SNOMED INTERNATIONAL

AI Position Statement

A Principled Approach to Artificial Intelligence in Clinical Terminology

Last Updated: September 2026

About this document

This position statement sets out SNOMED International's considered view on the responsible integration of artificial intelligence within clinical terminology and health information systems. It draws on documented implementations across 30+ health systems in 12 countries, and on SNOMED International's internal AI Strategy.

It does not constitute clinical or regulatory guidance.

EXECUTIVE SUMMARY

SNOMED International recognizes the growing adoption of artificial intelligence across global health systems. While this development carries meaningful potential, it also introduces risks that demand deliberate governance. This position statement sets out our considered view: AI technologies should be adopted cautiously, incrementally, and always within the constraints of validated clinical terminology standards.

Critically, SNOMED International does not position itself as a passive observer. We are actively investing in five structured AI initiatives — to support member implementation, improve clinical coding, enhance end user tools, embed SNOMED CT within AI technologies, and strengthen our own product quality assurance. In each case, the governing principle is the same: AI is a powerful set of technologies, and SNOMED CT is the guardrail that can make it safer in a healthcare context.

30+
Documented global implementations

12+
Countries with active deployments

THE PROBLEM

AI alone is insufficient for clinical environments

Large Language Models and probabilistic AI systems are designed to predict statistically likely outputs — not to guarantee clinical truth. In healthcare, the consequences of error are not abstract: a misclassified diagnosis, a missed contraindication, or a hallucinated drug code can directly harm patients.

Evidence from implementations across 12 countries confirms this risk. Unguided AI systems have produced nonsensical terminology mappings, invented clinical codes that do not exist, and assigned conditions to patients based on family history or negated findings. These are predictable failure modes of probabilistic systems operating in domains that require precision.

DOCUMENTED FAILURE — SINGAPORE, NATIONAL HEALTHCARE GROUP

When tasked with mapping ‘chronic kidney disease’ to a clinical terminology, ChatGPT initially produced nonsensical, clinically invalid terms. Only after a SNOMED CT Retrieval-Augmented Generation architecture was applied did accuracy rise to 80.8% — illustrating that general-purpose LLMs cannot be trusted in clinical settings without verified terminology constraints.

THE ACCURACY GAP — SOUTH KOREA, KAKAO HEALTHCARE

General LLM prompts achieved only a 39.8% match rate for Korean clinical terms. Routing GPT-4 through a SNOMED RAG pipeline raised exact match accuracy to 96.2%, a gap of nearly 57 percentage points that represents unacceptable risk in any clinical workflow.

OUR POSITION

SNOMED International's 6 principles for safe AI integration

Principle 1 **Determinism over probability** Clinical decisions must be grounded in verified, deterministic knowledge structures. SNOMED CT's formal logic, not statistical inference, provides this foundation.	Principle 2 **Clinician authority is non-negotiable** AI outputs should augment clinical judgement, not replace it. Every recommendation must remain auditable and overridable by a qualified clinician.	Principle 3 **Explainability as a prerequisite** Black-box AI has no place in clinical decision support. SNOMED CT's knowledge graph leaves an auditable trail for every AI decision.
Principle 4 **Incremental validation before deployment** No AI application should enter clinical workflows without prospective validation against SNOMED CT concepts and real-world outcome data.	Principle 5 **Hallucination is not acceptable** RAG architectures must constrain AI outputs to concepts that exist in verified SNOMED databases. Invented codes must be treated as system failures.	Principle 6 **Interoperability as a safety mechanism** SNOMED CT's multilingual, cross-institutional standardization ensures AI outputs are interpretable across care settings, a requirement, not a feature.

EVIDENCE BASE

Real-world deployments: where constrained AI has proven its value

The case for SNOMED CT as an AI trust layer is not theoretical. The following examples — drawn from hospitals, research institutions, and healthcare platforms across the globe — demonstrate both the power of constrained AI and the risks of deploying it without verified terminology governance.

Predictive and diagnostic AI

EARLY WARNING

Sepsis detection

NHS England deployed AI to monitor health records for early sepsis signs. SNOMED CT standardized varied clinical terms such as 'fever' and 'tachycardia', preventing dangerous ambiguity in automated alerts and enabling life-saving intervention at scale.

