International Journal of Education, Culture, and Society

e-ISSN : 3024-8973
p-ISSN : 3024-8981

Index: Harvard, Boston, Sydney University, Dimensions, Lens, Scilit, Semantic, Google, etc

https://doi.org/10.58578/IJECS.v4i3.12372

Culturally Responsive AI-Supported Cancer Pain Education: A Human-Centered Framework for Multilingual Oncology Care

Yahia El Karfa & Aicha Ziani
Ibn Tofail University, Morocco
yahia.elkarfa@uit.ac.ma; yahia.elkarfa@uit.ac.ma

Submitted:	Revised:	Accepted:	Published:
Sep 15, 2026	Oct 14, 2026	Oct 25, 2026	Oct 30, 2026

Abstract

Cancer pain education aims to help patients understand pain, communicate symptoms, and follow clinician-approved management plans; however, improved knowledge does not consistently translate into reduced pain intensity. Generative artificial intelligence (AI) is increasingly proposed for patient-facing education, while evidence supporting culturally and linguistically responsive oncology applications remains limited. This critical perspective develops a human-centered framework for AI-supported cancer pain education in multilingual oncology care. A focused, non-systematic critical synthesis integrated research on cancer pain education, health literacy, cultural adaptation, multilingual communication, and AI governance. No primary data were collected. The framework proposes six interdependent design domains: patient-led co-design and cultural humility; multilingual, accessible, dialogic learning; accurate pain science that avoids reductionist explanations; clearly bounded, evidence-grounded AI functions; explicit human oversight and escalation procedures; and lifecycle governance with equity-focused evaluation. These domains constitute conceptual recommendations requiring empirical evaluation. A staged evaluation pathway progresses from co-design and cognitive testing to offline AI benchmarking, supervised usability pilots, and comparative implementation studies. Initial evaluations should prioritize comprehension, communication, trust, safety, and access across language groups. Potential effects on pain interference and other clinical outcomes remain hypotheses for prospective testing. Illustrative examples involving

IJECS Journal is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License

Moroccan Darija highlight linguistic and cultural adaptation challenges without establishing cancer-specific validation. The proposed AI role is restricted to educational support and does not encompass diagnosis, prescribing, psychotherapy, professional interpretation, or replacement of oncology professionals. By connecting health education, cultural responsiveness, pain science, and responsible AI governance, the framework provides a basis for developing and evaluating multilingual educational support. Its contribution lies in specifying design and evaluation priorities while retaining clinical effectiveness as an open empirical question.

Keywords: Cancer Pain Education; Generative Artificial Intelligence; Cultural Adaptation; Multilingual Oncology Care; Health Literacy

INTRODUCTION

Cancer-related pain remains a consequential educational and care challenge. A systematic review and meta-analysis covering 444 studies estimated an overall pain prevalence of 44.5% among people with cancer and moderate-to-severe pain in 30.6%; these pooled estimates vary by treatment phase and should not be read as estimates for any single country or cancer population (Snijders et al., 2023). Pain is not merely a signal of tissue damage or a score to be reduced. The revised International Association for the Study of Pain definition describes pain as a personal experience influenced by biological, psychological, and social factors, and distinguishes pain from nociception (Raja et al., 2020). For education, this distinction matters: explaining that attention, emotion, expectations, sleep, social support, and nervous-system processing can influence pain must never imply that cancer pain is imaginary, self-caused, or evidence of psychological weakness. Cancer pain may reflect disease, treatment, procedures, co-occurring conditions, or more than one mechanism; assessment and treatment decisions therefore remain clinician-led (Fallon et al., 2018).

Patient education can support knowledge about pain, self-monitoring, communication with clinicians, and informed use of a care plan. It should not be framed as a substitute for analgesia, oncological treatment, rehabilitation, palliative care, or psychological support. Reviews of therapeutic and pain-neuroscience education in oncology report promising but heterogeneous findings. Improvements in pain-related knowledge and medication adherence appear more consistent than clinically meaningful

reductions in pain intensity or quality of life; apparent effects on pain vary with intervention content, comparison conditions, populations, and outcome measures (Galietta et al., 2025; González-Martín et al., 2023; Hernández-Hernández et al., 2024). A small randomized trial of an online programme combining pain neuroscience education with graded movement in breast-cancer survivors is informative about a particular rehabilitation package, but it does not isolate the effect of education, establish generalizability to other cancers, or provide evidence for AI-mediated teaching (Martínez-Miranda et al., 2024). A scoping review of adult cancer-pain self-management interventions likewise found education to be a common component, while delivery modes and intervention content varied (Sjattar et al., 2024). Digital psychosocial cancer interventions are a related but distinct evidence base: reviews examine engagement and efficacy, including whether interventions are guided, but do not establish the value of AI-mediated pain education (Akdemir et al., 2024; Zhang et al., 2024). The defensible educational aim is therefore to improve understanding and communication first; symptom relief is a plausible but unguaranteed downstream outcome.

Health literacy provides a useful educational foundation. Rather than treating literacy as an individual deficit in reading ability, contemporary models describe people's capacity to access, understand, appraise, and use information in interaction with the demands and resources of health systems (Nutbeam, 2000; Sørensen et al., 2012). A review of reviews in cancer care found that health literacy was often weakly theorized or operationalized and that organizational health literacy received limited attention, reinforcing the need to examine both the learner and the service (Jeitani et al., 2025). Patient education should consequently be dialogic and responsive: it should begin with what the person wants to know, recognize prior experience, use plain language without removing important nuance, and check whether the explanation made sense. Teach-back can support this process when it is framed as a test of the clarity of the explanation—not an examination of the patient. Reviews report benefits for comprehension and related outcomes, but the cancer-specific evidence base remains small and heterogeneous (Choi & Choi, 2021; Talevski et al., 2020). A patient's ability to repeat a phrase should not be mistaken for durable learning, informed choice, or access to appropriate care.

