




Bridging the patient gap: exploring generative AI to support meaningful patient involvement in health technology assessment

Journal of **Comparative Effectiveness Research**

Manuel Cossio¹  & Ramiro E Gilardino² 

¹ University of Barcelona, 08007 Barcelona, Spain

² Independent Health Executive, RG Consulting, Zurich, Zurich, 8600, Switzerland

Patient and public involvement in health technology assessment (HTA) has progressed from best practice to policy requirement, yet communication barriers persist. This perspective explores how plain language summaries (PLSs) and summaries of information for patients (SIPs) can enhance equity and transparency in HTA. Building on recent European regulatory developments and emerging research, it discusses the balance between accessibility, quality and feasibility. Generative artificial intelligence offers the potential to scale PLS and SIP production, but its responsible integration requires oversight, collaboration and a continued focus on equity and patient-centeredness within evolving HTA frameworks.

First draft submitted: 19 September 2025; Accepted for publication: 23 October 2025; Published online: 4 November 2025

Keywords: generative AI • health technology assessment • patient engagement • plain language summaries

The role of patients and the public in health technology assessment (HTA) has evolved over the past two decades from being considered best practice to becoming a policy requirement. In Europe and beyond, HTA agencies, regulators and policymakers increasingly recognize that patients contribute irreplaceable insights into the lived experience of disease, the outcomes that truly matter and the trade-offs that are acceptable in everyday life. Yet, a major obstacle persists: the language of HTA itself. Technical dossiers and assessment reports are often written in dense medical and methodological jargon, creating an effective barrier for nonspecialists.

This communication gap undermines meaningful participation. It risks reducing patient and public involvement to a procedural exercise rather than a genuine dialogue. It also reinforces inequities, since only the best-resourced patient groups are able to commission experts to interpret these complex documents.

These challenges contrast sharply with progress in global regulatory frameworks, which are moving steadily toward greater transparency and accessibility. For example, the European Union (EU) Clinical Trials Regulation requires sponsors to provide lay summaries of clinical trial results [1]. This sets an important precedent: plain language communication is no longer optional in certain domains, but mandated. Likewise, the new EU Regulation on HTA (EU-HTAR) calls for patient involvement throughout the Joint Clinical Assessment process [2]. If plain language is mandatory in clinical trials, then extending this principle to HTA, through summaries of information for patients (SIP) or plain language summaries (PLS), is both logical and necessary for equitable and transparent decision-making. Throughout this article, we use the terms PLS and SIP (and their plurals PLSs and SIPs) interchangeably to mention these documents.

This article reviews recent initiatives and considers whether generative artificial intelligence (AI) could help make the widespread adoption of such summaries feasible in practice.

The ongoing challenge of patient-accessible communication

Evidence consistently shows that without accessible communication, patient understanding is limited, which undermines genuine engagement and informed decision-making [3]. Simply being invited to the discussion table

is insufficient if the information presented there cannot be understood. For example, clinical trial registries still lack clear mechanisms to identify studies co-designed with patients, a shortcoming that constrains transparency. The same issue is evident in informed consent processes, where complex language continues to dominate. These examples highlight a structural problem: meaningful participation requires not only an invitation but also the tools that allow all stakeholders to fully engage with the conversation.

In HTA, highly technical language produces the same exclusionary effect. PLS were introduced precisely to address this gap. The Scottish Medicines Consortium pioneered their systematic use, requiring manufacturers to provide lay summaries alongside technical dossiers [4,5]. These templates were designed to ensure that patient organizations could review the same evidence base as committee members.

Building on this approach, the HTAi Patient and Citizen Involvement Interest Group developed an international template for SIPs [6]. The objective was to make submissions more consistent across countries and to lower barriers for patient organizations with different levels of HTA expertise. Pilots in England and Australia confirmed both the benefits and the challenges of this approach. Patient groups reported that SIP reduced preparation time and improved confidence, with 88% finding the format helpful. HTA bodies considered them a bridge between complex data and meaningful lay contributions. At the same time, issues such as maintaining neutrality, managing resource constraints and adapting to local contexts emerged as significant lessons [7].

SIPs are therefore more than a communication tool; they represent a structural innovation. By standardizing information in a patient-friendly format, they enable earlier, more confident and more informed engagement. Yet, despite progress, implementation remains uneven. A recent scoping review showed that reporting of patient and public involvement in Technology Appraisal and Assessment Reports is highly variable, making it difficult to trace how patient perspectives influence outcomes [8]. Standardized tools such as the international SIP template offer HTA bodies a mechanism to capture and report patient input more consistently.

