Chronic Heart failure detection medication titration recommendation System

Trajectory-Integrated Decision Engine for Heart Failure (Phase 1A)

Team: Pixel Minds

Team Members:
Hema Priya Balaji
Jayachandran Galda
Srinivasa Rao

Client: Dr. Ciprian Ionita, QAS.AI

MEETINGS

2/26 - Introductory Meeting
3/5 - Weekly Check-in and Scope Clarification
3/13 - Weekly Check-in and Dataset Discussion

Spring Break

3/26 - Weekly Check-in, Initial Dataset Findings and Titration Discussion
4/2 – Weekly check in on the GCP database and flowchart of logic presented
4/14 – Weekly check in on medical recommendation based on our logic
4/23 – Presented our final website and logic for the system built

Project Overview

- Our project focuses on building decision support systems for patients with chronic heart failure (HFrEF). Trajectory Integrated Decision Engine for Heart Failure mainly focuses on improvising how patients with chronic heart failure are monitored and treated before and after discharge. It uses wearable sensor data and clinical records to support safe and timely medication decisions.

- Currently most patients are monitored during hospital visits and medication adjustments are done every few weeks. This creates a gap where a patient's condition may worsen without timely intervention.

- The goal is to improve post discharge monitoring and reduce patient risk through continuous tracking. To address this, we are building a system that:

1. Continuously collects real-time physiological data
2. Applies clinical guidelines in a structured way
3. Provides safe medication recommendations
4. Alerts clinicians in case of risk

The goal is not to replace doctors, but to support them with better, continuous insights.

System Architecture Overview

Architecture Layers

1. **Data Collection Layer**
   Wearables + dataset inputs
   ECG, HR, BP, SpO_2, fluid
2. **Data Processing Layer**
   Cleaning and filtering (MIMIC-IV)
   Feature extraction
3. **Decision Engine**
   Rule-based clinical logic
   Medication titration rules

4. **Output Layer**
    Alerts
    Recommendations
    Reports for clinicians

Client Interaction Summary

We had recurring meetings with stakeholders to understand expectations and move in the problem statement's direction. The focus was on building a logic foundation before moving into model development in advance. We have aligned our work to ensure it remains clinically safe and easy to explain.

Project Description

The system is designed as a decision support engine that integrates:

- Wearable sensor data
- Clinical datasets (MIMIC-IV)
- Evidence based on medical logic

It works in a step-by-step decision process, like how a clinician would think, but continuously and consistently.

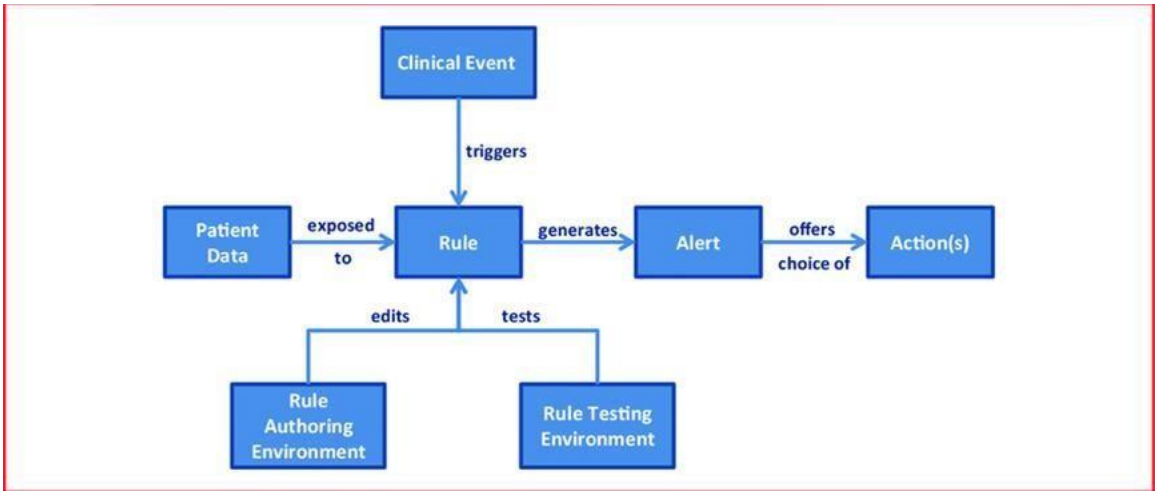

Step by Step Flow

1. Input Data
    - Sensor data + patient records

2. Safety Check
    - Detect emergencies (detection if they have low BP, low oxygen, arrhythmia)
3. Patient State Identification
    - Wet
    - Dry
    - Borderline
4. Medication Logic Execution
    - Apply rules for each drug
5. Monitoring and Feedback delivery
    - Track trends
    - Escalate if condition worsens

Project Overview Studied Congestive Heart Failure Mechanism and Summary Key concepts

Aim: To understand the clinical mechanism of congestive heart failure for our project research.

Source: YouTube Video

Congestive Heart failure - Clinical Medicine - Congestive Heart Failure | Clinical Medicine

Deliverables:

- Personal notes to understand the concept
- Key mechanism explained
- List of the symptoms and the causes
- Treatment overview

Basic Mechanism of heart attack

Heart failure happens when the heart cannot pump blood effectively and this leads to:

- Reduced blood flow to organs

- Fluid buildup in lungs and body
- Increased pressure inside the heart

Key Concept

BP = CO × SVR

- CO (Cardiac Output) decreasing → Heart is weak
- SVR (Resistance) increasing → Body tries to compensate
- This compensation worsens the condition over time.

Internal Conditions: What is happening in our body?

1. Reduced Cardiac Output

- Less oxygen reaches organs
- Body activates stress responses

2. Activation of RAAS System

- Kidneys retain salt + water
- Leads to fluid buildup

3. Fluid Backup

Left Heart Failure:

- Fluid goes into lungs
- Pulmonary edema
- Shortness of breath

Right Heart Failure:

- Fluid accumulates
- Leg swelling
- Abdominal swelling

Symptoms

Quantitative (measurable):

- High heart rate
- Low oxygen (SpO_2)
- Fluid retention (weight gain)

Qualitative:

- Shortness of breath
- Fatigue
- Swelling (in legs, abdomen)
- Cough (especially at night)

Complications:

- Pulmonary edema
- Cardiogenic shock
- Kidney failure
- Stroke / heart attack

Key Takeaway

Heart failure is a **cycle problem**:

- Weak heart → compensation → fluid buildup → worse heart function

Project Understand Heart Failure Drug Mechanism and Treatment Concepts

Aim: Understand major drug classes used in heart failure treatment and their clinical mechanisms to support domain knowledge for our healthcare project.

