INPLASY

Human-centered artificial intelligence in paediatric oncology: a rapid review protocol

INPLASY202690103
doi: 10.37766/inplasy2026.9.0103
Received: 23 September 2026
Published: 23 September 2026

Corresponding author:
Madhan Balasubramanian

madhan.balasubramanian@flinders.edu.au

Author Affiliation:
Flinders University.

Shivashankar, M; McMurray, A; Pereira, B; Balasubramanian, M.

ADMINISTRATIVE INFORMATION

Support - No support.

Review Stage at time of this submission - Preliminary searches.

Conflicts of interest - None declared.

INPLASY registration number: INPLASY202690103

Amendments - This protocol was registered with the International Platform of Registered Systematic Review and Meta-Analysis Protocols (INPLASY) on 23 September 2026 and was last updated on 23 September 2026.

INTRODUCTION

Review question / Objective This review addresses: "What is the nature and extent of health professional, child and family involvement in the problem definition, design, development, evaluation and implementation of artificial intelligence (AI) systems in paediatric oncology?" Three stakeholder groups are used throughout: health professionals (doctors, nurses, allied health professionals, analysed collectively with subgroups reported separately where studies permit); parents and carers; and children and young people. The primary objective is to map which participants were involved and at which AI lifecycle stage (problem definition, design, development, evaluation, implementation), and to characterise the depth of participation, from being informed through to collaborative decision-making, using a five-level scale adapted from the IAP2 Spectrum of Public Participation. By mapping current practice across the paediatric-oncology AI literature, the review will identify the lifecycle stages at which health professional, child and family participation remains absent, superficial, or under-reported, offering direction for clinicians, health service leaders and technology developers seeking to design AI systems around the people who use and are affected by them.

Background Artificial intelligence (AI) is increasingly applied across the paediatric cancer pathway, spanning tumour detection and segmentation, molecular classification, risk and treatment-response prediction, and clinical documentation. Paediatric oncology presents a distinct challenge for these technologies: childhood cancers are rare and biologically diverse, datasets are small and heterogeneous, imaging protocols are non-standardised across growing children, and clinical decisions carry consequences across a long survivorship. These characteristics make paediatric oncology both a demanding and an important setting for examining how AI is developed and deployed.

Translation of AI into routine paediatric cancer care remains limited, with barriers arising from the interaction of technical, organisational, regulatory and human factors. Reviews of clinical AI

repeatedly identify clinician trust, explainability, workflow integration and meaningful end-user involvement as principal barriers to adoption. Interviews with innovation experts in paediatric cancer specifically highlight legal and regulatory hurdles, validation standards, and communication gaps between clinicians and technical teams. A survey of healthcare workers across African paediatric-oncology units found awareness and readiness to be uneven, underscoring the equity dimension of who is engaged in AI development.

"Human-centred AI" (HCAI) has emerged as a response to these failures, arguing that AI systems should be designed with and around the people who use and are affected by them, rather than around the algorithm alone. In healthcare, this is operationalised through participatory co-design, attention to trust calibration and explainability, and evaluation against human-factors and experience outcomes, not technical accuracy alone. Yet whether, and how, this principle has been applied in paediatric oncology is unclear: most published work reports model performance, while far less is known about who was involved in designing these systems, at which stage, and with what degree of influence. For this review, human-centred practice is defined operationally: a study is characterised as demonstrating human-centred practice where it reports at least one documented instance of a stakeholder group contributing to the AI system at a level of consultation or above (level 2+ on the participation scale). Related constructs — explainability, trust calibration, workflow integration, and evaluation against human-factors or equity outcomes — are extracted and reported in their own right, not treated as criteria a study must meet to qualify.

Rationale No existing review has systematically mapped the extent and nature of human involvement in AI systems for paediatric oncology. Prior reviews concentrate on technical capability and diagnostic performance, or address adoption barriers in adult or general clinical settings. Where human-centred practice is discussed, it is typically presumed from a study's use of that vocabulary rather than assessed from what was actually done; studies that conducted usability testing, clinician workshops or family involvement frequently do not label themselves as "human-centred". The degree to which children, families and clinicians are genuine partners in the design of paediatric-oncology AI has therefore not been established.

This rapid review directly addresses that gap. Rather than restricting the search to studies using human-centred terminology, which would favour self-labelled studies, the review searches the broader paediatric-oncology AI literature and then assesses human involvement at full text, coding its depth and lifecycle stage across stakeholder groups. A rapid review design is appropriate given the timeliness of the question for health-service decision-making and the funded programme within which it sits, with methodological shortcuts declared transparently. Findings will give clinicians, health leaders and developers a clear account of where human-centred practice currently stands in childhood-cancer AI, and where earlier and deeper participation is required.