Generative AI in HTA: opportunities for patient insights

Producing high-quality SIP and PLS remains a resource-intensive task. Drafting, reviewing and adapting these documents requires substantial time, expertise and coordination, placing significant strain on both HTA agencies and patient organizations. Even when patient groups are well resourced, the manual process can create bottlenecks, delay meaningful involvement and reduce opportunities for broader engagement.

Generative AI offers a potential, though complex, solution. Large language models (LLMs) have demonstrated the ability to convert highly technical content into accessible, plain-language explanations. In oncology, for example, AI has been used successfully to transform dense trial reports into patient-friendly summaries, supporting comprehension and informed decision-making [9]. This capability not only reduces the workload for human authors but also has the potential to standardize the clarity and quality of content across multiple dossiers.

The potential for efficiency and scale is particularly relevant in the context of the EU-HTAR. With potentially dozens of Joint Clinical Assessment reports expected in the next year, HTA bodies will face increasing pressure to incorporate patient perspectives without overburdening either manufacturers or patient organizations. While SIPs are not yet formally required, they could serve as a practical bridge between highly technical dossiers and meaningful patient input. Relying solely on manual preparation for every assessment is unlikely to be feasible; AI could provide a pathway to generate initial drafts rapidly, which could then be refined by experts.

Beyond efficiency, AI also opens the door to personalization. SIPs and PLSs could be tailored to specific patient communities, literacy levels, and languages, supporting inclusivity and broader participation. Smaller or under-resourced organizations could thus gain access to the same quality of information as larger groups, mitigating inequities and expanding the reach of patient engagement initiatives. Additionally, AI may allow for faster updates when new evidence emerges, keeping summaries current and aligned with evolving clinical understanding.

A recent proof-of-concept study being presented at ISPOR Europe 2025 (PT21 – “Can generative AI deliver patient-friendly summaries? A case study using NICE guidance for spinal muscular atrophy”) provides early evidence on AI’s performance in a real-world HTA context. The study evaluated whether an LLM could generate a SIP from a full guidance document published by the National Institute for Health and Care Excellence on onasemnogene abeparvovec[®], a treatment for spinal muscular atrophy [10]. These documents exemplify the methodological and clinical complexity that frequently limits patient comprehension, making them an ideal test case.

Using a carefully structured prompt based on the Scottish Medicines Consortium template, the AI produced a multi-page, PLS within seconds ([Supplementary Material](#)). The output was then evaluated against critical criteria, including readability, accuracy, relevance, structure and overall patient-friendliness.

Findings were mixed. The AI-generated summary was well organized and largely free of jargon. Reviewers noted clear explanations of the treatment, strong readability and logical structure, confirming that AI can produce accessible content from complex sources. However, critical limitations were also identified. The AI struggled to clearly explain comparators and did not adequately address the concept of ‘unmet need’, a central component of HTA deliberations. These gaps are significant because they represent essential information that patients require to provide meaningful input on a treatment’s value. The study concluded that while AI can generate a rapid first draft, expert refinement, validation and human oversight remain essential to ensure quality and contextual nuance.

This case study is pivotal because it shifts the conversation from theoretical potential to practical application. It demonstrates that AI can accelerate the most labor-intensive steps of summary creation, but current outputs are not yet fit for autonomous use. In its present form, the technology appears best suited as a supportive tool for HTA bodies or manufacturers rather than as a standalone solution. If implemented thoughtfully, AI could become a powerful enabler of equitable, timely and patient-centered involvement in HTA decision-making.

Regulatory environment & implementation practices

Deploying AI to generate PLSs and SIPs in HTA is fraught with risk. Inaccurate or biased content could compromise not only the quality of assessments but also public trust in the process itself. These risks are not hypothetical. In the EU, they are explicitly addressed by a growing regulatory framework, which makes clear that ‘human-in-the-loop’ oversight is not optional but a legal and ethical necessity to ensure accountability and reduce bias. Within the EU, three complementary instruments shape how AI may be applied in HTA:

The AI Act: As the first comprehensive AI law, it classifies systems by risk. Any tool producing PLS or SIP would likely fall into the ‘high-risk’ category, triggering strict requirements for risk management, data quality and mandatory oversight [11]. Transparency obligations further require developers to disclose when content is AI-generated and summarize training data. Yet unresolved issues remain: much of the underlying evidence, such as trial datasets and manufacturer dossiers, constitutes proprietary intellectual property, creating conflicts between public accountability and commercial confidentiality.

