

About Open Clinical History

Overview, purpose, scope, safety, limitations and roadmap for Open Clinical History.

- [Clinical Safety & Limitations](#)
- [How Open Clinical History Works](#)
- [What is Open Clinical History?](#)
- [Why Open Clinical History Exists](#)

Clinical Safety & Limitations

Open Clinical History is a working proof of concept designed to organise fragmented clinical information into a structured longitudinal history.

It is not intended to replace the original medical record, clinical judgement, diagnosis, treatment decisions or direct review of source documentation.

The platform uses automated processing to extract, reconcile and classify information from clinical documents. Those processes can be useful, but they can also be wrong, incomplete or uncertain. Clinical information presented by Open Clinical History should therefore be treated as an aid to understanding the record, not as an unquestionable clinical fact.

Source documents remain authoritative

Open Clinical History creates structured clinical events from source material such as specialist letters, discharge summaries, pathology reports and other clinical documents.

The structured event is a derived representation of that source.

Where there is any discrepancy between an Open Clinical History event and the original clinical document, the original source document must be treated as the authoritative record.

For this reason, provenance and traceability are fundamental design principles of the platform.

Where possible, users should be able to understand:

- which document an event came from;
- what information supported the extracted event;
- when the source document was created;
- how the event was classified;
- and whether the result requires further review.

Automated extraction can be imperfect

Clinical documents are complex.

They may contain:

- historical information;
- current diagnoses;
- ruled-out diagnoses;
- family history;
- copied-forward information;
- uncertain or provisional findings;
- references to events that occurred years earlier;
- medication lists that are no longer current;
- and dates that relate to different parts of the clinical story.

An automated system may incorrectly interpret one of these as a current clinical event.

It may also:

- miss an important event;
- combine separate events;
- split one event incorrectly;
- assign the wrong date;
- misinterpret negation or uncertainty;
- misunderstand who a statement refers to;
- or lose important clinical context.

Open Clinical History therefore treats automated extraction as part of a processing pipeline that should remain reviewable and auditable.

SNOMED CT classification has limits

Open Clinical History uses SNOMED CT as its clinical terminology foundation.

A clinical phrase may map cleanly to a single SNOMED CT concept, but many real-world descriptions are ambiguous.

Classification can be affected by:

- incomplete clinical wording;
- abbreviations;
- local terminology;
- historical terminology;
- multiple possible SNOMED CT concepts;
- insufficient context;
- or differences between the wording in a document and the terminology model.

A technically valid SNOMED CT concept is not automatically the correct clinical interpretation.

Where classification is uncertain, the system should favour transparency and review rather than creating false confidence.

Anatomical mapping can be approximate

Open Clinical History can associate clinical events with anatomical structures and visual layers.

These mappings are intended to support understanding of the history.

They should not be interpreted as precise radiological, surgical or diagnostic localisation unless the underlying source information supports that level of detail.

A general anatomical relationship may sometimes be more appropriate than a highly specific one.

A timeline does not contain the whole clinical context

A longitudinal timeline can make a complex history much easier to understand, but reducing a clinical record to events inevitably removes some of the narrative context contained in the original documents.

Important information may exist in:

- surrounding clinical discussion;
- differential diagnoses;

- correspondence between clinicians;
- examination findings;
- risk discussions;
- treatment rationale;
- or information that was not selected as a discrete event.

The timeline should therefore provide a path into the original source material rather than becoming a replacement for it.

Missing documents create an incomplete history

Open Clinical History can only process the information it receives.

If source documents are missing, the longitudinal history may also be incomplete.

The absence of an event in Open Clinical History does not prove that the event did not occur.

Similarly, an incomplete set of records can create gaps that may make later events appear disconnected from their actual clinical context.

Dates require careful interpretation

Clinical documents often contain several different dates.

For example, a document may have:

- a document creation date;
- an admission date;
- a procedure date;
- a historical diagnosis date;
- a specimen date;
- a result date;
- or a reference to an event that happened years earlier.

Open Clinical History attempts to place events in the correct temporal context, but automated date interpretation can be uncertain.

Users should verify clinically important dates against the source record.