Culture and language shape how people interpret symptoms, communicate pain, assess medicines, involve family, and navigate services. They also interact with institutional practices, clinician assumptions, access to interpreters, and structural constraints. The

evidence does not support a simple claim that cancer pain education has never addressed culture. A mixed-methods systematic review identified 62 studies of patient-centred education for cancer pain; only six specifically focused on minority ethnic populations, and all six were conducted in the United States. That review also proposed an initial logic model for future resources (Dehpour et al., 2026). A separate review of culturally and linguistically diverse cancer-information interventions included 120 papers, but most addressed prevention rather than active treatment; community engagement was identified as important, while the review found little reporting of consumers' and carers' roles in the development of resources (Whitehead et al., 2025). Reviews of culturally incorporated psychosocial and behavioral interventions suggest potential benefits but cover different populations and outcomes and do not establish the effectiveness of AI-supported pain education (Arana-Chicas et al., 2025; Llave et al., 2024; Lui et al., 2023). The gap is therefore one of integration, context, and evaluation—not a total absence of prior cultural work.

Morocco offers a concrete, but not universal, multilingual setting in which to consider these design questions. In a 2026 cultural and linguistic adaptation of the Health Literacy Questionnaire, cognitive interviews with 28 primary-care users from Kenitra and Settat led to linguistic or cultural refinement of 19 of 44 Darija items. The authors describe this as an initial adaptation step, not a completed psychometric validation, and the study was not conducted in an oncology pain-education population (Boumendil et al., 2026). A national survey across five Moroccan pediatric oncology units reported that 145 of 155 children with cancer experienced pain and that most parents wanted further information and communication about pain management; these findings are relevant to family-centered education but cannot be generalized to adults (Khoubila et al., 2020). In another single-center Moroccan study of 100 patients, 38% reported needing more than three consultations to fully understand their cancer diagnosis and 79% felt their concerns were not adequately attended to. This cross-sectional study concerned diagnosis disclosure—not pain education—and is best treated as a local communication signal rather than a population estimate (Raoui et al., 2025). Together, these sources support locally co-designed communication, while leaving cancer-specific Darija pain education unevaluated.

Generative AI may help educators draft plain-language explanations, create question prompts, support teach-back, or make approved information easier to navigate. Reviews across medical education and patient care map a rapidly expanding set of

applications but emphasize the need for task-specific assessment rather than transfer of performance claims across settings (Aydin et al., 2024; Busch et al., 2025). In oncology, the implications for patient–physician communication, trust, and comprehension remain especially important (Thind & Tsao, 2025). However, technical fluency is not evidence of educational or clinical safety. A scoping review of generative AI in cancer patient education found that evaluations concentrated on accuracy, readability, and general quality; more than half of the reported evaluation metrics were custom scales, and patient-centered or culturally responsive frameworks were rarely applied (Leong & Ormsby, 2026). A scoping review of oncology chatbots included 60 studies, 25 of which reported bias or hallucination as a limitation (Chen et al., 2025). At the broader health-literacy interface, a systematic review found only 18 eligible AI studies and limited engagement with intended users (Abeo et al., 2025). Multilingual translation research suggests that human post-editing can improve or stabilize the quality of AI-generated discharge instructions, but findings across six languages do not validate Darija or cancer-specific education (Brewster et al., 2025). A randomized ePAL trial offers adjacent evidence for a digital pain-support package, but its symptom monitoring, AI-based triage, and education were bundled; it does not demonstrate the effectiveness of a generative-AI tutor or culturally adapted teaching (Kamdar et al., 2024). Likewise, OncoPainBot is an emerging clinical decision-support framework for pain assessment and analgesic recommendations, not a patient-education intervention (Liu et al., 2026).

This article addresses the intersection of patient education, culture, pain science, and responsible AI. It does not claim to offer the first logic model for minority-relevant cancer pain education. Its contribution is a design framework that makes multilingual adaptation, dialogic learning, boundaries on AI functions, human escalation, and equity-focused governance explicit in one educational architecture. The objectives are to: (a) synthesize what current evidence supports—and does not support—about cancer pain education, cultural adaptation, and AI; (b) specify a human-centered framework for safe AI-supported education in multilingual oncology; and (c) outline a staged, testable evaluation agenda, using Morocco as an illustrative context rather than a validated implementation site.

METHODS

Approach and scope

This article is a critical perspective based on a focused, non-systematic narrative synthesis. The search and source-checking process was undertaken during preparation of this draft, with a cutoff of October 17, 2026. It was designed to inform conceptual integration, not to estimate pooled effects, exhaustively identify every publication, or establish a definitive evidence base. No review protocol was registered, no formal systematic screening was conducted, no study counts were generated from database searches, and no formal risk-of-bias appraisal was performed. Accordingly, the article is not described as a systematic or scoping review, and no PRISMA flow diagram is reported.

Sources and selection logic

Targeted searches were conducted across PubMed/PubMed Central, Crossref, journal and publisher websites, and official guidance repositories. Search concepts included combinations of "cancer pain education," "pain neuroscience education," "health literacy," "cultural adaptation," "minority ethnic," "multilingual cancer information," "Darija," "Morocco," "generative AI," "large language model," "patient education," "translation," and "AI governance." Searches were supplemented by reference checking and direct verification of publication records, abstracts, and available full texts. Priority was given to recent systematic, scoping, or mixed-methods reviews; relevant randomized trials; Moroccan contextual studies; foundational health-literacy and pain concepts; and official AI-governance guidance. Older sources were retained only when they provided a foundational concept. The selection was relevance-based and interpretive; it should not be treated as a complete search of all databases or non-English literature.

Analytic procedure

Evidence was organized into four linked lenses: (1) pain science and cancer-pain education; (2) health-literacy and dialogic teaching; (3) cultural and linguistic responsiveness; and (4) generative-AI capability, risk, and governance. For each lens, the synthesis distinguished direct evidence from adjacent evidence and from conceptual recommendations. The six framework domains were derived by mapping educational objectives to sociocultural contexts, patient-facing AI functions, human responsibilities, and measurable safety and equity conditions. The evaluation pathway was then specified as a proposed sequence of future research stages, not as a tested intervention protocol. No

new participant data, patient records, experiments, or model-performance tests were collected or analyzed.

RESULTS

Conceptual synthesis and evidence boundaries

The synthesis produced a framework for designing—not yet validating—AI-supported cancer pain education. Table 1 summarizes the evidence streams most relevant to the framework and the limits on transferring their findings to multilingual oncology education. The evidence is strongest for the need to improve pain-related knowledge, communication, and access to information. It is weaker for the proposition that education alone produces consistent reductions in pain intensity, and it does not establish that generative AI improves outcomes for culturally and linguistically diverse oncology populations.