METHODS

Strategy of data synthesis Three bibliographic databases were searched: MEDLINE (Ovid), Scopus, and Web of Science Core Collection, from 1 January 2015 to the date of final search (searches run 6 August 2026). Database-specific strings combine controlled vocabulary (e.g., MeSH) and free-text terms across three concept domains combined with AND: (1) paediatric populations (paediatric, pediatric, child, adolescen, infant, neonat, juvenile, "young people"); (2) oncology (cancer, oncolog, neoplas, tumour, leukaemi, lymphom, sarcom, blastoma, malignan); and (3) artificial intelligence ("artificial intelligence", "machine learning", "deep learning", "neural network", "large language model", "clinical decision support", radiomic), plus named model/vendor terms and algorithmic-decision terms. Episodes of care (diagnosis, imaging, treatment, monitoring, survivorship, supportive care) are not a separate search domain, since the disease domain already retrieves records across the full pathway; pathway stage is instead coded at extraction. The search is not restricted to human-centred terminology (e.g., "co-design", "participatory", "user-centred", "usability"), since studies reporting usability testing or family involvement frequently do not use such terms; these terms may be used only as an optional supplementary layer. Only peer-reviewed, English-language publications are eligible. The strategy was validated against five known-relevant seed papers spanning the cancer pathway; four are MEDLINE-indexed and retrieved, the fifth (published in EJC Paediatric Oncology) is retrieved via Scopus. No publication-type filter was applied in any database, to avoid inconsistency across databases with differing publication-type taxonomies; ineligible publication types are instead excluded at screening. Forward and backward citation tracking of included studies will be undertaken via Scopus and Google Scholar, and reference lists of relevant reviews will be screened.

Searches retrieved 10,966 records in total: 2,709 from MEDLINE, 5,967 from Scopus and 2,290 from Web of Science. All records were imported into Nested Knowledge (AutoLit), which identified 3,967 duplicates, leaving 6,999 unique records for title and abstract screening. This protocol is reported in accordance with PRISMA-P, and the completed review will be reported in accordance with PRISMA 2020. Rapid-review shortcuts adopted to accelerate the review (and to be addressed in the limitations of the completed review) include: search restricted to three databases with no grey literature; single-reviewer title/abstract screening after a 50-record calibration pilot with a second reviewer; single-reviewer full-text screening with second-reviewer verification of a sample; single-reviewer data extraction with second-reviewer verification of a random 25% sample; and no formal risk-of-bias assessment, since the review maps human involvement rather than estimating intervention effectiveness. Software-assisted (Nested Knowledge AutoLit) screening prioritisation and extraction are used, with every automated output confirmed by a human reviewer before acceptance.

Eligibility criteria Population: Children and adolescents aged 0–18 years with cancer, and/or their families and carers, and/or the health professionals (doctors, nurses, allied health professionals) who care for them. Mixed-age studies are eligible where paediatric findings are reported separately or where participants aged 0–18 years constitute more than 50% of the sample; where this cannot be established, corresponding authors will be contacted, and studies excluded if unresolved within two weeks.

Intervention/Concept: Any artificial intelligence or machine-learning system applied to any part of the paediatric cancer pathway — detection, diagnosis, segmentation, classification, risk or response prediction, prognostication, treatment planning, monitoring, documentation, or psychosocial/supportive care. Systems may be rule-based, machine-learning, deep-learning, or foundation-model (e.g., large language model) based.

Comparator: None required; studies with and without a comparator are eligible, since the review maps human involvement rather than estimating an intervention effect.

Outcomes: (1) barriers and enablers to human-centred practice reported by each study; (2) higher-order implementation challenges derived across the corpus by pattern-coding those findings; (3) depth and lifecycle stage of involvement of health professionals, parents/carers, and children/young people, coded at full text, distinguishing documented absence of involvement from absence of reporting; (4) categories of outcome each study measures — technical performance, human-factors/usability, process, patient/family experience, clinical, and equity.

Study designs/Context: Any empirical design is eligible — model development and validation studies, qualitative studies, surveys, mixed-methods, usability and implementation studies, and controlled trials. Purely technical or simulation papers are eligible (coded according to what they report, including a not-reported code). Editorials, commentaries, letters and reviews are excluded, though reference lists of relevant reviews will be screened. Full peer-reviewed conference papers are eligible; conference abstracts without an accompanying full paper are excluded, since a substantial share of co-design and participatory research relevant to this review is published in full conference proceedings rather than journals. Articles must be peer-reviewed, in English, published from 1 January 2015 onwards.