Source: YouTube Video - Drugs for Heart Failure

Deliverables:

Summary notes

Drug classification table

Mechanism summary

Key clinical insights relevant to our project

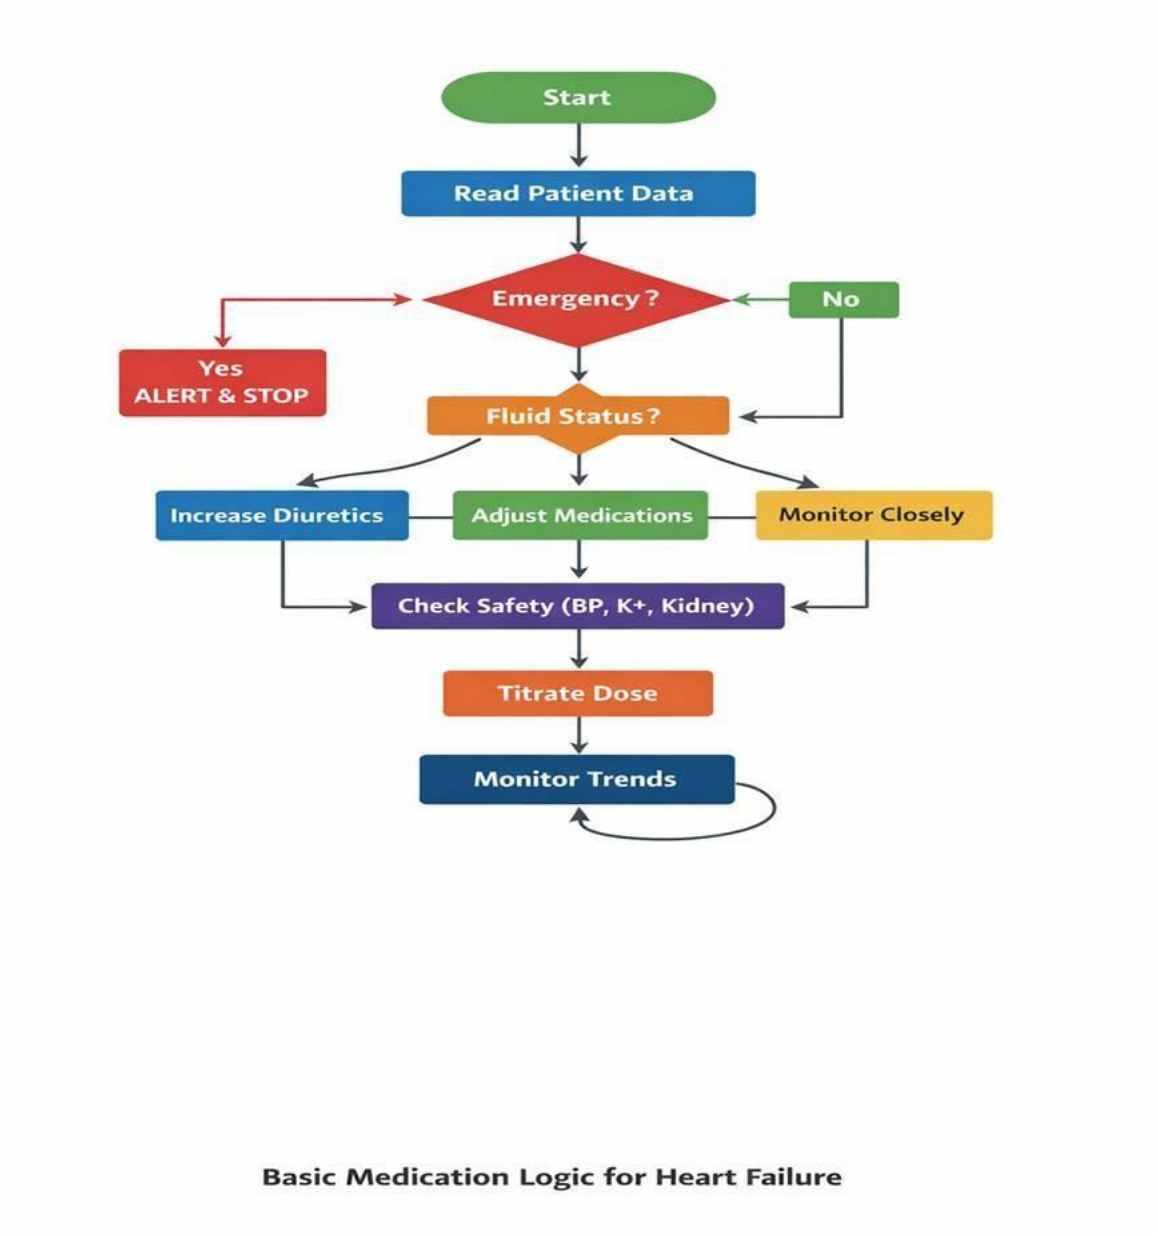

Basic Medication Logic for Heart Failure

Core Idea

Each drug class targets a **different part of the problem**:

- Reduce workload on heart
- Remove excess fluid
- Improve pumping efficiency

Drug Classes & Mechanism

1. Beta Blockers

- Reduce heart rate
- Reduce oxygen demand

Helps heart work more efficiently

Initial effect: slight worsening, but improves long-term

2. RAAS Inhibitors (ACEI / ARB / ARNI)

- Block harmful hormonal system (RAAS)
- Reduce: Blood pressure, Fluid retention, Heart stress

Core drug for heart failure treatment

3. Diuretics

- Remove excess fluid Helps with:
- Swelling
- Breathing issues

Does NOT improve survival, but improves symptoms

4. MRA (Aldosterone Antagonist)

- Reduce sodium and water retention
- Protect heart

5. SGLT2 Inhibitors

- Originally diabetes drugs
- Now used in heart failure
- Fluid balance
- Heart protection

Simple Drug Table

Drug Class	Purpose
Beta Blockers	Reduce heart workload
RAAS Inhibitors	Improve heart function
Diuretics	Remove fluid
MRA	Control electrolytes & fluid
SGLT2	Stabilize patient

Basic Medication Logic

Simple Flow of our project in stepwise format:

Step 1: Check Emergency

- Low BP
- Low oxygen
- Arrhythmia

If yes STOP & alert

Step 2: Check Fluid Status

- Wet (Increase diuretics)
- Dry (Start titration)
- Borderline (Monitor)

Step 3: Apply Drug Logic

Start with:

- RAAS inhibitors
- Beta blockers

- Add: SGLT2 o MRA

Step 4: Monitor Response

- Check trends
- Adjust doses slowly

Key Clinical Insights for Your Project

- Always check safety first
- Fluid status drives most decisions
- Medications are not given together blindly
- Titration is slow and step by step
- Trends matter more than single readings

Explanation in simple terms for our understanding:

We studied how heart failure progresses and how medications interact with the body. Based on that, we understood that treatment mainly focuses on reducing fluid overload, stabilizing the heart, and gradually optimizing medications. This helped us design our logic system in a step by step, safety-first manner.

