy lives and promote well-being for all at all ages). Furthermore, by focusing on economically disadvantaged groups, we hope that our project will make a contribution to UN SGD 2 (End hunger, achieve food security and improved nutrition and promote sustainable agriculture). In addition to that, better healthcare outcomes will directly lead to better educational outcomes for pediatric cancer patients, thus contributing to UN SDG 4 (Ensure inclusive and equitable quality education and promote lifelong learning opportunities for all).

2 Overview and Research Challenges

To tackle these challenges, we propose a novel framework called **`SaludASR`**. It consists of four major steps: 1) receives spoken input using voice recognition tools, making use of existing phone lines, sidestepping the technological challenges; 2) analyzes information for structured medical data extraction using an LLM; 3) generates questions for missing medical information using an LLM and synthesizes them into speech; 4) stores the obtained data in a privacy-respecting electronic medical record system. Figure 1 shows an overview of the proposed solution.

Our approach is data-driven, linguistically, and culturally contextualized and designed for low-resource settings, integrating social determinants of health (SDoH) to enhance caregiver support and pediatric cancer care.

2.1 Healthcare and Community engagement

The team behind this proposal has a long history of equal and fruitful partnership with HIMFG, focusing on utilizing cutting-edge technologies developed within the Mexican healthcare context to help pediatric cancer patients [Schnur *et al.*, 2024; Avilés-Robles *et al.*, 2022]. This partnership is dedicated to addressing the barriers to healthcare access and quality of care arising from the fragmentation of Mexico's healthcare system. We will achieve this objective through engagement with caregivers or families to understand their needs and barriers in using this technology, ensuring the deployed solution is both inclusive and responsible.

2.2 Focus groups for caregivers

Engaging deeply with the local community, we conducted a preliminary study involving caregivers, healthcare personnel, and professors at HIMFG, ensuring that our research was grounded in the real-world experiences and needs of families affected by childhood cancer.

Study design and participants: A mixed-methods study was conducted within a prospective cohort study from July to August 2024. The study aimed to track complications after chemotherapy in children and adolescents with cancer using the SaludConectaMx system [Schnur *et al.*, 2024], centered around a mobile application. The dynamic cohort is ongoing and has 355 enrolled participants at HIMFG

Semi-structured interviews and focus groups were conducted with key informants, including caregivers of children under six years and healthcare personnel. Thirty-one semi-structured interviews were conducted with caregivers of children undergoing scheduled chemotherapy in both the external consultation area and the ambulatory chemotherapy room.

The question matrix was developed based on the Nurturing Care Framework [Black *et al.*, 2017; WHO, 2018] to ensure comprehensive coverage of the five essential components for child development: good health, adequate nutrition, responsive care, safety and protection, and early learning opportunities. A literature review was conducted to identify key thematic areas relevant to the concerns of caregivers, educators, and healthcare professionals in pediatric oncology. Additionally, a participatory approach was adopted, incorporating feedback from these stakeholders to refine the wording of the questions for accessibility and clarity. For example, terms like “NGOs” were replaced with “foundations” to align with the language commonly used by caregivers. The final set of questions underwent three rounds of refinement based on content and thematic analysis [Braun and Clarke, 2021] to ensure that the matrix effectively prioritized the most recurrent concerns of caregivers.

The questions were organized into a three-tier hierarchical structure. At the first level, broad analytical categories such as Adequate Nutrition, Good Health, and Responsive Care were established. The second level included subcategories; for example, Adequate Nutrition was further divided into Normal Diet, Ultra-Processed Foods, and Lack of Appetite. At the third level, specific topics were assigned to each subcategory. A structured coding process was implemented, linking questions to their frequency of mention and variations in phrasing.

Data Analysis: For data analysis, a combined approach of thematic and axial analysis was applied [Patton, 2014; Creswell and Creswell, 2017]. An iterative coding process refined categories and subcategories, enabling the identification of patterns in caregivers' and healthcare professionals' concerns. A participatory ranking technique was used to enhance prioritization, employing a double-entry prioritization matrix to compare and rank topics by relative importance. The final questionnaire reflects the most frequently reported and thematically relevant concerns, ensuring that participants' knowledge and experiences are systematically integrated into the data collection procedures.

Results: We have uncovered 4 broad categories of concerns that pediatric patients' caregivers have. The first category deals with the lack of functionality in existing systems. For example, caregivers have pointed out that the current system they use lacks a nutritional guidance module, appointment and medication reminders, as well as a lack of tools to monitor health indicators over time. They have also noted having difficulty with data entry, a lack of ability to communicate with medical staff, and a general lack of emotional and mental health support resources.