Table 1. Evidence streams and transfer limits relevant to the proposed framework

Evidence stream	What the evidence supports	Transfer limit for this article's question
Cancer-pain education	Reviews indicate that knowledge and some self-management outcomes can improve; pain-intensity effects are variable (Galietta et al., 2025; González-Martín et al., 2023; Hernández-Hernández et al., 2024).	Heterogeneous interventions and populations; no basis for promising analgesia from education or AI.
Minority-relevant pain education	A mixed-methods review found 62 studies, including six focused on minority ethnic populations, all based in the United States, and proposed an initial logic model (Dehpour et al., 2026).	Evidence is geographically concentrated; the review does not evaluate generative AI or Moroccan settings.
Culturally and linguistically diverse cancer information	Reviews emphasize cultural relevance and community engagement, but many studies focus on prevention and consumer roles in resource development are underreported (Llave et al., 2024; Whitehead et al., 2025).	Culture-incorporated psychosocial or information interventions are not equivalent to cancer-pain education.
Moroccan language and communication context	Darija adaptation required changes to 19/44 HLQ items; Moroccan oncology surveys identify information and communication needs (Boumendil et al., 2026; Khoubila et al., 2020; Raoui et al., 2025).	These sources do not validate an adult oncology pain-education tool or an AI system; the Darija HLQ is not yet fully validated.

Evidence stream	What the evidence supports	Transfer limit for this article's question
AI in cancer education and multilingual communication	Cancer-AI education reviews show inconsistent evaluation and limited cultural frameworks; human post-editing can support multilingual translation quality (Brewster et al., 2025; Chen et al., 2025; Leong & Ormsby, 2026).	Translation and chatbot benchmarks do not establish comprehension, trust, safety, or clinical benefit in a target population.
Digital or clinical AI pain tools	Digital symptom support and clinical decision-support systems are being evaluated (Kamdar et al., 2024; Liu et al., 2026).	Monitoring, triage, and analgesic recommendations are not the same as patient education; bundled effects cannot be attributed to education or AI alone.

Proposed framework

The framework treats the AI system as an optional educational aid embedded within a human-led learning relationship. Its central purpose is to help patients and, where they choose, caregivers understand approved information and formulate questions for their care team. It is not a diagnostic, prescribing, or autonomous triage system. Figure 1 displays six interdependent design domains; none is sufficient on its own.

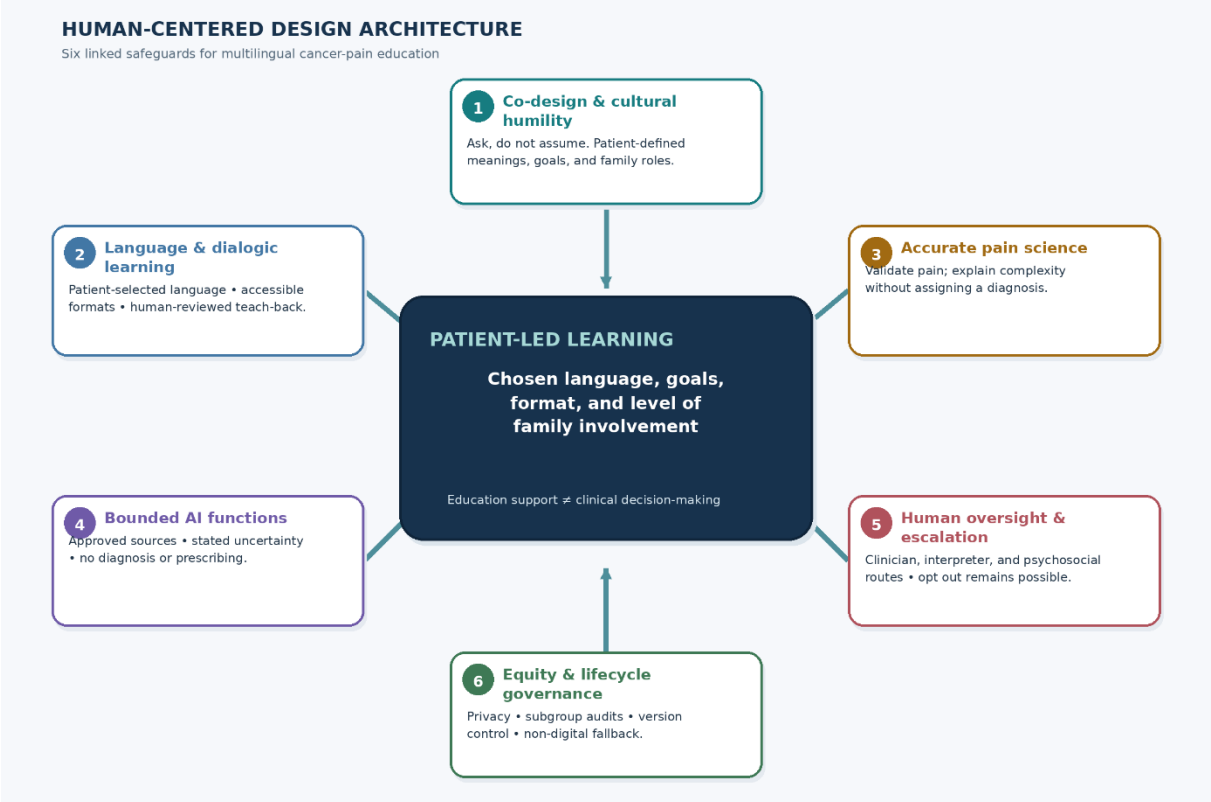


Figure 1. *Human-centered framework for culturally responsive AI-supported cancer pain education.*

Note. The patient's priorities and chosen language anchor the framework. Cultural adaptation, educational quality, AI scope, human oversight, and governance operate together; the non-digital and human-interpreter routes remain available.

