Journal of Integrative Nursing and Palliative Care (JINPC)

Original Article

Volume 7 | Issue 2 | Page 134-145
Copyright CC BY NC SA 4.0

# Whose Observation is Recorded in an AI-Generated Nursing Note? A Methodological Framework for Preserving Source Attribution, Temporal Context, and Uncertainty in Palliative Care Documentation

**Klaus Schmidt[1]*, Anna Fischer[2], Dieter Müller[3]**

[1]*Department of Nursing Informatics, University of Heidelberg, Heidelberg, Germany.*
[2]*Department of Palliative Care, University of Freiburg, Freiburg, Germany.*
[3]*Department of Nursing Science, University of Munich, Munich, Germany.*

**Abstract**

To develop a source-aware methodology for AI-assisted nursing documentation that preserves who supplied a clinically meaningful proposition, when the underlying event or statement occurred, whether the proposition represents observation or inference, what uncertainty remains, and who is responsible for correction. Generative artificial intelligence can reduce documentation workload and produce fluent clinical summaries, yet linguistic fluency can conceal changes in epistemic status. This problem is especially important in palliative care, where a note may combine patient self-report, relatives' accounts, direct nursing observations, information inherited from earlier records, and algorithmically inferred statements across a rapidly changing clinical trajectory. Original non-empirical methodological framework developed through critical synthesis of contemporary nursing, palliative-care, clinical-documentation, natural-language-processing, generative-AI, temporal-reasoning, human-factors, and governance scholarship. A critically selected 2017–2026 evidence corpus was used to examine clinical text provenance, serious-illness documentation, temporal information, AI-generated summaries, uncertainty, human editing, nursing workflow, and accountability. The evidence set was used as a conceptual foundation and was not treated as a systematic-review sample. The central methodological unit is the clinically meaningful proposition within a nursing note. Each proposition should retain an identifiable source class, temporal anchor, epistemic status, uncertainty state, and correction pathway. Patient self-report, relatives' interpretation, clinician observation, imported historical information, and algorithmic inference therefore remain distinguishable even when they are compressed into a single generated note. Trustworthy AI-assisted palliative-care documentation depends on preserving the provenance and current interpretability of clinical statements. A generated sentence should not acquire greater authority merely because an AI system has rendered it concise, coherent, or definitive.

**Keywords:** Artificial intelligence, Nursing documentation, Palliative care, Information provenance, Temporal reasoning, Clinical uncertainty

## Introduction

Artificial intelligence is beginning to alter how clinical information is captured, organized, and converted into nursing documentation. Contemporary nursing scholarship has already envisaged systems in which documentation is produced with substantially less manual burden, potentially allowing nurses to redirect time from record production toward clinical work [1]. Ambient documentation systems provide a more immediate demonstration of this trajectory: spoken encounters can be transformed into draft notes, with measurable effects

**Corresponding author:** Klaus Schmidt
**Address**: Department of Nursing Informatics, University of Heidelberg, Heidelberg, Germany.
**E-mail:** ✉ klaus.schmidt@uni-heidelberg.de
**Received:** 19 May 2026; **Revised:** 6 July 2026; **Accepted:** 9 July 2026

**How to Cite This Article:** Schmidt K, Fischer A, Müller D. Whose Observation Is Recorded in an AI-Generated Nursing Note? A Methodological Framework for Preserving Source Attribution, Temporal Context, and Uncertainty in Palliative Care Documentation. J Integr Nurs Palliat Care. 2026;7(2):134-145. https://doi.org/10.51847/veo8wBPcDX

on documentation burden and clinician experience [2]. These developments make the question of documentation efficiency increasingly practical. They also expose a second question that efficiency metrics cannot answer: when an AI-generated sentence appears in a nursing note, whose observation, report, interpretation, or inference does that sentence represent?

Clinical language models can summarize complex source material at a level that performs strongly against expert comparators in selected tasks [3]. Such performance supports legitimate interest in using generative systems to condense lengthy records, organize encounter content, and improve information accessibility. Healthcare frameworks for generative AI nevertheless emphasize that usefulness depends on the relationship among the model, intended task, users, data environment, and mechanisms for oversight [4]. For nursing documentation, source attribution is part of that relationship. A sentence can be factually plausible while leaving unclear whether the information originated from the patient, a family member, a nurse's direct observation, an earlier record, or an inference produced during automated synthesis.

This distinction matters because clinically fluent output can contain unsupported content, omissions, or unjustified certainty [5]. Rapid healthcare adoption also means that hybrid human–AI documentation practices are developing while evidence on outcomes, governance, and appropriate controls remains incomplete [6]. The central methodological problem is therefore narrower than whether generative AI can produce readable notes. It concerns whether a generated note preserves the epistemic status of the information from which it is assembled.

This paper develops Source-Aware Nursing Documentation Methodology for that problem. Its unit of analysis is the clinically meaningful proposition: a statement in the record that may influence interpretation of symptoms, preferences, goals, needs, responses, or care decisions. The proposed methodology asks five linked questions of such a proposition: Who supplied or generated it? When was it true, observed, reported, or recorded? Is it an observation, report, interpretation, or inference? What uncertainty remains? Who has responsibility and authority to correct it? These questions are particularly consequential in palliative care because information is frequently relational, preferences may evolve, family interpretations may legitimately differ from patient statements, symptom meaning can change over time, and documentation is repeatedly reused by professionals who were not present when the original information was produced.

