Bridging the Gap: Inclusive Artificial Intelligence for Patient-Oriented Language Processing in Conversational Agents in Healthcare

Kerstin Denecke
Bern University of Applied Sciences
Quellgasse 21, 2502 Biel/Bienne
Switzerland
kerstin.denecke@bfh.ch

Abstract

Conversational agents (CAs), such as medical interview assistants, are increasingly used in healthcare settings due to their potential for intuitive user interaction. Ensuring the inclusivity of these systems is critical to provide equitable and effective digital health support. However, the underlying technology, models and data can foster inequalities and exclude certain individuals. This paper explores key principles of inclusivity in patient-oriented language processing (POLP) for healthcare CAs to improve accessibility, cultural sensitivity, and fairness in patient interactions. We will outline, how considering the six facets of inclusive Artificial Intelligence (AI) will shape POLP within healthcare CA. Key considerations include leveraging diverse datasets, incorporating gender-neutral and inclusive language, supporting varying levels of health literacy, and ensuring culturally relevant communication. To address these issues, future research in POLP should focus on optimizing conversation structure, enhancing the adaptability of CAs' language and content, integrating cultural awareness, improving explainability, managing cognitive load, and addressing bias and fairness concerns.

1 Introduction

Conversational agents (CAs) in healthcare - intelligent systems that enable natural language interaction - have the potential to improve access to healthcare services, enhance patient literacy (Wynia and Osborn, 2010), and empower individuals to make informed healthcare decisions. These systems are increasingly being used in a variety of applications, including medical history taking (Denecke et al., 2024), blended psychotherapy, and delivery of cognitive behavioural therapy (e.g., WoeBot (Sackett et al., 2024)). While early healthcare CAs were predominantly rule-based, the emergence of transformer-based models has opened up new possibilities for more dynamic, flexible and engaging conversations enabled through patient-oriented language processing (POLP) (Sarker et al., 2021).

A critical component of effective healthcare CAs is the ability of these systems to tailor communication to individual patient needs, taking into account factors such as health literacy, cultural context and linguistic diversity that is realized by artificial intelligence (AI), natural language processing (NLP), ideally incorporating POLP. While AI-driven CAs hold promise for improving access to healthcare, they also pose risks, particularly in terms of exacerbating existing inequalities. The design and implementation of these systems, including the AI models and datasets on which they rely, often fail to adequately represent diverse user populations, leading to biased outcomes and interaction barriers (Cross et al., 2024). Marginalised communities, already disproportionately affected by structural inequalities in healthcare, may be further excluded if AI fails to process and deliver health information in a way that meets their linguistic and cognitive needs. In addition, if patient-facing AI systems present excessive, irrelevant or poorly prioritised information, users may feel overwhelmed, hindering their ability to derive meaningful insights for their health concerns.

To develop inclusive POLP within healthcare CAs, the methods must be designed with fairness, adaptability, and user-centred communication strategies in mind. Inclusive AI involves integrating diverse human attributes and perspectives throughout the entire lifecycle of AI systems—from data collection and model training to implementation and governance (Zowghi and Bano, 2024). Nadarzynski et al. proposed a 10-phase roadmap for the design and implementation of inclusive CAs in healthcare (Nadarzynski et al., 2024), primarily focusing on system development and evaluation. In contrast, this paper takes a more specific approach, focusing on the linguistic and technical aspects required to achieve inclu-

Proceedings of the Second Workshop on Patient-Oriented Language Processing (CL4Health), pages 456–462
May 4, 2025 ©2025 Association for Computational Linguistics

sive POLP within healthcare CAs. While some general aspects of usability and accessibility will be discussed, our primary focus lies in ensuring that language processing within CAs effectively accommodates diverse patient populations.

In the era of digitisation and AI-driven healthcare, it is imperative to ensure that AI-driven conversational systems are not only functionally efficient, but also linguistically and culturally inclusive. Despite increasing discussions on AI ethics and responsible AI practices, there is still a lack of practical strategies to bridge the technology divide, especially for underserved populations.

