Check for updates

Patient Artificial Intelligence Literacy as a Component of Responsible Artificial Intelligence Implementation

Isabel Preeshagul, MD[1] and Eleonora Teplinsky, MD, FASCO[2]

DOI https://doi.org/10.1200/CCI-26-00206

ACCOMPANYING CONTENT

Article, 10.1200/CCI-26-00114

Accepted July 8, 2026
Published July 31, 2026

JCO Clin Cancer Inform
10:e2600206
© 2026 by American Society of Clinical Oncology

In 2026, medical education is increasingly happening outside the examination room as a growing number of patients turn to artificial intelligence (AI) for cancer information. This shift is occurring alongside the rapid integration of AI into oncology, with applications spanning screening, diagnosis, pathology, treatment planning, clinical documentation, and patient communication.[1] Simultaneously, patients are now independently engaging with large language models (LLM) and other AI platforms to better understand their cancer diagnosis, treatment options, and prognosis.[2] Although considerable attention has been devoted to developing frameworks for the responsible implementation of AI within health care systems, substantially less focus has been paid to how patients understand, interpret, and use these technologies.

Subramanian et al[3] evaluate the educational resources individuals are likely to find when searching online for information about AI in cancer care. The authors examined indexed web content rather than chatbot-generated responses to characterize the information available to patients seeking to understand the role of AI in cancer care.

In a cross-sectional analysis using Google Trends-derived search terms, the authors screened 170 Google webpages and 150 YouTube videos, totaling 320 combined, of which only 82 (25%) were considered both patient-facing and relevant. Of the available resources, many were considered low quality based on DISCERN, a tool for evaluating quality of written health information,[4] with only 33% of webpages and 23% of videos meeting the threshold for high quality. Resources were written above recommended reading levels and inconsistent in their discussion of key AI safety concepts. The median reading level, assessed across three readability indices, corresponded to a college level, which is well above the sixth- to eighth-grade reading levels recommended by the American Medical Association and National Institutes of Health.

Although most webpages discussed clinician oversight and transparency, only about half addressed bias and equity and just 15% discussed hallucinations or the risk of misinformation, arguably among the most significant limitations and risks associated with generative AI. Because AI-generated misinformation is personalized, conversational, and often delivered with a high degree of confidence, patients who are unaware that LLMs can fabricate information may interpret these responses as authoritative, making inaccuracies more difficult to recognize.

The strengths of this study include dual independent review, the use of validated instruments, reflection of real-world patient search behavior, and evaluation of both text- and video-based resources. However, important limitations should also be recognized. DISCERN was not designed to evaluate AI-specific educational content. In addition, the authors limited their analysis to Google and YouTube, excluding social media platforms such as Instagram, Facebook, and TikTok, where many patients increasingly seek health information. The YouTube videos analyzed also had relatively few views, limiting generalizability. Finally, the findings do not address what patients are told when they ask chatbots directly about AI and cancer care.

This study naturally raises questions about patient outcomes. Does exposure to this content help patients understand key AI safety concepts? After viewing this information, are patients able to distinguish hallucinated or inaccurate cancer information from reliable information and does this education about AI and cancer care influence patient decision making, behavior, trust, or clinical outcomes?

Downloaded from ascopubs.org by University of Toronto on October 5, 2026 from 099.245.175.177
Copyright © 2026 American Society of Clinical Oncology. All rights reserved.

The implications of this study extend well beyond the quality of internet resources and raise broader questions about how patients navigate an increasingly AI-mediated information ecosystem. AI has transformed the health information supply chain. Patients no longer interact solely with information created by identifiable authors or institutions but increasingly with synthesized responses that aggregate and reinterpret multiple sources, often making it more difficult to interpret and judge their reliability. In addition, in a repeat search, the authors identified substantial turnover in online content over just 9 months, with 9% of webpages and < 1% of YouTube videos appearing in both search periods. This finding highlights the dynamic nature of the digital information environment and underscores the challenge of maintaining accurate, up-to-date patient education resources.

Current discussions surrounding responsible AI implementation in oncology have largely focused on clinician oversight, transparency, bias mitigation, and governance. However, these frameworks are primarily directed toward health care professionals and institutions, whereas comparable patient-centered guidance remains largely absent.

This leads to the question of responsibility. Currently, accountability for patient AI education remains poorly defined and without clear ownership. In this environment, educational efforts are likely to remain fragmented. Addressing these gaps will require coordinated efforts from multiple stakeholders. Professional societies and patient advocacy organizations can develop evidence-based educational resources and patient-facing guidelines, whereas AI developers should prioritize transparency and clearly communicate the limitations of their tools. Health systems and clinicians also play an important role. Embedding AI education into oncology care could facilitate conversations about the appropriate use and limitations of AI and may eventually become part of standard patient education to help patients engage with these technologies more safely and critically.

