Journal of Integrative Nursing and Palliative Care (JINPC)

Original Article

Volume 7 | Issue 2 | Page 146-156
Copyright CC BY NC SA 4.0

From Observable Behavior to Inferred Suffering: A Conceptual Framework for Defining the Limits of Artificial Intelligence in Symptom Assessment among Nonverbal Adults Receiving Palliative Care

Giulia Romano[1]*, Francesco Esposito[2]

[1]*Department of Nursing Ethics, University of Rome, Rome, Italy.*
[2]*Department of Palliative Care, University of Naples, Naples, Italy.*

Abstract

To define the epistemic limits of artificial intelligence (AI) when symptom-related information from nonverbal adults receiving palliative care is transformed from observable behavior into claims about subjective suffering. Patients who cannot reliably communicate may express distress through facial activity, movement, vocalization, physiological change and contextual alterations recognized by caregivers or clinicians. Structured observational instruments and emerging AI systems can improve detection, consistency and information synthesis. Yet observed signals, proxy interpretations, algorithmic classifications and lived suffering are not interchangeable forms of knowledge. Original non-empirical conceptual boundary analysis informed by critically selected literature on nonverbal symptom assessment, observational pain measures, proxy reporting, machine-learning inference, clinical decision support, nursing judgment and responsible healthcare AI. The analysis distinguishes five epistemically different positions: observable evidence, human or proxy interpretation, model classification, probabilistic clinical inference and unobserved subjective suffering. AI may strengthen detection and support clinical reasoning when its inputs, target construct, validation population and uncertainty are explicit. Its output does not independently establish the content, intensity or meaning of an individual patient's subjective suffering. Human interpretation is also fallible and should not be treated as an infallible comparator. Responsible use of AI in nonverbal palliative care requires calibrated claims rather than either technological overconfidence or categorical rejection. The strength of any statement about suffering should remain proportionate to the evidence source, model target, validation context and unresolved uncertainty, while nursing accountability and relational interpretation remain clinically consequential.

Keywords: Artificial intelligence, Nonverbal communication, Palliative care, Pain assessment, Symptom assessment, Clinical inference

Introduction

Pain is formally understood as a personal sensory and emotional experience, and an inability to communicate does not negate the possibility that pain is present [1]. That principle becomes especially consequential in palliative care, where cognitive impairment, neurological disease, profound weakness, delirium, reduced consciousness, language disturbance and the final stages of illness can limit a person's ability to provide a stable verbal account. In such circumstances, the absence of conventional self-report changes the available evidence; it does not convert an internal experience into a directly observable object. For nursing practice, this creates an epistemic problem as well as an assessment problem: clinicians must often act under uncertainty while remaining careful about what their observations actually establish.

Corresponding author: Giulia Romano
Address: Department of Nursing Ethics, University of Rome, Rome, Italy.
E-mail: ✉ giulia.romano@uniroma1.it
Received: 20 May 2026; **Revised:** 7 July 2026; **Accepted:** 11 July 2026

How to Cite This Article: Romano G, Esposito F. From Observable Behavior to Inferred Suffering: A Conceptual Framework for Defining the Limits of Artificial Intelligence in Symptom Assessment among Nonverbal Adults Receiving Palliative Care. J Integr Nurs Palliat Care. 2026;7(2):146-156. https://doi.org/10.51847/xMuOMrgYIE

Structured approaches to nonverbal pain assessment can improve the consistency with which behavioral and contextual cues are recognized in palliative populations [2]. Such approaches are clinically valuable because untreated pain and distress are serious harms, and a requirement for verbal confirmation before responding would itself be unsafe. The relevant question is therefore not whether indirect evidence should be used. It must be used. The harder question is how far a clinician, instrument or computational model may legitimately move from an observable pattern to a statement about the patient's current subjective experience.

AI intensifies this question because it can combine quantities and forms of data that exceed unaided human processing. A recent mapping of AI in palliative care found a growing field spanning prognostication, symptom-related applications, communication and decision support, while also identifying substantial deficits in prospective validation, transparency and clinically grounded evaluation [3]. Parallel work on AI for pain assessment in older adults shows expanding use of machine learning, deep learning, computer vision, natural-language processing and other computational approaches, but also continuing problems of limited datasets, contextual transfer and validation [4]. These findings make a categorical claim that AI is incapable of contributing to symptom assessment untenable. They also make premature claims that AI can determine suffering difficult to defend.

Cancer-pain research further demonstrates that computational methods are already being applied to clinically meaningful pain-related tasks [5]. Nursing research similarly shows growing use of AI-enabled decision support, although implementation evidence, reproducibility, prospective testing and meaningful nursing involvement remain uneven [6]. The central problem is therefore one of claim calibration. A system may detect a facial configuration, classify a segment of video, estimate a model-defined pain state, flag a clinical note, predict a care need or recommend reassessment. Each output can be useful. None should acquire a stronger experiential meaning merely because it was produced computationally.

This article examines the boundary between what can be observed and what may be inferred about suffering in nonverbal adults receiving palliative care. Its contribution is an Observable-to-Inferred Suffering Claim-Boundary Framework that separates observable behavior, proxy interpretation, model classification, probabilistic clinical inference and unobserved subjective experience. The analysis does not privilege clinicians simply because they are human, nor algorithms because they are computationally consistent. Instead, it asks what evidentiary warrant supports each form of claim and where additional inference enters.