The General Data Protection Regulation: processing health data must comply with General Data Protection Regulation principles, requiring a clear legal basis such as explicit consent or a defined public-interest task. Non-negotiable safeguards like data minimization and pseudonymization remain central to protecting patient privacy [12].

The European Health Data Space: this emerging framework provides mechanisms for the secure reuse of health data for secondary purposes, including policymaking [13]. For HTA, it represents both an opportunity for more robust evidence generation and a set of procedural requirements that will directly shape AI deployment.

Together, these instruments establish strict boundaries: AI can support the scaling of plain-language summaries, but only within systems that guarantee oversight, transparency and data protection.

Safeguards in practice

Standardization is a critical safeguard. Templates such as the international SIP provide structural anchors to guide AI outputs, ensuring that summaries remain complete, relevant and balanced. Readability standards, explicit neutrality and the use of clear visual aids should be embedded from the outset. Equally important is human oversight: validation by patient representatives, HTA staff or independent reviewers is indispensable to guarantee accuracy and contextual nuance. While trust in automated outputs may grow over time, early implementation must be tightly controlled.

Yet regulatory compliance alone is insufficient. Broader governance and equity challenges must also be addressed:

Accountability and shared oversight: regulatory agencies should define ethical baselines, HTA agencies adapt them to local contexts and patient groups act as validators of readability, accuracy and neutrality. Such co-governance models anchor accountability in practice, ensuring that outputs remain both technically reliable and contextually meaningful [14].

Equity concerns: AI systems reflect the data they are trained on. If SIPs rely on materials from high-income or well-resourced groups, less-represented patient communities risk exclusion. Deliberate safeguards, such as diverse training datasets, equity requirements in governance and low-cost accessible tools, are essential to prevent GenAI from widening participation gaps in HTA [15,16]. **The digital language divide:** LLMs are disproportionately trained on high-resource languages like English, which weakens performance in underrepresented languages. Without deliberate inclusion of diverse linguistic data and robust multilingual frameworks, many populations may remain excluded from patient-facing summaries [17].

Transparency trade-offs: The key confidentiality challenge lies not in algorithm design but in using manufacturer dossiers and patient-level clinical data. A tiered transparency model, disclosing metadata about training sources (e.g., dossier type, trial demographics and time frames) and performance benchmarks, without releasing raw data, could balance auditability with protection of commercial and patient interests.

Tools such as the international SIPs and PLSs are emerging across HTA systems, but their format and quality vary. Harmonizing core principles internationally, such as how benefits, risks and uncertainties are presented in plain language, would improve comparability while still allowing cultural and linguistic adaptation [7]. Long-term implementation will depend on structured collaboration. Public–private partnerships could combine the policy expertise of HTA agencies with the technical infrastructure of industry. Co-developed or open-source tools, subject to joint oversight, would ensure AI systems remain technically sound, ethically grounded and accessible to patient groups with varying capacity.

Toward a responsible & equitable future

PLSs and SIPs are now widely recognized as fundamental to equitable patient involvement [6,7]. Pioneering initiatives have demonstrated their value, but they have also exposed the significant resource burden they impose [3–7]. Generative AI offers a compelling solution to this problem of scale. When paired with standardized templates and rigorous oversight, it could accelerate the production of summaries and lower barriers to entry. Yet this promise introduces an equity paradox: while AI could democratize access to information, it also risks widening the gap between organizations with the digital literacy and resources to manage advanced tools and those without. Addressing this tension, potentially through HTA body-led initiatives or public–private partnerships, will be essential.

The coming years, shaped by the implementation of the EU-HTAR and international collaborations by many HTA agencies outside EU, present a critical opportunity to align regulatory mandates, patient advocacy and technological innovation. The central question is no longer whether we should use PLSs, but how we can deploy them at scale in ways that are high-quality, ethical and legally sound. Generative AI may serve as a catalyst, but it is not a panacea. Its contribution will depend entirely on our commitment to keeping human experience, oversight and equity at the very center of the process.

