

B

Competency Standards in Clinical Artificial Intelligence

A draft for consideration by KMPDC and the Nursing Council of Kenya

Dr Neal Aggarwal

Written in a personal capacity

AUGUST 2026

Annex B to Another Arrow in the Quiver.

I write here in a personal capacity. This is a working draft offered for comment. It has no regulatory status.

1. Purpose

This document proposes a competency standard for the safe use of artificial intelligence by registered health professionals in Kenya. It is written so that either Council may adopt it whole, adopt it in part, or use it as a starting point for its own drafting.

It exists because of a specific gap. Registrants are already using these systems. No Council currently states what competent use looks like, which means no registrant can be held to a standard, no employer can specify one in a job description, and no training provider can be held to account for producing it.

The standard describes **what a registrant must be able to do**. It does not prescribe which tools they use, and it should survive the replacement of any particular system.

2. Scope and application

2.1 To whom this applies

Council	Cadres	Proposed application
KMPDC	Medical practitioners, dentists	All registrants in clinical practice
Nursing Council of Kenya	Nurses, midwives	All registrants in clinical practice

The standard is written to be cadre-neutral at Foundation level and cadre-specific above it. A ward nurse and a consultant surgeon should meet the same Foundation standard; what differs is what they then do with it.

2.2 What counts as “artificial intelligence” for this standard

Any computational system that generates clinical content, recommendations, classifications or predictions which a registrant may act upon, where the system’s output is not fully determined by an explicit, inspectable rule set.

This deliberately includes general-purpose conversational systems used informally on personal devices, which is where the majority of current unregulated use sits. It deliberately excludes deterministic calculators and rule-based alerts, which are already covered by existing standards of practice.

2.3 What this standard does not do

It does not authorise any use of AI that would otherwise be outside a registrant's scope of practice. It does not transfer any part of a registrant's accountability to a system or its supplier. It does not require any registrant to use AI.

3. Structure of the standard

Four **domains**, derived from the AI Fluency Framework of Dakan and Feller (see §10 for attribution and licensing).

Code	Domain	The question it answers
D1	Delegation	What may I hand to a system, and what must never leave my hands?
D2	Description	How do I state a clinical problem so that the answer is worth reading?
D3	Discernment	How do I judge what comes back, before it reaches a patient?
D4	Diligence	How do I remain accountable for the result?

Each domain contains numbered **performance criteria** at three tiers:

Tier	Expectation	Who
F Foundation	Every registrant in clinical practice	All
P Practitioner	Registrants who use AI in direct patient care	Most
A Advanced	Registrants who supervise, teach, procure or evaluate	Some

Criteria are cumulative: a Practitioner must also meet all Foundation criteria.

4. Domain D1 — Delegation

Statement of the domain. The registrant decides, before acting, which parts of a clinical task may appropriately be performed by an AI system and which may not, and can justify that boundary.

Foundation

D1.F1 Distinguishes the three modalities of use — *automation* (the system performs a defined task on instruction), *augmentation* (the registrant and system work a problem together), and *agency* (the system is configured to act on future cases, possibly for others) — and states which they are operating in before they act.

D1.F2 Maintains and can articulate a personal **non-delegable list**: clinical acts they will not hand to a system under any circumstances. At minimum this must include obtaining consent, breaking significant news, the decision to undertake an invasive procedure, the final diagnosis committed to the record, the prescription, and the signature.

D1.F3 Recognises that a task's suitability for delegation depends on the reversibility and severity of an error in it, not on the convenience of delegating it.

D1.F4 Declines to use AI where the registrant lacks the underlying clinical knowledge to evaluate the output.

Practitioner

D1.P1 Decomposes a clinical workflow into component tasks and assigns each to registrant, system, or joint work, with a stated reason.

D1.P2 Identifies, in their own service, tasks where AI use would introduce risk disproportionate to benefit, and does not use it there.