*Background*

Nursing documentation is simultaneously a clinical record, a professional communication medium, and a relational practice. When patients gain direct access to nursing notes, nurses describe changes in how they select language and represent sensitive information, demonstrating that documentation practices are influenced by anticipated readership and by the relationship among nurse, patient, and others who may read the record [7]. Documentation is also a major component of nursing workload. Interventions aimed at redesigning documentation processes can reduce time spent producing records, confirming that documentation burden is neither trivial nor immutable [8]. AI-assisted note generation therefore addresses a genuine operational problem. The methodological challenge is to reduce that burden without stripping statements of information about their origin.



Provenance problems did not begin with generative AI. Electronic health records have long combined newly authored language with copied, imported, or reused material, making the origin and freshness of note text clinically relevant [9]. Serious-illness documentation adds another layer because information can be distributed differently across structured fields and narrative text. Content analysis of serious-illness conversations shows that structured and free-text representations do not preserve identical information [10]. A generated nursing note may therefore inherit existing documentation limitations before any generative transformation occurs.

Semantic ambiguity also matters. A systematic review and discourse analysis found substantial heterogeneity in how "goals of care" was used in healthcare literature, illustrating that familiar clinical terms can compress different meanings and purposes [11]. At the encounter level, patient, clinician, and electronic-record accounts of whether a goals-of-care discussion occurred can disagree [12]. Such discordance should not automatically be interpreted as evidence that one source is erroneous. The patient may recall the conversation according to its meaning, the clinician according to its intended function, and the record according to what was documented. Source identity is therefore part of the clinical information itself.

The same principle applies to time. Digital tools can improve the retrievability and visibility of serious-illness documentation across an electronic record [13], yet accessible information is not necessarily current information. Longitudinal examination of documented goals among hospitalized patients at high risk of mortality found that goals could be conditional, concurrent, changing, or insufficiently clear for straightforward classification [14]. A historical statement such as "patient wishes to return home" may remain clinically important without constituting a present preference after deterioration, new information, or a change in capacity. An AI system that incorporates the statement into a current note without its temporal context can transform accurate historical information into a misleading present-tense proposition.

These findings support an attribution requirement that is broader than identifying who typed a sentence. Source-aware documentation needs to preserve the provenance of the clinical claim.

The source distinctions that establish this requirement are summarized in **Table 1**.

**Table 1.** Provenance classes for clinically meaningful propositions in AI-assisted nursing documentation

| Proposition source class | What the recorded statement represents | Attribution requirement | Temporal requirement | Epistemic boundary that must remain visible | Principal documentation risk |
|---|---|---|---|---|---|
| Patient self-report | The patient's own account of symptom, experience, preference, goal, concern, or meaning | Identify the patient as the reporting source | Preserve when the patient reported it and, where clinically relevant, the period to which it referred | A report is evidence of the patient's expressed experience or preference; it is not automatically an externally verified observation | AI rephrasing may convert subjective language into an apparently objective clinical fact |
| Relative or caregiver account | Information supplied by a family member, caregiver, proxy, or other relational informant | Identify the informant and their relationship where relevant | Preserve when the account was supplied and whether it concerns a current or earlier state | The informant's interpretation must remain distinguishable from the patient's own words or preferences | Speaker substitution may cause a relative's interpretation to appear as the patient's statement |
| Direct nursing or clinician observation | Findings directly observed, examined, measured, or witnessed by a professional | Identify the observation as clinician-derived and distinguish measurement from interpretation | Anchor the observation to the relevant encounter or observation period | Observation and subsequent clinical interpretation are related but analytically distinct | Generated prose may merge what was seen with what was inferred |
| Imported or inherited record information | Information originating in earlier documentation, another clinician's note, or another part of the record | Retain traceability to prior documentation instead of presenting inherited text as newly established | Preserve original event/report time when available and identify uncertain current validity | Historical documentation is evidence of what was previously recorded; it may not describe the present state | Stale information may be linguistically refreshed and appear current |
| Algorithmic inference or synthesis | A proposition generated through model-based combination, classification, completion, or interpretation | Identify that the proposition is inferred or synthesized by the system unless independently verified | Relate the inference to the source material and time interval on which it depends | Algorithmic inference is not a witnessed observation and does not acquire clinical authority through fluent wording | Unsupported certainty, invented linkage, or silent transformation of source meaning |
| Explicitly corroborated multi-source statement | A proposition supported by more than one independently identifiable source | Preserve each materially relevant source or state that corroboration occurred | Confirm that the sources concern compatible time periods | Agreement strengthens corroboration but does not erase the identity of contributing sources | Compression may remove clinically important disagreement or timing differences |

*Data sources*

The data entering an AI-supported nursing note are heterogeneous before generation begins. Natural-language-processing research demonstrates that symptoms and other clinically meaningful phenomena can be extracted from free-text electronic health records, while also showing substantial variation in corpora, annotation practices, target concepts, and evaluation methods [15]. Broader clinical-NLP scholarship similarly emphasizes that validity depends on data quality, annotation, context, and external evaluation [16]. Computational accessibility therefore does not make a narrative statement source-neutral. A phrase available for extraction may be a patient quotation, a clinician's description of behaviour, a copied historical assertion, or an earlier interpretation.

Time constitutes a second data layer. Contemporary methods can extract temporal relationships from clinical narratives and combine textual and structured information to improve temporal reasoning [17]. Yet reviews of clinical timeline extraction identify persistent difficulty with implicit dates, relative expressions, cross-sentence relationships, and events whose timing must be inferred from context [18]. These limitations are directly relevant to palliative care, where expressions such as "recently more breathless," "previously declined admission," or "family reports worsening overnight" contain clinically meaningful but different temporal structures. When exact timing cannot be established, the unresolved temporal status should remain visible.