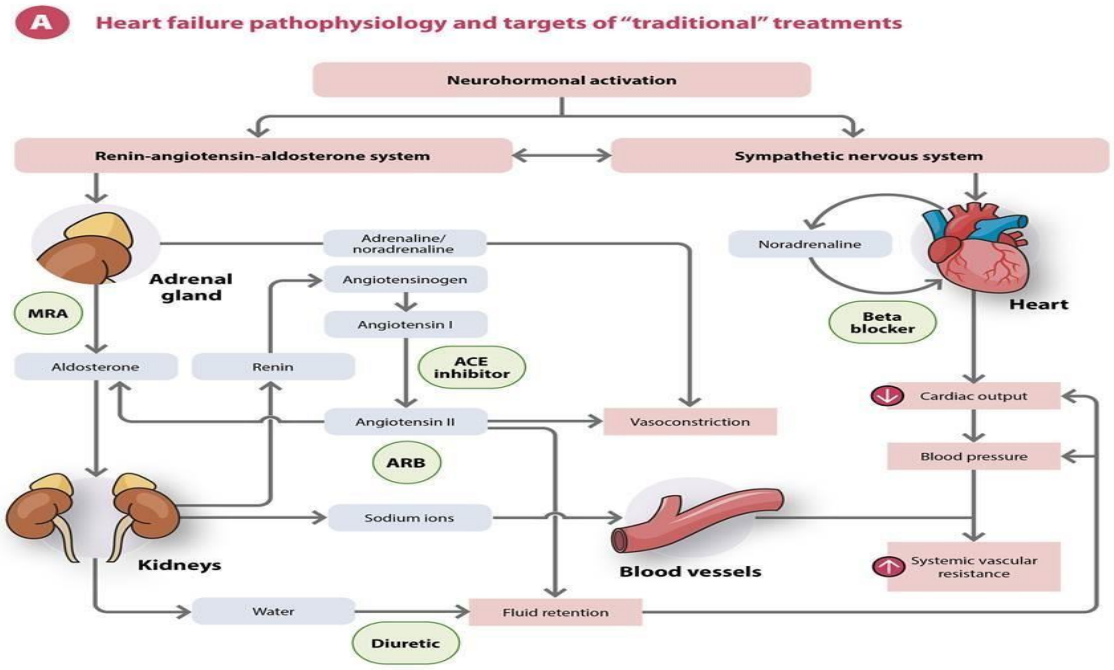

Problem Understanding

Heart failure patients are typically monitored at intervals, which can delay critical interventions. This creates a gap between patient condition changes and treatment updates. Our system aims to bridge this gap using continuous monitoring and decision support.

Clinical Background

We studied how heart failure affects the body and how medications help stabilize patients. This helped us understand how fluid retention, blood pressure, and heart rate influence treatment decisions.

System Design

The system collects real-time data from sensors such as ECG, BP, HR, oxygen levels, and fluid indicators. It first checks emergency conditions before making any recommendations.

Patients are classified into states such as wet, dry, or borderline to guide treatment decisions. Medication adjustments are based on structured rules derived from clinical guidelines.

Work Completed as of March 3rd week:
Completed literature review and domain understanding. Finalized the MIMIC-IV dataset and began filtering. Developed initial logic framework and flowchart. Prepared materials for client discussion.

Next Steps for April month:
Clean and structure dataset for use. Map data fields for required inputs. Convert logic into implementation-ready format. Begin testing and validation.

Literature Review & Research Background

The development of our Trajectory-Integrated Decision Engine for Heart Failure is grounded in a body of peer-reviewed clinical and computational research. The following six sources collectively justify the clinical framework we adopted, the

safety rules embedded in our logic, the medications we recommend, and the machine learning approach we chose. Each study contributed a specific and necessary piece of evidence to our system design.

Research paper 1:

AHA/ACC 2022 Heart Failure Guidelines

The 2022 guidelines from the American Heart Association and American College of Cardiology are basically the rulebook for how heart failure gets treated in the US right now. The big thing this edition did was formally lock in the four pillars framework the idea that every HFrEF patient who can tolerate it should be on four specific drug classes. Those four are the ARNI, a beta blocker, an MRA, and an SGLT2 inhibitor. Before this update, the approach was messier and varied a lot depending on which cardiologist you were seeing.

The guidelines also spell out how heart failure progresses in stages. Stage A is someone who does not have heart failure yet but is at risk someone with hypertension or diabetes, for example. Stage B is where structural changes in the heart have started, but the person is not yet feeling symptoms. Stage C is when symptoms show up the breathlessness, fatigue, the swelling. Stage D is the most advanced stage where standard treatments are no longer keeping things under control.

For our project, these stages and the four drug classes are what the whole logic engine is built around. The system looks at where a patient is clinically and what medications they are already on and figures out what should be started, increased, or held exactly the way these guidelines lay it out.

Link: -https://www.ahajournals.org/doi/10.1161/CIR.0000000000001063

Research paper 2:

Machine Learning–Based Evaluation of Guideline-Directed Medical Therapy in Heart Failure Patients

This study out of Frontiers in Cardiovascular Medicine looked at data from over 3,800 heart failure patients and checked how many of them were getting the

medications the guidelines say they should be on. The numbers were not good. Despite ARNI being a first line of medication, only 8.4 percent of eligible patients in this study were taking it. The other drug classes showed similar gaps.

The reason the researchers pointed to was straightforward most heart failure patients only see their doctor a few times a year. Between those visits, nobody is actively checking whether medications need adjustment. A patient can be getting slowly worse for weeks, and nobody knows until the next appointment.

That 8.4 percent is basically the whole reason we built this system. The gap between what patients should be getting and what they are getting is enormous and a big part of why it exists is the lack of continuous monitoring and regular adjustment.

Link: - https://pmc.ncbi.nlm.nih.gov/articles/PMC10321403/

Research paper 3:

Frontiers Meta-Analysis (2025) - MRA Safety Evidence

This was a systematic review that pooled data from nine separate clinical trials looking at MRA use in heart failure patients. The headline finding was that MRAs cut overall mortality risk by around 22 percent compared to placebo, which is a strong result and backs up why they are in the four pillars.

