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- medication lists that are no longer current;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.

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