Background

Observational pain instruments illustrate why this distinction is necessary even before AI is introduced. Systematic psychometric evaluations of pain tools for people with dementia show that observational measures vary substantially in reliability, validity, responsiveness and the quality of evidence supporting their use [7,8]. These instruments operationalize patterns such as facial behavior, vocalization, body movement, consolability or changes from usual behavior. Their scores are therefore clinically structured interpretations of observable phenomena. A well-supported observational measure can provide stronger evidence than unstructured impression, but the numerical or categorical output still represents a measured construct rather than direct access to the person's phenomenological experience.

Facial behavior is among the most studied observable channels. A systematic review of facial muscle activity associated with pain identified recurrent actions across experimental and clinical research while also showing meaningful heterogeneity in their occurrence and combination [9]. This matters because the clinical usefulness of a signal does not require it to be universal, but its absence cannot automatically function as evidence that pain is absent. Nonverbal symptom expression is shaped by individual variability, illness, motor capacity, fatigue, medication, social context and potentially culture. A cue therefore acquires meaning through context rather than through appearance alone.

Direct observational evidence makes the same point from another direction. In residents with dementia and chronic pain, some facial expressions were associated with observational pain whereas other expressions were not [10]. Conversely, structured technology-supported approaches such as PainChek have demonstrated measurable agreement and validity against established observational comparators in hospitalized people with dementia [11]. These findings are important counterweights to excessive skepticism. Indirect assessment can be systematic, reproducible and useful. The relevant limitation is narrower: validation against behavioral or clinician-observed comparators does not transform a detected behavioral pattern into direct measurement of subjective suffering.

Human interpretation also introduces uncertainty. Nurses' beliefs and approaches to pain assessment in people living with dementia are associated with factors such as knowledge, education and practice context [12]. Implementation research in acute care similarly shows that perceived staff capability and actual documented pain-management practice do not necessarily align [13]. Clinical judgment is therefore not an epistemically neutral gold standard against which algorithmic inference can simply be judged. Human interpretation can integrate biography, trajectory, relational knowledge and situational subtleties unavailable to a model, yet it can also be inconsistent, incompletely documented or shaped by training and organizational conditions.

Proxy reporting adds another layer. Patient–caregiver comparisons show that agreement differs by symptom and circumstance and that proxies can either understate or overstate a patient's report [14]. A family member who

recognizes an individual's usual expressions may possess highly valuable longitudinal knowledge, while a clinician may recognize physiological deterioration or behavioral deviation that a family member cannot interpret clinically. Neither form of proxy knowledge is equivalent to direct self-report. Their importance lies precisely in their partial and situated access.

These distinctions create a hierarchy of different evidence types, not a hierarchy of human worth or professional authority. Direct self-report, when meaningful and available, provides access to the patient's own account. Observable behavior provides externally detectable evidence. Physiological signals provide bodily correlates. Family members contribute relationally accumulated knowledge. Nurses contribute clinical observation, trajectory recognition and contextual judgment. Instruments structure selected cues. Each can become more or less informative depending on the patient and setting, but none should be silently relabeled as another type of evidence.

The evidence-source distinctions that will govern the subsequent analysis are summarized in **Table 1**.

Table 1. Symptom-information sources and the inferential limits attached to each evidentiary form

Information source	What enters the clinical information field	Principal interpretive strength	Dependence on context or prior knowledge	What the source can reasonably support	What remains epistemically inaccessible from that source alone
Patient self-report	Patient-selected words, ratings, gestures or other intentional communication	Direct account of experienced symptoms when communication is meaningful	Language, cognition, communication capacity, timing and willingness to disclose	Presence, quality, intensity or meaning as reported by the patient	Independent verification that every aspect of experience has been fully communicated
Observable behavior	Facial activity, movement, posture, vocalization, guarding, withdrawal or change from baseline	Detectable evidence associated with discomfort or pain-related states	Motor capacity, illness, medication, individual expression pattern, environment	Presence of behavior compatible with a symptom-related state	Exact subjective intensity, personal meaning or broader suffering
Physiological signal	Autonomic, cardiovascular, respiratory or other measurable bodily changes	Continuous or objective measurement of bodily response	Comorbidity, medication, arousal, infection, exertion and competing causes	Physiological change compatible with a pain- or distress-related state	Unique attribution of the change to pain or suffering
Family or caregiver proxy interpretation	Longitudinal recognition of usual behavior, preferences and deviations	Person-specific relational knowledge that may be unavailable to short-contact clinicians	Relationship history, emotional burden, knowledge of patient baseline	Plausible interpretation of behavior and change in context	Direct access to the patient's current subjective state
Nurse or clinician interpretation	Bedside observation integrated with history, examination, trajectory and care context	Clinical synthesis across multiple information sources	Expertise, workload, continuity, expectations, local practice and documentation quality	Clinically actionable judgment with explicit uncertainty	Infallible determination of private experience
Structured observational instrument	Predefined combination of selected behaviors or signs	Standardization and potentially improved reproducibility	Construct definition, validation population and observer quality	Instrument-defined score or category supported within validated use	A broader experiential claim not represented by the instrument's construct