But the other finding matters just as much for our project MRA use was linked to a 2.19 times higher risk of hyperkalemia. That is a condition where potassium in the blood goes too high and at dangerous levels it can trigger life threatening heart rhythm problems. The meta-analysis did not say stop using MRA because of this the mortality benefit is real and meaningful. But it made it very clear that any system managing MRA dosing must have a hard stop when potassium gets too high.

That 2.19 times risk figure is directly why our logic engine stops MRA dosing when potassium climbs above 5.0 and escalates to a clinician. That rule is not something we made up of it comes directly from this evidence.

Link: -https://www.thelancet.com/journals/lanepe/article/PIIS2666-7762(25)00302-3/fulltext

Research paper 4:

EMPEROR-Reduced Trial - SGLT2 Inhibitors in Heart Failure

This was a randomized controlled trial published in the New England Journal of Medicine. It enrolled 3,730 HFrEF patients and gave half of them empagliflozin and half a placebo. The patients on empagliflozin had a 25 percent reduction in the combined risk of cardiovascular death or heart failure hospitalization. That is a genuinely large effect for a drug in this space.

There were two things from this trial that shaped how we built our system. First the kidney protection finding. Empagliflozin cut the risk of serious kidney complications by about 50 percent on top of the heart benefits. That matters because a lot of HFrEF patients already have some degree of kidney impairment and having a drug that protects both organs at once is significant.

Second these benefits showed up in patients without diabetes too. SGLT2 inhibitors started as diabetes drugs and for a while there was genuine uncertainty about whether the heart failure benefit applied to non-diabetic patients or whether it was something specific to the metabolic effects in diabetes. EMPEROR-Reduced settled that. The benefit was there regardless of diabetes status.

This is why our logic recommends SGLT2 for all eligible HFrEF patients and not just the ones with diabetes.

Link: - https://www.frontiersin.org/journals/cardiovascular-medicine/articles/10.3389/fcvm.2025.1667236/full

Research paper 5:

Machine Learning–Driven Identification of Heart Failure Subtypes Using Electronic Health Record

This came out in Lancet Digital Health, and it is one of the bigger machine learning studies done on heart failure. The researchers ran four different unsupervised machine learning methods on EHR data from over 313,000 patients across three

large UK datasets. They used 645 clinical variables and let the model find natural groupings in the data rather than telling it what to look for.

What came out were five distinct subtypes. Early onset patients who tended to be younger with fewer risk factors. Late onset which skewed older and more female with lower cardiovascular disease burden. An atrial fibrillation related group with strong links to valve disease and arrhythmias. A metabolic group overweight patients with moderate risk factors. And a cardiometabolic group which was the highest risk with widespread cardiovascular disease and multiple conditions at once.

The survival differences between these groups were striking. One year mortality ranged from 0.11 in the metabolic group all the way to 0.61 in the atrial fibrillation related group. The model validated well externally with c-statistics, reaching 0.94. They even built a working clinical prototype to show how subtype classification could be used during actual patient consultations.

For our project, this study matters because it confirms that heart failure is not one thing it is several distinct conditions that happen to share a name. Machine learning can find those differences from routine health records. That is the same kind of approach we are taking with MIMIC-IV.

Link: -https://www.nejm.org/doi/full/10.1056/NEJMoa2022190

Research paper 6:

SHAP-Based Explainability Study - Why Transparency Matters in Clinical AI

This study tested several machine learning models including neural networks, gradient boosting, and random forests for predicting heart disease outcomes. The more complex models performed well in terms of accuracy. But the study made an important point those models are black boxes. They give you an answer, but they do not tell you why. For a doctor, making a clinical decision about a real patient is a serious problem. An unexplained recommendation from a computer is very hard to act on with confidence.

The researchers used SHAP values to break down what each model was doing and show exactly how much each input variable blood pressure, creatinine, potassium,

age was contributing to any given prediction. That turned the model from something opaque into something that could be read and understood by a clinician.

This shaped how we think about our own system. When our logic engine makes a recommendation, it does not just say increase the diuretic or hold the RAAS. It records which specific value triggered that decision whether it was the fluid status, the blood pressure reading, the potassium level, or the creatinine trend. A clinician looking at the output can see exactly what drove it. That transparency is not a nice thing to have in a clinical system it is necessary for anyone to trust and use it safely.

Link: -https://www.thelancet.com/journals/landig/article/PIIS2589-7500(23)00065-1/fulltext

Overview of Literature Survey:

These six sources are not just a reading list. Each one fed directly into a specific part of how the system was built. The guidelines gave us the drug framework. The prescribing gap study told us why continuous monitoring matters. The MRA metaanalysis gave us the potassium safety threshold. The EMPEROR-Reduced trial justified including SGLT2 for all HFrEF patients. The Banerjee clustering study validated using machine learning on EHR data for this kind of problem. And the SHAP research shaped why we built explainability into every output the system produces.

1. DRUGS FOR HEART FAILURE:

Heart failure with reduced ejection fraction means the heart muscle is too weak to pump blood properly. The left ventricle the main pumping chamber cannot squeeze the way it should, so the rest of the body ends short on blood flow. Doctors treat this using four types of medication that each targets a different part of the problem. These four together are called GDMT guideline directed medical therapy.