This paper aims to address this gap by exploring concrete strategies for ensuring inclusivity in AI-driven healthcare CAs. Specifically, we explore aspects for improving POLP for ensuring that healthcare CAs account for linguistic diversity, varying levels of health literacy, and cultural sensitivities. The World Health Organization's 2021 Ethical Guidelines emphasise inclusivity, transparency and accountability as core principles for AI in healthcare (World Health Organization, 2024). Building on this framework, our study aims to contribute to the development of more equitable AI-driven healthcare solutions that prioritise inclusivity and accessibility for all patient demographics.

To achieve this, we address the following research questions (RQs):

- RQ1: How does inclusive AI impact POLP in healthcare CAs?
- RQ2: What NLP or other technologies are required to achieve inclusive POLP in healthcare CAs?
- RQ3: What are the most important future research directions for achieving inclusive POLP in healthcare CAs?

By answering these questions, this paper aims to contribute actionable insights for the design and implementation of fair, effective, and inclusive POLP in healthcare CAs. It is intended as starting point of research towards inclusive POLP for CA in healthcare.

2 Methods

To answer our research questions, we apply two steps. First, we use the facets of inclusive AI that have been collected in a recent review (Bokolo et al., 2025) and study how inclusive AI impacts on POLP within healthcare CA based on our experience in developing CAs.

In a previously conducted review (Bokolo et al., 2025), we retrieved papers that included the keyword "Inclusive AI" in their abstract from six databases (PubMed, PsycINFO, CINAHL, Academic Search Premier, IEEE Xplore, and Scopus). The included research studies should address inclusive AI in the context of healthcare. Out of 1377 papers, 18 were included with information extracted on strengths, weaknesses, opportunities and threats (SWOT). From this SWOT analysis, six facets of inclusive AI in healthcare were concluded: 1) Accessibility, 2) Equity, 3) Usability and Navigability, 4) Diversity and Cultural Sensitivity, 5) Mitigation of Disparities, and 6) Skill Development and Literacy. In more detail, accessibility asks AI technologies in healthcare to be designed in a useable manner by individuals with diverse needs (Accessibility). They should offer fair access and outcomes for everyone (Equity) and must be designed in a user-friendly manner (Usability and Navigability). AI technologies in healthcare should recognize and accommodate socio-demographic diversity (Diversity and Cultural Sensitivity), should mitigate existing disparities (Mitigation of Disparities). Users should be equipped with relevant skills to engage efficiently with AI technologies when used for healthcare purposes (Skill development and literacy).

We will consider these six facets and assess how inclusive AI impacts on POLP within CA in healthcare. In a second step, we suggest a research agenda for technologies and methodologies to address the identified impact factors.

3 Impact of Inclusive AI on POLP within Healthcare CA

This section is structured along the 6 facets of inclusive AI described in the section before. We will outline how inclusive AI shapes POLP within CA with regard to these specific facets.

3.1 Accessibility

Accessibility takes into account users with disabilities and different abilities such as visual, hearing, motor or cognitive impairments (Henni et al., 2022). An accessible CA must provide multimodal interaction to ensure that users with different abilities can effectively engage with it. This includes speech-to-text and text-to-speech capabilities for visually impaired or low-literacy users, keyboard-only navigation for those with motor impairments, and support for screen readers. In addition, the CA should provide high-contrast visual options, adjustable font sizes, and easy navigation to accommodate users with cognitive or visual impairments.

From a POLP perspective, the CA should:

- Use short, clear sentences to break down complex medical information.
- Offer step-by-step explanations for processes such as measuring blood pressure, ensuring better comprehension.
- Enable adaptive communication styles, allowing users to choose between brief responses and detailed explanations based on their needs.
- Provide direct answers with optional elaboration, offering additional details upon request.

To effectively implement these or similar features into healthcare CAs, adaptive language generation could be applied such as text simplification models to adjust the complexity of responses. User profiling and context-aware interactions consider the user's preferences for adapting answer length, style, etc.