Table 2. Six design domains and minimum safeguards

Domain	Educational purpose	Minimum safeguard and example of evaluation
1. Cultural humility and co-design	Begin with the patient's own account, explanatory model, goals, and preferences; recognize variation within communities.	Paid or otherwise meaningful participation by patients, caregivers, clinicians, interpreters, and community partners; cognitive interviews; documentation of whose feedback changed the content.
2. Multilingual access and dialogic learning	Make information understandable in the patient's chosen language, register, and format; support questions and teach-back.	Human-reviewed language versions; audio, visual, captioned, printable, and low-bandwidth options; comprehension and teach-back assessed by language and literacy group.
3. Accurate, non-reductionist pain science	Explain pain as real and multidimensional; distinguish general education from individualized assessment.	Clinician-approved, dated content; no psychologizing, minimization, or claim that a pain score explains its cause; expert review for clinical accuracy.
4. Bounded, evidence-grounded AI	Adapt approved information to a question, simplify wording, generate questions, or rehearse teach-back.	Retrieval from a controlled knowledge base; source and date shown; uncertainty stated; abstention and referral when the question is outside scope; no diagnosis or prescribing.
5. Human oversight and relational care	Preserve clinician, interpreter, psycho-oncology, and caregiver roles according to patient choice.	Visible AI disclosure; clear routes to a qualified person; urgent escalation based on a clinician-approved protocol; consent before involving family; no AI-only care pathway.
6. Lifecycle governance and equity	Protect privacy, assess performance across groups, and maintain accountability over time.	Data minimization, consent, access controls, version logs, incident reporting, dialect- and subgroup-stratified audits, and an accessible non-digital fallback.

Domain 1: Cultural humility and co-design

Cultural responsiveness begins with asking rather than inferring. A system should not presume that a patient's ethnicity, religion, language, gender, or nationality predicts beliefs about pain, opioids, family roles, disclosure, or mental health. It should invite the patient to choose what to discuss, who may be involved, what language or register feels comfortable, and whether a particular explanation or metaphor is helpful. This approach follows cultural humility's emphasis on lifelong self-reflection and attention to power rather than a checklist of presumed group traits (Tervalon & Murray-García, 1998). Co-design

should influence the questions, examples, audio, visuals, and referral pathways—not merely review a finished translation. Where feasible, patients and caregivers should be compensated and should be able to reject, revise, or withhold personal narratives.

Domain 2: Multilingual access and dialogic learning

Language availability is not the same as linguistic access. The education pathway should ask the patient's preferred language and mode, then support validated versions in the languages a service can safely provide. In a Moroccan context this may involve Darija, Modern Standard Arabic, Amazigh varieties, French, or another language selected by the patient; the list is an implementation question, not an assumption about any individual. Translation should be followed by cognitive testing for meaning, tone, local usage, and actionability. Patient-facing content should be available in short modules and more than one format, with audio and visual explanations for people who prefer them or have reading, vision, or digital-access barriers. Teach-back should be phrased as a request for feedback on the explanation—"I want to make sure I explained that clearly; what would you tell someone at home?"—and followed by re-teaching where needed. It should never be used to shame, grade, or deny care.

Domain 3: Accurate, non-reductionist pain science

Approved educational content can explain that pain is an individual experience shaped by sensory, emotional, cognitive, biological, and social factors, and that pain does not map perfectly onto tissue injury or a single test. It can distinguish general concepts—such as nociceptive, neuropathic, and treatment-related pain—without assigning a mechanism to a patient. Pain neuroscience education should validate the symptom, explain that the nervous system participates in pain processing, and present coping or movement strategies only when they are part of a clinician-approved plan. It should not imply that "the brain creates the pain" in a way that dismisses cancer or treatment effects. Any new, severe, rapidly changing, or otherwise concerning symptom should be directed to the oncology team using a service-approved escalation protocol, not interpreted by the AI system.

Domain 4: Bounded, evidence-grounded AI

The safest initial functions are educational and administrative: retrieving a reviewed passage, explaining a term in simpler language, generating patient questions, offering a short comprehension check, or directing the user to a relevant human service. A retrieval-

augmented system may constrain answers to a controlled source set, but retrieval does not guarantee correctness. Responses should identify the source and update date, distinguish general information from personal advice, and abstain if the source set does not answer the question. The system should not diagnose the cause of pain, recommend or change a medicine or dose, infer urgency from free text, provide psychotherapy, or independently decide whether a patient needs emergency care. If a service later considers symptom triage, that is a separate clinical function requiring its own intended-use specification, validation, governance, and authorization; it should not be smuggled into an "education" product.

Domain 5: Human oversight and relational care

The interface should disclose that the user is interacting with AI, explain what it can and cannot do, and provide a visible route to an oncology professional, interpreter, pharmacist, nurse, psycho-oncology clinician, or other appropriate person. Patients should be able to use education without the AI, decline digital support, or switch to a human conversation. Caregiver access should depend on patient preference and consent. Distress, fear, family conflict, uncertainty about medicines, or a request for individualized treatment should lead to human support rather than increasingly persuasive automated dialogue. Clinical guidance on cancer pain and psychosocial care remains the responsibility of qualified professionals and their institutions (Andersen et al., 2023; Fallon et al., 2018; World Health Organization [WHO], 2021, 2024).

Domain 6: Lifecycle governance and equity

Governance must begin before deployment and continue after updates. A service should record the intended use, approved source set, model and prompt versions, language variants, human reviewers, limitations, and route for reporting errors. It should minimize data collection, avoid sending identifiable patient records to unapproved external systems, explain retention and access, and offer a non-digital alternative. Performance should be tested separately across languages, dialects, literacy levels, ages, disability access needs, and connectivity conditions; pooled accuracy can conceal failures for a smaller group. WHO guidance emphasizes autonomy, safety, transparency, accountability, inclusion, and sustainability, while the NIST AI Risk Management Framework offers a lifecycle structure for identifying and managing risks (Tabassi, 2023; WHO, 2021, 2024). Equity is a design and monitoring requirement, not a one-time translation check.

Design propositions for future testing

The framework yields six propositions for future evaluation. First, co-design and cultural humility should improve perceived relevance and reduce stereotyping compared with translation-only adaptation; this requires direct testing. Second, patient-selected language and format should improve proximal comprehension only when the materials are linguistically and conceptually validated. Third, teach-back and question-generation functions may support communication, but their value depends on the human response to the patient's questions. Fourth, retrieval from an approved corpus may reduce unsupported statements but cannot eliminate hallucinations, omissions, or misleading emphasis. Fifth, visible human escalation and an offline alternative should support trust and equitable access, especially for users who are distressed, digitally excluded, or uncertain about the system. Sixth, improvements in knowledge or communication should not be presumed to reduce pain intensity; that claim requires controlled clinical evaluation with adequate follow-up and appropriate comparators.

Staged evaluation pathway

Figure 2 and Table 3 translate the conceptual framework into a research sequence. This pathway is a proposal, not an existing validated protocol. The first studies should test the educational content and safety of the system before any claim of clinical utility. A human-led comparator is essential: an appropriate later trial would compare standard human-led education with the same education augmented by optional AI support, not AI-only education against no care.