Computational feature extraction or classification	Digitized behavioral, physiological, textual or multimodal features processed by an algorithm	Consistency, scalability and detection of complex statistical patterns	Training labels, dataset composition, target definition and model transportability	Model-defined classification or estimate within its validated domain	Independent proof of what the patient subjectively feels or how that experience is valued

Data sources

Computational pain assessment makes the transformation from evidence source to inferential output especially visible. Computer-vision systems trained on pain-related facial or behavioral labels can recognize patterns that are difficult to monitor continuously by human observers. Work in older adults with dementia has shown that pairwise and contrastive learning strategies can support unobtrusive pain monitoring from observed behavior [15]. In perioperative settings, machine-learning approaches have also been developed to estimate pain-related states quantitatively from clinical data [16]. These studies demonstrate technical feasibility. What they establish, however, is performance against a defined target or reference label. The epistemic content of the output remains constrained by how that target was constructed.

Physiological signals create a parallel problem. Contemporary synthesis shows that heart-rate-related measures, electrodermal activity, respiratory changes and other physiological markers have been investigated as inputs for pain assessment across multiple forms of pain [17]. Such data may provide information when overt behavioral expression is minimal or difficult to observe. Their interpretive limitation is non-specificity: autonomic or physiological change can reflect pain, anxiety, dyspnea, fever, medication effects, exertion, fear or other processes. A model may improve statistical discrimination among predefined states without converting a non-specific physiological response into a uniquely identifiable subjective experience.

Textual data introduce yet another level of mediation. Natural-language processing has been used to identify uncontrolled symptoms in clinical documentation for patients with cancer [18]. This can increase the recoverability of clinically meaningful information at scale. Nevertheless, an NLP system analyzes what has been documented. If an experience was not verbalized, noticed, interpreted or recorded, the model cannot recover that missing experience merely through greater computational sophistication. The distinction between patient state and recorded representation of patient state therefore remains essential.

Patient-reported outcome measures occupy a different epistemic position because they deliberately preserve the patient's own account as a data source. Work on intelligent palliative care based on patient-reported outcome measures illustrates how patient-generated information can be incorporated into data-driven care [19]. Yet that route becomes less available precisely in the population considered here. When meaningful self-report cannot be obtained, a system must rely more heavily on behavioral, physiological, proxy or record-based signals. The inferential distance between input and subjective experience therefore increases.

Clinical decision-support systems demonstrate why model usefulness and experiential validity should be evaluated separately. A scoping review of decision support in palliative care found applications spanning referral, medication, symptom management and other clinical decisions [20]. Such systems can organize information and support action without claiming to measure suffering itself. Similarly, prediction models have been developed to identify palliative-care needs among people with advanced cancer receiving chemotherapy [21]. The endpoint in that setting is an operationally defined care need. Recasting the same prediction as evidence that a patient is suffering would change the construct being claimed without necessarily changing the underlying model.

The same distinction appears in implementation studies. Machine-learning targeting can be used to identify patients for advance-care-planning activity [22]. AI-based decision support has also increased palliative-care referral in a randomized clinical setting [23]. These are clinically meaningful outcomes because they show that algorithmic information can alter care delivery. Their value does not depend on the algorithm having direct access to an internal experience. A system may therefore be effective at directing attention while remaining epistemically limited concerning what the patient is experiencing.

Generative AI introduces a further complication because its outputs resemble human reasoning linguistically. In standardized nursing vignettes, generative AI has shown an ability to approximate aspects of clinical decision-making performed by nurses and nursing students [24]. Such findings are relevant to future symptom-assessment support, but linguistic plausibility or concordance with a structured answer does not establish bedside sensitivity to embodied, relational and temporally changing cues. The model receives a representation of the situation rather than the patient's experience itself.

Across these data sources, four transformations recur. First, some aspect of the patient or care environment is rendered observable or recordable. Second, the observation is assigned a label, interpretation or operational target. Third, a model learns or applies a relationship between inputs and that target. Fourth, a clinician may use the output to make a further inference about the patient. The greatest conceptual risk occurs when these stages are compressed so completely that the final output is described as though the system had directly observed the subjective state.

Figure 1 makes these different epistemic positions visible while preserving legitimate clinical routes from observation to action.

