ReconMed

One correct medicine list, before the patient goes home.

Team Phoenix Forge 0, MedHack UAE, August 2026

The prototype

Built, deployed and open. No login and no setup.

Try it https://maryammeda.github.io/ReconMed/

Check our work https://maryammeda.github.io/ReconMed/tests.html

Read the code https://github.com/maryammeda/ReconMed

Choose Healthcare Worker, open the first patient, press Begin reconciliation. The first finding is the one described below.

The observation

A woman of 68 is admitted to hospital. She takes eleven different medicines, for her heart, her diabetes and her kidneys. Her medicine list exists in three places at once, and the three do not match. The national health record holds one version. Her clinic letter holds another. The photograph she took of the boxes in her kitchen holds a third.

Before she can go home, one person has to turn those three lists into one correct list. That person is the ward pharmacist, and today they are covering thirty other patients.

The problem, stated properly

Nobody has a single correct list of what a patient actually takes. Building one by hand takes fifteen to thirty minutes per patient, so most patients never get the check. They go home on whichever list happens to be printed.

62 in 100 against 9 in 100

Patients leaving hospital with a clinically important medication error, without a pharmacist check and with one. Jošt and colleagues, 2024.

> The check already works. There is simply nobody with the time to run it on everyone.

Why it is a UAE problem

The UAE has already built the hard part. Malaffi, Nabidh and Riayati bring a patient’s records together from different hospitals onto one screen. But putting records next to each other is not the same as comparing them. A human still has to notice that two lists disagree, and that the dose on all of them is wrong for this particular patient’s kidneys.

Patients here also move between emirates and between private and public care, so one person's history is spread across more systems than usual. And the people who would do the checking often cannot reach the records at all.

68 in 100

UAE hospital pharmacists who agreed that lack of access to patients' medical notes or shared information is a barrier to doing this check. National survey of 342 pharmacists.

Who is affected

The patient. Usually older, usually on many medicines. She goes home with a list she cannot check herself, often in a language she does not read comfortably. If a medicine has quietly disappeared, she will not know until she is unwell again.

The pharmacist. One person covering a whole ward. This is the task that gets dropped when the ward is busy, because nothing visibly goes wrong when it is skipped.

The hospital. Carries the readmission, the incident report and the accreditation finding, all of them preventable.

The constraint that shapes everything

Two approaches are already known to fail, and both of them shaped what we built.

You cannot ask a hospital to replace the record system it already runs on. Even where the factors that make these projects succeed are well documented, a majority still fail in some sense, and the largest fail hardest.

And you cannot solve it by warning about everything. Clinicians already ignore between 49 and 96 of every 100 automated drug warnings, because those systems cry wolf. A warning nobody reads is not a safety feature.

> So it has to work with what the hospital already has, and raise very few findings, each one worth a clinician's time.

The solution shape

ReconMed reads the three lists in whatever form they arrive, including a photograph. It works out that "Eliquis" and "apixaban" are the same drug, and finds where the lists disagree: something listed twice, something missing, two different doses.

Then it checks each medicine against the patient's own test results, which is what lets it catch a dose that is wrong even when all three lists agree perfectly.

The pharmacist gets a short ranked list, each finding showing the exact lines it came from and the reasoning behind it, and decides on every one. The tool never changes a medicine itself, and never writes to the hospital record.

What the prototype proves

For our demonstration patient it finds six things. Five are the disagreements you would expect. The sixth is the reason this project exists.

All three lists agree that she takes 1000 mg of metformin twice a day, so there is no disagreement to find. But a kidney test from six days earlier shows her kidneys working at roughly a third of normal, and at that level the recommended daily maximum is half what she is on.

> No amount of comparing lists finds this, because the lists were never the problem. It is found by reading the patient, not the paperwork.

A second patient in the prototype takes the same drug at the same dose with better kidney function, and is correctly not flagged. Same rule, different patient, different answer. There is also a box where you can type a medicine line of your own and watch the system read it. Nothing is pre-recorded, and 93 automated tests run from the second link above.

What it changes

For the pharmacist. A thirty minute manual job becomes a review of a few ranked findings, with the riskiest patient first. If only three get checked today, they are the right three.