This study raises several important questions regarding AI health literacy, AI-specific quality standards, management of AI-clinician conflicts, and patient-facing AI guidelines. Traditional health literacy emphasizes the ability to obtain, understand, and apply health information, whereas digital health literacy focuses on navigating online resources. The widespread availability of generative AI suggests the emergence of other essential competencies, including the ability to critically evaluate AI-generated information, recognize uncertainty, understand the limitations of these systems, and appropriately integrate AI-generated content into discussions with health care professionals. Patients must understand that AI responses are probabilistic rather than authoritative, that recommendations may vary depending on how questions are asked, and that outputs may reflect outdated information, embedded biases, or fabricated citations presented with confidence. Equally important is knowing when AI can be a useful educational tool and when clinical questions require individualized discussion with a health care professional. As patients increasingly bring AI-generated recommendations into the clinic, guidance on navigating situations in which AI advice conflicts with clinician recommendations will also become increasingly important to support shared and informed decision making while maintaining trust.

The broader questions are even more challenging: Who is responsible when AI-generated misinformation causes harm? How do we improve AI literacy without widening equity gaps? Where should trustworthy AI education reside and ultimately, is more internet content the answer? Expanding the availability of high-quality online resources is necessary but unlikely to be sufficient, particularly when existing materials exceed recommended readability standards and the digital landscape evolves rapidly. Simply creating more content risks perpetuating disparities if that content remains inaccessible to many patients.

The study also highlights important priorities for future research. Existing instruments such as DISCERN provide a valuable framework for evaluating consumer health information but were not designed to assess AI-specific educational content. Development of AI-specific quality metrics may be warranted, incorporating domains such as hallucination disclosure, uncertainty communication, data privacy, and management of discordance between AI-generated recommendations and clinician advice. Ultimately, future studies should move beyond content evaluation and examine whether AI literacy influences patient outcomes. Does AI literacy influence treatment decisions, adherence, health care utilization, trust in clinicians, or patient-clinician communication? Does improving AI literacy lead to safer and more effective use of these technologies?

Subramanian and colleagues have identified important deficiencies in the current patient-facing AI information landscape, but their work also highlights a broader challenge. Responsible AI implementation must extend beyond algorithms and governance structures to include the patients who increasingly rely on these technologies. Developing evidence-based, patient-centered AI literacy frameworks should be considered an essential component of responsible AI implementation. The question is no longer whether patients will use AI to inform their cancer care. They already are. The challenge now is whether health care systems, professional societies, patient advocacy organizations, technology developers, and policymakers can work together to ensure that patients have the knowledge and skills necessary to use these tools safely, critically, and effectively.

Downloaded from ascopubs.org by University of Toronto on October 5, 2026 from 099.245.175.177
Copyright © 2026 American Society of Clinical Oncology. All rights reserved.

AFFILIATIONS

[1]Memorial Sloan Kettering Cancer Center, New York, NY
[2]Valley Health System, Paramus, NJ

CORRESPONDING AUTHOR

Eleonora Teplinsky, MD, FASCO; Social Media Handle: @drteplinsky; e-mail: teplel@valleyhealth.com.

AUTHOR CONTRIBUTIONS

Conception and design: All authors
Collection and assembly of data: All authors
Data analysis and interpretation: All authors
Manuscript writing: All authors
Final approval of manuscript: All authors
Accountable for all aspects of the work: All authors

AUTHORS' DISCLOSURES OF POTENTIAL CONFLICTS OF INTEREST

The following represents disclosure information provided by authors of this manuscript. All relationships are considered compensated unless otherwise noted. Relationships are self-held unless noted. I = Immediate Family Member, Inst = My Institution. Relationships may not relate to the subject matter of this manuscript. For more information about ASCO's conflict of interest policy, please refer to www.asco.org/rwc or ascopubs.org/go/authors/author-center.

Open Payments is a public database containing information reported by companies about payments made to US-licensed physicians (Open Payments).

Isabel Preeshagul
Consulting or Advisory Role: Takeda, Medscape, Regeneron, Boehringer Ingelheim, OncLive/MJH Life Sciences, AstraZeneca, OncLive/MJH Life Sciences, Plexus, IDEOlogy Health, Jazz Pharmaceuticals, AbbVie, OncLive/MJH Life Sciences, Nuvalent, Inc, Decera Clinical Education, Pfizer, EMD Serono/Merck

Eleonora Teplinsky
Honoraria: Sermo, MJH Life Sciences, Total Health Conferencing
Consulting or Advisory Role: Novartis, Daiichi Sankyo/Astra Zeneca, Pfizer, Immunogen, AbbVie
Travel, Accommodations, Expenses: OncLive/MJH Life Sciences, Novartis, Total Health Conferencing, Sermo

No potential conflicts of interest were reported.

REFERENCES

1. Kolla L, Parikh RB: Uses and limitations of artificial intelligence for oncology. Cancer 130:2101-2107, 2024
2. Yun HS, Bickmore T: Online health information-seeking in the era of large language models: Cross-sectional web-based survey study. J Med Internet Res 27:e68560, 2025
3. Subramanian P, Shyamsunder S, Eshaghi K, et al: Gaps in patient-facing information on artificial intelligence in cancer care: A cross-sectional analysis. JCO Clin Cancer Inform 10.1200/CCI-26-00114
4. Charnock D, Shepperd S, Needham G, et al: DISCERN: An instrument for judging the quality of written consumer health information on treatment choices. J Epidemiol Community Health 53: 105-111, 1999

Downloaded from ascopubs.org by University of Toronto on October 5, 2026 from 099.245.175.177
Copyright © 2026 American Society of Clinical Oncology. All rights reserved.