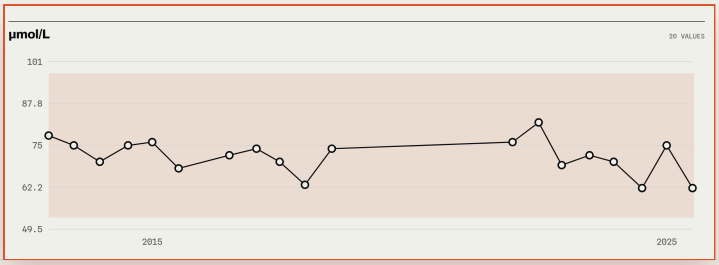
The folder the patient brought, read before the visit.
Every value one click from its page.

Epicrisis turns the paper a patient carries in — other clinics' letters, other laboratories' forms, in five languages — into a structured archive on your own servers. Your clinician reads it in one place, can question it through an AI assistant, and every answer points back to the scan. **The archive shows what the documents say. Your clinicians decide.**



DEMO DATA, INVENTED PATIENT. ONE TEST OVER THIRTEEN YEARS, DRAWN FROM FORMS SOMEBODY ELSE PRINTED. THE SHADED BAND IS THE RANGE PRINTED ON EACH FORM, HELD UNTIL ANOTHER FORM PRINTED A DIFFERENT ONE: IT MOVES BECAUSE LABORATORIES DIFFER, NOT BECAUSE ANYTHING WAS CALCULATED. EVERY VALUE KEEPS ITS DATE, ITS INSTITUTION AND THE SCAN IT CAME FROM.

01 WHAT THE CLINICIAN GETS

One timeline per test across laboratories, countries and languages. Units are kept as printed; mg/dL and µmol/L never land on one chart.

The range of that form, on that day. The band on a chart is what the laboratory printed beside the value, not a table of ours — where a form printed no range, there is no band.

Every value one click from its scan. The page is the authority, and a line read wrongly is corrected by hand in seconds.

Search over every word of every document, matching across Cyrillic and Latin spellings of one test.

Checks with no model at all: the same result filed three times, a stored number that differs from the print, a date nobody could read.

02 WHY IT FITS A CLINIC

Three modes, strictest by default. *As printed*: shown as the forms wrote it, nothing compared. *May also read the values*: a model may say what a value means and how it moved. *No limits set here*: the model answers under no rule of ours, and only here does the application itself compare a number with the range printed beside it. The last two were written for a person reading their own records — the model is told that the person asking is the person the records are about — so the mode that fits a clinic reading somebody else's file is the first. Your instance, your choice, on its own Settings page.

Inside your perimeter. Self-hosted, no vendor cloud, no telemetry. Model calls go out with your own key, or not at all.

One archive per patient, isolated by construction, and every page carries the name of the one whose records are open.

It reads and cites; it does not diagnose, triage or decide. Not a medical device, and nothing here changes who is responsible for a clinical decision.

It holds what was brought and could be read, and says so: a test missing from this archive is not a test that was never taken.

Source-available. Your IT and clinical governance read every line before it is installed.

01	02	03	04
A new patient with a folder	**Referral and second opinion**	**Before a procedure**	**Patients from abroad**
Years of somebody else's paper, searchable before the visit ends.	The whole history in one place, with the form behind every number.	Which tests exist, when they were taken, and what that form printed.	RU, UK, EN, ES and EL forms on one timeline, each in its own words.

600+	6,000+	2063	1989–2026	150+
DOCUMENTS READ	VALUES AS PRINTED	TESTS IN THE SUITE	YEARS COVERED	CLINIC NAMES · 5 LANGUAGES

Looking for two or three clinics or departments for a pilot.

Your own files, on your own server. We learn what your clinicians actually need; you keep everything. A data processing agreement under Article 28 GDPR comes with the licence.

Artem Berman
EPICRISIS.BERMANA.NET
LINKEDIN.COM/IN/ARTEMBERMAN

NOT A MEDICAL DEVICE UNDER REGULATION (EU) 2017/745: IT STORES AND DISPLAYS WHAT THE FORMS PRINTED (MDCG 2019-11). NOT FOR DIAGNOSIS, TREATMENT OR CLINICAL DECISIONS. BUSINESS SOURCE LICENSE 1.1 — FREE FOR A PERSON AND THEIR FAMILY, A COMMERCIAL LICENCE FOR ORGANISATIONS.