Source of evidence screening and selection All records will be imported into Nested Knowledge (AutoLit; Nested Knowledge Inc., St Paul, MN, USA) for de-duplication, screening, tagging and extraction, which supports both single and dual screening modes. The first author (MS) will act as primary reviewer, screening all titles and abstracts against the eligibility criteria; an initial pilot of approximately 50 records will be screened independently by the primary and second reviewers to calibrate the inclusion threshold before full screening proceeds. Where limited information is available at title/abstract stage, full text will be obtained to determine eligibility. Full-text screening is undertaken by the primary reviewer, with the second reviewer verifying a sample and adjudicating uncertain cases. Disagreements between reviewers will be resolved through discussion; where consensus is not reached, a senior arbiter will be involved and the decision recorded. The selection process will be presented in a PRISMA 2020 flow diagram. Nested Knowledge provides AI-assisted functions including screening prioritisation and qualitative extraction; every AI-generated decision or extraction will be reviewed and confirmed by a human reviewer before acceptance, and the stages at which AI assistance was used will be reported for transparency and replicability.

Data management Data extraction will use a piloted, standardised template configured within Nested Knowledge (AutoLit), whose hierarchical tagging structure applies the participation coding frame and framework-based barrier/enabler categories consistently and exports them for analysis. The primary reviewer will perform extraction; the second reviewer will independently verify a random 25% sample, including participation coding and NASSS-based coding of barriers and enablers. Extracted items include bibliographic details, country/setting, sample and cancer type, AI system task/method/data modality, cancer-pathway stage addressed, and — for five stakeholder categories (doctors, nurses, allied health professionals, parents/carers, children/young people) — depth of involvement at each of five lifecycle stages (problem definition, design, development, evaluation, implementation), using a six-level scale (0/None to 4/Partner, plus NR/Not reported) adapted from the IAP2 Spectrum of Public Participation. The coding frame will be piloted on ten studies independently by both reviewers before full extraction. Where information is unclear, corresponding authors will be contacted.

Reporting results / Analysis of the evidence A descriptive synthesis will be followed, structured around the four outcomes in priority order. Barriers and enablers will be coded deductively to the seven NASSS framework domains and summarised by domain and frequency across studies (interpreted descriptively, not as effect size or importance). Study-level barriers/enablers will then be pattern-coded into higher-order implementation challenges, tested by systematic comparison across NASSS domain, stakeholder group and lifecycle stage, with negative cases actively sought and an audit table linking each challenge to its source codes and studies. The central output is a stakeholder × lifecycle matrix showing depth of involvement, with proportions reported against two denominators (all included studies, and studies with sufficient reporting) as a sensitivity analysis reflecting reporting quality. Inter-rater agreement (Cohen's kappa, weighted kappa, percentage agreement) will be reported for dual-screened/extracted samples. No meta-analysis is planned, given expected heterogeneity and the descriptive mapping aim.

Presentation of the results Planned outputs, pre-declared against each outcome: (1) a table of barriers/enablers by NASSS domain, with number and proportion of studies reporting each; (2) a narrative synthesis of higher-order implementation challenges organised by NASSS domain, accompanied by an audit table linking each challenge to its constituent domains, codes and studies; (3) a stakeholder × lifecycle matrix — three stakeholder groups (health professionals; parents/carers; children and young people) by five lifecycle stages — with each cell showing relevant studies, studies with sufficient reporting, number/proportion reaching level 3 (Involve) or above, and number coded NR, plus a supplementary matrix disaggregating health professionals into doctors, nurses and allied health professionals; (4) a frequency table of studies by outcome tier, cross-tabulated against cancer-pathway stage. These will be accompanied by a PRISMA 2020 flow diagram and a table of study characteristics.

Language restriction English language only; non-English-language publications are excluded.

Country(ies) involved Australia.

Other relevant information Conducted within the AI4NCD Innovation Collaborative at Flinders University. Protocol reported per PRISMA-P; completed review will follow PRISMA 2020.

Keywords artificial intelligence; human-centred AI; paediatric oncology; childhood cancer; co-design; stakeholder engagement; participation; implementation; rapid review.

Dissemination plans Findings will be submitted for peer-reviewed publication (target journal to be confirmed) and presented at relevant paediatric oncology and health-services research conferences. Results will also be shared with the AI4NCD Innovation Collaborative at Flinders University and with paediatric oncology services and health service leaders in South Australia, to directly inform where human-centred AI practice needs strengthening.

Contributions of each author

Author 1 - Mahalakshmi Shivashankar - Design; Protocol; Search strategy; Screening; Data extraction; Synthesis; Writing first draft.
Email: mahalakshmi.shivashankar@flinders.edu.au
Author 2 - Adela McMurray - Design; Protocol; Synthesis.
Email: adela.mcmurray@flinders.edu.au
Author 3 - Bruno Pereira - Design; Protocol; Synthesis.
Email: bruno.pereira@flinders.edu.au
Author 4 - Madhan Balasubramanian - Design; Protocol; Screening; Synthesis.
Email: madhan.balasubramanian@flinders.edu.au