A reverse Chain-of-Thought prompting, where the CA explicitly guides users through stepwise instructions (e.g., breaking a process into incremental, explainable steps for better user comprehension) could better guide through processes. Depending on the purpose of the CA, the conversation could be implemented as progressive disclosure where information is revealed gradually,

3.2 Equity

Studies show that limited health literacy is linked to poor health outcomes, increased healthcare costs, and health disparities (Gibney et al., 2020). Digital communication tools in healthcare, including CAs, have the potential to improve health literacy and empower individuals to take a more active role in managing their health (Fitzpatrick, 2023). In this context, CAs can play a critical role in providing equitable responses tailored to different socio-economic backgrounds, ensuring that all individuals - regardless of location, income or education - receive accurate and relevant health information formulated and presented in a way that addresses their reading skills, health literacy and data literacy (Nadarzynski et al., 2024).

Linguistic and culture inclusivity could be achieved by multilingual support, but needs also additional aspects such as cultural-appropriate health recommendations (see section 3.4). Trauma-informed conversational strategies (Berring et al., 2024) could be applied to address specific needs of users with trauma: NLP models should be designed to recognize distress and provide gentle, supportive responses. For example, if a user expresses suicidal thoughts, the CA should prioritize crisis intervention resources over generic health advice.

To address these and similar aspects related to equity, healthcare CAs should use simple, jargon-free language, integrating explanatory visuals, and providing localised health advice based on regional medical practices. Underlying AI models need to be trained on datasets that include diverse user groups to avoid biases that could lead to misinformation or exclusion of marginalised populations.

3.3 Usability and Navigability

Previous research also showed that user interfaces must be designed with consideration of the information requirements, cognitive capabilities, and limitations of end users in healthcare environments (Patel and Kushniruk, 1998). Therefore, healthcare CAs should be designed with an intuitive, patient-friendly interface that prioritizes clarity, guidance, and responsiveness (Denecke, 2023). For example, guidance would mean that the CA guides the user through the conversation, supports when the user has no idea what to write or say. Also structuring the dialogue could help or summarizing previously said aspects from time to time when the interaction gets long.

CAs should provide clear fallback options, such as the ability to speak with a human operator or access a help menu when the CA fails in recognizing user intent. They should maintain a consistent tone throughout the conversation which could be formal, friendly or empathetic. Proactive engagement would make the interaction more user-centric and intuitive. By anticipating what the user might need next, CAs can offer relevant information or actions before the user even asks. For example, if a user has been discussing symptoms of a cold, the CA might proactively suggest remedies or ask if they need a doctor's appointment. Based on previous interactions, the CA can suggest next steps or related information, making the user feel understood and supported.

Aspects mentioned for accessibility or equity could also support usability (e.g. clear, concise language).

3.4 Diversity and Cultural Sensitivity

Cultural factors have been identified to affect access to and uptake of digital health technologies among culturally and linguistically diverse populations (Davies et al., 2024; Whitehead et al., 2023). An inclusive patient-facing healthcare CA must be culturally competent and linguistically adaptable. This requires multilingual support with real-time translation capabilities to communicate in regional dialects and under-represented languages. In addition, the CA should be able to adapt health recommendations based on cultural beliefs, dietary restrictions and traditional medical practices. Avoiding gender bias in language, respecting gender identity pronouns, and acknowledging religious sensitivities in healthcare (e.g. fasting during Ramadan) are critical to making the CA more inclusive. By incorporating cultural nuances and linguistic diversity, healthcare CA can foster trust, improve engagement and increase the effectiveness of interactions, ultimately leading to better health outcomes for all communities (Davies et al., 2024).

3.5 Mitigation of Disparities

Individuals may struggle with limited health literacy, so it is essential for inclusive POLP in healthcare CA to simplify medical language and ensure that critical health information is easy to understand. Strategies to achieve this include dynamic simplification, where the CA adjusts its complexity based on the user's familiarity with medical terms, and interactive learning features such as visual aids, audio explanations and quizzes that reinforce understanding. The CA should proactively identify and clarify misunderstood terms and offer alternative explanations in simpler language to bridge gaps in understanding. Again, underlying data has to accurately represent different demographics, preventing the reinforcement of harmful stereotypes and the exclusion of marginalised groups.