The second category deals with the burden that the existing app-based system places on the caregivers. They have noted that the current system is difficult to use, lacks an engaging interface, and requires manual data entry, which can be time consuming. In general, this leads to a reduction of the app use over time, especially when the patients get better. We have also uncovered that some caregivers were unaware of the system's existence.

The last two categories deal with the lack of resources faced by vulnerable populations. Caregivers have noted a lack of access to even the simplest medical tools, such as thermometers, and a lack of ways to receive real-time guidance from medical staff in emergency situations, when they cannot afford travel to the hospital. They have also noted a general lack of access to support network and other resources.

Finally, caregivers have outlined the fact that they often face logistical challenges, exacerbated by their inability to afford long-distance travel. This often leads to emotional and physical strain, for which they find in difficult to get support.

We have then taken advantage of this insight to create an initial database of questions for caregivers, which we have utilized in early prototypes of **`SaludASR`**.

Consent letters were obtained from participants before they were interviewed. This study was performed in line with the principles of the Declaration of Helsinki.

3 Concept of Operations and Research Questions

The data collected and analyzed during our preliminary study leads us to the following project concept and main research questions, which will direct the intellectual contributions behind the development and deployment of this project. Figure 2 shows an overview of the operation of **`SaludASR`**.

To address the issue of low internet penetration and the burden placed on the caregivers by complex mobile applications, we propose utilizing either landlines or cell phones belonging to caregivers. By employing automated voice recognition tools, their spoken input can be effectively routed into an LLM-based solution, thereby sidestepping these challenges.

The combined LLM and ASR system will rely on contextual knowledge we have collected during our preliminary mixed-methods study. Our contextualized knowledge base includes over 200 questions tailored to caregivers' language and lived realities in the most vulnerable contexts, address-