Provenance must also be considered at more than one level. Responsible health-AI practice increasingly emphasizes documentation of dataset suitability and intended use because missing provenance information

constrains safe reuse [19]. At the clinical-document level, automated methods can identify serious-illness conversations in electronic records with useful performance [20]. The successful detection of documentation, however, answers a different question from whether the detected proposition faithfully represents what a patient meant. Dataset provenance, document provenance, statement provenance, and speaker provenance are nested but non-equivalent.

Generative models add a further source category because they can produce clinically plausible language from learned representations and supplied context. Large language models have demonstrated substantial medical knowledge and reasoning capability while retaining clinically important limitations under expert evaluation [21]. An AI-generated proposition is therefore not a direct observation merely because it is derived from clinical data. Where the source text does not explicitly support a generated statement, the resulting content belongs to an inferential layer. Even when an inference is useful, its status needs to remain recoverable.

For source-aware documentation, a clinically meaningful proposition can consequently carry at least four independent dimensions: source, time, epistemic status, and current applicability. Correction responsibility becomes a fifth dimension once the proposition enters a draft note. These dimensions cannot be reconstructed reliably if generation collapses them into undifferentiated prose.

**Figure 1** depicts this proposition-centred provenance structure and the points at which meaning can be altered.

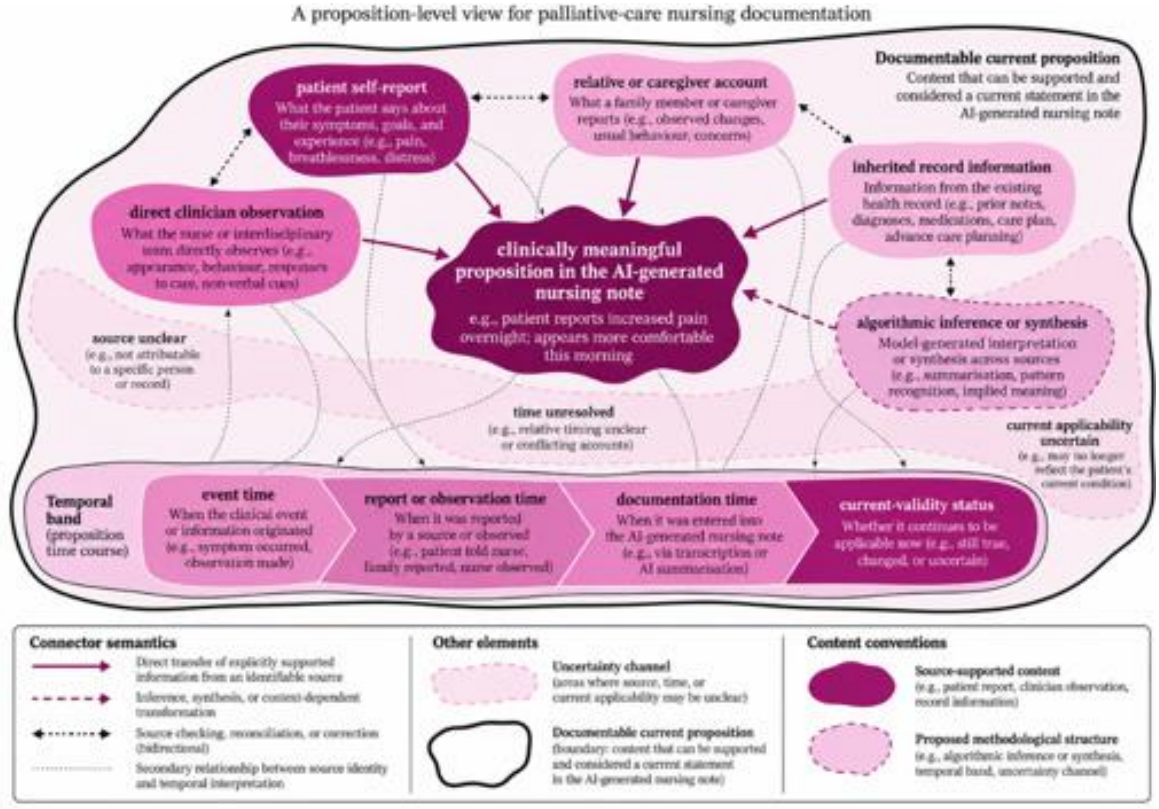



**Figure 1.** Tracing source, time, and inference around a clinically meaningful proposition in an AI-generated nursing note

*Overview of the issues*

*Transformation of source material during generation*

The principal danger in AI-generated documentation is not limited to fabrication. Generation can also alter a statement while leaving its surface plausibility intact. Early real-world evaluation of generative inpatient summarization shows why produced summaries require explicit quality assessment before clinical reliance [22]. More directly, clinicians describing their edits to ambient AI-generated notes have identified problems involving transcription, speaker attribution, context, omissions, and unwarranted certainty [23]. These editing decisions reveal a class of errors that conventional factuality measures can overlook: the underlying clinical idea may have some basis in the encounter while its documented source or epistemic status has changed.

Consider a hypothetical encounter in which a daughter says that her father "seemed confused overnight." A source-preserving note can represent this as a relative's account. A generated statement such as "Patient was confused overnight" changes two elements simultaneously: the reporter disappears, and an interpreted observation becomes an apparently established patient state. If the nurse later observes the patient to be oriented, the distinction

becomes clinically consequential. The AI did not necessarily invent the concept of confusion; it transformed its provenance.