3.6 Skill Development and Literacy

From a POLP perspective, skill development and CA literacy are essential to ensure that patients can effectively interact with healthcare CA, understand medical information, and make informed health decisions. To support skill development, healthcare CA should incorporate strategies to teach users how to engage with the CA, increase health literacy, and build confidence in using digital health tools. When users first interact with a CA, it should provide a guided onboarding experience that explains its capabilities, how to ask questions, and how to navigate responses and inform about the possibilities and shortcomings of the CA. Offering simple scenarios or guided exercises (e.g., "Try asking me about your symptoms!") can help users become comfortable with the interactions. Users may be unsure of how to phrase health-related questions effectively and what could the CA be asked. To address this, a CA can guide users by offering question templates (e.g. "You can ask me: 'What are the symptoms of diabetes?'").

4 Research Agenda for Inclusive Patient-oriented Language Processing in CA

The previous sections described the characteristics that inclusive POLP within a healthcare CA should provide. Considered from multiple facets, we can recognize that some aspects are of relevance to support multiple facets (e.g. concise language supports accessibility, equity and usability). Some technologies are already available to realize these aspects, while others still require research efforts. In this section, we are outlining possible research directions recommended for future research towards inclusive patient-oriented language processing within healthcare CAs. Table 1 lists some possible research questions for the future.

Conversation structure. Research is needed to determine how conversations in healthcare CA should be structured and how to implement these structures in POLP. A well-structured conversation flow ensures that information is delivered in a clear, understandable and digestible way. This minimises confusion and allows users to focus on key health information without unnecessary complexity. To achieve this, it is essential to analyse the linguistic and cultural barriers that affect communication with patients. These barriers can include low health literacy, non-native speakers and regional dialects. In addition, understanding how patients with disabilities - such as blindness, hearing loss or cognitive impairment - interact with health care CAs is crucial. Incorporating participatory methods during the design phase can help gather input from these user groups to ensure accessibility and usability.

Research area	Examples for possible research questions
Conversation structure	How should language simplification be implemented in healthcare CAs to ensure user comprehension of complex medical concepts?
Adaptability	How can CAs dynamically adapt to user's reading level, health literacy or cognitive load?
Cultural awareness	How can CA responses become culturally sensitive?
Explainability	How can explainability be included in the conversation flow without disturbing it?
Cognitive load analysis	How can a CA analyse a user's cognitive load in real time?
Bias and fairness-awareness	Which bias mitigation techniques could be implemented into healthcare CAs in general or into POLP specifically?

Table 1: Research areas and examples for research questions towards inclusive patient-oriented language processing in healthcare CA

Considering established rules of communication or best practices from patient-doctor interactions can help in designing effective conversation flows (Denecke, 2023).

Research could also explore how patients interpret and respond to medical terminology. Identifying areas where NLP-based language simplification or explanation of medical concepts is needed can improve comprehension. However, simplification must be carefully balanced, as excessive reduction of medical terminology may result in the loss of critical health information.

Cognitive load analysis. Interacting with a CA could be overwhelming when the conversation gets long and comprehensive. It could be studied whether NLP or other techniques can be used for real-time cognitive load detection (Zayim et al., 2023). This would mean signs of frustration, stress or cognitive fatigue could be recognized while the interaction takes place which in turn would allow to adjust CA responses accordingly. For example, inclusion of multimodal AI could offer an opportunity to address detected cognitive load by allowing for various modes of communication.

Cultural awareness. Another research direction regarding inclusive POLP within CA is how to consider cultural aspects in the health dialogue. Research could explore ways to incorporate cultural sensitivity into language models used in healthcare CAs. This may involve adapting conversation flows to align with cultural norms, addressing variations in health-related beliefs, or ensuring that medical terminology is explained in ways that resonate with different communities.

Adaptability. The previous three areas already indicate another important direction of research. While rule-based CAs capture the flow of conversations through predefined rules, LLMs and other AI methods offer greater flexibility and adaptability. Future healthcare CA can adapt their responses to the user's health literacy, reading level, cognitive abilities, or culture. Future research could explore adaptive models that dynamically adjust text complexity based on the user's level of comprehension or reading skills. Additionally, research could consider personalised AI for health literacy growth, enabling CAs to dynamically adapt their responses based on a patient's evolving comprehension of medical concepts. This will require the development of adaptive NLP models that assess a user's level of comprehension in real time and adjust explanations accordingly - offering simpler definitions for beginners, while gradually introducing medical terminology for more advanced users. Related to this, it could be studied how to realize a closed-loop-communication that ensures and verifies patient's comprehension by methods such as teach-back (Kreps, 2018).