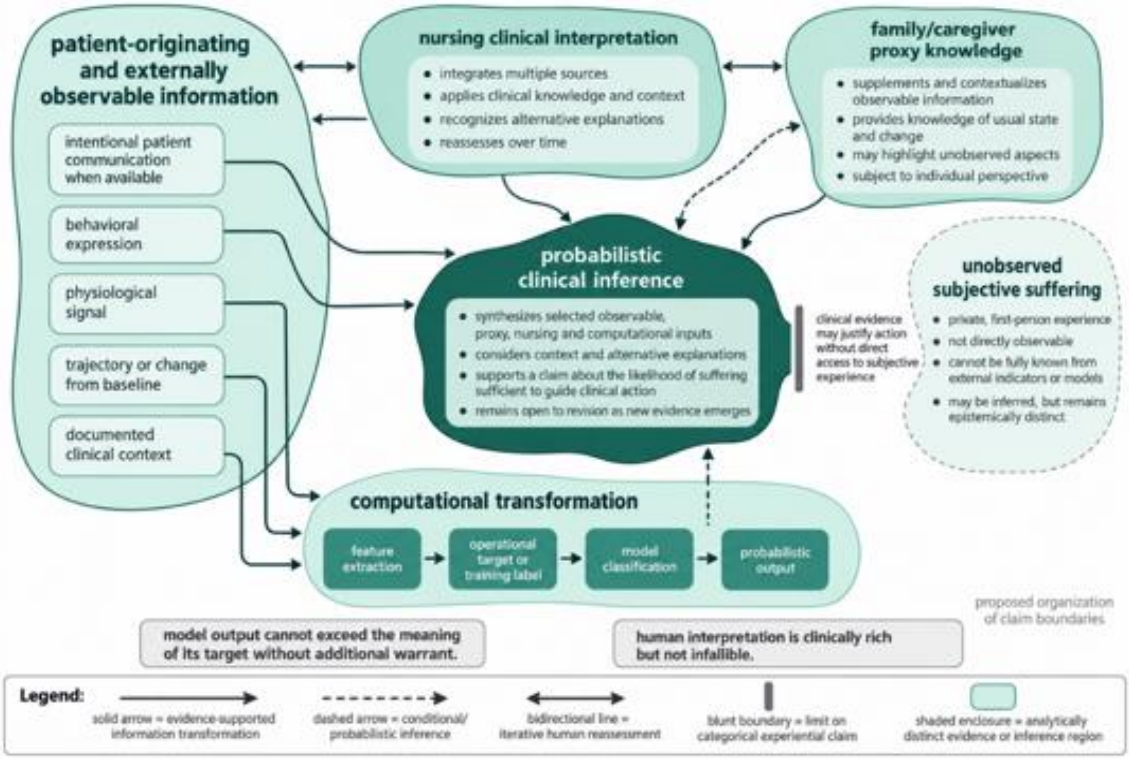


Figure 1. Where symptom evidence changes epistemic status before a claim about suffering is made

Overview of the issues

The central problem is not scarcity of cues but ambiguity in what cues mean. Contemporary synthesis of pain cues in people with dementia identifies facial expressions, movement, vocal behavior and other observable changes across a heterogeneous evidence base [25]. This multiplicity can improve sensitivity to possible pain because clinicians need not wait for one canonical sign. It also makes deterministic interpretation less defensible. A person may express pain in a familiar way on one occasion and differently later as illness, medication, energy, consciousness or social context changes. The absence of an expected cue therefore has weak exclusionary value unless the cue's sensitivity for that individual is well established.

AI can help manage this multiplicity by maintaining longitudinal attention, identifying recurrent patterns or integrating multimodal data. The risk appears when computational consistency is treated as epistemic superiority. An integrative review of generative AI in clinical nursing describes potential gains alongside concerns involving trust, communication, workflow, equity and professional role boundaries [26]. These issues become particularly important when a patient cannot readily correct an inaccurate interpretation. An erroneous model output may acquire persuasive force precisely because the person whose experience is being inferred has limited ability to contest it.

Palliative clinicians' responses to AI prognostication show that this concern is not abstract. Physicians report conditional interest in AI while also identifying responsibility and liability questions surrounding its use [27]. Prognostication differs from symptom assessment, but the accountability issue is analogous: computational contribution does not automatically relocate responsibility for interpretation. When an AI system suggests that a nonverbal patient is unlikely to be in pain, someone still has to decide whether that output is sufficiently trustworthy to influence observation, medication, examination, family consultation or escalation.

Explainability does not eliminate this problem. A systematic review of explainable AI in healthcare found that explanations can increase trust, decrease it or have little effect depending on how explanations and tasks are designed [28]. An explanation can help a clinician interrogate a system, but it can also make an output appear more coherent than its underlying target or validation justifies. The relevant question is consequently not only whether a model can explain *why* it produced a result, but whether the result itself concerns the clinical construct that users believe it concerns.

This distinction can be expressed through five progressively stronger forms of statement. An observational statement reports what is detectable: for example, a change in facial activity or movement. A proxy-interpretive statement attributes probable meaning to that observation using relational or clinical knowledge. A model-

classification statement reports assignment to an operational class learned from specified labels. A probabilistic clinical inference integrates several sources and expresses a qualified judgment about the likelihood of a symptom-related state. A subjective-experience claim states what the patient is actually feeling or how that experience is affecting the person. The transitions between these statements require different kinds of warrant.

This hierarchy does not imply that nurses should suspend treatment until subjective certainty becomes possible. Clinical action and epistemic certainty are separate questions. When a patient shows repeated distress-related behavior, a family member recognizes the pattern as unusual, a nurse identifies a plausible pain-generating cause and a validated tool or model yields concordant information, precautionary analgesia, reassessment or diagnostic evaluation may be entirely reasonable. What should remain qualified is the language used to describe what has been established. Clinical responsibility may require action under uncertainty rather than elimination of uncertainty.

Table 2 translates this distinction into claim-boundary and nursing-oversight principles without converting them into unvalidated clinical thresholds.