Artificial intelligence is not the clinical authority

AI-assisted processing can help identify and structure information contained in large volumes of unstructured clinical material.

It must not be treated as the authority on what happened clinically.

The authority remains the underlying clinical evidence and the healthcare professionals responsible for interpreting that evidence.

Open Clinical History should therefore be designed so that AI-generated or AI-assisted results are:

- inspectable;
- traceable;
- reviewable;
- correctable;
- and distinguishable from the original source information.

Clinical review remains essential

Open Clinical History is intended to reduce the effort required to understand a fragmented history.

It does not remove the need for clinical review.

For information that could influence diagnosis, treatment, medication, surgery or other clinical decisions, the relevant source records should be reviewed by an appropriately qualified healthcare professional.

Privacy and sensitive information

Clinical histories contain highly sensitive personal information.

Any deployment of Open Clinical History must consider appropriate controls for:

- authentication;

- authorisation;
- secure transmission;
- data storage;
- backups;
- logging;
- access auditing;
- document retention;
- and protection of API credentials or other secrets.

Real patient information should only be stored or processed in environments that are appropriate for that information and the obligations of the organisation operating the system.

Proof-of-concept status

Open Clinical History currently demonstrates a technical approach to:

- ingesting fragmented health information;
- extracting clinical events;
- applying terminology classification;
- associating events with anatomical structures;
- preserving source provenance;
- and presenting a longitudinal clinical history.

Further work is required before the platform should be considered suitable for use as part of a production clinical workflow.

Areas requiring ongoing validation include clinical accuracy, terminology governance, security, privacy, interoperability, usability and formal safety controls.

Safety principle

The core safety principle of Open Clinical History is:

Make the history easier to understand without hiding uncertainty or disconnecting information from its source.

Structured information should help the user navigate the clinical record.

It should never make the original evidence harder to find, imply certainty where certainty does not exist, or replace appropriate clinical judgement.

How Open Clinical History Works

Open Clinical History converts fragmented clinical documents into a structured longitudinal history through a staged ingestion, extraction, classification and presentation process.

The system is designed so that the original source material remains important throughout the process. Structured clinical events are created to make the history easier to understand, but those events remain linked back to the documents from which they were derived.

1. Clinical documents are ingested

The process starts with raw clinical material.

This can include documents such as:

- specialist letters;
- discharge summaries;
- pathology reports;
- imaging reports;
- hospital correspondence;
- scanned historical records;
- other unstructured clinical documents.

The source document is stored and registered for processing.

Rather than assuming that one document represents one clinical event, Open Clinical History treats the document as a source that may contain many separate clinically meaningful events.

2. The document is processed into clinical events

The ingestion pipeline analyses the document and identifies discrete clinical events.

Examples may include:

- diagnoses;
- procedures;
- investigations;
- admissions;
- discharges;
- medication changes;
- complications;
- follow-up events;
- significant clinical findings.

Each event is treated as an individual part of the patient's longitudinal history.

Where possible, the system also identifies relevant dates, clinical context and supporting evidence from the source document.

3. Events are reconciled and normalised

Clinical documents are rarely written in a perfectly structured way.

The same event may be described in different sections, dates may be expressed indirectly, and important context may be spread across several sentences.

Open Clinical History therefore performs reconciliation before committing extracted information into the longitudinal model.

The aim is to reduce duplication, preserve the intended clinical meaning and produce a cleaner representation of the patient's history.

4. Clinical concepts are classified using SNOMED CT

Once an event has been identified, the system attempts to associate it with an appropriate SNOMED CT concept.

Open Clinical History maintains a local SNOMED CT terminology database that is used to validate terminology and concept relationships.

This means the classification process is not simply accepting an arbitrary text label. The selected concept can be checked against the local terminology foundation.

The original clinical wording is still important and is retained alongside structured terminology where appropriate.

5. Anatomical relationships are identified

Where a clinical event relates to a body structure, Open Clinical History can associate the event with an anatomical concept.

The system maintains an anatomical catalogue and mappings between anatomical image layers and relevant clinical terminology concepts.

This allows clinical events to be connected to a visual representation of the body.

For example, an event involving the brain, heart, spine or another anatomical structure can contribute to the corresponding visual layer in the patient view.