For the patient. A list a person has actually checked, in plain English and Arabic, with every change explained.

For the hospital. A sealed record of who decided what, when and why.

The patient's screen

The same check that clears her medicines for discharge also gives her one screen to keep, in place of the three lists that never agreed.

Her medicines, as they stand now. The reconciled list, name, dose and time, shown in plain English or Arabic.

What is going on, in plain words. Her most recent diagnosis, re-written without the jargon, so she knows what she is taking and why (A recommendation from a 3rd year medical student in the UAE.)

A nudge for each dose. A reminder by name and time for every medicine, so nothing is left to memory.

A small word of advice. A short, kind advice, because staying on a plan is easier when someone is in her corner.

A reminder to check back in. A prompt to book the next follow-up with her doctor, timed to when the plan says she is due.

Keeping the record safe

- **The approved list is sealed.** When the doctor signs, the system takes a fingerprint of the exact list. If anyone changes it afterwards, the fingerprint no longer matches, and the system says so on every screen and records it. This is a short non-cryptographic digest suitable for a prototype; a real deployment would use a server-side cryptographic hash with a write-once store.
- **Every look is logged.** Opening a record is recorded with who opened it, even when nothing is changed, because who read a record is itself auditable.
- **The screen locks itself.** Ward workstations are shared, so a record left open hides after five minutes.

Known limitations

- The risk score orders a worklist. It is not a validated clinical instrument, and it never changes a medicine.
- Artificial intelligence reads the photograph and translates the patient copy after sign-off. It makes no clinical decision. Every clinical rule is ordinary code.
- Forty medicines, not a national registry. Two patients carry full records; the rest of the ward carries summary figures, and the screen says so.
- It is a prototype, has not been used on real patients, and all data in it is invented. Real accuracy needs a clinical pilot.
- There is no real authentication; the role switch is a demonstration, and a deployment would need identity and enforced role-based access.

What would come next

Reading directly from Malaffi or Nabidh instead of an exported file, and writing the approved list back so the pharmacist retypes nothing. After that the work is data rather than engineering: growing the drug dictionary is data entry, not redesign, which is what turns one ward into a hospital. UAE law requires health data to stay in the country, so a real deployment would run on in-country infrastructure such as Med42. Also, true English ⇆ Arabic translation backed by translators and research for honest bilingual accommodation.

A professional in the field complained about lack of medicine compliance and patient ghosting of their follow ups thus, there is also a loyalty loop worth testing: once a patient has kept to their plan for a defined stretch, the dashboard could offer ten percent off vitamins at the pharmacy that filled their prescription. A small thank you for staying on track, and a reason to keep filling prescriptions and booking follow-ups on time.

Where the numbers come from

Jošt M, Kerec Kos M, Kos M, Knez L. Frontiers in Pharmacology, 2024.

Clinically important medication errors at discharge affected 9.3 per cent of patients in the pharmacist led group and 61.9 per cent in the control group.

https://pmc.ncbi.nlm.nih.gov/articles/PMC11007427/

Farajallah A, Zainal H, Palaian S, Alomar M. Scientific Reports, 2024.

National survey of 342 UAE hospital pharmacists. 68 per cent agreed that lack of access to patients' medical notes or shared information is a barrier.

https://doi.org/10.1038/s41598-024-64605-4

van der Sijs H, Aarts J, Vulto A, Berg M. J Am Med Inform Assoc, 2006.

Clinicians override 49 to 96 per cent of drug safety alerts.

https://pubmed.ncbi.nlm.nih.gov/16357358/

Kaplan B, Harris-Salamone K. J Am Med Inform Assoc, 2009.

Despite research identifying what makes health IT projects succeed, a majority still fail in some sense.

https://pubmed.ncbi.nlm.nih.gov/19261935/

World Health Organization. Medication Without Harm.

The global cost of medication errors is estimated at 42 billion US dollars a year.

https://www.who.int/initiatives/medication-without-harm

Dose ceilings follow the European Medicines Agency summary of product characteristics. The United States FDA label sets no numeric daily cap for this drug, so that figure is deliberately not attributed to it.