Table 2. Proposed claim-boundary and nursing-oversight principles for AI-supported assessment in nonverbal palliative care

Claim escalation point	Minimum evidentiary condition	Acceptable formulation	Nursing oversight function	Condition requiring caution	Statement that should not be made from that evidence alone
From signal to observation	Signal is detectable, technically credible and attributable to the patient or clinical record	"The system detected a change in facial activity, movement, physiology or documented symptom language."	Confirm data quality, clinical plausibility and competing explanations	Artifact, poor recording, altered motor function, medication effect, documentation gap	"The patient is suffering."
From observation to proxy interpretation	Observable change is considered with baseline, trajectory, illness context and relational knowledge	"The observed change is compatible with pain or distress and warrants further assessment."	Integrate bedside examination, prior pattern, family knowledge and care goals	New illness, delirium, fear, dyspnea or other plausible causes	"This behavior uniquely means pain."
From data to model classification	Model target is explicit and the system has been evaluated for the intended type of data	"The model classified the available input within its defined pain-related category."	Determine whether the target and training labels correspond to the present clinical question	Target mismatch, dataset shift, missing modality, unfamiliar population	"The algorithm measured the patient's suffering."
From classification to probabilistic clinical inference	Relevant validation, target-population fit, contextual concordance and uncertainty are available	"Taken with other evidence, the output increases or decreases concern about a symptom-related state."	Weigh model output against patient trajectory, examination, proxy information and alternative explanations	Poor calibration, subgroup uncertainty, contradictory bedside information	"The model has established the patient's internal state."
From convergent evidence to clinical action	Several sources support concern or potential harm of non-treatment exceeds harm of reassessment or cautious intervention	"The available evidence justifies reassessment, treatment trial, investigation or escalation."	Retain decision responsibility and evaluate response over time	Irreversible intervention, high treatment burden, discordant goals or uncertain cause	"Action proves the original interpretation was correct."

From repeated response patterns to stronger individualized inference	Within-person trajectory shows reproducible relationships among cues, context, intervention and subsequent change	"For this patient, this recurring pattern has become more informative over time."	Preserve individualized baseline and re-test assumptions when condition changes	Disease progression, sedation, neurological change, new environment or loss of prior behavioral capacity	"Past pattern guarantees the meaning of the current episode."
From probabilistic inference to subjective suffering	No indirect evidentiary configuration can provide direct phenomenological access when meaningful self-report remains unavailable	"The patient may be experiencing pain, distress or broader suffering; the degree and personal meaning remain uncertain."	Protect the patient through proportionate action while maintaining interpretive humility	Pressure to convert uncertainty into categorical documentation or automated decisions	"AI has directly determined what this patient feels."

The distinctive boundary is therefore not between "human" and "machine." Both human and algorithmic interpretations can be informative, mistaken or incomplete. The more important boundary separates what the available evidence directly represents from what clinicians or systems infer from that evidence. In nonverbal palliative care, this distinction is ethically significant because patients may be especially dependent on others to interpret their experience without either ignoring indirect evidence or overstating what it proves.

Conceptual framework

The central rule of the Observable-to-Inferred Suffering Claim-Boundary Framework is that the strength of a clinical statement should not exceed the epistemic reach of the evidence supporting it. Ethical analyses of healthcare AI have emphasized the importance of epistemic humility when computational outputs are used to make claims about complex human states [29]. In nonverbal palliative care, this principle applies equally to algorithmic and human interpretation. A facial movement may be observed with high confidence, yet its meaning remains interpretive. A model classification may be statistically reproducible, yet its clinical meaning remains limited by the target against which the model was trained. A nurse may possess rich contextual knowledge, yet still be uncertain about what a patient subjectively experiences.

Accordingly, progression from classification toward stronger clinical inference requires transparent target definition, appropriate validation and trustworthy deployment. Reporting and evaluation standards for prediction models require explicit specification of outcomes, predictors, data handling and validation, while contemporary trustworthy-AI guidance additionally emphasizes fairness, traceability, usability, robustness and explainability [30–32]. These requirements are necessary because a technically accurate model can still answer the wrong clinical question. If training labels represent observer-rated pain behavior, successful prediction establishes correspondence with that operational construct. It does not independently establish correspondence with the patient's unobserved phenomenological suffering.

Clinical admissibility also depends on what occurs after an output is generated. DECIDE-AI emphasizes evaluation of AI systems within real clinical workflows, including interactions between clinicians, technology and decision processes [33]. For nonverbal palliative assessment, the relevant unit is therefore not the algorithm alone but the patient–nurse–family–technology configuration in which evidence is interpreted. Human oversight has substantive meaning only when clinicians can interrogate the target, notice discordance, integrate trajectory and context, and disregard or override an output when the patient's condition makes it implausible.

Transfer across populations introduces another boundary. International recommendations concerning health datasets emphasize representation, provenance and transparency because model behavior can change when demographic, clinical or institutional conditions differ from those represented during development [34]. Nonverbal expression may also be altered by neurological impairment, weakness, medication, culture and disease trajectory. A system validated in postoperative adults, for example, cannot be assumed to retain equivalent construct meaning in a sedated patient approaching the end of life merely because both settings involve pain-related data.

Finally, probabilistic outputs should remain probabilistic. Prediction research increasingly emphasizes uncertainty around individualized estimates rather than presentation of point predictions as though they were certain [35]. For symptom assessment, at least two kinds of uncertainty must be distinguished: uncertainty about the model's estimate and uncertainty about whether the modeled target adequately represents the patient's lived state. Better calibration can narrow the first without eliminating the second. The strongest permissible claim is therefore determined jointly by what was observed, how the target was defined, how well the system was validated in the relevant population, whether other evidence converges, and what uncertainty remains.