Similar transformations can occur when a patient reports breathlessness, a nurse observes rapid breathing, and an algorithm generates "worsening respiratory distress." Each proposition may be clinically related, but they do not have the same source or evidentiary status. Source-Aware Nursing Documentation Methodology treats such transformation as a documentation event that should remain recoverable.

*Efficiency and evidentiary fidelity are separate outcomes*

Ambient documentation can produce substantial workflow benefits. Real-world evaluations have reported favorable effects on clinician experience and selected documentation measures [24]. Prospective implementation research has also demonstrated meaningful uptake and reductions in time spent documenting, although use varies across clinicians and contexts [25]. These are important outcomes because excessive documentation burden can compete with direct care, contribute to after-hours work, and reduce the practical sustainability of detailed records. Efficiency, however, measures a different property from fidelity to source. A note can be produced more rapidly while still obscuring who made a statement, when it applied, or whether the final wording represents observation or inference. Conversely, a source-aware system may remain worthwhile even if it adds a small verification step, provided that the step protects clinically material distinctions. Evaluation should therefore avoid using reduced documentation time as indirect evidence that the resulting record preserves patient voice, temporal context, or uncertainty.

The methodological implication is that source fidelity should be assessed independently of note-generation speed, completeness, readability, or clinician satisfaction. These outcomes may interact, yet they should remain analytically separable.

*Correction occurs inside a nursing work system*

Human review is frequently treated as the answer to generative-documentation error. In nursing practice, this assumption is incomplete because review occurs inside a work system with competing demands, interface constraints, interruptions, handovers, role expectations, and variable authority to amend documentation. Research on nursing EHR workarounds demonstrates that clinicians adapt when system design and clinical work do not align [26]. A workflow that depends on nurses to detect every source-attribution or temporal error can therefore create new invisible work, especially if provenance information is no longer visible in the generated draft.



Questions of correction also concern professional authority. Critical nursing scholarship on AI warns that technological efficiency can alter how professional work and judgment are organized [27]. In documentation, responsibility becomes especially problematic if an AI system performs a transformation that the nurse did not request, while the signed note places full evidentiary responsibility on the nurse. A meaningful correction process must consequently provide enough traceability for the reviewer to identify what requires checking.

Source-aware correction should address more than whether a sentence is linguistically accurate. The reviewer may need to ask whether the speaker has been preserved, whether two sources were silently combined, whether an earlier statement has been presented as current, whether uncertainty has been hardened into certainty, or whether an algorithmic inference has been expressed as an observation. These failure states are summarized in **Table 2**.

**Table 2.** Provenance, temporal, and uncertainty failure states in AI-drafted nursing documentation

| Failure state | Source material before generation | Transformation in the draft note | Information lost or altered | Clinically relevant interpretation risk | Required status before acceptance |
|---|---|---|---|---|---|
| Speaker substitution | A patient, relative, or clinician is the identifiable origin of a statement | The generated sentence removes or changes the speaker | Source identity | A relative's interpretation may appear to be the patient's report, or a clinician interpretation may appear to be patient speech | Restore attributable source or mark source as unresolved |
| Observation–inference collapse | A clinician observes a sign and separately interprets its possible meaning | Observation and interpretation are merged into one factual statement | Epistemic status | Readers cannot determine what was directly observed and what was inferred | Separate observation from interpretation |

| | | | | | |
|---|---|---|---|---|---|
| Temporal migration | A statement is clearly historical or linked to an earlier encounter | The generated note expresses it in an unqualified present tense | Event time/current applicability | Historical information may be acted on as though current | Restore temporal anchor and current-validity status |
| Uncertainty hardening | Source text uses tentative, conditional, approximate, or conflicting language | Generated text adopts definitive wording | Degree or type of uncertainty | Clinical ambiguity may disappear from the record | Reinstate uncertainty explicitly |
| Multi-source compression | Patient, family, and clinicians provide related but non-identical accounts | Generated note produces one synthesized statement | Perspective differences | Legitimate disagreement or contextual variation is erased | Preserve materially different sources and disagreements |
| Unsupported inferential completion | Available information suggests, but does not establish, a proposition | Model supplies a plausible connecting statement | Boundary between evidence and inference | Readers may treat algorithmic completion as witnessed fact | Identify inference and verify independently before factual use |
| Stale-information refresh | Earlier documentation contains a previously valid statement | Generated note restates it without historical qualification | Original timing and validity | Old goals, symptoms, or circumstances appear newly confirmed | Confirm present applicability or retain historical status |
| Correction without provenance retention | Reviewer edits inaccurate generated wording | Final note contains corrected text but no recoverable indication of the relevant source distinction | Auditability of correction | Future readers cannot determine why the correction was necessary or which source prevailed | Preserve the corrected source/temporal status in the clinical proposition |



Taken together, these issues define the problem that the next section must solve. AI-assisted nursing documentation can generate useful and efficient records while simultaneously changing the provenance, timing, or certainty of clinically meaningful propositions. The challenge is therefore to design a method in which those dimensions remain visible during generation, review, correction, and subsequent reuse.

*Methodological framework*

Source-Aware Nursing Documentation Methodology treats the clinically meaningful proposition, rather than the whole note, as the primary unit requiring provenance control. A generated note may contain dozens of sentences assembled from different evidentiary origins, and reliability at note level can conceal transformation at proposition level. Human evaluation frameworks for healthcare language models similarly indicate that assessment must extend beyond a single performance dimension and incorporate clinically meaningful judgments about generated output [28]. For nursing documentation, this means that each proposition capable of influencing care interpretation should retain enough information to determine its source, temporal relevance, epistemic status, uncertainty, and correction pathway.