NHS England, UK

RISK PREDICTION

Diabetic foot — amputation risk

14 million clinical notes were mapped to SNOMED clusters at Barts Health NHS Trust. A SNOMED Graph Processor grouped 90 codes into meaningful features, enabling transparent and auditable prediction of amputation risk in diabetic foot disease.

Barts Health NHS Trust, UK

PATIENT TIMELINE

Digital twin forecasting

King's College and South London Hospitals trained a GPT model on SNOMED-coded patient histories to act as a 'digital twin' — safely forecasting the next likely diagnosis, procedure, or drug without relying on unstructured clinical text.

King's College & South London Hospitals, UK

CARDIOLOGY

Heart failure under-diagnosis

At University Hospital Aachen, AI mined unstructured discharge letters to extract HFpEF risk parameters. SNOMED standardization enabled early-detection cluster analysis to identify patients at risk of end-stage hospitalisation.

University Hospital Aachen (DARIO project), Germany

Data extraction and automation

CLINICAL CODING

74% automated admission coding

Hospital Clínic de Barcelona deployed NLP with SVM to generate structured problem lists at admission. By limiting AI output to valid SNOMED nodes and handling negations correctly, the hospital achieved a 74.2% automatic coding rate.

Hospital Clínic de Barcelona, Spain

REVENUE RECOVERY

£5.5M recovered via SNOMED NLP

MedCAT, a SNOMED-based NLP engine, structured 22 million records at Guy's and St Thomas' NHS Trust, automating outpatient coding and flagging armed forces veterans for benefits entitlement, recovering £5.5 million in missed revenue.

Guy's and St Thomas' NHS Trust — CogStack, UK

PATHOLOGY

300,000 pathology reports labelled

SVM and LSTM AI models assigned SNOMED codes to 300,000 unstructured pathology reports in Italy. An explainability module (LIME) confirmed the AI genuinely distinguished benign from malignant lesions — not just pattern-matched text.

Italian Digital Pathology Network

RARE DISEASE

48% more rare disease data recovered

AI extracted rare neurodevelopmental disease features from Spanish clinical text. Using SNOMED CT to bridge to the Human Phenotype Ontology recovered 48% more data compared to unstructured extraction — without hallucinating false positives.

Hospital Clínic de Barcelona & Bitac, Spain

LLM guardrails and platform governance

ALLERGY SAFETY

Anaphylaxis risk alerts

Allergy errors cause 10% of fatal adverse drug events. An agentic AI system at Hospital Sírio-Libanês extracts allergens, grounds them in a SNOMED vector store, and structures data into FHIR resources — enabling deterministic decision-support alerts at point of care.

Hospital Sírio-Libanês, Brazil

CLINICAL TRIAL MATCHING

Subsumption-powered trial eligibility

At UCLH, trial criteria were annotated with SNOMED CT. Because SNOMED understands subsumption, a search for 'lung cancer' automatically linked to 'non-small cell lung cancer' — without AI inference or manual synonym management.

University College London Hospitals, UK

KNOWLEDGE AT SCALE

80 billion clinical objects standardized

Babylon Health uses SNOMED CT to standardize 80 billion objects from diverse clinical sources, forming the foundation of a unified Knowledge Graph used to power proactive AI care interventions across its global patient population.

Babylon Health, Global

SAFE CONTENT GENERATION

Preventing toxic medical advice

Elsevier's Clinical Key AI uses a strict RAG architecture that prevents the AI from drawing on general internet data. It can only retrieve facts from trusted medical journals and SNOMED-backed taxonomies — with human oversight at every stage.

Elsevier — Clinical Key AI, Global

SNOMED INTERNATIONAL IN PRACTICE

How we are applying these principles to SNOMED CT

These principles are not simply positions we hold about others' use of AI. SNOMED International has committed to five structured AI initiatives of its own, each designed to demonstrate the responsible application of AI in service of the terminology and its members. Governance may sit with a proposed dedicated AI Advisory Group, comprising internal stakeholders, member representatives, and external AI specialists.

INITIATIVE 1

AI-supported implementation

Using AI to lower the cost and complexity of SNOMED CT implementation, covering translation into low-resource languages, local extensions, reference set construction, and legacy code mapping.

