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

---

Revision #1
Created 2026-08-25 09:11:29 UTC by Admin
Updated 2026-08-25 09:11:29 UTC by Admin