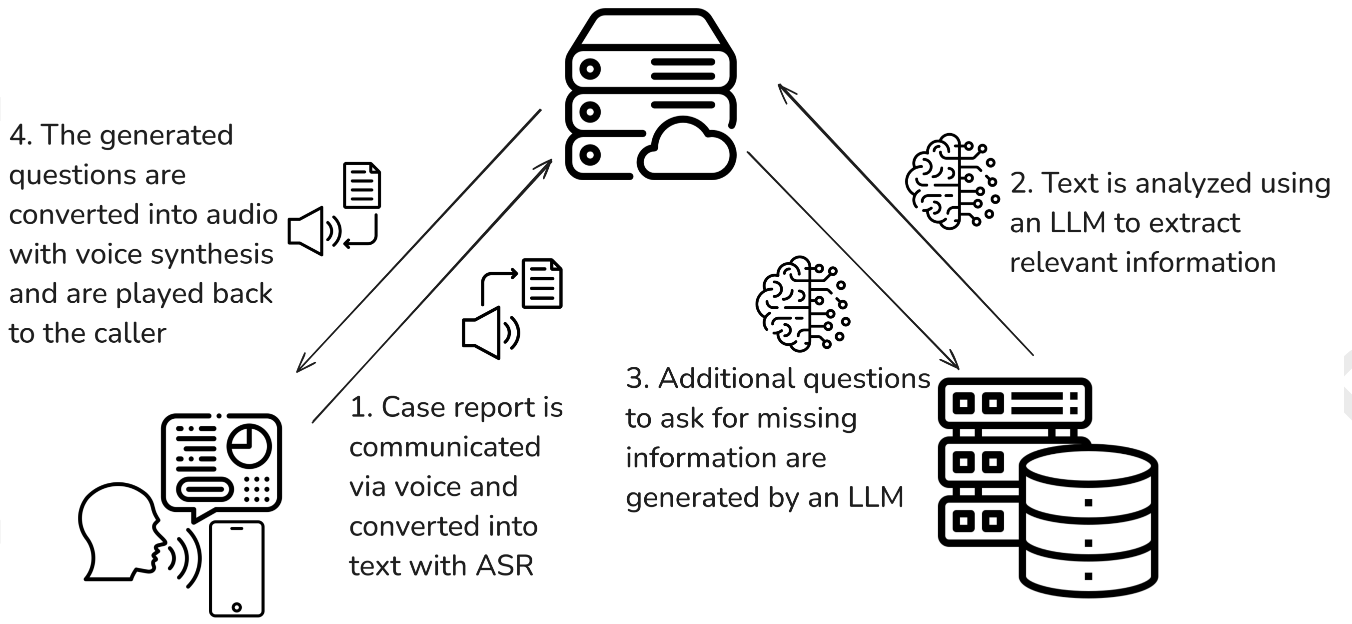

Figure 2: Graphical representation of the proposed LLM-ASR pipeline

ing critical SDoH such as poverty, education, healthcare access, and social support networks. This allows us to effectively collect data about the most pressing concerns faced by caregivers. It also enables medical staff to take timely action based on that data, such as providing medical advise and guidance. Furthermore, is allows them to connect caregivers to the relevant support organizations and groups, mitigating the effects of the financial and emotional burden felt by some caretakers. Finally, this solution allows us to largely reuse existing systems for reminding the caregivers of important dates and times.

RQ1 [Reducing Burden] How can the combination of LLMs and ASR reduce the burden the existing applications place on the caregivers?

RQ2 [Addressing Community Needs] How can the combination of LLMs and ASR be used to address a broad range of caregiver needs, such as receiving nutritional guidance and monitoring patients' symptoms?

RQ3 [Addressing Material Challenges] How can the combination of LLMs and ASR be used to address the material challenges faced by the community, such as the lack of adequate nutrition or access to medical instruments?

RQ4 [Addressing Logistical Challenges] How can the combination of LLMs and ASR be used to respond to logistical, financial and emotional needs of patients' caregivers and families?

Through adaptive AI interaction, **`SaludASR`** enables natural conversations [Ma *et al.*, 2024] in Spanish, reducing literacy barriers and empowering caregivers to express concerns while receiving AI-powered responses aligned with pediatric oncology guidelines. By facilitating interaction between caregivers, the AI assistant, and local support figures—such as health personnel —the tool strengthens social support mechanisms, mitigating the effects of healthcare fragmentation [Zhai *et al.*, 2023].

Designed for scalability and generalizability, the AI-based solution integrates indicators from UNICEF MICS surveys [UNICEF, 2025b] to align with global child health standards and is adaptable to low-resource settings where disparities in healthcare access, financial insecurity, and limited infrastructure challenge pediatric cancer care. By addressing the SDoH, our approach ensures equitable, accessible, and contextually relevant support for families navigating complex health systems [Mercado *et al.*, 2024].

4 Method: Integrating Research and Practice

We introduce the overall operation of **`SaludASR`**. We then describe our objectives, the way they will be achieved, a strategy for continual improvement of the system.

Our aim is to develop an AI-based system that engages with patients' caregivers over existing landlines and mobile networks using voice communication. This system will offer human-like conversations and guidance, answering questions posed by its users.

The proposed system, denoted as **`SaludASR`**, will merge the latest advancements in LLMs (e.g., tool calls and advanced Retrieval-Augmented Generation) with our existing knowledge of SDoH and additional information gathered during RO1. The workflow involves automatic calls from the system to users, greeting them upon answering (via LLM text generation and voice synthesis). Authentication will be conducted by inquiring about information known only to the patient's caregiver and hospital, such as the last digits of their medical record number. If the authentication fails, or the caregiver refuses to talk, hospital staff are notified of this event so they can take relevant action.

Following successful authentication, the system will prompt the caregiver to provide a basic report on both the patient's condition and their current SDoH status. The LLM will generate questions based on a list of variables and potential questions collected during the preliminary mixed-methods study and refined further in RO1. Upon receiving the report, ASR will convert voice into text, which is then processed by an LLM to extract structured information, including communicated values of variables of interest and their confirmation by the caregiver. If the caregiver indicates that they are in an emergency, or if some of the indicators communicated by them show that the patient is in need of urgent care, the system will automatically connect the caregiver to the relevant medical professional, who will be provided with all the collected data to aim them in providing medical guidance to the caregiver.

The LLM will be fine-tuned and prompted to comprehend subtle culturally-specific cues about SDoH when expressed by the caregiver. This approach enhances the human-like nature of conversations, allowing for more comprehensive information extraction. Communicated variable values will then be securely stored in a HIPAA-compatible database. For data extraction, the system will utilize constrained generation [Geng *et al.*, 2023], ensuring that the output of the LLM corresponds to a format that is acceptable by the database.

For phone network integration, HIMFG already employs the Twilio platform for interactions with patients and caregivers; the pilot version of **`SaludASR`** will leverage this existing integration for phone functionality. This will enable us to effectively reuse the hospital's resources, while also ensuring regulatory compliance for this part of the system.