*Attribution at proposition level*

The first rule is source retention. A proposition should remain attributable to the patient, a relative or caregiver, a clinician's direct observation, an earlier record, an external data source, or an algorithmic inference whenever that distinction changes its clinical interpretation. Where several sources contribute, the system should not erase a material disagreement simply to produce fluent prose. Attribution does not require cumbersome repetition of provenance labels in every sentence. The requirement is that provenance remain recoverable and visible when clinically relevant.

The second rule is observation–inference separation. A nurse may observe reduced oral intake, hear a relative describe increasing fatigue, and encounter previous documentation of progressive decline. An AI system may synthesize those inputs into an inference about deterioration. That inference can be clinically useful while

remaining categorically different from any individual observation. The methodology therefore resists converting synthesis into witnessed fact.

*Temporal anchoring and current validity*

Source identification is insufficient when the time represented by a proposition is unclear. The methodology separates event time, report or observation time, documentation time, and current-validity status. These dimensions may coincide, but they frequently do not in longitudinal palliative care.

A previously documented preference, for example, can be accurately reproduced while being inappropriate as an unqualified statement of present preference. Temporal anchoring should therefore preserve historical truth without silently converting it into current truth. When timing remains uncertain, the record should preserve that uncertainty instead of generating a false degree of temporal precision.

*Preserving uncertainty*

Generative systems can produce language that sounds more certain than the information supporting it. Probability-generation research demonstrates that model confidence and probabilistic outputs are not inherently reliable representations of clinical uncertainty [29]. Conversely, experimental work shows that AI systems can be designed to recognize situations in which available evidence is insufficient for confident inference [30]. The relevant principle for nursing documentation is not to create a new numerical uncertainty score. It is to prevent unresolved uncertainty from disappearing during generation.

Uncertainty may concern the source, timing, content, relationship among accounts, or present applicability of a proposition. Expressions such as “patient reports,” “family believes,” “appears,” “possibly,” “timing unclear,” or “current preference not reconfirmed” can contain clinically important epistemic information. Automated rewriting should not remove such qualifiers simply because a definitive sentence is stylistically cleaner.

*Verification and correction responsibility*

Explainable output alone does not establish trustworthy healthcare AI. Explanations can be persuasive without demonstrating that the underlying inference is clinically valid [31]. Verification should therefore test the generated proposition against its relevant source rather than relying only on a plausible model rationale.

Responsibility also belongs to the combined human–AI work system. Experimental evidence from a safety-critical nursing task demonstrates that AI support may improve or degrade human performance depending on the circumstances in which the recommendation is encountered [32]. A nurse should consequently not be treated as an unlimited final safety barrier for an opaque generative process. Effective correction requires access to the relevant source information, sufficient time to review it, a usable mechanism for amendment, and clear authority over the final record.

The proposed methodology therefore defines five linked controls for a clinically consequential generated proposition**:** source identification, temporal anchoring, observation–inference separation, uncertainty preservation, and traceable correction. These controls form a methodological specification for documentation design and evaluation, not a validated scoring instrument or mandatory clinical standard.



**Results and Discussion**

The methodological significance of source-aware documentation lies in what it prevents readers from assuming. Conventional clinical notes often encourage a reader to treat declarative prose as though its propositions occupy the same evidentiary level. Generative systems can intensify that tendency by producing stylistically consistent language from heterogeneous inputs. A patient’s own statement, a family interpretation, an earlier clinician’s note, a current nursing observation, and an algorithmic inference may therefore become grammatically indistinguishable even though they have different clinical meanings.

For palliative nursing, preserving those distinctions is particularly important because clinical knowledge is often relational and temporally unstable. Symptoms may fluctuate, preferences may change, family accounts may complement or conflict with the patient’s account, and previous documentation may remain clinically relevant without remaining current. Source awareness does not require declaring one contributor universally more authoritative than another. It requires preventing a transformation in which the identity and status of an important contributor disappear.

*Validation should test proposition fidelity*

Evaluation of AI-assisted documentation should therefore ask more than whether the note is readable, concise, complete, or judged generally accurate. A source-aware validation programme would examine whether clinically consequential propositions preserve the correct source, whether temporal relations remain interpretable, whether inferred content is distinguishable from observation, whether uncertainty is retained where the source material is uncertain, and whether reviewers can detect and correct transformations efficiently.

Such evaluation must also include realistic human interaction. Clinicians already report correcting attribution, context, certainty, and factual problems in AI-drafted notes [23]. Human evaluation frameworks reinforce the need for multidimensional assessment [28], while human–AI safety research demonstrates that system quality cannot be inferred solely from isolated model performance [32]. Validation should therefore examine the combined documentation system: model output, interface, reviewer behaviour, workflow constraints, correction process, and final signed record.

*Accountability cannot stop at the individual nurse*

Organizational governance is necessary because provenance failure is not always detectable from the final sentence alone. A recent review of healthcare AI governance frameworks found considerable variability in lifecycle coverage and comparatively limited operational oversight in many frameworks [33]. For AI-generated documentation, governance should clarify who evaluates the technology before deployment, what provenance information the interface exposes, how material errors are reported, who monitors recurrent transformation patterns, and how system changes are reassessed.

This allocation of responsibility matters ethically. Requiring the nurse to sign a generated note while withholding the source trace needed for efficient verification transfers accountability without equivalent control. Conversely, retaining extensive technical provenance that clinicians cannot practically interpret would not solve the problem. Source-aware design therefore requires proportional visibility: enough provenance for clinically meaningful verification without turning documentation into an unusable technical audit log.