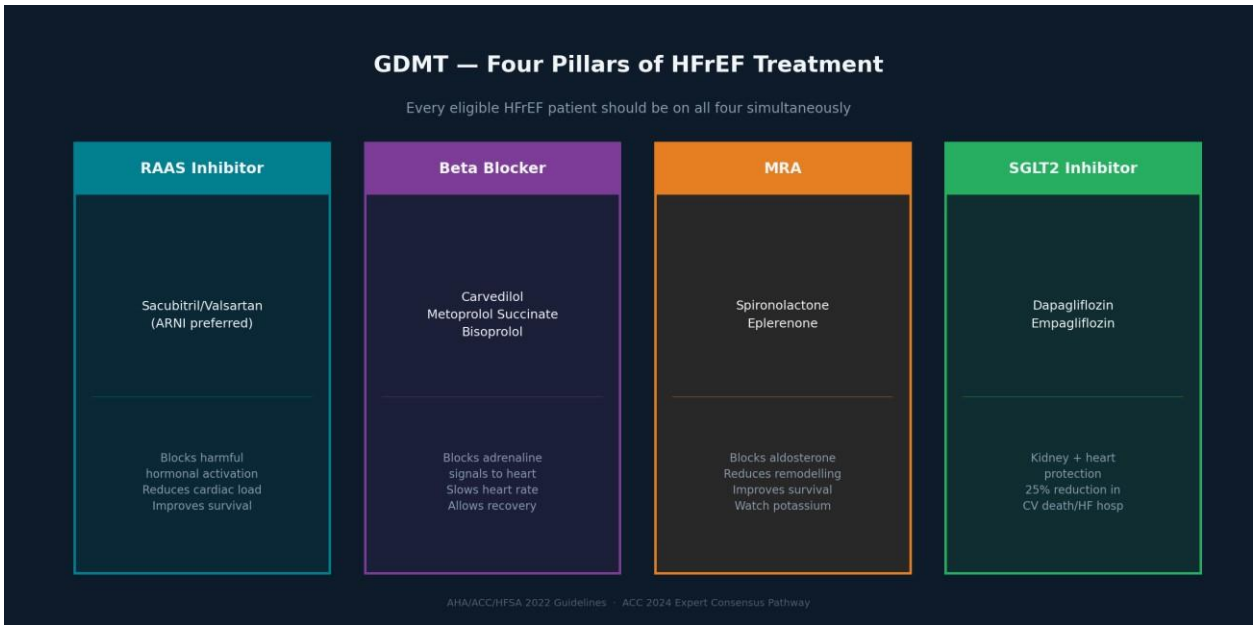

RAAS Inhibitors

Drugs in this group include lisinopril, enalapril, ramipril, valsartan, losartan, candesartan, and sacubitril/valsartan which is sold under the name Entresto. These drugs block something called the renin-angiotensin-aldosterone system. When the heart starts struggling, that system kicks in and the body thinks it is helping but what happens is blood vessels tighten up, the body holds onto more fluid, and the heart gets put under even more strain. These drugs stop that chain reaction. Sacubitril/valsartan the ARNI is the preferred one now because the survival data behind it is better than the older options.

Beta Blockers

Three specific beta blockers have been proven in heart failure carvedilol, metoprolol succinate, and bisoprolol. When the heart is struggling, the body sends out adrenaline signals to push it harder and faster. That helps in the short term but after months and years of it the heart muscle takes real damage. Beta blockers block those adrenaline signals from getting through. The heart rate comes down, the load drops, and some patients see their heart function improve over time. Starting them can feel rough at first and patients sometimes feel worse before they feel better. But the long-term benefit is there if they stay on it.

Diuretics

Furosemide and torsemide are the main ones here. These get rid of excess fluids. When the heart cannot pump well enough, fluid backs up into the lungs, the legs, the belly. Diuretics push the kidneys to produce more urine, which pulls that fluid out. They do not extend life the way the other drugs do but they matter a lot for day to-day comfort. Someone badly fluid overloaded cannot breathe lying down, cannot walk far, and feels drained constantly. Getting rid of that fluid changes how they feel quickly.

MRA and SGLT2 Inhibitors

MRA stands for mineralocorticoid receptor antagonists spironolactone and eplerenone are the two main ones. They block a hormone called aldosterone that makes the body hold onto sodium and fluid and slowly causes structural damage to the heart muscle. Blocking it slows damage and improves survival. The main thing to watch is that MRA pushes potassium up, so blood tests need to happen regularly while someone is on it.

SGLT2 inhibitors dapagliflozin and empagliflozin started as diabetes drugs. Researchers were doing routine safety studies on diabetic patients and noticed these patients were having far fewer heart failure events than expected. That was surprising enough that dedicated heart failure trials followed and both drugs showed real benefit even in patients without diabetes. Both are now approved specifically for heart failure. Why exactly they help the heart so much is still being studied, but the trial evidence is strong, and they are now part of the standard four-drug treatment plan.

2. Medications and Chronic Treatment

Managing heart failure long-term is genuinely difficult. The challenge is not just knowing which drugs to use that part is well established. The challenge is getting every patient onto all four classes at the right doses and keeping them there over months and years while the drugs interact with each other and with the patient's other health conditions.

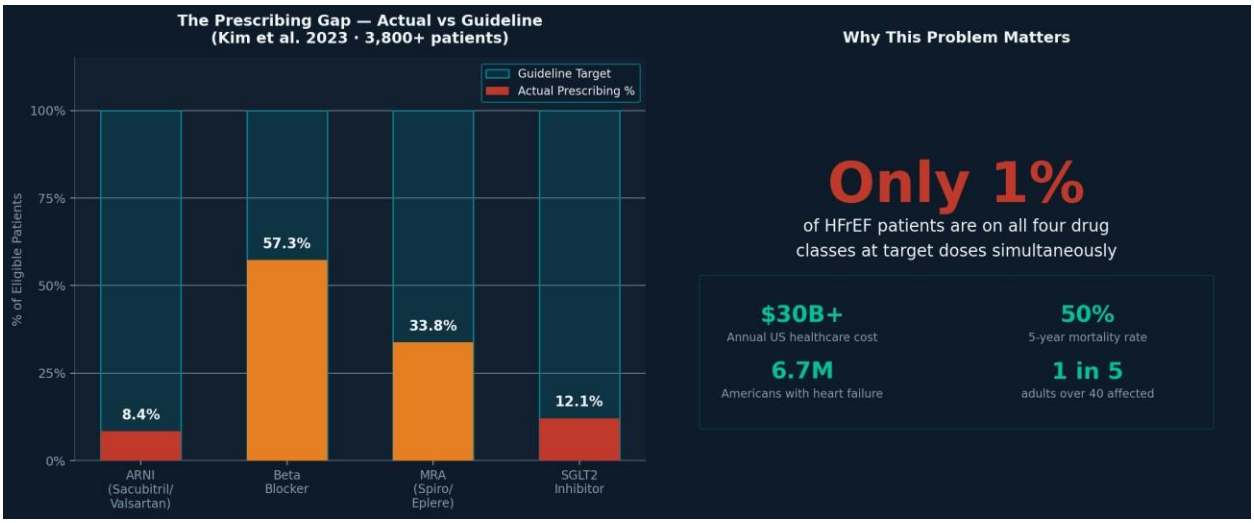

The process of gradually increasing each drug to its target dose is called titration. In an ideal world this would happen quickly over a few weeks after diagnosis or after a hospital discharge. It often takes much longer or never gets completed at all. The CHAMP-HF registry a large study that looked at real-world heart failure care found that only about one percent of HFrEF patients were on all four drug classes at guideline recommended doses at the same time. That number is genuinely shocking given what we know these drugs can do.

The reason titration so hard comes down to the fact that every single drug in this combination has conditions that can stop you from increasing or even maintaining the dose. Beta blockers slow heart rate and can cause it to drop too low. They also lower blood pressure, which can become a problem when the patient is already on a RAAS inhibitor that also lowers blood pressure. RAAS inhibitors can push potassium up to dangerous levels, especially when combined with MRA, which also raises potassium. They can also worsen kidney function in patients whose kidneys were already struggling. Diuretics can cause the opposite potassium problem dropping it too low and can also cause dehydration and kidney stress if the dose goes too high. MRA raises potassium and requires careful monitoring, especially in patients with any kidney impairment. SGLT2 inhibitors are the most straightforward to manage but still cannot be used if kidney function drops below a certain threshold.