During the initial stages of deployment, our system will store direct patient quotes as text strings in the database, ensuring accurate data retention and facilitating error diagnosis through manual reviews conducted by hospital staff. Following deployment, we plan to enhance the system's information extraction capabilities, allowing it to retrieve numerical data when required by healthcare personnel for efficient information delivery.

Next, the system will identify missing information for the patient in the database. Depending on this assessment, it will either conclude the call (if no information is missing) or generate additional questions to the caregiver to solicit the missing data. These questions will undergo voice synthesis before being relayed back to the user, thereby completing the interaction loop.

The questions will be retrieved from the database and processed using an adaptive RAG approach [Gao *et al.*, 2024] that combines both checking for missing information, and synthesizing new questions to the caregivers that allow for more effective solicitation of medical information. This strategy allows the system to retrieve only the information that is missing, and to do it in the most efficient way possible, ensuring efficient communication with the caregivers.

This system design ensures that caregivers can convey crucial medical information over the phone, even in areas with unreliable or absent internet access. It also enables patients to communicate this information using simple language, as the LLM can accurately interpret precise medical statements from their speech. This approach facilitates effective communication of important medical data without requiring specialized training to understand complex medical terminology.

Furthermore, this design does not necessitate application use (or smartphone access), thereby offering more opportunities for communication to individuals without smartphones or computers. Voice communication allows for multitasking (e.g., while driving) and eliminates the need for lengthy, time-consuming forms, thus streamlining communication and improving response rates.

To address concerns regarding the processing of protected health information (PHI) through a cloud-based AI system, we will employ specialized Client Cloud Computing for PHI (C3PHI) secure storage and computing environment. Our solution will be initially hosted in a C3PHI cloud provided by Amazon AWS, offering access to a HIPAA-compatible AI API based on Anthropic Claude [Anthropic, 2024], a leading AI model, as well the capability to run the agentic system in an EC2 environment. The environment provided by Amazon additionally allows for deployment of custom LLMs, making it possible to use specialized models and complex pipelines for medical information extraction and parsing. [1]

For the initial deployment, HIMFG will expand their current use of Twilio, an ASR/voice synthesis/phone calling system, by leveraging its capabilities to connect caregivers with the LLM component of the proposed system.

4.1 Research Objectives

In this section we introduce the research objectives of our study, which correspond to our research questions.

RO1 Reducing Burden: The voice-based interface provides a more natural avenue for caregivers to interact with the system through the use of their voice. This system will allow them to easily communicate important medical information to the hospital without the need to fill out long, complex, and time-consuming forms.

It will also enable human-like conversations, making it more dynamic and engaging for caregivers. Furthermore, since the system can now become an active party by calling the caregivers first, it will effectively address the issue of them reducing the use of the existing applications over time.

RO2 Addressing community needs: Our bank of questions collected during the preliminary study will enable the system to collect a wide gamut of relevant information, enabling near-realtime symptom monitoring. Furthermore, we will collect information on nutrition, thus enabling medical staff (and, in the future, an automated system) to provide personalized nutritional advice.

Furthermore, the system will allow for direct connections between medical staff and caregivers. This will allow for effective communication in emergencies, by providing caregivers with a direct link to hospital staff, and by providing hospital staff with the latest data collected via the tool.

RO3 Addressing Material Challenges: The system will be able to log and track the caregivers' material and financial concerns. This will enable healthcare personnel to connect

[1] A prototype of this system (using a text interface) is available at https://himfg-llm-test.grisha.xyz/

caregivers with the relevant NGOs to ensure that their needs are met. Having this communication channel will also enable the hospital to schedule treatments in a way to minimize the burden to the most vulnerable caregivers.

RO4 Addressing Logistical Challenges: `SaludASR` will include additional questions concerning the caregivers themselves, including their emotional state and their financial situation. This information will be given to healthcare personnel, who will then be able to refer the caregivers to relevant support groups. This will allow the caregivers to receive timely support, both in the financial, and in the emotional sense.

4.2 Continual Evaluation

We will perform continual evaluation of the performance of **`SaludASR`** a two-pronged approach, which we will employ after deploying the system at HIMFG and allowing for a 2-week period of initial use. In the first prong, we will conduct follow-up interviews with caregivers to gather their experiences with the system. This feedback will be invaluable in identifying and addressing any consistent patterns of inadequate behavior exhibited by the system.