6. Provenance is retained

A structured clinical history is only useful if users can understand where the information came from.

Open Clinical History therefore retains source provenance for extracted events.

The user can move from the longitudinal representation back towards the supporting clinical record rather than being presented with an unexplained AI-generated summary.

This traceability is an important design principle of the platform.

7. The longitudinal history is built

Once events have been extracted, reconciled and classified, they can be placed into chronological context.

Instead of displaying only a list of source documents, Open Clinical History presents the clinical events contained within those documents as a longitudinal history.

This allows the user to see:

- what happened;
- when it happened;
- how events relate over time;
- which anatomical areas were involved;
- which clinical concepts were identified;
- and which source material supports each event.

8. The patient history is visualised

The resulting history is presented through a combination of timeline information, event details and anatomical visualisation.

The intent is to allow a user to first understand the broad clinical story and then progressively drill into the detail.

A clinician, patient or carer should be able to move from:

the overall history → an individual event → the supporting source

without needing to reconstruct the entire history manually from a collection of unrelated documents.

The key principle

Open Clinical History does not attempt to replace the original clinical record.

It creates a structured layer over fragmented source information so that the patient's history can be understood more quickly, while retaining the source material needed to verify and interpret that history.

The overall flow is:

Clinical documents → ingestion → event extraction → reconciliation → SNOMED CT classification → anatomical association → longitudinal history → visualisation

What is Open Clinical History?

Open Clinical History is a working proof-of-concept platform for turning fragmented, unstructured health information into a coherent longitudinal clinical history.

Health information is often spread across scanned letters, discharge summaries, pathology reports, specialist correspondence and records from different hospitals or clinical systems. Even when the information exists, it can be difficult for a clinician or patient to quickly understand the complete history.

Open Clinical History is designed to take those source documents and transform them into a structured, navigable clinical timeline while preserving the relationship back to the original source material.

The platform combines document ingestion, clinical event extraction, terminology classification and anatomical visualisation to make a patient's history easier to understand.

What the platform does

At a high level, Open Clinical History:

- ingests raw and unstructured clinical documents;
- identifies individual clinical events from those documents;
- records dates, event types and supporting source information;
- classifies clinical concepts using SNOMED CT;
- associates relevant events with anatomical structures where possible;
- presents the resulting information as a longitudinal clinical history;
- provides anatomical visualisation to help show where relevant events occurred;
- retains provenance so extracted information can be traced back to its source.

The goal is not simply to summarise documents. It is to create a structured clinical history from information that was never originally designed to work together.

A longitudinal view instead of a document list

Traditional record repositories often present health information as a collection of documents.

A clinician may need to open multiple PDFs, letters or reports and reconstruct the patient's story manually.

Open Clinical History instead focuses on the clinical events contained inside those documents.

A discharge summary, specialist letter or pathology report may contain several clinically meaningful events. These can be extracted, dated, classified and positioned within the patient's history.

The result is a timeline that can show how the patient's health has changed over months, years or decades.

SNOMED CT classification

Open Clinical History uses SNOMED CT as its clinical terminology foundation.

Rather than relying only on free-text labels, the platform attempts to associate clinical events with recognised SNOMED CT concepts.

This provides a consistent terminology layer that can support:

- clinical classification;
- terminology searching;
- anatomical relationships;
- future interoperability;
- more reliable analysis of longitudinal health information.

The platform maintains a local SNOMED CT terminology database and uses that local terminology as the authoritative source when validating classifications.

Anatomical visualisation

Where a clinical event relates to a part of the body, Open Clinical History can associate that event with an anatomical structure.

The platform maintains an anatomical image catalogue linked to clinical terminology concepts. Relevant anatomical layers can then be highlighted when viewing a patient's history.

This provides an additional visual way to understand the patient's clinical story rather than relying entirely on text.

Source traceability

Clinical information should not become disconnected from the document from which it originated.

Open Clinical History therefore retains provenance information linking extracted events back to their source material.

The structured event is intended to help a user understand the history quickly, while the original source remains available for verification and clinical context.

A working proof of concept

Open Clinical History is currently a working proof of concept.