*Applicability and limitations*

Generalist medical foundation models can perform an expanding range of clinical-language tasks, but broad capability increases the number of contexts in which task-specific validation remains necessary [34]. The proposed methodology should therefore not be interpreted as evidence that a particular commercial or experimental system is safe for palliative nursing documentation. It also does not establish the optimal interface, provenance data structure, or division of correction labour.

The current argument is constrained by the literature from which it was developed. Much direct evidence on ambient generative documentation comes from physician and ambulatory settings rather than palliative nursing. Research on serious-illness documentation establishes relevant semantic and temporal problems but largely predates widespread generative note production. Technical studies of uncertainty and temporal reasoning show what computational systems can and cannot represent, yet they do not validate the proposed documentation method in clinical nursing practice.



Empirical evaluation is consequently required. Future work should test whether proposition-level provenance can be preserved without recreating the documentation burden that generative tools are intended to reduce. Studies should examine usability for nurses, comprehensibility for patients and families, source-identification accuracy, handling of changing preferences, correction behaviour, effects of time pressure, and performance across languages, cultures, care settings, model versions, and documentation systems. Qualitative research will also be necessary because acceptable attribution may depend on relational context; for example, patients may differ in how they want family interpretations represented alongside their own accounts.

The safeguards and validation targets arising from these considerations are consolidated in **Table 3**.

**Table 3.** Safeguards and validation targets for source-aware AI-assisted nursing documentation

| Methodological safeguard | What must remain recoverable in the final documentation process | Validation question | Human responsibility | System or organizational responsibility | Failure that would invalidate the safeguard |
|---|---|---|---|---|---|
| Proposition-level source attribution | The materially relevant origin of a clinical statement | Can a reviewer determine whether the proposition came from the patient, a relative, a clinician, prior documentation, or algorithmic inference? | Confirm or correct clinically consequential attribution | Preserve usable provenance and expose it where verification is needed | Generated prose makes the source materially ambiguous |
| Temporal anchoring | Relevant event, report/observation, documentation, and current-validity relationships | Can a reader determine whether the proposition is current, historical, or temporally unresolved? | Reconfirm information when current applicability matters | Preserve source timestamps and avoid silent tense conversion | Historical information appears newly established |

| | | | | | |
|---|---|---|---|---|---|
| Observation–inference separation | What was directly reported or observed versus what was synthesized or inferred | Can an inferred proposition be distinguished from its supporting observations? | Judge whether inference is clinically acceptable | Label or trace transformations that introduce inference | Model-generated interpretation appears as witnessed fact |
| Uncertainty preservation | Material ambiguity, conditionality, disagreement, or missing information | Does generation retain clinically relevant uncertainty present in the source material? | Resolve uncertainty when possible; retain it when unresolved | Avoid automatic conversion to unwarranted certainty | Tentative information becomes definitive without verification |
| Source-visible correction | The information needed to identify and repair an attribution, timing, or inference error | Can the reviewer efficiently compare a questionable proposition with its relevant source? | Correct the clinical record when authorized and appropriate | Provide usable source access and correction functionality | Reviewer receives responsibility without practical verification access |
| Correction traceability | The clinically relevant provenance status after amendment | Does the corrected record accurately represent the source and temporal status that justified the correction? | Ensure the final proposition is clinically defensible | Maintain appropriate auditability without burdening routine readers | Correction changes wording but leaves provenance misleading |
| Lifecycle monitoring | Recurring provenance-related failure patterns across system use | Are systematic attribution, temporal, or uncertainty transformations detected after deployment? | Report clinically meaningful recurrent problems | Monitor, investigate, update, and revalidate the system | Known transformation patterns persist without accountable response |
| Use-case-specific validation | Performance under the intended patient population, workflow, setting, and note type | Has the complete human–AI documentation process been evaluated in the environment where it will be used? | Participate in realistic evaluation | Prevent unsupported transfer from unrelated tasks or settings | General model capability is treated as documentation validation |

A source-aware documentation system therefore requires correction to function as an accountable relationship among generated text, underlying evidence, the reviewing professional, and organizational oversight. **Figure 2** represents that relationship without assuming that human review can compensate for every upstream design failure.