To optimize communication with caregivers, these interviews will be conducted via phone and designed to be brief, thereby improving response rates and minimizing time commitments for both caregivers and interviewers.

In the second prong, we will record random conversations with caregivers along with ASR and LLM interaction logs. These recordings will then undergo human verification to ensure accuracy in extracting information from the system's interactions. This process will enable us to promptly detect and rectify instances where the system misinterprets or overlooks relevant information.

In all processes involving human feedback, we will implement automatic deidentification pipelines based on LLMs [Liu *et al.*, 2023] to remove Personally Identifiable Information (PII) and Protected Health Information (PHI) from data sent to human reviewers. This strategy guarantees high-quality human feedback for effectively diagnosing system issues while safeguarding patients' and caregivers' privacy.

These strategies will foster continuous enhancements in the system's ability to extract information and engage in effective conversations, ultimately ensuring superior care for pediatric cancer patients. Furthermore, these approaches will facilitate the system's adaptation to diverse settings (e.g., other hospitals) by helping us anticipate and mitigate potential failure modes.

4.3 Project Timeline

The proposed study is structured within a 18-month timeline, incorporating all our research objectives. We initiate with the initial deployment if **`SaludASR`**, which will happen concurrently with HIMFG staff training the use of the new system. This initial phase is anticipated to conclude within three months (months 1-3).

This initial phase will be followed by a 12-months period of the use of the system in the hospital setting. During this period, the system will be used by actual HIMFG patients, their caregivers, and hospital staff. We anticipate that this phase of our study will take 12 months (months 4-15).

After this, the final phase of the study will start, where we will collect patient, caregiver and staff feedback on their use of the system. We will use the same data collection methodology as we did for our preliminary study. This will enable us to compare the feedback before and after the introduction of **`SaludASR`**.

Throughout the study's duration, we commit to ongoing observation and improvement of the system as deployed in HIMFG, commencing with its initial deployment in month 10 and continuing while the system is in use at HIMFG.

5 Scalability and Economic Sustainability

For the initial deployment of **`SaludASR`**, we have established a partnership with HIMFG, with it serving as a platform for pilot testing and study. In addition, the University of Notre Dame's partnership with Amazon provides us access to their HIPAA-compliant computing environments, which we can utilize to host the system.

We are currently planning to scale up our project within Mexico and expand into other countries in Latin America. In Mexico, we are working on establishing a partnership with an another large pediatric care organization. Specializing in pediatric oncology, hematology, and neonatal surgery, the hospital boasts a 65.7% occupancy rate. It caters to children from both Mexico City and other parts of Mexico, offering advanced diagnostic and treatment options. Notably, the hospital employs a fully digital medical record system, enabling us to integrate **`SaludASR`** within its fully digitized environment.

Outside of Mexico, we are focused on building partnerships with leading oncology centers in Colombia. These institutions represent advanced public oncology centers providing cutting-edge treatment options to pediatric cancer patients. Furthermore, they serve as major research hubs, contributing to advancements in pediatric oncology care within Colombia and the wider Latin American context.

Beyond the pilot phase, successful outcomes from the study and clinical trial pave the way for commercial viability through further collaborations with hospitals throughout Latin America. A potential startup will assume responsibility for the system's development and maintenance, enabling its adaptation and implementation across hospitals in the region.

6 Conclusions

Pediatric oncology patients who are living in social and economic marginalization face severe healthcare access barriers due to fragmented systems and absent social health protection programs [Chalfant *et al.*, 2023]. Caregivers struggle with navigating healthcare processes and receiving timely guidance, exacerbated by limited internet penetration and reliance on outdated technology in large Mexican hospitals. Success depends on combining AI, natural language processing for literacy-independent interactions, and evidence-based guidance with deep understanding social determinants of health to enhance caregiver support and improve pediatric cancer care in low-resource settings. **`SaludASR`** brings all these elements together in a novel way.

We formulated hypotheses regarding the potential benefits of our proposed innovations: (1) AI, specifically the combination of LLM and ASR, can enhance healthcare outcomes for marginalized pediatric cancer patients; (2) these AI-based systems can be continually refined through user feedback and community input; (3) these AI systems can be effectively scaled up and transferred across countries, at least within Latin America, while maintaining success despite cultural and technological differences. To test these hypotheses, we proposed a comprehensive research framework that integrates social science methodologies (interviews and focus groups), computer science techniques (audio and text analysis, AI research), and established strategies from healthcare research (clinical trials). These approaches will allow us to validate our hypotheses and inform the future development and commercialization of **`SaludASR`**.