In practice: The Netherlands' NICTIZ used Gemini LLM to translate 15,000 SNOMED concepts into plain patient-friendly language — with outputs manually validated and tethered

INITIATIVE 2

Machine-assisted clinical coding

SNOMED is driving improvements in Entity Linking, the automatic mapping of unstructured clinical text to SNOMED concepts, through the Entity Linking Benchmark, and real-world field tests with healthcare providers.

In practice: Barcelona Supercomputing Center used deep learning to extract 14,000+ surgical procedures from case reports, using SNOMED's

to exact SNOMED concept IDs to prevent unsafe advice.

Cost: $150k–$300k Risk: Medium Member benefit: High

post-coordination grammar to assemble anatomical parts, actions, and tools into accurate codes without losing clinical detail.

INITIATIVE 3

Enhanced end user experience

LLM-powered interfaces will be demonstrated and integrated into Snowstorm, and Bahmni, along with others, enabling natural language querying, improved concept search, and intelligent post-coordination suggestions for clinicians.

In practice: CSIRO's OntoGPT already translates plain English clinical questions into SNOMED Expression Constraint Language (ECL), turning natural language into deterministic, executable queries that return verified facts for cohort building.

INITIATIVE 4

Embedding SNOMED CT in AI models

SNOMED will partner with clinical LLM researchers to demonstrate how RAG architectures and production guardrails constrain AI outputs to verified terminology, and publish thought leadership to position SNOMED CT as the essential trust layer.

In practice: 3M Health Information Systems replaced 11,000 hard-coded text strings with SNOMED CT mappings, using post-coordination to dynamically combine concepts like 'chest pain', 'dyspnea', and 'acute' into single universally understood atomic concepts.

INITIATIVE 5

AI-assisted modelling and quality assurance

Machine learning models from academic literature will be re-implemented and run against SNOMED CT to support terminologists in identifying QA issues, such as unusual fully specified names or missing relationships. This directly supports SNOMED's ability to manage an increasingly complex terminology without proportional cost increases to members, as well as ensuring that SNOMED CT is the high quality product needed by AI technologies to fulfill their potential.

In practice: Severance Hospital in South Korea deployed an agentic AI that queried SNOMED CT to resolve ambiguous nursing actions, locking terms like 'raise bedside rail' and 'side rail up' into a single deterministic concept, eliminating inconsistency across ICU records.

RECOMMENDATIONS

What SNOMED International calls for

FOR HEALTH SYSTEMS

Require SNOMED CT compliance as a procurement condition for any AI clinical decision support system. The evidence from NHS England, Barts Health, and UCLH shows that constrained AI is deployable, and that unconstrained AI is not safe.

FOR AI DEVELOPERS

Adopt SNOMED CT as the semantic validation layer for clinical NLP outputs. The Kakao Healthcare and Singapore NHG cases demonstrate that RAG architectures anchored to SNOMED CT can close accuracy gaps that would otherwise be clinically unacceptable.

FOR REGULATORS	FOR THE RESEARCH COMMUNITY
Mandate explainability and terminology traceability for AI systems used in diagnosis, coding, or prescribing. Hallucination events, such as those documented in Singapore, should be reportable as clinical near-misses.	Prioritise prospective studies measuring patient outcomes, not just coding accuracy, when AI operates with and without SNOMED CT constraints. Engage with SNOMED's Entity Linking Benchmark to advance the field on shared, validated datasets.

CONCLUSION

The deployment of AI in clinical settings is not a question of whether, but of how. The implementations documented in this paper, from sepsis detection in NHS England to rare disease phenotyping in Barcelona, from anaphylaxis alerts in Brazil to surgical code mapping in Switzerland, share a common architecture: AI generates; SNOMED CT governs.

SNOMED International holds itself to the same standard. Our five internal AI initiatives are each built on the principle that AI accelerates, but verified terminology is what makes acceleration safe. Without that anchor, speed becomes a risk.

The goal is not to slow progress. It is to ensure that as AI is placed in clinical environments, it operates within boundaries that protect patients. SNOMED CT provides those boundaries. Their use should not be optional, including within our own operations.

SNOMED International

snomed.org | September 2026