Cultural adaptations can take multiple dimensions. The provided content can be tailored to align with the user's cultural context (e.g., dietary suggestions should respect cultural norms). Additionally, research could explore how the tone and structure of interactions adapt to different cultures. However, careful design is essential to avoid reinforcing stereotypes or generating biased responses.

Explainability. When transitioning from rule-based CAs in healthcare to LLM-based CAs, ensuring patient safety is crucial to maintaining control over the information provided. Research on explainability in conversational AI is essential for enhancing transparency and trust. AI-generated explanations of medical information should be interpretable, contextually relevant, and aligned with user expectations, while also ensuring they do not pose any risk to patients. Some approaches are already available for explainable CA such as the one presented by Nguyen et al. (Nguyen et al., 2023) or Garofalo et al. (Garofalo et al., 2023) .

Bias and fairness-awareness. Existing NLP models still have problems regarding fairness and come along with bias (Hovy and Prabhumoye, 2021). Therefore, research is necessary regarding bias detection and mitigation frameworks for gender, racial, disability, and socio-economic biases to be integrated in POLP within healthcare CA. Such advances in diverse and bias-aware dataset curation, along with fairness-driven fine-tuning of medical NLP models, are essential to mitigate model biases in POLP. Models capturing the peculiarities of specific user groups could help in handling local languages and developing specifically focussing solutions. This would require developing low-resource NLP models for underserved communities, integrating local dialects and indigenous languages.

5 Conclusions

This paper explored how patient-facing language processing in healthcare CAs should be designed to achieve inclusivity. Ensuring that healthcare CAs are accessible, equitable, usable, and useful to all individuals - regardless of their social or socio-economic background, cultural identity, health literacy, digital literacy, or cognitive abilities - is critical to their effectiveness as digital health interventions. These topics gain in relevance when moving from rule-based CAs to LLM-based systems as they allow for more flexibility.

We identified several key research directions for future work, including optimising conversation structure, improving the adaptability of CA language and content, integrating cultural awareness, improving explainability, managing cognitive load, and addressing bias and fairness concerns. These aspects are particularly important in healthcare settings, where CAs are used by a diverse patient population and must effectively support users with different needs.

Inclusive POLP is essential to prevent the unintentional exclusion of certain user groups, which could exacerbate existing health disparities and inequalities in healthcare. By prioritising inclusivity in the design of healthcare CAs, research can contribute to a more equitable and patient-centred digital health landscape.

6 Limitations

This work comes along with some limitations. While the facets of inclusive AI have been collected in a literature review, the impact of inclusive AI on POLP within healthcare CA was only reflected based on the experiences in CA development of the author. In future work, this should be verified by input from other experts in the field.

References

Lene Lauge Berring, Tine Holm, Jens Peter Hansen, Christian Lie Delcomyn, Rikke Søndergaard, and Jacob Hvidhjelm. 2024. Implementing trauma-informed care—settings, definitions, interventions, measures, and implementation across settings: a scoping review. In *Healthcare*, volume 12, page 908. MDPI.

Anthony Junior Bokolo, Kerstin Denecke, and Elia Gabarron. 2025. A literature review on inclusive ai in healthcare- a user-centered approach to potential benefits, challenges, and recommendations.

James L Cross, Michael A Choma, and John A Onofrey. 2024. Bias in medical ai: Implications for clinical decision-making. *PLOS Digital Health*, 3(11):e0000651.

Godson Kofi Davies, Martin Luther King Davies, Esther Adewusi, Kenechukwu Moneke, Olwaseun Adeleke, Lateefat Abiodun Mosaku, Abdulbasit Oboh, Damilola Sherifat Shaba, Isa Aisha Katsina, Joshua Egbedimame, et al. 2024. Ai-enhanced culturally sensitive public health messaging: A scoping review. *E-Health Telecommunication Systems and Networks*, 13(4):45–66.

