

Self Project

2026-10-10

HealthBook AI: Privacy-First Family Health Workspace for Longitudinal Records, Human- Verified Extraction, and Selective Doctor Sharing



A privacy-first family health workspace for longitudinal medical records, human-verified document extraction, daily health tracking, clinical timelines, secure PDF exports and selective doctor sharing across web and Android.

Project Snapshot

DATE

2026-10-10

CATEGORY

Self Project

SHORT DESCRIPTION

A privacy-first family health workspace for longitudinal medical records, human-verified document extraction, daily health tracking, clinical timelines, secure PDF exports and selective doctor sharing across web and Android.

TAGS / TECH STACK

Health AI

Healthcare AI

Full-Stack Platform

Web + Android

Personal Health Records

Document Intelligence

OCR

Human-in-the-Loop

Privacy Engineering

Project Links

GitHub Link

GitHub Link

Project Link

Full Project Content

HealthBook AI is a privacy-first family health workspace. It keeps original medical documents encrypted and private, turns them into structured records only after a person has verified them, separates what was prescribed from what was actually given, and lets a family export or share exactly the part of the history a clinician needs, from one common web and Android architecture.

HealthBook AI is a personal health record, not a medical device, and it never prescribes or changes a dose.

Project Links

- [Technical Wiki](#): architecture, data model, security boundaries and every feature pipeline, with diagrams
- [GitHub Repository](#): sanitized public source code, tests and deployment templates
- [Live Application](#): the deployed web application (sign-in required; no public data)

GitHub Repository: [dranubhaparashar/HealthBook-AI](#)

One-Line Idea

Give every family one private place that answers a practical question: *what is actually documented about this person's health, and where is the original that proves it?*

Why This Project Exists

A family's medical history usually lives in many places at once:

- paper and PDF lab reports;
- prescriptions and discharge notes;

- vaccination cards;
- phone photos of documents;
- memory of what was given, when, and for how long.

When a doctor asks a simple question, such as which antibiotic courses a child has had or when a vaccine was actually given, the answer is often reconstructed from memory. HealthBook AI keeps the original evidence, turns it into structured history only after human review, and keeps every statement linked to the document page it came from.

Project at a Glance

Area	Description
Project type	Self-built full-stack health platform
Domain	Personal and family health records
Clients	Web application and Android application
Backend	One FastAPI API that owns every secret and every rule
Data	Supabase PostgreSQL with encrypted clinical fields; private encrypted object storage for originals
Extraction	PyMuPDF text layer with local Tesseract OCR fallback; field-level human verification
Outputs	Designed PDF dossiers and scoped, expiring doctor-share links
Governance boundary	Nothing becomes a structured record until a person confirms it
Public showcase	Sanitized repository and technical wiki using synthetic data only

Key Capabilities

- **Family health profiles:** each person (infant to senior) is a separate profile with its own timezone, modules and records
- **Medical document library:** PDF, PNG, JPEG and HEIC originals stored encrypted, with SHA-256 duplicate detection
- **Human-verified extraction:** OCR produces a *candidate* transcription that a guardian checks field by field before anything is saved
- **Medicine history:** draft and verified prescriptions, revisions and separate administration events

- **Vaccination history:** card entries and a reference chart that is never auto-activated; due is never treated as given
- **Labs and growth trends:** observations, cultures, lab series, growth entries and vitals
- **Daily Health:** a daily journal for fluids, sleep, activity, symptoms, vitals and medicines given
- **Care Calendar:** day, week and month views over the unified timeline
- **Longitudinal timeline** and **Lifetime Health:** illness episodes, fever episodes and antibiotic courses linked to evidence
- **Secure PDF export:** a designed dossier with the original documents appended and linked
- **Selective doctor sharing:** preview, choose exactly which records to share, and send an expiring, revocable link with an optional PIN
- **Android + Web:** one backend serving a web application and an Expo React Native Android app
- **Private encrypted storage:** Fernet encryption for clinical text and every stored file

Clinical Safety Rules

These rules are enforced by the data model and the API, not only by the interface:


Rule	Meaning
Prescribed $\neq$ scheduled $\neq$ given	A planned dose is never counted as an administered dose
OCR extracted $\neq$ verified	Extraction output stays unverified until a person confirms it
Due date $\neq$ administered	A vaccination reminder is not a vaccination record
Reminder $\neq$ proof	A notification is never evidence that a dose was given
Provenance is kept	Structured records keep a link to the original document page

Architecture

Web application (Firebase Hosting)

Android application (Expo React Native)





The [Technical Wiki](#) documents each layer in detail, including deployment, authentication and session security, family profile isolation, the OCR and human-verification workflow, PDF export, doctor sharing, encryption and the API design.

Technology Stack

Layer	Technologies
Backend	Python 3.12, FastAPI, SQLAlchemy, Alembic, Pydantic
Database and storage	PostgreSQL, Supabase (PostgreSQL and private storage bucket)
Documents	PyMuPDF, Tesseract OCR, Pillow, ReportLab
Security	cryptography (Fernet), hashed session tokens, CSRF protection
Web	Firebase Hosting, no-build JavaScript client with a service worker
Android	Expo, React Native, TypeScript
Testing	pytest, Playwright

Testing and Validation

The public repository includes a pytest suite of **130** test functions and **20** Playwright browser tests. All tests use synthetic data only (fictional names, `example.invalid` emails, generated PDFs and images). Many tests encode the clinical safety rules directly, for example that a scheduled medicine is never counted as given, and that one family can never read or confirm another family’s records.

Privacy and Scope

- The public repository and wiki are a **sanitized showcase**. Production infrastructure, credentials, source medical documents and patient records are maintained separately and are not public.

- Every example in the documentation uses synthetic data.
- HealthBook AI is a family-maintained record. It is not a medical device and makes no claim of clinical validation or regulatory clearance.

For limitations and the roadmap, see the [Technical Wiki](#).

Source: src/content/posts/HealthBook-AI/HealthBook-AI.md