D1.P3 Recognises when a colleague or junior has delegated inappropriately, and intervenes.

Advanced

D1.A1 Constructs delegation boundaries for a whole service or department, documents them, and reviews them at defined intervals.

D1.A2 Distinguishes a system deployed in agency mode from one deployed in augmentation mode, and ensures the governance applied matches the modality actually in use.

5. Domain D2 — Description

Statement of the domain. The registrant can specify a clinical problem to an AI system with sufficient structure and local context that the response is clinically meaningful, without disclosing information they are not entitled to disclose.

Foundation

D2.F1 Does not enter identifiable patient data into any system the registrant or their employer does not control. This criterion is absolute and is assessed as a conjunctive requirement (see §7.4).

D2.F2 States clinical context adequate to the question: the level of facility, the investigations actually available, the formulary in use, and the relevant local epidemiology.

D2.F3 Structures a query as they would a clinical handover, using a recognised structure such as SBAR, rather than as a search term.

D2.F4 Requests reasoning, differentials, and the case against a proposition — not only a conclusion.

Practitioner

D2.P1 Iterates a query in response to an inadequate answer, and recognises when further iteration is not going to help.

D2.P2 Specifies the population to which a question refers, and does not accept guidance framed for a population materially different from the patient in front of them.

D2.P3 Where a history was taken in a language other than that of the query, states this, and accounts for what was lost in translation.

Advanced

D2.A1 Designs and documents standard query structures for a service, and evaluates whether they improve output quality.

D2.A2 Specifies the context requirements for a system being procured or configured for local use.

6. Domain D3 — Discernment

Statement of the domain. The registrant evaluates AI output critically and independently, detects clinically significant error, and does not allow the fluency or confidence of a response to substitute for its verification.

This is the domain on which the standard chiefly rests. A 2025 randomised controlled trial found that physicians who had completed twenty hours of AI-literacy training nonetheless deferred to deliberately erroneous model output. Familiarity with these systems does not produce the capacity to catch them being wrong. That capacity must be trained and assessed directly, and it is the reason this standard specifies error detection as a conjunctive requirement rather than a compensatable one.

Foundation

D3.F1 Forms and records an independent clinical impression before consulting an AI system, on any question where the registrant is capable of forming one. The order matters and is assessable.

D3.F2 Reads the reasoning offered before the conclusion offered.

D3.F3 Identifies output that is plausible, fluent and wrong, in a clinical vignette, at the threshold set for the assessment.

D3.F4 States that confidence of expression carries no information about accuracy, and demonstrates this in practice rather than only in recall.

D3.F5 Verifies any citation, dose, threshold or guideline reference against an authoritative source before acting on it.

Practitioner

D3.P1 Detects and corrects seeded clinical error in a simulated encounter, at or above the standard-set threshold.

D3.P2 Recognises the characteristic failure modes of these systems in the Kenyan context, including but not limited to:

1. under-weighting of locally prevalent conditions and over-weighting of conditions prevalent in the training corpus;
2. degraded performance on darker skin in dermatological and wound assessment;
3. recommendations presupposing investigations, medicines or referral pathways unavailable at the facility level in question;
4. drug nomenclature that resolves differently in the local supply chain;
5. degraded performance on paediatric, obstetric and geriatric presentations;
6. fabricated references presented in correct citation format.

D3.P3 Maintains calibrated trust: neither refuses assistance that would benefit the patient, nor accepts it without verification. Both directions of miscalibration are failures of this criterion.

D3.P4 Escalates appropriately when an AI-supported decision by a colleague appears wrong, irrespective of the seniority of that colleague.

Advanced

D3.A1 Designs error-detection exercises for others and can justify the seeded-error rate used.

D3.A2 Appraises published evidence on the performance of a clinical AI system, including its validation population, and states whether it transfers to the local case mix.

D3.A3 Monitors their own and their team's unassisted clinical performance for evidence of skill decay.