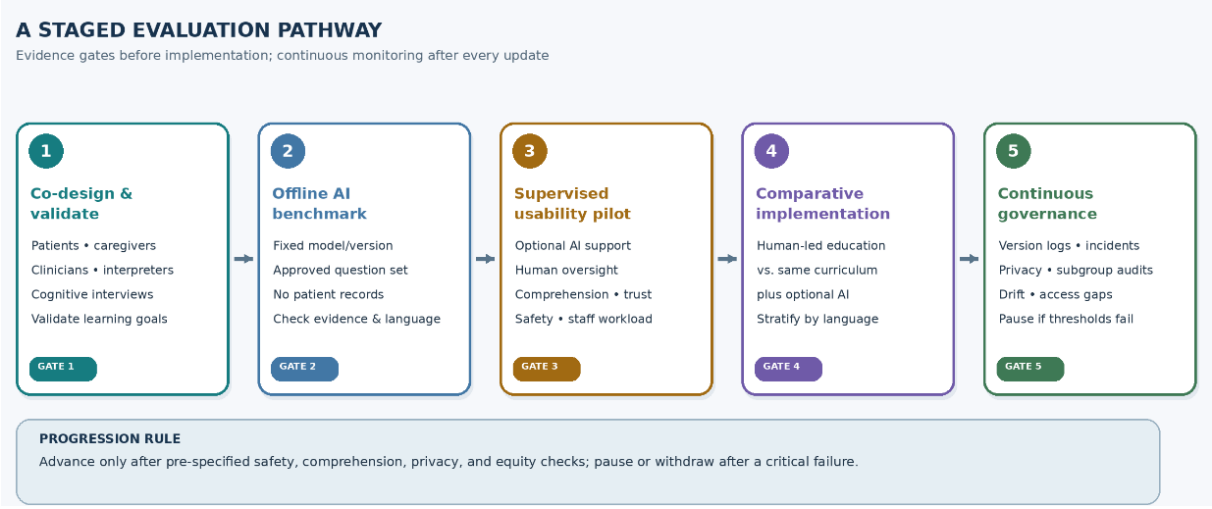


Figure 2. *Proposed staged pathway for evaluating an AI-supported education system.*

Note. Progression between stages requires documented safety and comprehension thresholds. Post-deployment monitoring is continuous; a material safety, privacy, translation, or subgroup-performance failure should trigger review and, where appropriate, suspension.

Table 3. Proposed evaluation stages, measures, and progression criteria

Stage	Main activity	Priority measures	Example progression condition
1. Co-design and content validation	Define intended use; develop approved pain-education content with patients, caregivers, clinicians, educators, and interpreters.	Relevance, cultural acceptability, cognitive-interview findings, terminology and comprehension across language versions.	No unresolved high-severity misunderstanding; patient partners endorse the learning goals and format.
2. Offline AI benchmark	Test a fixed model and version against a clinician-approved question set, without identifiable patient data.	Factual support, completeness, citation fidelity, reading level, refusal, escalation, translation quality, and performance by language/dialect.	No critical unsafe recommendations; failures are documented and corrected; results meet pre-specified subgroup thresholds.
3. Supervised usability pilot	Offer the tool as an optional supplement under human oversight; retain a non-digital route.	Comprehension, teach-back, usability, trust, perceived agency, distress, completion, help-seeking, staff workload, and incidents.	Users understand the system's limits; no unresolved safety event; acceptable access across the intended groups.
4. Comparative implementation study	Compare human-led education with the same human-led programme plus optional AI support.	Primary: comprehension and communication. Secondary: knowledge, self-efficacy, pain interference, distress, service navigation, and implementation outcomes.	Pre-registered protocol; adequate power; language-stratified analysis; independent safety monitoring.
5. Continuous governance	Reassess after model, source, workflow, or language updates and during routine use.	Error reports, subgroup performance, privacy events, drift, referral completion, usage inequities, staff burden, and user feedback.	Version-controlled approval; transparent corrective action; pause or withdrawal when safety or equity thresholds are breached.

DISCUSSION

The principal contribution is a set of design requirements for linking health education, culture, pain science, and AI governance. The framework addresses a practical

gap between two bodies of work: cancer pain education has evidence on knowledge, self-management, and communication but variable effects on pain; cancer-AI education studies increasingly evaluate accuracy and readability but rarely assess cultural responsiveness, patient experience, or equity. Recent reviews have already developed logic models and cultural-adaptation guidance for cancer-pain or psychosocial interventions (Costas-Muniz et al., 2023; Dehpour et al., 2026). The present contribution is narrower and complementary: it specifies how an AI-supported educational layer might be bounded, co-designed, linguistically tested, and evaluated without treating generative AI as a clinical decision-maker.

The educational premise is that accessible information is not enough. Health literacy is produced in an interaction among people, communication practices, and institutional demands; an interface that rewrites a leaflet into shorter sentences may still fail if its examples are culturally alien, if the patient cannot ask a follow-up question, or if the clinic has no interpreter or clinician available. The framework therefore treats patient agency, comprehension, and navigation as educational outcomes. Teach-back is useful as a dialogic check, but evidence from a small cancer-specific literature should not be inflated into a guarantee of improved adherence or clinical outcomes (Choi & Choi, 2021; Talevski et al., 2020). In practice, the professional remains responsible for adjusting the explanation and responding to what the patient says.

Cultural humility is equally important. "Culturally responsive" should not mean assigning beliefs to a nationality, ethnicity, faith, or language group. Such essentialism can reproduce the very inequities that education is meant to address. Instead, adaptation should examine whose knowledge is treated as authoritative, whether family involvement is desired, what terms a patient uses for pain, and whether structural barriers—transport, cost, time, interpreter access, stigma, literacy, or digital connectivity—make a recommended action feasible. Reviews of culturally incorporated cancer interventions show that community involvement and deeper cultural adaptation are relevant, but also that these strategies are inconsistently operationalized and are often studied in populations unlike Moroccan oncology patients (Arana-Chicas et al., 2025; Llave et al., 2024; Lui et al., 2023). The local Moroccan sources in this synthesis are useful for designing questions and partnerships, not for claiming that one language version or cultural model fits all patients.