So, what the doctor has to do at every visit is look at blood tests and vitals and decide whether each of the four medications can be increased, needs to stay the same, or needs to come down. This juggling act requires frequent clinic visits with

blood tests, and it requires a clinician who knows the patient well. Even then it is genuinely time-consuming and easy to leave incomplete especially when patients have multiple other conditions, and the appointment time is limited.

What makes this even harder is that the consequences of getting the timing wrong go both ways. Moving too slowly means the patient spends months on lower doses than they need, which means worse outcomes and a higher readmission risk. Moving too quickly or in the wrong order can cause a deterioration that sends the patient back to the hospital. The right order matters too diuretics first to get rid of fluid, then RAAS, then beta blockers only once the patient is stable. Starting a beta blocker in a patient who is still fluid overloaded can tip them into acute decompensation.

All of this is what our project is trying to address. The manual, periodic, clinic-visit dependent approach to titration is the root cause of most patients never reaching optimal therapy.

3. Existing Solutions and Titration Research

When we started looking at what was already out there to help with this problem, we found a few different categories of existing solutions.

The most advanced hardware-based solution is CardioMEMS, a small sensor implanted in the pulmonary artery that measures pressure continuously and sends readings to a monitoring centre. When pressure starts rising which is one of the early signs that fluid is accumulating the care team gets an alert and can adjust

diuretics before the patient becomes symptomatic. Studies have shown it reduces heart failure hospitalizations significantly. But it is invasive, requires a procedure to implant, costs a lot, and only measures one signal. It tells you fluid is building up, but it does not tell you what to do about each of the four drug classes.

There are also wearable monitoring platforms from companies like Biofourmis that collect data from wrist-worn devices and use machine learning to predict when a patient is likely to deteriorate. These systems are good at early warning, but they are still primarily alerting systems they tell a clinician to pay attention to a particular patient; they do not actually guide the medication decision. The clinician still must look at the data and decide what to do.

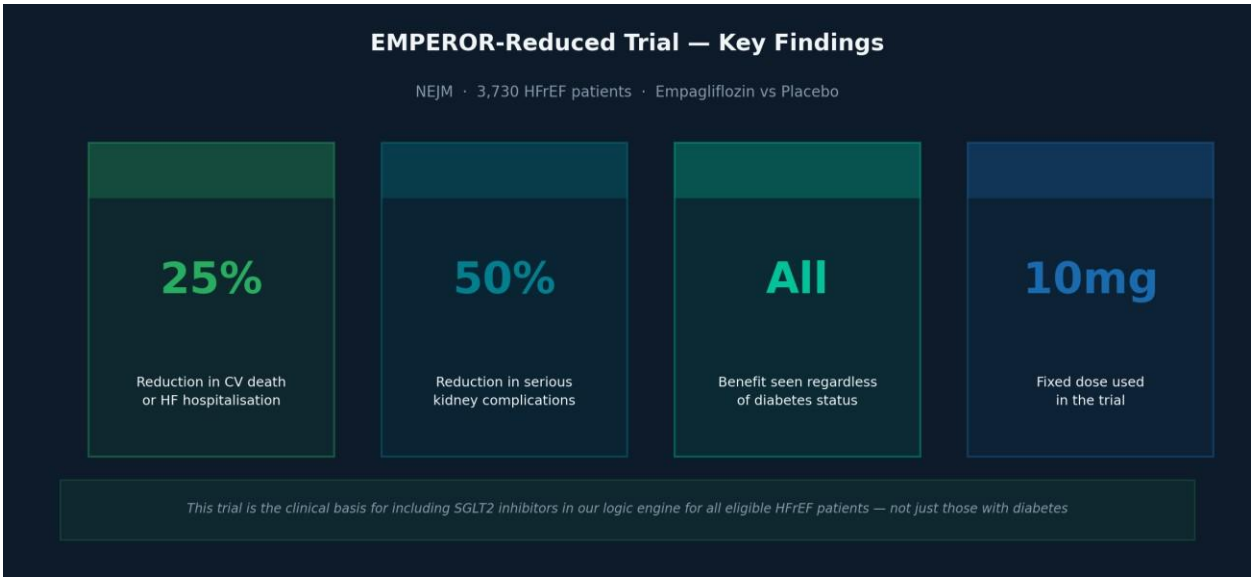

The STRONG-HF trial showed rapid up titration under close supervision checking blood tests and vitals every two weeks and pushing doses up quickly significantly reduced readmission rates compared to standard care. The results were compelling enough that they influenced the most recent heart failure guidelines. But the approach requires frequent clinic visits and a lot of clinician time, which makes it hard to scale.

On the technology side, there are also clinical decision support tools built into electronic health record systems. These typically give doctors reminders or alerts things like flagging that a patient with heart failure is not on a beta blocker, or that their potassium has risen, and the MRA dose might need reviewing. These are helpful nudges, but they are reactive and still require the clinician to do the

thinking. They do not have a built-in understanding of how all the factors interact or what the right next step is for a given combination of vitals and lab values.