7. Domain D4 — Diligence

Statement of the domain. The registrant documents, discloses and takes responsibility for AI-assisted decisions, and complies with Kenyan law governing health data.

Foundation

D4.F1 Documents in the clinical record where AI materially informed a decision, and how.

D4.F2 States the requirements of the Data Protection Act 2019 and the Digital Health Act 2023 as they bear on the entry, storage and transfer of health data, and complies with them.

D4.F3 Discloses AI involvement to a patient where a reasonable patient would wish to know, and can conduct that conversation.

D4.F4 States plainly that accountability for the decision remains wholly with the registrant.

Practitioner

D4.P1 Reports, through the local incident system, any occasion on which AI use contributed or may have contributed to patient harm or a near miss.

D4.P2 Takes deliberate steps to maintain the clinical skills that AI use might otherwise erode.

D4.P3 Declares any personal or financial interest in a system they use, recommend or evaluate.

Advanced

D4.A1 Establishes local governance for AI use: documentation standards, incident routes, and periodic audit.

D4.A2 Conducts or commissions audit of AI-assisted practice and acts on the findings.

D4.A3 Ensures that procurement decisions are informed by evaluation evidence relevant to the local population.

8. Assessment specification

8.1 Principle

Competence is inferred from performance, not from attendance. No award should be made on the basis of having been present.

8.2 Acceptable evidence, by tier

Tier	Minimum acceptable evidence
Foundation	Invigilated knowledge assessment from a blueprinted item bank, plus one observed simulated encounter
Practitioner	The above, plus a simulation-based performance assessment with seeded error, plus a workplace logbook of not fewer than 20 supervised encounters countersigned by a supervisor of appropriate seniority
Advanced	The above, plus a substantial piece of work — an evaluation, an audit, a taught programme — assessed against published criteria

8.3 Standard setting

Pass standards should be set by a formal method — modified Angoff for knowledge assessment, borderline regression for performance assessment — by a panel of practising registrants, and published. A pass mark set at an arbitrary percentage, or a pass rate decided in advance, is not acceptable.

8.4 Conjunctive requirements

Two criteria may **not** be compensated by strong performance elsewhere:

- **D2.F1** — entry of identifiable patient data into an uncontrolled system;
- **D3.P1** — detection of seeded clinical error at the standard-set threshold.

A candidate failing either fails the assessment, whatever their aggregate score. The precedent is established: a candidate does not pass a conventional OSCE by compensating for a fatal drug error with excellent communication.

8.5 Psychometric obligations on providers

Any provider accredited to assess against this standard should be required to conduct item analysis after every sitting, withdraw items with negative discrimination and recalculate affected candidates' scores, and publish reliability for each instrument. Providers unable to do this should not be accredited to assess, whatever their competence to teach.

8.6 Currency

Competence in this domain decays and the technology changes. **Recertification at not more than two years** is proposed, and recertification should be assessed, not attested.

9. Mapping to continuing professional development

9.1 KMPDC

KMPDC requires 50 CPD points per calendar year for retention on the register. Indicative mapping:

Activity	Proposed points
Foundation certificate (12 contact hours + assessment)	12
Practitioner certificate (36–40 hours + logbook)	40
Advanced certificate (60 hours + project)	60
Certified instructor track	40
Recertification at two years	10

Accreditation as a CPD provider is by application to the Council, with a non-refundable application fee (KES 15,000 at the time of writing), disclosure of sponsorship, a calendar of activities, two referees, and inspection of premises, activities, facilitators and resources. Providers receive a unique identification number which must appear on all activity documentation.

9.2 Nursing Council of Kenya

The Council operates a points requirement for licence renewal, administered through its online services portal, with provider registration via the same portal.

A point on which I could not obtain a reliable figure. Published secondary sources disagree on the annual points requirement for nurses — some state 20, others 40. I have not been able to resolve this from primary Council documentation and have therefore not proposed a mapping. **This must be confirmed directly with the Council before any mapping is adopted.** I would rather leave the cell blank than fill it with a number I cannot stand behind.