It demonstrates that fragmented clinical information can be:

1. ingested from unstructured source material;
2. converted into discrete clinical events;
3. classified against a recognised clinical terminology;
4. associated with anatomical structures;
5. presented as a coherent longitudinal history.

The project is continuing to evolve, particularly around clinical validation, interoperability, terminology governance and the safe use of AI-assisted extraction and classification.

What Open Clinical History is not

Open Clinical History is not intended to replace a clinician, electronic medical record, diagnostic system or source clinical documentation.

It is a tool for organising and presenting existing clinical information in a more understandable form.

Any extracted or classified information must remain traceable to its source and should be reviewed in the context of the original clinical record.

Why Open Clinical History Exists

Open Clinical History began with a simple problem: a patient's medical history may exist, but that does not mean the person treating them can easily see or understand it.

A complex health history can span many years, hospitals, specialists, imaging providers, pathology services and clinical systems. Important information may be contained in scanned letters, discharge summaries, old reports and documents created by organisations that do not share a common clinical record.

The result is fragmentation.

At a new specialist appointment, the patient or their carer may once again be asked to explain years of medical history from memory, even though much of that history has already been documented somewhere.

Open Clinical History was created to explore a better way of presenting that information.

The problem is not always missing information

A significant challenge in healthcare is that information can be present but difficult to use.

A folder containing hundreds of clinical documents is technically a health record, but it is not necessarily an understandable clinical history.

A clinician may need to determine:

- what happened;
- when it happened;
- which conditions are ongoing;
- which procedures were performed;
- which investigations were significant;
- how one event relates to another;
- which part of the body was involved;

- and which source document supports each piece of information.

Doing this manually takes time.

For patients with long or complicated histories, the cognitive burden can be substantial for both clinicians and patients.

The patient should not be the integration layer

In fragmented healthcare environments, the patient or carer often becomes the person responsible for connecting the systems together.

They carry documents between providers, remember previous diagnoses, explain earlier procedures and try to reconstruct the sequence of events for each new clinician.

That is an unreliable way to manage a complex longitudinal history.

People forget dates. Clinical terminology can be difficult to remember. Important details can be lost when a complicated history is reduced to a short verbal explanation during an appointment.

Open Clinical History is based on the idea that software should help perform that integration work.

From documents to clinical events

The project therefore focuses on converting source material into a structured sequence of clinical events.

Instead of treating a discharge summary or specialist letter as a single object, Open Clinical History examines the clinical information contained within it.

A single document may describe:

- a diagnosis;
- a procedure;
- an investigation;
- a medication change;
- an admission;
- a complication;
- a follow-up;

- or a clinical outcome.

These events can then be placed into the wider longitudinal history.

Consistent clinical terminology

Free text is essential in healthcare, but free text alone makes it difficult to consistently classify and connect information from different sources.

Open Clinical History therefore uses SNOMED CT as a terminology foundation.

The goal is to preserve the original clinical meaning while also associating events with recognised clinical concepts where appropriate.

This provides a common vocabulary across records that may have been created years apart and by different organisations.

Making the history understandable

Open Clinical History is also designed around the idea that a clinical history should be understandable at more than one level.

A user may first want to see the broad story:

- major diagnoses;
- admissions;
- procedures;
- complications;
- important investigations;
- and changes over time.

They can then drill into the details and, where necessary, return to the original source document.

Anatomical visualisation provides another way of understanding the history by showing which parts of the body are associated with particular events.

Supporting, not replacing, clinical judgement

The purpose of Open Clinical History is not to make clinical decisions on behalf of a healthcare professional.

Its purpose is to organise fragmented information so that the existing history can be understood more quickly and reviewed more effectively.

Where automated extraction or classification is used, the result should remain auditable and traceable to the underlying source.

The broader goal

The longer-term goal of Open Clinical History is to demonstrate a practical model for converting fragmented health information into a structured, standards-aware longitudinal history that can be useful to both clinicians and patients.

The project brings together:

- unstructured document ingestion;
- clinical event extraction;
- SNOMED CT terminology;
- anatomical classification;
- provenance;
- longitudinal visualisation;
- and future interoperability standards.

The central idea is simple:

A patient's history should be more than a pile of documents. It should tell the clinical story.