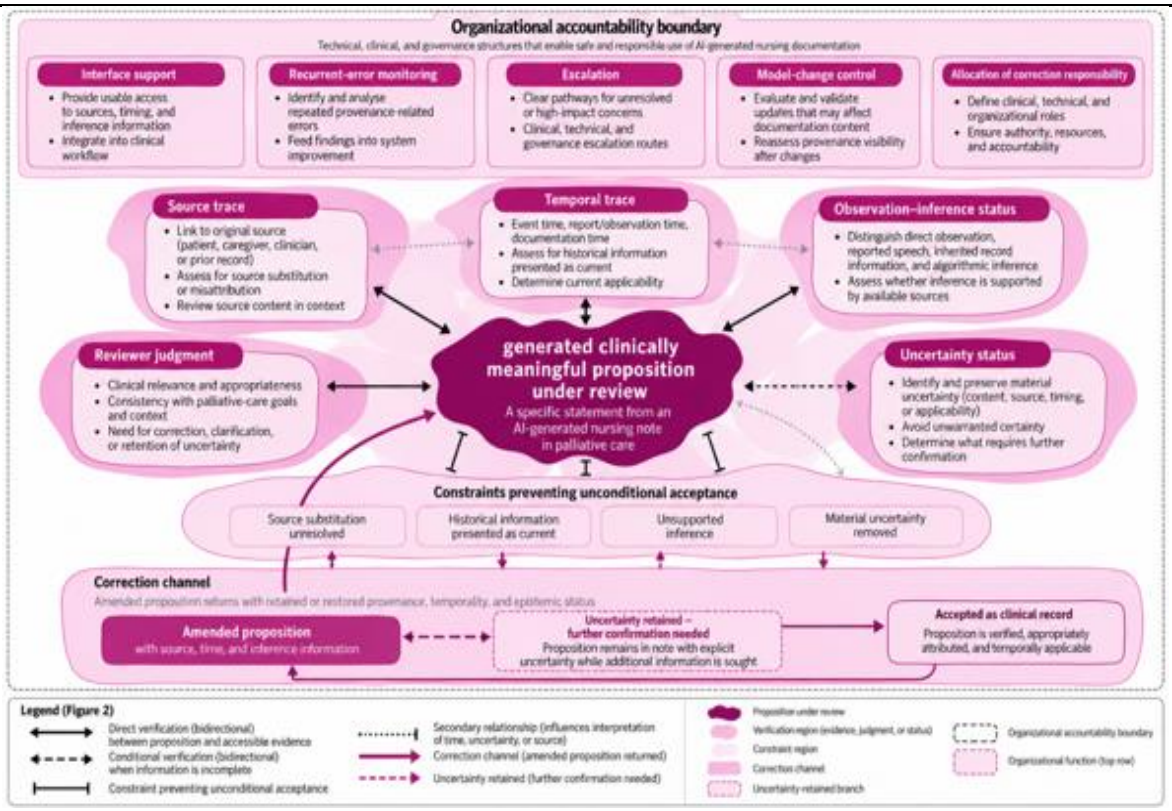



**Figure 2.** Verification and correction of provenance before an AI-generated nursing note becomes the clinical record

## Conclusion



An AI-generated nursing note should not be treated as though every sentence has a single, self-evident source. In palliative care, clinically meaningful propositions may originate in patient self-report, relatives' accounts, direct nursing observation, inherited documentation, or algorithmic inference, and their relevance may change over time. Source-Aware Nursing Documentation Methodology therefore places provenance inside the meaning of the recorded statement. Clinically consequential propositions should retain identifiable source, temporal context, separation of observation from inference, material uncertainty, and a workable correction pathway. These requirements do not reject generative documentation or its potential efficiency benefits. They define conditions under which efficiency can coexist with professional accountability and interpretable patient-centred records.

The methodology now requires empirical testing in real nursing workflows. Until such validation is available, its strongest contribution is a bounded principle: generated language should not become more certain, more current, or more authoritative than the information from which it was produced.

**Acknowledgments:** None

**Conflict of interest:** None

**Financial support:** None

**Ethics statement:** None

## References

1. Michalowski M, Topaz M, Peltonen LM. An AI-enabled nursing future with no documentation burden: A vision for a new reality. J Adv Nurs. 2026;82(1):907-12. doi:10.1111/jan.16911
2. Shah SJ, Devon-Sand A, Ma SP, Jeong Y, Crowell T, Smith M, et al. Ambient artificial intelligence scribes: Physician burnout and perspectives on usability and documentation burden. J Am Med Inform Assoc. 2025;32(2):375-80. doi:10.1093/jamia/ocae295