10. Proposed regulatory wording

Offered as drafting material, not as settled text.

Use of artificial intelligence in clinical practice

1. A registrant who uses an artificial intelligence system in the course of clinical practice remains wholly accountable for any decision taken, and may not delegate that accountability to the system, its supplier, or any other person.
2. A registrant shall not enter information by which a patient may be identified into any system which the registrant or the registrant's employer does not control.
3. A registrant shall record in the clinical record where an artificial intelligence system has materially informed a clinical decision, and the nature of that contribution.
4. A registrant shall not act upon the output of an artificial intelligence system in a matter falling outside the registrant's competence to evaluate that output.
5. A registrant shall report, through the relevant incident reporting system, any occasion on which the use of an artificial intelligence system contributed, or may have contributed, to patient harm.
6. The Council may specify competencies in the use of artificial intelligence, and may require evidence of them as a condition of retention on the register.

Clause 6 is the operative one. The remainder can be adopted immediately as standards of practice; clause 6 creates the hook for competency requirements once training capacity exists to meet them.

11. Implementation and transition

Imposing a competency requirement before training capacity exists creates an unmeetable obligation and brings the standard into disrepute. A phased approach:

Phase	Timing	Status of the standard
1	On adoption	Clauses 1–5 in force as standards of practice. Competencies published as guidance
2	+12 months	Foundation competencies expected of registrants using AI in patient care; provider accreditation open
3	+24 months	Foundation competence required for new registrants; CPD credit available at all levels
4	+36 months	Practitioner competence required of registrants using AI in direct patient care in accredited facilities

Phase 4 should not commence until the Councils are satisfied that accredited providers can meet demand — including at county facilities, and not only in Nairobi.

12. Review

This standard should be reviewed **not less often than every eighteen months**. A standard in this field that has not been reviewed in two years should be presumed out of date.

Review should be informed by published outcome data from accredited providers, including error-catch rates and any incidents in which a certificated registrant's use of AI contributed to harm.

13. Attribution and licensing

The four domains adapt the **AI Fluency Framework** of Prof. Rick Dakan (Ringling College of Art and Design) and Prof. Joseph Feller (Cork University Business School, University College Cork), elaborated into an open course series in partnership with Anthropic PBC with support from Ireland's Higher Education Authority.

Two distinct licences apply and they are not interchangeable:

- the open **course materials** are released under Creative Commons **BY-NC-SA 4.0**, which permits adaptation provided the authors are credited and the derivative carries the same licence;
- the authors' *Framework for AI Fluency (Practical Summary Document)*, v1.1, is released under Creative Commons **BY-NC-ND 4.0**, which permits redistribution but **not** adaptation.

This draft adapts from the former and cites the latter. Any Council adopting this standard should take its own view on the licensing position of the resulting instrument, and should note that a regulatory standard derived from BY-NC-SA material may carry share-alike obligations.

Suggested citation for the framework: Dakan, Rick and Feller, Joseph. "Framework for AI Fluency (Practical Summary Document)," Version 1.1, Ringling.edu/ai, 2025.

14. Note on the evidence base

The single most important source behind this standard is:

Automation Bias in Large Language Model-Assisted Diagnostic Reasoning among Physicians Trained in AI Literacy — A Randomized Clinical Trial, NEJM AI, 2025.

Its finding — that prior AI-literacy training did not prevent deference to erroneous output — is the reason this standard makes error detection conjunctive rather than compensatable, and the reason it specifies the independent-impression rule as an assessable behaviour rather than an exhortation.

Councils considering adoption should read it directly rather than relying on my summary of it.

Prepared by Dr Neal Aggarwal, August 2026, in a personal capacity, as a draft for comment. Drafted with AI assistance; the judgements and any errors are mine.