Results and Discussion

The framework proposed here should be treated as conceptual synthesis requiring empirical evaluation rather than as a validated clinical instrument. Prospective studies of AI-supported symptom assessment should be reported with AI-specific methodological transparency so that training, updating, human interaction and clinically relevant outcomes can be appraised [36]. Validation should determine not merely whether an algorithm reproduces an existing score but whether its output adds useful, safe information for nonverbal palliative patients without creating false certainty.

Nursing oversight is central because current evidence does not support replacing relational clinical interpretation with autonomous algorithmic attribution. Reviews of AI-supported nursing decisions and generative AI in clinical nursing indicate potential benefits alongside unresolved questions concerning workflow, professional roles, trust and real-world validation [6]. Bedside nurses can integrate temporality, response to prior interventions, family observations, examination findings and changes from an individual's baseline—information that may be absent from a model input [26]. Human involvement, however, should not be romanticized as inherently error-free. The framework therefore calls for complementary scrutiny of both algorithmic output and human interpretation.

Explainability and professional acceptance should likewise be separated from validity. Explanations can influence clinician trust in different directions [28], while prediction-model methodology requires validation beyond the data used for development [31]. A persuasive explanation cannot compensate for a poorly chosen target, population shift or invalid proxy. Conversely, an accurate but opaque system may still have limited clinical usability. Responsible assessment consequently requires both epistemic adequacy and usable communication of uncertainty.

Clinical usefulness also does not require a system to identify subjective suffering directly. Machine-learning targeting has altered advance-care-planning processes, and AI decision support has increased palliative-care referral in prospective clinical evaluation [22, 23]. These examples show that algorithms can improve attention allocation or prompt beneficial action by predicting operationally defined outcomes. Similar logic can apply to nonverbal symptom care: a system might legitimately identify a pattern warranting reassessment without claiming that it has established the patient's internal experience.

Figure 2 summarizes the resulting admissibility logic: stronger AI-supported claims require progressively stronger evidence, while nursing accountability and uncertainty remain present rather than disappearing at higher levels of computational sophistication.

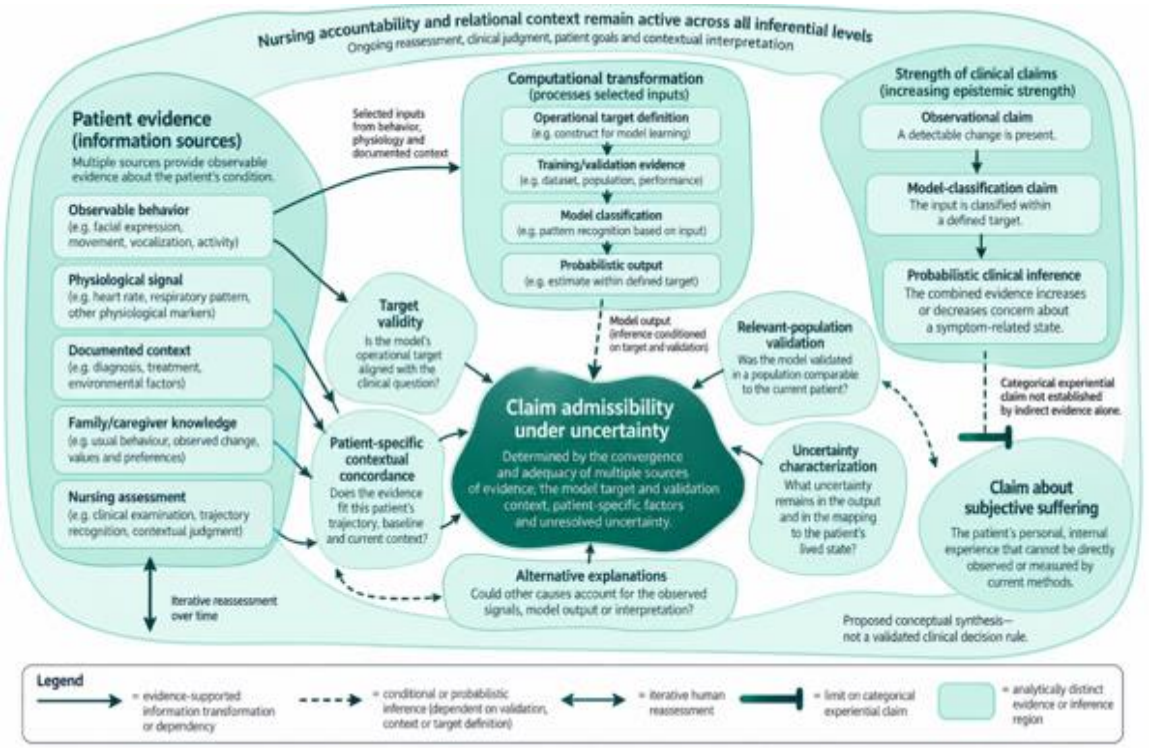


Figure 2. Evidence admissibility and uncertainty across AI-supported claims about nonverbal suffering

A future validation program should therefore examine technical, clinical and relational performance together. Recent palliative-AI evidence, nursing decision-support research and trustworthy-AI guidance collectively indicate that deployment quality depends on more than predictive performance alone [3,6,32]. Studies should

assess within-person trajectories, subgroup and cultural performance, disagreement among model, nurse and family interpretations, communication of uncertainty and consequences for treatment, reassessment and missed suffering. Such work would test whether the proposed claim boundaries improve clinical calibration without discouraging appropriate action when evidence is incomplete.