Ultimately, we aim to develop a data-driven, linguistically and culturally contextualized AI tool that integrates SDoH to enhance caregiver support and improve pediatric cancer care in low-resource settings and for marginalized groups. This tool will facilitate natural, literacy-independent interactions in Spanish, provide reliable, evidence-based guidance, and strengthen community support networks by connecting caregivers with local health workers and educators. By ensuring scalability and alignment with global pediatric health standards, we seek to reduce healthcare disparities, improve access to timely care, and address structural barriers affecting families in vulnerable communities.

7 Directions

We rely heavily on participatory design to continually improve our service; however, we will also consider more sophisticated methods that engage the public to guide further development . The project may be affected by inherent biases in LLM and ASR technologies [Kotek *et al.*, 2024]. These include the potential inability of LLMs and ASR to understand certain Spanish dialects [Kantharuban *et al.*, 2023] and their tendency to ignore cultural aspects of communication [Qiu *et al.*, 2024]. Potential mitigation strategies include fine-tuning models on texts from specific dialects, or using in-context learning to improve understanding. One additional avenue for improvement would be to utilize multimodal LLMs with native voice capability in the pipeline, as a way to enhance or replace ASR.

Additionally, **`SaludASR`** faces challenges related to potentially inappropriate model outputs caused by users, both inadvertently and deliberately. To combat these issues, we can introduce additional postprocessing steps [Inan *et al.*, 2023] into the pipeline that detect and address such outputs when generated by the model. We also plan to invite community members to review anonymized model outputs and help improve both the LLM and the inappropriate output detector.

Ensuring patient privacy and HIPAA compliance remains a significant challenge for any project combining healthcare and AI. In our project, we will mitigate these concerns by using a compliant AI API provider in a secure environment compatible with HIPAA [Amazon, 2018] and by utilizing only data that has been anonymized and cleared for release by hospital staff.

Hospitals like HIMFG, along with patients and caregivers, demonstrate a strong sense of commitment to their mission. We witnessed this firsthand during our interactions and pilot study. By engaging with these stakeholders, we will continue refining our research direction, testing the applicability of our innovations beyond HIMFG, and disseminating our findings and developments in both academia and industry.

Ethical Statement

One potential risk with any healthcare application is a potential for leaks of protected health information, which can have potentially devastating consequences to patients and their caregivers. **`SaludASR`** faces this risk as well. Furthermore, during the pilot stages of the system deployment, certain de-identified conversations will be revealed to human annotators, who are part of a third party certified to process PHI. The risk in this case is comparable to any other healthcare application utilizing a similar process. **`SaludASR`** faces an additional privacy risk where PHI can inadvertently end up in an LLM training or fine-tuning dataset, and then be revealed by the model. We will mitigate this risk by having an agreement with the API provider forbidding them to use the data for training their models, and by manually verifying our own model fine-tuning datasets.

The main new ethical risk introduced by **`SaludASR`** is that our system might inadvertently perpetuate societal biases that were imprinted on LLMs by their training datasets, therefore negatively affecting people already vulnerable populations. We mitigate this risk by performing a continual review of the questions asked and information collected by the system, to ensure early detection and correction of such biases. For correction of these biases, we will use both prompt engineering techniques and model fine-tuning, to make the correction of model biases both timely and complete. Moreover, the caregivers will always be able to indicate that the questions asked to them by the system are inappropriate, and they will automatically be flagged for review.

An additional risk that is common to all LLM-based solutions is the risk of them generating inappropriate outputs. Existing LLMs employ a variety of train-time strategies to mitigate this risk. We plan to address this further, by introducing additional stages in the pipeline that will detect and remove any potentially harmful content generated by the model. Obviously, this needs to be balanced with the ability of the system to process medical information.

This project is in compliance and will continue to comply with the IRB review process at the University of Notre Dame.

Contribution Statement

Grigorii Khvatskii: Writing – original draft, Software, Conceptualization, Methodolody; **Angelica Garcia Martinez**: Writing – original draft, Methodology, Investigation, Conceptualization, Project administration, Resources; **Deng Pan**: Software, Investigation, Methodology Matthew Belcher: Software, Investigation, Methodology, Project administration; **Gerónimo Medrano Loera**: Data curation, Investigation; **Dayana Pineda Pérez**: Data curation, Investigation; **Juan Emmanuel Ferrari Muñoz-Ledo**: Data curation, Investigation; **Horacio Márquez-González**: Funding acquisition, Resources, Supervision, Project administration, Conceptualization; **Nuno Moniz**: Conceptualization, Funding acquisition, Methodology, Resources, Supervision, Writing - review and editing; **Nitesh V Chawla**: Conceptualization, Funding acquisition, Methodology, Resources, Supervision, Writing - review and editing;

References

[Amazon, 2018] Amazon. Architecting of hipaa security and compliance on amazon web services, 2018.