Pain science also requires careful pedagogy. A biopsychosocial account can help explain why pain varies and why multiple approaches may be used, but it should not be used to psychologize an individual's cancer pain or imply that reframing beliefs is a substitute for medical assessment. The distinction between pain education and treatment is particularly important when AI can produce authoritative-sounding individualized recommendations. A patient-facing tool should explain what is known, acknowledge uncertainty, and support a conversation with the care team. It should not infer a pain mechanism, suggest a medication change, or minimize a report because it conflicts with a model's expectation. For the same reason, pain intensity is not the only appropriate measure of educational value: comprehension, question-asking, reporting, self-efficacy, and timely connection to care are nearer to the intervention's intended mechanism.

The AI boundary is not simply a technical feature. A system may appear empathetic or linguistically fluent while producing incomplete, unsupported, or culturally inappropriate information. In cancer education studies, accuracy and readability assessments coexist with inconsistent tools, custom scales, and limited patient-centered frameworks (Chen et al., 2025; Leong & Ormsby, 2026). Human-in-the-loop translation offers a useful precedent for language quality assurance, but it does not remove the need to test each context, register, and intended use (Brewster et al., 2025). Similarly, ePAL and OncoPainBot illustrate different digital and clinical-AI pathways; neither establishes that generative AI improves culturally responsive pain education (Kamdar et al., 2024; Liu et al., 2026). A narrow educational scope, explicit abstention, evidence-linked content, human escalation, and version control are therefore minimum requirements rather than optional features. WHO and NIST guidance further support lifecycle oversight and accountability (Tabassi, 2023; WHO, 2021, 2024).

For education policy and practice, the framework favors a hybrid model. Oncology services could co-produce a small library of reviewed multilingual modules, audio explanations, visual aids, and question prompts, then test them with patients before adding an AI interface. Educators and clinicians should receive training in health-literacy communication, uncertainty, AI limitations, and escalation. Patients should retain access to human interpreters and non-digital materials. In Morocco, a locally designed pathway could begin with patient-selected language and cognitive testing of key terms; Darija is a plausible priority for some users, but availability should be based on expressed preference and service capacity, with Amazigh varieties, Modern Standard Arabic, French, or other

languages considered as appropriate. The 2026 HLQ adaptation demonstrates why literal translation is inadequate, while also underscoring that language adaptation alone does not validate an oncology education intervention (Boumendil et al., 2026). The pediatric pain survey and adult diagnosis-disclosure study provide contextual signals of information and communication needs, but adult cancer-pain education must be studied directly (Khoubila et al., 2020; Raoui et al., 2025).

The evaluation sequence is designed to prevent premature deployment and can be adapted to implementation-context frameworks that attend to intervention characteristics, outer and inner settings, individuals, and implementation processes (Damschroder et al., 2022). Early work should test content and language; technical benchmarking should include error types, source support, comprehension, and subgroup performance rather than accuracy alone; supervised pilots should document adverse events and staff workload; and comparative studies should use the same human-led curriculum in both arms so the added value of AI is interpretable. Core outcomes should be selected with patients and clinicians in advance. A reasonable primary endpoint for an initial education trial is comprehension or communication quality, with knowledge, navigation, and self-efficacy as additional proximal outcomes. Pain interference, distress, adherence, and service use may be important secondary outcomes, but they should not be treated as guaranteed consequences. Results should be stratified by language and literacy where sample size and privacy permit; small groups should not be exposed through over-detailed reporting. A successful average score cannot justify deployment if a lower-resource language group receives unsafe or unusable outputs.

This article has important limitations. First, its literature synthesis was focused and non-systematic; it may have missed relevant work, particularly non-English and regional publications. Second, no formal quality appraisal was conducted, and the conceptual framework has not been validated by patients, clinicians, interpreters, or educational specialists. Third, several Moroccan sources concern primary care, pediatric oncology, or diagnosis disclosure rather than adult cancer pain education; they are contextual examples, not direct evidence for the proposed intervention. Fourth, the AI evidence is evolving rapidly, and studies of translation, clinical decision support, or digital symptom monitoring cannot be treated as evidence for a patient-facing educational model. Finally, no model was tested and no participant data were collected. The proposed design principles, thresholds,

and outcomes must therefore be refined through participatory research and prospective evaluation before implementation.

CONCLUSION

Culturally responsive AI-supported cancer pain education should be designed as a patient-led, human-governed educational service—not as an automated substitute for oncology care. The proposed framework links cultural humility, multilingual access, dialogic health-literacy teaching, accurate pain science, bounded AI functions, human escalation, and lifecycle equity governance. Existing evidence supports attention to pain knowledge, communication, and access, but does not justify promises that generative AI or education alone will reduce cancer pain. Moroccan Darija and oncology examples make the need for local adaptation visible while remaining distinct from cancer-specific validation. The immediate research priority is to co-design and test educational content and AI safety across languages, then evaluate comprehension, communication, trust, and equitable access before studying clinical outcomes. Any future deployment should retain clinician and interpreter access, protect privacy, disclose AI use, and pause when safety or subgroup performance is inadequate.

Declarations

Ethics approval and consent to participate: Not applicable. This critical perspective involved no human participants, identifiable records, or new patient data.

Data availability: No new dataset was generated or analyzed for this article.

Funding: No external funding was received.

Competing interests: The authors declare no competing interests.

Author contributions: To preserve double-blind review, the author-specific CRediT roles are provided on the separate title page. All authors reviewed and approved the final manuscript.

Originality and exclusive submission: The authors confirm that this manuscript is original, unpublished, and not under consideration elsewhere.

Use of generative AI and AI-assisted tools: Paperpal was used solely for grammar correction, language polishing, and stylistic refinement. It was not used to

generate the theoretical argument, develop the propositions, or interpret the literature. The authors reviewed and verified the final manuscript and take responsibility for its content.