3. Van Veen D, Van Uden C, Blankemeier L, Delbrouck JB, Aali A, Bluethgen C, et al. Adapted large language models can outperform medical experts in clinical text summarization. Nat Med. 2024;30(4):1134-42. doi:10.1038/s41591-024-02855-5
4. de Vere Hunt IJ, Jin KX, Linos E. A framework for considering the use of generative AI for health. NPJ Digit Med. 2025;8(1):297. doi:10.1038/s41746-025-01695-y
5. Lee P, Bubeck S, Petro J. Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. N Engl J Med. 2023;388(13):1233-9. doi:10.1056/NEJMsr2214184
6. Angus DC, Khera R, Lieu T, Liu V, Ahmad FS, Anderson B, et al. AI, health, and health care today and tomorrow: The JAMA Summit report on artificial intelligence. JAMA. 2025;334(18):1650-64. doi:10.1001/jama.2025.18490
7. Lerbæk B, Kusk KH, Jørgensen L, Laugesen B. Navigating the complexities of nursing documentation when patients have access to the content: A qualitative study. J Adv Nurs. 2025;81(5):2604-16. doi:10.1111/jan.16502
8. Cooper AL, Albrecht MA, Kelly S, Eccles SP, Brown JA. A pre-post interventional study to reduce time spent on clinical documentation by nurses and midwives. J Adv Nurs. 2024;80(4):1452-63. doi:10.1111/jan.15931
9. Wang MD, Khanna R, Najafi N. Characterizing the source of text in electronic health record progress notes. JAMA Intern Med. 2017;177(8):1212-3. doi:10.1001/jamainternmed.2017.1548
10. Durieux BN, Zupanc SN, Tarbi EC, Manz CR, Lakin JR, Lindvall C. Content analysis of serious illness conversation documentation: Structured vs. free-text information. J Pain Symptom Manage. 2024;68(4):369-81.e4. doi:10.1016/j.jpainsymman.2024.06.025
11. Secunda K, Wirpsa MJ, Neely KJ, Szmuilowicz E, Wood GJ, Panozzo E, et al. Use and meaning of "goals of care" in the healthcare literature: A systematic review and qualitative discourse analysis. J Gen Intern Med. 2020;35(5):1559-66. doi:10.1007/s11606-019-05446-0
12. Modes ME, Engelberg RA, Downey L, Nielsen EL, Curtis JR, Kross EK. Did a goals-of-care discussion happen? Differences in the occurrence of goals-of-care discussions as reported by patients, clinicians, and in the electronic health record. J Pain Symptom Manage. 2019;57(2):251-9. doi:10.1016/j.jpainsymman.2018.10.507
13. Hachem Y, Lakin J, Teuteberg W, Cullinan A, Gonzales MJ, Lindvall C, et al. Electronic health record serious illness conversation dashboards: An implementation case series. J Pain Symptom Manage. 2025;69(2). doi:10.1016/j.jpainsymman.2024.10.032
14. Auriemma CL, Song A, Walsh L, Han JJ, Yapalater SR, Bain A, et al. Classification of documented goals of care among hospitalized patients with high mortality risk: A mixed-methods feasibility study. J Gen Intern Med. 2024;39(10):1839-49. doi:10.1007/s11606-024-08773-z
15. Koleck TA, Dreisbach C, Bourne PE, Bakken S. Natural language processing of symptoms documented in free-text narratives of electronic health records: A systematic review. J Am Med Inform Assoc. 2019;26(4):364-79. doi:10.1093/jamia/ocy173
16. Velupillai S, Suominen H, Liakata M, Roberts A, Shah AD, Morley K, et al. Using clinical natural language processing for health outcomes research: Overview and actionable suggestions for future advances. J Biomed Inform. 2018;88:11-9. doi:10.1016/j.jbi.2018.10.005
17. Knez T, Žitnik S. Multimodal learning for temporal relation extraction in clinical texts. J Am Med Inform Assoc. 2024;31(6):1380-7. doi:10.1093/jamia/ocae059
18. Olex AL, McInnes BT. Review of temporal reasoning in the clinical domain for timeline extraction: Where we are and where we need to be. J Biomed Inform. 2021;118:103784. doi:10.1016/j.jbi.2021.103784
19. Heinke A, Huang L, Simpkins KU, Kalaw FGP, Karsolia A, Singh K, et al. Dataset documentation for responsible AI: Analysis of suitability and usage for health datasets. NPJ Digit Med. 2026;9(1):547. doi:10.1038/s41746-026-02714-2
20. Lee RY, Brumback LC, Lober WB, Sibley J, Nielsen EL, Treece PD, et al. Identifying goals of care conversations in the electronic health record using natural language processing and machine learning. J Pain Symptom Manage. 2021;61(1):136-42.e2. doi:10.1016/j.jpainsymman.2020.08.024
21. Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. Nature. 2023;620(7972):172-80. doi:10.1038/s41586-023-06291-2
22. Peters SG, Boyum JP, Legler SR, Heise KJ, Griffin A, Heaton HA. Generative artificial intelligence for inpatient documentation summarization: Mixed-methods quality assessment and early real-world experience. J Am Med Inform Assoc. 2026;33(8):1466-73. doi:10.1093/jamia/ocag076
23. Guo Y, Hu D, Yang Z, Chow E, Tam S, Perret D, et al. Clinicians' rationale for editing ambient AI-drafted clinical notes: Persistent challenges and implications for improvement. J Am Med Inform Assoc. 2026;33(7):1345-53. doi:10.1093/jamia/ocag059

24. Stults CD, Deng S, Martinez MC, Wilcox J, Szwerinski N, Chen KH, et al. Evaluation of an ambient artificial intelligence documentation platform for clinicians. JAMA Netw Open. 2025;8(5). doi:10.1001/jamanetworkopen.2025.8614
25. Ma SP, Liang AS, Shah SJ, Smith M, Jeong Y, Devon-Sand A, et al. Ambient artificial intelligence scribes: Utilization and impact on documentation time. J Am Med Inform Assoc. 2025;32(2):381-5. doi:10.1093/jamia/ocae304
26. Nguyen OT, Carey A, Steege LM, Cho H. Characteristics and contributors to workarounds in the electronic health record system among inpatient nurses: Secondary analysis of a mixed-methods study. Int J Nurs Stud. 2026;173:105261. doi:10.1016/j.ijnurstu.2025.105261
27. Richards DA. Artificial intelligence and nursing: Freedom from the mundane or subservience to a new nobility. J Adv Nurs. 2026. doi:10.1111/jan.70726
28. Tam TYC, Sivarajkumar S, Kapoor S, Stolyar AV, Polanska K, McCarthy KR, et al. A framework for human evaluation of large language models in healthcare derived from literature review. NPJ Digit Med. 2024;7(1):258. doi:10.1038/s41746-024-01258-7
29. Gu B, Desai RJ, Lin KJ, Yang J. Probabilistic medical predictions of large language models. NPJ Digit Med. 2024;7(1):367. doi:10.1038/s41746-024-01366-4
30. Zhou S, Wang J, Xu Z, Wang S, Brauer D, Welton L, et al. Uncertainty-aware large language models for explainable disease diagnosis. NPJ Digit Med. 2025;8(1):690. doi:10.1038/s41746-025-02071-6
31. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9
32. Morey DA, Rayo MF, Woods DD. Empirically derived evaluation requirements for responsible deployments of AI in safety-critical settings. NPJ Digit Med. 2025;8(1):374. doi:10.1038/s41746-025-01784-y
33. Wang A, Freeman S, Magrabi F. Governance for safe and responsible AI in healthcare organisations: A scoping review of frameworks. NPJ Digit Med. 2026;9(1):516. doi:10.1038/s41746-026-02679-2
34. Moor M, Banerjee O, Abad ZSH, Krumholz HM, Leskovec J, Topol EJ, et al. Foundation models for generalist medical artificial intelligence. Nature. 2023;616(7956):259-65. doi:10.1038/s41586-023-05881-4