Several limitations of the present analysis remain. Much of the measurable literature concerns pain rather than the broader construct of suffering, and substantial supporting evidence derives from dementia, oncology, acute-care or perioperative settings. These sources are informative because they expose the structure of indirect symptom inference, but they cannot establish performance in nonverbal adults receiving palliative care. The framework also does not identify a universal evidentiary threshold for treatment or escalation. Thresholds would depend on intervention burden, potential harm of non-treatment, patient goals and local context and therefore require empirical and ethical evaluation rather than conceptual invention.

The principal implication is consequently one of disciplined language and responsibility. AI may detect, classify, estimate, flag and support inference. Clinicians may then integrate those outputs with examination, trajectory, relational knowledge and therapeutic response. As evidence converges, confidence may reasonably increase. Yet confidence and direct access are different concepts. Maintaining that distinction protects both innovation and the patient whose subjective experience remains the object of care.

Conclusion

Artificial intelligence can extend symptom assessment in nonverbal palliative care by detecting patterns, integrating complex information and supporting probabilistic clinical reasoning. Its legitimate contribution depends on the evidentiary meaning of its inputs, the construct represented by its target, validation in the relevant population and transparent treatment of uncertainty.

Observable behavior, proxy interpretation, model classification, probabilistic inference and subjective suffering therefore should not be treated as interchangeable states of knowledge. Strong indirect evidence may justify compassionate action even when subjective certainty is impossible. The appropriate limit is not a prohibition on inference but a requirement that claims remain proportionate to what the evidence can establish. Within that boundary, AI can support nursing judgment without displacing clinical accountability or overstating access to the patient's lived experience.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Raja SN, Carr DB, Cohen M, Finnerup NB, Flor H, Gibson S, et al. The revised International Association for the Study of Pain definition of pain: Concepts, challenges, and compromises. Pain. 2020;161(9):1976-82. doi:10.1097/j.pain.0000000000001939
2. Larson S, Laures E, Seo M, Cox M, Wagner M. Evidence-based pain assessment in nonverbal palliative care patients. Pain Manag Nurs. 2024;25(2):152-9. doi:10.1016/j.pmn.2023.12.005
3. Bozkurt S, Fereydooni S, Kar I, Diop Chalmers C, Leslie SL, Pathak R, et al. AI in palliative care: A scoping review of foundational gaps and future directions for responsible innovation. J Pain Symptom Manage. 2025;70(6). doi:10.1016/j.jpainsymman.2025.08.009
4. Ilenwabor R, Mohammed M, Martini N. Artificial intelligence in pain assessment and management for older adults: A scoping review. Artif Intell Med. 2026;179:103455. doi:10.1016/j.artmed.2026.103455
5. Salama V, Godinich B, Geng Y, Humbert-Vidan L, Maule L, Wahid KA, et al. Artificial intelligence and machine learning in cancer pain: A systematic review. J Pain Symptom Manage. 2024;68(6). doi:10.1016/j.jpainsymman.2024.07.025
6. Cant R, Ryan C, Chugh R. Artificial intelligence technologies in nursing clinical decision-making: An umbrella review. J Adv Nurs. 2026;82(10):8968-83. doi:10.1111/jan.70579
7. McLennan AIG, Winters EM, Gagnon MM, Hadjistavropoulos T. Pain in dementia: A systematic review of the measurement properties of observational assessment tools. Clin Psychol Rev. 2026;124:102704. doi:10.1016/j.cpr.2026.102704
8. Smith TO, Harvey K. Psychometric properties of pain measurements for people living with dementia: A COSMIN systematic review. Eur Geriatr Med. 2022;13(5):1029-45. doi:10.1007/s41999-022-00655-z