REFERENCES

Abeo, A. N. A., Armstrong, S., Scriney, M., & Goss, H. (2025). Artificial intelligence techniques and health literacy: A systematic review. *Mayo Clinic Proceedings: Digital Health, 3*(4), Article 100269. https://doi.org/10.1016/j.mcpdig.2025.100269

Akdemir, A., Smith, A. B., Wu, V. S., Rincones, O., Russell, H., Lyhne, J. D., Kemp, E., David, M., & Bamgboje-Ayodele, A. (2024). Guided versus non-guided digital psychological interventions for cancer patients: A systematic review and meta-analysis of engagement and efficacy. *Psycho-Oncology, 33*(1), Article e6290. https://doi.org/10.1002/pon.6290

Andersen, B. L., Lacchetti, C., Ashing, K., Berek, J. S., Berman, B. S., Bolte, S., Dizon, D. S., Given, B., Nekhlyudov, L., Pirl, W., Stanton, A. L., & Rowland, J. H. (2023). Management of anxiety and depression in adult survivors of cancer: ASCO guideline update. *Journal of Clinical Oncology, 41*(18), 3426–3453. https://doi.org/10.1200/JCO.23.00293

Arana-Chicas, E., Zhang, Y., Chávez-Iñiguez, A., Lin, P.-J., Mattick, L. J., Kamen, C., Clark, V., Cartujano-Barrera, F., & Mustian, K. M. (2025). Use of cultural appropriateness strategies and behavioral frameworks in behavioral interventions for Black and Hispanic cancer survivors: A systematic review. *BMC Cancer, 25*, Article 835. https://doi.org/10.1186/s12885-025-14182-0

Aydin, S., Karabacak, M., Vlachos, V., & Margetis, K. (2024). Large language models in patient education: A scoping review of applications in medicine. *Frontiers in Medicine, 11*, Article 1477898. https://doi.org/10.3389/fmed.2024.1477898

Boumendil, K., Iderdar, Y., Al Wachami, N., Arraji, M., Najmi, S., Hawkins, M., Osborne, R., & Chahboune, M. (2026). Translation and cultural adaptation of the health literacy questionnaire (HLQ) in Morocco. *Journal of Patient-Reported Outcomes, 10*, Article 149. https://doi.org/10.1186/s41687-026-01111-3

Brewster, R. C. L., Tse, G., Fan, A. L., Elborki, M., Newell, M., Gonzalez, P., Hoq, A., Chang, C., Chowdhury, M., Geeti, A., Hana, M., Hassan, H., Ibrahim, O., Keseyan, L., Li, Q., Mamoon, M., Nageye, M., Ohannessian, A., Eisenberg, I. R., … Khan, A. (2025). Evaluating human-in-the-loop strategies for artificial intelligence-enabled translation of patient discharge instructions: A multidisciplinary analysis. *npj Digital Medicine, 8*, Article 629. https://doi.org/10.1038/s41746-025-02055-6

Busch, F., Hoffmann, L., Rueger, C., van Dijk, E. H. C., Kader, R., Ortiz-Prado, E., Makowski, M. R., Saba, L., Hadamitzky, M., Kather, J. N., Truhn, D., Cuocolo, R., Adams, L. C., & Bressem, K. K. (2025). Current applications and challenges in large language models for patient care: A systematic review. *Communications Medicine, 5*, Article 26. https://doi.org/10.1038/s43856-024-00717-2

Chen, D., Avison, K., Alnassar, S., Huang, R. S., & Raman, S. (2025). Medical accuracy of artificial intelligence chatbots in oncology: A scoping review. *The Oncologist, 30*(4), Article oyaf038. https://doi.org/10.1093/oncolo/oyaf038

Choi, S., & Choi, J. (2021). Effects of the teach-back method among cancer patients: A systematic review of the literature. *Supportive Care in Cancer, 29*(12), 7259–7268. https://doi.org/10.1007/s00520-021-06445-w

Costas-Muniz, R., Torres-Blasco, N., Gany, F., Gonzalez, C. J., Galindo-Vazquez, O., Bergerot, C. D., Rocha-Cadman, X., Lui, F., Lichtenthal, W. G., Velazquez, A. I., Tergas, A. I., & Castro-Figueroa, E. M. (2023). Cultural adaptation process of cancer-related interventions: A step-by-step guide. *Psycho-Oncology, 32*(1), 179–186. https://doi.org/10.1002/pon.6071

Damschroder, L. J., Reardon, C. M., Widerquist, M. A. O., & Lowery, J. (2022). The updated Consolidated Framework for Implementation Research based on user feedback. *Implementation Science, 17*, Article 75. https://doi.org/10.1186/s13012-022-01245-0

Dehpour, T., Wakefield, D., Luckett, T., Hussain, J., & Koffman, J. (2026). Patient-centred educational interventions for assessing and managing cancer-related pain relevant to minority ethnic populations: A mixed-methods systematic review. *Palliative Care and Social Practice, 20*, Article 26323524261421936. https://doi.org/10.1177/26323524261421936

Fallon, M., Giusti, R., Aielli, F., Hoskin, P., Rolke, R., Sharma, M., & Ripamonti, C. I. (2018). Management of cancer pain in adult patients: ESMO Clinical Practice Guidelines. *Annals of Oncology, 29*(Suppl. 4), iv166–iv191. https://doi.org/10.1093/annonc/mdy152

Galietta, E., Donati, C. M., Bazzocchi, A., Sassi, R., Zamfir, A. A., Hovenier, R., Bos, C., Hendriks, N., Boomsma, M. F., Huhtala, M., Blanco Sequeiros, R., Grüll, H., Ferdinandus, S., Verkooijen, H. M., & Morganti, A. G. (2025). Impact of pain education on pain relief in oncological patients: A narrative review of systematic reviews and meta-analyses. *Cancers, 17*(10), Article 1683. https://doi.org/10.3390/cancers17101683

González-Martín, A. M., Aguilera-García, I., Castellote-Caballero, Y., Rivas-Campo, Y., Bernal-Suárez, A., & Aibar-Almazán, A. (2023). Effectiveness of therapeutic education in patients with cancer pain: Systematic review and meta-analysis. *Cancers, 15*(16), Article 4123. https://doi.org/10.3390/cancers15164123

Hernández-Hernández, S., Heredia-Ciuró, A., Martín-Núñez, J., Calvache-Mateo, A., Navas-Otero, A., López-López, L., & Valenza, M. C. (2024). The impact of pain education interventions for cancer survivors and caregivers: A systematic review with meta-analysis. *Cancers, 16*(13), Article 2468. https://doi.org/10.3390/cancers16132468

Jeitani, C., Van den Broucke, S., & Leemans, C. (2025). Interventions addressing health literacy in cancer care: A systematic review of reviews. *International Journal of Environmental Research and Public Health, 22*(2), Article 212. https://doi.org/10.3390/ijerph22020212