[Anthropic, 2024] Anthropic. The Claude 3 Model Family: Opus, Sonnet, Haiku. https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf, 2024. [Accessed 07-02-2025].

[Avilés-Robles *et al.*, 2022] Martha Avilés-Robles, Jennifer J Schnur, Elisa Dorantes-Acosta, Horacio Márquez-González, Luis A Ocampo-Ramírez, and Nitesh V Chawla. Predictors of septic shock or bacteremia in children experiencing febrile neutropenia post-chemotherapy. *Journal of the Pediatric Infectious Diseases Society*, 11(11):498–503, 08 2022.

[Black *et al.*, 2017] Maureen M Black, Susan P Walker, Lia C H Fernald, Christopher T Andersen, Ann M DiGirolamo, Chunling Lu, Dana C McCoy, Günther Fink, Yusra R Shawar, Jeremy Shiffman, Amanda E Devercelli, Quentin T Wodon, Emily Vargas-Barón, and Sally Grantham-McGregor. Early childhood development coming of age: science through the life course. *The Lancet*, 389(10064):77–90, 2017.

[Braun and Clarke, 2021] Virginia Braun and Victoria Clarke. *Thematic analysis: A practical guide*. sage, 2021.

[Chalfant *et al.*, 2023] Victor Chalfant, Carlos Riveros, Scott M Bradfield, and Andrew A Stec. Impact of social disparities on 10 year survival rates in paediatric cancers: a cohort study. *The Lancet Regional Health–Americas*, 20, 2023.

[Creswell and Creswell, 2017] John W Creswell and J David Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications, 2017.

[Dellinger, 2022] Jolynn Dellinger. *Pandemic ethics: the intersection of technology, trust, and privacy, and implications for marginalized communities*, page 204–215. Edward Elgar Publishing, November 2022.

[Elango *et al.*, 2015] Sneha Elango, Jorge Luis García, James J Heckman, and Andrés Hojman. Early childhood education. In *Economics of Means-Tested Transfer Programs in the United States, Volume 2*, pages 235–297. University of Chicago Press, 2015.

[Gao *et al.*, 2024] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. Retrieval-augmented generation for large language models: A survey, 2024.

[Geng *et al.*, 2023] Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. Grammar-constrained decoding for structured NLP tasks without finetuning. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Singapore, December 2023. Association for Computational Linguistics.

[HIMFG, 2022] HIMFG. Anuario Estadistico de Servicios Medicos. http://himfg.com.mx/descargas/documentos/planeacion/AnuariosEst/Anuario_Estadistico_HIMFG_2022.pdf, 2022. [Accessed 10-02-2025].

[Inan *et al.*, 2023] Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabsa. Llama guard: Llm-based input-output safeguard for human-ai conversations, 2023.

[Kantharuban *et al.*, 2023] Anjali Kantharuban, Ivan Vulić, and Anna Korhonen. Quantifying the dialect gap and its correlates across languages, 2023.

[Knaul *et al.*, 2023] Felicia Marie Knaul, Hector Arreola-Ornelas, Michael Touchton, Tim McDonald, Merike Blofield, Leticia Avila Burgos, Octavio Gómez-Dantés, Pablo Kuri, Adolfo Martinez-Valle, Oscar Méndez-Carniado, et al. Setbacks in the quest for universal health coverage in mexico: polarised politics, policy upheaval, and pandemic disruption. *The Lancet*, 402(10403):731–746, 2023.

[Kotek *et al.*, 2024] Hadas Kotek, David Q. Sun, Zidi Xiu, Margit Bowler, and Christopher Klein. Protected group bias and stereotypes in large language models, 2024.

[Liu *et al.*, 2023] Zhengliang Liu, Yue Huang, Xiaowei Yu, Lu Zhang, Zihao Wu, Chao Cao, Haixing Dai, Lin Zhao, Yiwei Li, Peng Shu, Fang Zeng, Lichao Sun, Wei Liu, Dinggang Shen, Quanzheng Li, Tianming Liu, Dajiang Zhu, and Xiang Li. Deid-gpt: Zero-shot medical text de-identification by gpt-4, 2023.