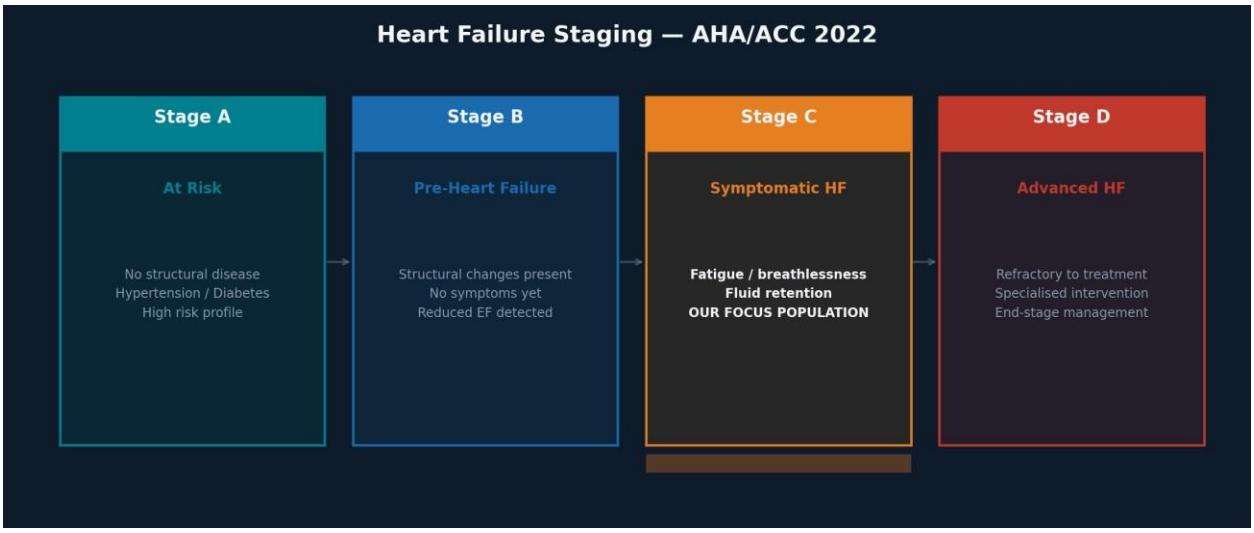

What all these solutions share is that they either monitor without acting, or they alert without guiding, or they guide without a complete picture of all five drug classes together. None of them put continuous sensor data, multi-drug clinical rules, safety gates, and a structured decision output into one system. That is the space we are working in.

4. Titration Research and Algorithm Design

The clinical rules that form the backbone of our logic engine came from seven sources. Two were the main ones the AHA/ACC/HFSA 2022 Heart Failure Guidelines and the ACC 2024 Expert Consensus Decision Pathway. The remaining five covered specific aspects of individual drug classes and monitoring approaches.

From the AHA/ACC guidelines we got the foundational framework the four drug classes, the idea that titration should happen quickly and simultaneously rather than sequentially, the monitoring parameters to watch, and the target doses for each drug. The 2024 Expert Consensus Pathway added more detail on how to handle specific situations like low blood pressure, high potassium, and declining kidney function during up titration.

From the diuretic research we learned that creatinine rising during diuresis does not always mean the kidneys are being damaged sometimes it is just

hemoconcentration as fluid is removed, and it resolves once the patient is euvolemic. But a sharp rise above certain thresholds, especially combined with continued fluid overload and no improvement in congestion, is a signal to stop and reassess. The creatinine threshold of 2.0 we use for escalating rather than just increasing the diuretic came from this literature.

From the beta blocker papers, we confirmed the dry-before-you-try principle and found specific heart rate targets 70 beats per minute for sinus rhythm patients and a lenient rate control target of around 110 for atrial fibrillation patients where strict rate control does not add much benefit. The restriction to bisoprolol and metoprolol only in COPD patients came from these papers too carvedilol is a non-selective beta blocker that can constrict airways in COPD patients.

From the RAAS inhibitor research we got the three safety gates that must all pass before any up titration systolic BP at or above 100, potassium below 5.5, and eGFR at or above 30. These are not arbitrary numbers. Each one represents a point below which the drug starts creating more risk than benefit. We also got the 36-hour washout rule for switching from ACE inhibitors to ARNI there is a risk of serious angioedema if the switch happens too quickly because both drugs affect bradykinin metabolism and overlapping them can cause a dangerous accumulation.

From SGLT2 and MRA research we got two important practical insights. First, the early eGFR dip after starting an SGLT2 inhibitor is expected and does not mean the drug should be stopped it is a hemodynamic effect of the drug reducing glomerular pressure, and it stabilizes. Stopping the drug at this point means the patient loses the benefit unnecessarily. Second, SGLT2 inhibitors tend to lower potassium slightly which provides some protective buffer when managing MRA patients on both drugs can tolerate the MRA a bit longer before potassium becomes a concern.

From the wearable monitoring research, we got the framework for thinking about which signals matter most and how they relate to each other. The concept that impedance-based fluid monitoring is more clinically meaningful than weight alone for detecting fluid accumulation early was central to how we designed the fluid classification step. HRV dropping before other vitals deteriorate came from this literature too.

Taking all this together, the algorithm we designed follows a strict sequence. Emergency safety comes first because no automated recommendation is worth making if the patient is in a life-threatening state. Fluid status comes second because it shapes every drug decision that follows. Diuretics come third because fluid removal is the immediate priority when a patient is overloaded. RAAS comes fourth. Beta blockers fifth with fluid state as a mandatory gate. SGLT2 and MRA sixth with their own independent safety checks. Trajectory lasts because looking at trends rather than isolated readings is what catches slow deterioration before it becomes a crisis.

5. Dataset

So, for the data side of this project, we went with MIMIC-IV. It is a hospital database that MIT put together with Beth Israel Deaconess Medical Center in Boston and you can get access to it through PhysioNet after going through their credentialing process. It goes from 2008 to 2022 and covers around 500,000 hospital admissions. Pretty much everything that got recorded during a patient stay is in there vitals, labs, medications, diagnoses, all with timestamps.

We ended up building our dataset by running five scripts in order. Each one did one thing, and the output fed into the next.