Kerstin Denecke. 2023. How to design successful conversations in conversational agents in healthcare? In *International Conference on Human-Computer Interaction*, pages 39–45. Springer.

Kerstin Denecke, Daniel Reichenpfader, Dominic Willi, Karin Kennel, Harald Bonel, Knud Nairz, Nikola Cihoric, Damien Papaux, and Hendrik von Tengg-Kobligk. 2024. Person-based design and evaluation of mia, a digital medical interview assistant for radiology. *Frontiers in Artificial Intelligence*, 7:1431156.

Patrick J Fitzpatrick. 2023. Improving health literacy using the power of digital communications to achieve better health outcomes for patients and practitioners. *Frontiers in Digital Health*, 5:1264780.

Marco Garofalo, Alessia Fantini, Roberto Pellugrini, Giovanni Pilato, Massimo Villari, and Fosca Giannotti. 2023. Conversational xai: Formalizing its basic design principles. In *Joint European Conference* *on Machine Learning and Knowledge Discovery in Databases*, pages 295–309. Springer.

Sarah Gibney, Lucy Bruton, Catherine Ryan, Gerardine Doyle, and Gillian Rowlands. 2020. Increasing health literacy may reduce health inequalities: evidence from a national population survey in ireland. *International journal of environmental research and public health*, 17(16):5891.

Silje Havrevold Henni, Sigurd Maurud, Kristin Skeide Fuglerud, and Anne Moen. 2022. The experiences, needs and barriers of people with impairments related to usability and accessibility of digital health solutions, levels of involvement in the design process and strategies for participatory and universal design: a scoping review. *BMC public health*, 22(1):35.

Dirk Hovy and Shrimai Prabhumoye. 2021. Five sources of bias in natural language processing. *Language and linguistics compass*, 15(8):e12432.

Gary L Kreps. 2018. Promoting patient comprehension of relevant health information. *Israel Journal of Health Policy Research*, 7(1):56.

Tom Nadarzynski, Nicky Knights, Deborah Husbands, Cynthia A Graham, Carrie D Llewellyn, Tom Buchanan, Ian Montgomery, and Damien Ridge. 2024. Achieving health equity through conversational ai: A roadmap for design and implementation of inclusive chatbots in healthcare. *PLOS Digital Health*, 3(5):e0000492.

Van Bach Nguyen, Jörg Schlötterer, and Christin Seifert. 2023. From black boxes to conversations: Incorporating xai in a conversational agent. In *World Conference on Explainable Artificial Intelligence*, pages 71–96. Springer.

Vimla L Patel and Andre W Kushniruk. 1998. Interface design for health care environments: the role of cognitive science. In *Proceedings of the AMIA Symposium*, page 29.

Casey Sackett, Devin Harper, and Aaron Pavez. 2024. Do we dare use generative ai for mental health? *IEEE Spectrum*, 61(6):42–47.

Abeed Sarker, Mohammed Ali Al-Garadi, Yuan-Chi Yang, Jinho Choi, Arshed A Quyyumi, Greg S Martin, et al. 2021. Defining patient-oriented natural language processing: a new paradigm for research and development to facilitate adoption and use by medical experts. *JMIR Medical Informatics*, 9(9):e18471.

Lara Whitehead, Jason Talevski, Farhad Fatehi, and Alison Beauchamp. 2023. Barriers to and facilitators of digital health among culturally and linguistically diverse populations: qualitative systematic review. *Journal of medical Internet research*, 25:e42719.

WHO World Health Organization. 2024. Who releases ai ethics and governance guidance for large multimodal models.

Matthew K Wynia and Chandra Y Osborn. 2010. Health literacy and communication quality in health care organizations. *Journal of health communication*, 15(S2):102–115.

Neşe Zayim, Hasibe Yıldız, and Yılmaz Kemal Yüce. 2023. Estimating cognitive load in a mobile personal health record application: A cognitive task analysis approach. *Healthcare Informatics Research*, 29(4):367–376.

Didar Zowghi and Muneera Bano. 2024. Ai for all: Diversity and inclusion in ai. *AI and Ethics*, 4(4):873–876.