[Ma *et al.*, 2024] Yingyi Ma, Zhe Liu, and Ozlem Kalinli. Effective text adaptation for llm-based asr through soft prompt fine-tuning, 2024.

[Mercado *et al.*, 2024] Carla I Mercado, Kai McKeever Bullard, Michele LF Bolduc, Courtni Alexis Andrews, Zoe RF Freggens, Grace Liggett, Desmond Banks, Shanice Battle Johnson, Ana Penman-Aguilar, and Rashid Njai. A shift in approach to addressing public health inequities and the effect of societal structural and systemic drivers on social determinants of health. *Public Health Reports®*, page 00333549241283586, 2024.

[Nurturing Care, 2025] Nurturing Care. Nurturing Care Framework for Early Childhood Development — nurturing-care.org. https://nurturing-care.org/, 2025. [Accessed 08-02-2025].

[Patton, 2014] Michael Quinn Patton. *Qualitative research & evaluation methods: Integrating theory and practice*. Sage publications, 2014.

[Qin *et al.*, 2024] Ruiyang Qin, Dancheng Liu, Gelei Xu, Zheyu Yan, Chenhui Xu, Yuting Hu, X. Sharon Hu, Jinjun Xiong, and Yiyu Shi. Tiny-align: Bridging automatic speech recognition and large language model on the edge, 2024.

[Qiu *et al.*, 2024] Haoyi Qiu, Alexander R. Fabbri, Divyansh Agarwal, Kung-Hsiang Huang, Sarah Tan, Nanyun Peng, and Chien-Sheng Wu. Evaluating cultural and social awareness of llm web agents, 2024.

[Schnur *et al.*, 2024] Jennifer J. Schnur, Angélica Garcia-Martínez, Patrick Soga, Karla Badillo-Urquiola, Alejandra J. Botello, Ana Calderon Raisbeck, Sugana Chawla, Josef Ernst, William Gentry, Richard P. Johnson, Michael Kennel, Jesús Robles, Madison Wagner, Elizabeth Medina, Juan Garduño Espinosa, Horacio Márquez-González, Victor Olivar-López, Luis E. Juárez-Villegas, Martha Avilés-Robles, Elisa Dorantes-Acosta, Viridia Avila, Gina Chapa-Koloffon, Elizabeth Cruz, Leticia Luis, Clara Quezada, Emanuel Orozco, Edson Serván-Mori, Martha Cordero, Rubén Martín Payo, and Nitesh V. Chawla. Saludconectamx: Lessons learned from deploying a cooperative mobile health system for pediatric cancer care in mexico. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing*, CSCW Companion '24, page 316–322, New York, NY, USA, 2024. Association for Computing Machinery.

[Serván-Mori *et al.*, 2023] Edson Serván-Mori, Octavio Gómez-Dantés, David Contreras, Laura Flamand, Diego Cerecero-García, Héctor Arreola-Ornelas, and Felicia M Knaul. Increase of catastrophic and impoverishing health expenditures in mexico associated to policy changes and the covid-19 pandemic. *Journal of Global Health*, 13, 2023.

[UNICEF, 2025a] UNICEF. Early Childhood Development. https://data.unicef.org/topic/early-childhood-development/overview/, 2025. [Accessed 07-02-2025].

[UNICEF, 2025b] UNICEF. UNICEF MICS. https://mics.unicef.org/surveys, 2025. [Accessed 07-02-2025].

[WHO, 2018] WHO. Nurturing care — who.int. https://www.who.int/teams/maternal-newborn-child-adolescent-health-and-ageing/child-health/nurturing-care, 2018. [Accessed 08-02-2025].

[World Bank, 2025] World Bank. Early Childhood Development — worldbank.org. https://www.worldbank.org/en/topic/earlychildhooddevelopment, 2025. [Accessed 08-02-2025].

[Yang *et al.*, 2024] Ziqi Yang, Xuhai Xu, Bingsheng Yao, Ethan Rogers, Shao Zhang, Stephen Intille, Nawar Shara, Guodong Gordon Gao, and Dakuo Wang. Talk2care: An llm-based voice assistant for communication between healthcare providers and older adults. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 8(2), May 2024.

[Zhai *et al.*, 2023] Shumenghui Zhai, Frances Chu, Minghui Tan, Nai-Ching Chi, Teresa Ward, and Weichao Yuwen. Digital health interventions to support family caregivers: An updated systematic review. *Digital health*, 9:20552076231171967, 2023.