9. Kunz M, Meixner D, Lautenbacher S. Facial muscle movements encoding pain—a systematic review. Pain. 2019;160(3):535-49. doi:10.1097/j.pain.0000000000001424
10. Pu L, Coppieters MW, Smalbrugge M, Jones C, Byrnes J, Todorovic M, et al. Associations between facial expressions and observational pain in residents with dementia and chronic pain. J Adv Nurs. 2024;80(9):3846-55. doi:10.1111/jan.16063
11. Sampson EL, Davies N, Vickerstaff V. Evaluation of the psychometric properties of PainChek in older general hospital patients with dementia. Age Ageing. 2025;54(2). doi:10.1093/ageing/afaf027
12. Gamage MWK, Pu L, Todorovic M, Moyle W. Factors related to nurses' beliefs regarding pain assessment in people living with dementia. J Clin Nurs. 2024;33(11):4367-80. doi:10.1111/jocn.17093
13. Mahama F, Harkin D, Hayes N, Doherty M, Coates V. Exploring pain management for people living with dementia in acute care settings using the COM-B framework to guide baseline data gathering. J Clin Nurs. 2026:1-17. doi:10.1111/jocn.70508
14. Kroenke K, Stump TE, Monahan PO. Agreement between older adult patient and caregiver proxy symptom reports. J Patient Rep Outcomes. 2022;6(1):50. doi:10.1186/s41687-022-00457-8
15. Rezaei S, Moturu A, Zhao S, Prkachin KM, Hadjistavropoulos T, Taati B. Unobtrusive pain monitoring in older adults with dementia using pairwise and contrastive training. IEEE J Biomed Health Inform. 2021;25(5):1450-62. doi:10.1109/JBHI.2020.3045743
16. Ryu G, Choi JM, Seok HS, Lee J, Lee EK, Shin H, et al. Machine learning based quantitative pain assessment for the perioperative period. NPJ Digit Med. 2025;8(1):53. doi:10.1038/s41746-024-01362-8
17. Camacho-Navas C, Li L, Poply K, Mehta V, Kyriacou P. Pain assessment using physiological responses/markers in different types of pain: A scoping review. NPJ Digit Med. 2026;9(1):64. doi:10.1038/s41746-025-02241-6
18. DiMartino L, Miano T, Wessell K, Bohac B, Hanson LC. Identification of uncontrolled symptoms in cancer patients using natural language processing. J Pain Symptom Manage. 2022;63(4):610-7. doi:10.1016/j.jpainsymman.2021.10.014
19. Sandham MH, Hedgecock EA, Siegert RJ, Narayanan A, Hocaoglu MB, Higginson IJ. Intelligent palliative care based on patient-reported outcome measures. J Pain Symptom Manage. 2022;63(5):747-57. doi:10.1016/j.jpainsymman.2021.11.008
20. Dos Santos FC, Snigurska UA, Keenan GM, Lucero RJ, Modave F. Clinical decision support systems for palliative care management: A scoping review. J Pain Symptom Manage. 2023;66(2). doi:10.1016/j.jpainsymman.2023.03.006
21. Kawashima A, Furukawa T, Imaizumi T, Morohashi A, Hara M, Yamada S, et al. Predictive models for palliative care needs of advanced cancer patients receiving chemotherapy. J Pain Symptom Manage. 2024;67(4):306-16.e6. doi:10.1016/j.jpainsymman.2024.01.009
22. Patel MN, Mara A, Acker Y, Gollon J, Setji N, Walter J, et al. Machine learning for targeted advance care planning in cancer patients: A quality improvement study. J Pain Symptom Manage. 2024;68(6):539-47.e3. doi:10.1016/j.jpainsymman.2024.08.036
23. Wilson PM, Ramar P, Philpot LM, Soleimani J, Ebbert JO, Storlie CB, et al. Effect of an artificial intelligence decision support tool on palliative care referral in hospitalized patients: A randomized clinical trial. J Pain Symptom Manage. 2023;66(1):24-32. doi:10.1016/j.jpainsymman.2023.02.317
24. Saban M, Dubovi I. A comparative vignette study: Evaluating the potential role of a generative AI model in enhancing clinical decision-making in nursing. J Adv Nurs. 2025;81(11):7489-99. doi:10.1111/jan.16101
25. Smrke U, Milošič A, Mlakar I, Kadiš M, Mulej Bratec S. Pain cues in people with dementia: Scoping review. JMIR Ment Health. 2025;12. doi:10.2196/75671
26. Watson AL, Bond C, Aveyard H, Smith GD, Jackson D. Generative AI at the bedside: An integrative review of applications and implications in clinical nursing practice. J Clin Nurs. 2026;35(10):4022-37. doi:10.1111/jocn.70151
27. Fischer SM, Fink RM, Alasmar AY, Campbell EG, DeCamp M. Palliative care physicians' perceptions about using artificial intelligence for prognostication. J Pain Symptom Manage. 2026;72(4). doi:10.1016/j.jpainsymman.2026.07.005
28. Rosenbacke R, Melhus Å, McKee M, Stuckler D. How explainable artificial intelligence can increase or decrease clinicians' trust in AI applications in health care: Systematic review. JMIR AI. 2024;3. doi:10.2196/53207
29. McCradden M, Hui K, Buchman DZ. Evidence, ethics and the promise of artificial intelligence in psychiatry. J Med Ethics. 2023;49(8):573-9. doi:10.1136/jme-2022-108447
30. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385. doi:10.1136/bmj-2023-078378
31. Collins GS, Dhiman P, Ma J, Schlussel MM, Archer L, Van Calster B, et al. Evaluation of clinical prediction models (part 1): From development to external validation. BMJ. 2024;384. doi:10.1136/bmj-2023-074819

32. Lekadir K, Frangi AF, Porras AR, Glocker B, Cintas C, Langlotz CP, et al. FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388. doi:10.1136/bmj-2024-081554
33. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
34. Alderman JE, Palmer J, Laws E, McCradden MD, Ordish J, Ghassemi M, et al. Tackling algorithmic bias and promoting transparency in health datasets: The STANDING Together consensus recommendations. Lancet Digit Health. 2025;7(1). doi:10.1016/S2589-7500(24)00224-3
35. Riley RD, Collins GS, Kirton L, Snell KIE, Ensor J, Whittle R, et al. Uncertainty of risk estimates from clinical prediction models: Rationale, challenges, and approaches. BMJ. 2025;388. doi:10.1136/bmj-2024-080749
36. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: The CONSORT-AI extension. Lancet Digit Health. 2020;2(10). doi:10.1016/S2589-7500(20)30218-1