Summary points

- Patient involvement in health technology assessment (HTA) is evolving from best practice to a policy requirement, underscoring the need for transparent and accessible communication.
- Plain language summaries (PLSs) and summaries of information for patients (SIPs) enable equitable understanding of complex evidence.
- Producing high-quality PLSs/SIPs remains resource-intensive and requires technical and linguistic expertise.
- Generative AI can accelerate summary production but introduces risks of bias, inaccuracy, and data misuse.
- Ethical frameworks and ‘human-in-the-loop’ oversight are crucial for safeguarding trust and accountability.
- International harmonization of PLS/SIP standards would strengthen comparability and patient engagement.
- Addressing linguistic and digital divides is critical to prevent inequities in AI-enabled communication.
- Collaborative, publicly guided models and current developments like the EU Regulation on HTA can ensure sustainable, responsible deployment of AI in HTA communication.

Supplementary data

To view the supplementary data that accompany this paper please visit the journal website at: <https://becarispublishing.com/doi/epdf/10.57264/cer-2025-0150>

Financial disclosure

The authors received no financial and/or material support for this research or the creation of this work.

Competing interests disclosure

M Cossio is an employee of Cytel, Inc. The authors have no other competing interests or relevant affiliations with any organization or entity with the subject matter or materials discussed in the manuscript apart from those disclosed.

Writing disclosure

No funded writing assistance was utilized in the production of this manuscript.

Open access

This work is licensed under the Attribution-NonCommercial-NoDerivatives 4.0 Unported License. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>

References

1. European Parliament and Council of the European Union. Regulation (EU) No 536/2014 on clinical trials on medicinal products for human use, and repealing Directive 2001/20/EC (2014).
2. European Parliament and Council of the European Union. Regulation (EU) 2021/2282 of the European Parliament and of the Council of 15 December 2021 on health technology assessment and amending Directive 2011/24/EU. *Off. J. Eur. Union L.* 458, 1–32 (2021).
3. Barnes A, Patrick S. Lay summaries of clinical study results: an overview. *Pharmaceut. Med.* 33(4), 261–268 (2019).
4. Scottish Medicines Consortium. Guidance to manufacturers for completion of summary information for submitting patient groups form (Version 4). (2017). Available from: <https://scottishmedicines.org.uk/media/2771/guidance-on-summary-information-for-patient-groups.pdf>
5. Dickson J, Taylor L, McCormack J, Downes N, Lockhart L, Walker L. OP79 improving public understanding of Scottish Medicines Consortium advice. *Int. J. Tech. Assess. Health Care* 35(Suppl. 1), 20 (2019).
6. Cook N, Livingstone H, Dickson J *et al.* Development of an international template to support patient submissions in health technology assessments. *Int. J. Tech. Assess. Health Care* 37(1), e50 (2021).
7. Coombes M, Morgan K, Diaz J *et al.* Plain language summaries supporting patient involvement: lessons and guidance from HTAi patient and citizen involvement interest group (PCIG). *Int. J. Tech. Assess. Health Care* 41(1), e59 (2025).
8. Johnson EE, Uteh CO, Belilios E, Pearson F. Reporting of patient and public involvement in technology appraisal and assessment reports: a rapid scoping review. *Patient* 18(2), 109–114 (2025).
9. Waters M. AI meets informed consent: a new era for clinical trial communication. *JNCI Cancer Spectr.* 9(2), pkaf028 (2025).
10. National Institute for Health and Care Excellence (NICE). *Guidance on Onasemnogene Apeparvovec*. Available from: <https://www.nice.org.uk/guidance/hst15>
11. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) (2024).
12. European Parliament and Council of the European Union. Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation) (2016).
13. European Parliament and Council of the European Union. Regulation (EU) 2025/327 on the European Health Data Space (2025).
14. Carapinha JL, Botes D, Carapinha R. Balancing innovation and ethics in AI governance for health technology assessment. *J. Med. Econ.* 27(1), 754–757 (2024).
15. Zemplényi A, Tachkov K, Balkanyi L *et al.* Recommendations to overcome barriers to the use of artificial intelligence-driven evidence in health technology assessment. *Front. Public Health* 11, 1088121 (2023).
16. Fleurence RL, Bian J, Wang X *et al.* Generative artificial intelligence for health technology assessment: opportunities, challenges, and policy considerations—an ISPOR Working Group report. *Value Health* 28(2), 175–183 (2025).
17. Hendy A, Abdelrehim M, Sharaf A *et al.* How good are GPT models at machine translation? A comprehensive evaluation. arXiv preprint arXiv:2302.09210 (2023). Available at: <https://arxiv.org/abs/2302.09210>