Notes from the Village

Building the trust network around your health data, one update at a time.

TrustMyOwn.Health 014 — Sorting the pile, not practicing medicine

TMOH (TrustMyOwn.Health) refers to our company.
PHDB (Personal Health Data Bank) is our product.

A Villager asked me four questions this week. Where does our AI get its training data? Are 32 use cases enough? Does it know when an herb acts like a medication? Does PHDB use Gemini, Claude, or ChatGPT? I shared them with my co-founders, Tomas Moran and Dr. Marianne Hudgins.

Tomas started with a line we should draw: organizing your records versus interpreting them. Picture a pile of paper on my desk. I sort it into buckets: labs, imaging, appointments, medications. When a clinician asks, "Any cardiac labs?" I say, "Right here." That is organizing. Deciding what those labs mean and what to do about them is medicine. PHDB is built for the first job.



Figure 1: Created with OpenArt AI

Which AI, and what was it trained on? Claude Sonnet 5.5 from Anthropic. Our technical team tested several. It is a general-purpose model that understands language. It recognizes "asthma" as a diagnosis and files a blood pressure reading in the right place. It is not a doctor, and no one trained it to be one. It is also not one of the frontier models you read about in the news. Sorting your records is a far simpler task than what those models tackle.

Does my information leave my control? Last week, I said I did not have the full answer. Now I do. Your vault is processed by a private AI agent running in a private virtual machine. Your records do not leave the vault, and nothing in them trains any external system. Private AI by default, cloud only by permission. If you want a big-name AI to help interpret your records, PHDB can prepare a summary and strip identifying details first. That is your decision to make, not ours.

Are 32 use cases enough? I thought they were mostly about user interface, what each person sees on screen. Tomas set me straight: the sorting is basic AI, driven by prompts, and the cases do not teach it diseases. They ensure PHDB covers every kind of record, such as bone density imaging, which men rarely have. Marianne put it another way. The structure is already built. The use cases tell us what else an Owner would find useful once their records are organized.

Does it know when an herb acts like a medication? If a document in your vault says you take an herb or supplement as medicine, it is treated that way. It will not add it to your medication list on its own. Something in your records has to say it belongs there.

Here is the part I did not expect to care about so much. Because the sorting reads the actual text of your documents, you can click an element to see the exact chunk it came from. And because it reads the narrative, not just the codes, it can surface a diagnosis a clinician mentioned in a note that never made it into an ICD-10 code. I have lived in that gap. Things about me that were true, written down, and invisible to the next person in the room.

The line is still blurrier than it sounds. Say two records list the same medication, one at 25 milligrams and the other at 50. Showing both is organizing. Choosing which is current is a clinical call. The line runs through hundreds of small choices, and the Data Sovereignty Covenant is where we start: certain decisions are yours from the beginning, not ours.

Where do you think the line belongs? Is organizing the right place for PHDB to stop? Hit reply and tell me.

Notes from the Village

Building the trust network around your health data, one update at a time.

Onward, Danny Health Hats