Kamdar, M., Jethwani, K., Centi, A. J., Agboola, S., Fischer, N., Traeger, L., Rinaldi, S., Strand, J., Ritchie, C., Temel, J. S., Greer, J. A., Kvedar, J., El-Jawarhi, A., & Jackson, V. (2024). A digital therapeutic application (ePAL) to manage pain in patients with advanced cancer: A randomized controlled trial. *Journal of Pain and Symptom Management, 68*(3), 261–271. https://doi.org/10.1016/j.jpainsymman.2024.05.033

Khoubila, N., Bendari, M., Benmiloud, S., ElHoudzi, J., Maani, K., Elmouden, L., Alzemmouri, M., Hachim, J., Kili, A., Kababri, M., Khattab, M., ElBouri, H., Hassoune, S., Nani, S., Cherkaoui, S., Quessar, A., Madani, A., & Hessissen, L. (2020). Pain management in children with cancer: National surveys of practices and perceptions in Morocco. *Open Access Macedonian Journal of Medical Sciences, 8*(B), 723–730. https://doi.org/10.3889/oamjms.2020.5252

Leong, A., & Ormsby, K. (2026). Generative AI for patient education in cancer care: A scoping review of evaluation practices and emerging trends. *Technical Innovations & Patient Support in Radiation Oncology, 37*, Article 100373. https://doi.org/10.1016/j.tipsro.2025.100373

Liu, H., Hu, Y., Li, D., Shi, B., Niu, Y., Zhang, X., Zheng, G., Li, C., Wang, L., & Bao, Y. (2026). LLM-driven collaborative framework for knowledge-enhanced cancer pain assessment and management. *npj Digital Medicine, 9*, Article 180. https://doi.org/10.1038/s41746-026-02362-6

Llave, K., Cheng, K. K., Ko, A., Pham, A., Ericson, M., Campos, B., Perez-Gilbe, H. R., & Kim, J. H. J. (2024). Promising directions: A systematic review of psychosocial and behavioral interventions with cultural incorporation for advanced and metastatic cancer. *International Journal of Behavioral Medicine, 31*(6), 848–870. https://doi.org/10.1007/s12529-024-10264-8

Lui, F., Lewicka, M., Bao, G. C., Moyer, A., Boyce, L., & Leng, J. (2023). A systematic review and meta-analysis of psychosocial interventions for immigrant and limited English proficient cancer patients. *Psycho-Oncology, 32*(4), 516–557. https://doi.org/10.1002/pon.6110

Martínez-Miranda, P., Jiménez-Rejano, J. J., Muñoz-Fernández, M. J., García-Muñoz, C., & Casuso-Holgado, M. J. (2024). Effectiveness of an interactive online group intervention based on pain neuroscience education and graded exposure to movement in breast cancer survivors with chronic pain: A randomised controlled trial. *Supportive Care in Cancer, 32*(10), Article 705. https://doi.org/10.1007/s00520-024-08887-4

Nutbeam, D. (2000). Health literacy as a public health goal: A challenge for contemporary health education and communication strategies into the 21st century. *Health Promotion International, 15*(3), 259–267. https://doi.org/10.1093/heapro/15.3.259

Raja, S. N., Carr, D. B., Cohen, M., Finnerup, N. B., Flor, H., Gibson, S., Keefe, F. J., Mogil, J. S., Ringkamp, M., Sluka, K. A., Song, X.-J., Stevens, B., Sullivan, M. D., Tutelman, P. R., Ushida, T., & Vader, K. (2020). The revised International Association for the Study of Pain definition of pain: Concepts, challenges, and compromises. *Pain, 161*(9), 1976–1982. https://doi.org/10.1097/j.pain.0000000000001939

Raoui, A., El Ghalib, L., & Agoub, M. (2025). Psychological and social impact of announcing cancer diagnosis in a Moroccan university hospital oncology unit: A cross-sectional study. *BMC Psychology, 13*, Article 1417. https://doi.org/10.1186/s40359-025-03690-0

Sjattar, E. L., Arafat, R., & Ling, L. W. (2024). Cancer pain self-management interventions in adults: Scoping review. *BMJ Supportive & Palliative Care, 14*(4), 411–415. https://doi.org/10.1136/spcare-2024-004893

Snijders, R. A. H., Brom, L., Theunissen, M., & van den Beuken-van Everdingen, M. H. J. (2023). Update on prevalence of pain in patients with cancer 2022: A systematic literature review and meta-analysis. *Cancers, 15*(3), Article 591. https://doi.org/10.3390/cancers15030591

Sørensen, K., Van den Broucke, S., Fullam, J., Doyle, G., Pelikan, J., Slonska, Z., Brand, H., & (HLS-EU) Consortium Health Literacy Project European. (2012). Health literacy and public health: A systematic review and integration of definitions and models. *BMC Public Health, 12*, Article 80. https://doi.org/10.1186/1471-2458-12-80

Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

Talevski, J., Wong Shee, A., Rasmussen, B., Kemp, G., & Beauchamp, A. (2020). Teach-back: A systematic review of implementation and impacts. *PLOS ONE, 15*(4), Article e0231350. https://doi.org/10.1371/journal.pone.0231350

Tervalon, M., & Murray-García, J. (1998). Cultural humility versus cultural competence: A critical distinction in defining physician training outcomes in multicultural education. *Journal of Health Care for the Poor and Underserved, 9*(2), 117–125. https://doi.org/10.1353/hpu.2010.0233

Thind, B. S., & Tsao, C.-K. (2025). Artificial intelligence in oncology: Promise, peril, and the future of patient–physician interaction. *Frontiers in Digital Health, 7*, Article 1633577. https://doi.org/10.3389/fdgth.2025.1633577

Whitehead, L., Kirk, D., Chejor, P., Liu, W., Nguyen, M., Balczer, C., Lan, C., & Evans, M. (2025). Interventions, programmes and resources that address culturally and linguistically diverse consumer and carers' cancer information needs: A mixed methods systematic review. *BMC Cancer, 25*, Article 599. https://doi.org/10.1186/s12885-025-13931-5

World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. https://iris.who.int/handle/10665/341996

World Health Organization. (2024). *Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models*. https://iris.who.int/handle/10665/375579

Zhang, Y., Flannery, M., Zhang, Z., Underhill-Blazey, M., Bobry, M., Leblanc, N., Rodriguez, D., & Zhang, C. (2024). Digital health psychosocial intervention in adult patients with cancer and their families: Systematic review and meta-analysis. *JMIR Cancer, 10*, Article e46116. https://doi.org/10.2196/46116