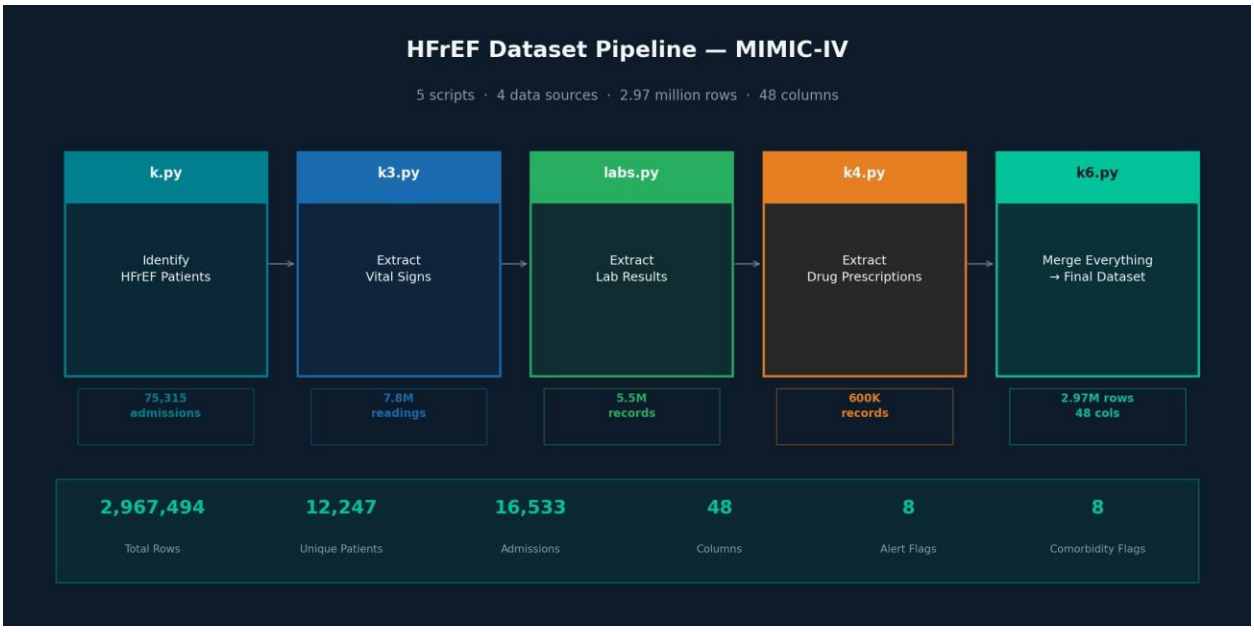

First script - finding the right patients

We went through all the diagnosis records and kept only the admissions where the patient had HFrEF. To do that, we used ICD code. For records from 2015 onwards we used ICD-10 systolic heart failure I5020 to I5023, combined heart failure I5040 to I5043, left ventricular failure I501 and I509, end stage and biventricular failure codes, hypertensive heart failure I110 I130 I132, and rheumatic heart failure I0981. For anything older than 2015 we used ICD-9 instead systolic failure 42820 to 42823, combined failure 42840 to 42843, and hypertensive variants.

We also added eight comorbidity flags to each admission while we were at it — atrial fibrillation, CKD, Type 1 diabetes, Type 2 diabetes, COPD, cardiomyopathy, hypertensive heart disease, and coronary artery disease. Those flags matter later because they change what the logic is allowed to do. Type 1 diabetes means no SGLT2. COPD means only specific beta blockers. CKD means being careful on pretty much every drug that touches kidney function.

After this step we had 75,315 admissions.

Second script -vital signs

This one went into the chartevents table in the ICU data. That is where every single bedside monitor of reading ends up. We pulled heart rate on item ID 220045, systolic BP from cuff readings on 220179 and arterial line on 220050, oxygen saturation on 220277, respiratory rate on 220210, and body weight on 226512. We only kept readings that had actual value and made clinical sense. Ended up with 7,844,613 readings across 12,247 patients.

Third script - lab results

This one hit the labevents table in the hospital data. We pulled creatinine, potassium, sodium, BUN, eGFR, pH, anion gap, hematocrit, platelets, and white blood cell count. Each of those feeds into at least one part of the logic — creatinine and eGFR for the kidney safety checks on diuretics and RAAS, potassium for the RAAS and MRA gates. Came out to 5,537,380 records.

Fourth script -medications

This one went into the prescriptions table and searched by drug name. We grouped everything into seven classes.

Diuretics - furosemide, lasix, torsemide, bumetanide, metolazone, hydrochlorothiazide, chlorthalidone, amiloride, triamterene, acetazolamide.

Beta blockers - carvedilol, metoprolol, bisoprolol, nebivolol, bystolic.

RAAS - lisinopril, enalapril, ramipril, captopril, perindopril, fosinopril, quinapril, trandolapril, benazepril, losartan, valsartan, candesartan, irbesartan, olmesartan, azilsartan, sacubitril, entresto.

SGLT2 - dapagliflozin, empagliflozin, canagliflozin, ertugliflozin.

MRA - spironolactone, eplerenone.

Plus, ivabradine and hydralazine/nitrate combinations.

That gave us 600,723 records across 28,325 patients.

Fifth script - putting it all together

This was the merge step. The sensor data was in long format one row per reading, so we pivoted it to a wide format where each row is one patient at one timestamp with all the sensor values in their own columns. Same thing for the lab data except we rounded the lab timestamps to the nearest hour first so they would line up with the sensor timestamps when joining. Drug records came in on admission ID, so every row knew what medications were being given during that stay. Comorbidity flags joined the same way.

Then we added eight pre-calculated alert flags directly to every row SpO2 below 90, SBP below 90, potassium above 5.5, potassium above 6, creatinine above 3.5, heart rate below 50, heart rate above 100, and SpO2 below 95. Those sit there ready to go, so the logic engine does not have to recalculate them on every row.

What came out

2,967,494 rows. 12,247 unique patients. 16,533 admissions. 48 columns.

Each row is one patient at one specific moment in time, and everything needed to make a medication decision is sitting right there in that row sensor readings, lab values, drug doses, comorbidity flags, and alert flags all in one place.

6. Final Logic and Code

The logic engine is a Python function called run logic that takes one row from the dataset as input and returns a dictionary of decisions one for each medication class, plus an emergency flag, a fluid state classification, and a trajectory assessment.

The function starts by reading all the relevant values from the row. Heart rate, systolic blood pressure, SpO2, respiratory rate, creatinine, potassium, eGFR, and the comorbidity flags for atrial fibrillation, Type 1 diabetes, and COPD. A helper function handles missing values so that if a reading is absent, the system uses a safe default rather than failing.

Step 1 checks all six emergency conditions. If SpO2 is below 90, or systolic BP is below 90, or potassium is above 6, or creatinine is above 3.5, or heart rate is below 40, or eGFR is below 15 any one of those being true causes the function to set all medication steps to HOLD_EMERGENCY, record the specific values that triggered it, set the alert flag to true, and return immediately. Nothing else runs. This is the most important design decision in the whole system safety must be able to stop everything at any point.

Step 2 classifies fluid status using respiratory rate as a proxy since thoracic impedance data is not available in MIMIC-IV. Respiratory rate above 22 breaths per minute means WET. Below 16 means DRY. Between 16 and 22 is BORDERLINE. This classification gets stored and passes into every subsequent step.

Step 3 decides the diuretic action based on the fluid state and kidney safety. If the patient is wet and creatinine is above 2.0, the decision is ESCALATE the patient needs more diuresis, but the kidneys are struggling, and this needs clinical oversight. If wet and potassium is below 3.5, the decision is REDUCE the diuretic is pulling too much potassium out. If wet with neither of those problems, the decision is INCREASE. If dry and creatinine is above 1.5, the decision is REDUCE suggesting over-diuresis. If borderline or stable dry, it is HOLD.

Step 4 decides the RAAS action using three gates. Systolic BP must be at or above 100. Potassium must be below 5.5. eGFR must be at or above 30. If a value is missing, that gate is treated as passing because we cannot penalize missing data. If all three pass, the decision is UPTITRATE. If any fail, the decision is HOLD with the specific failing values recorded in the output, so it is completely clear why the decision was made.

Step 5 decides the beta blocker action. Fluid state is the first check if the patient is wet or borderline, the decision is immediately SKIP with the dry-before-you-try reason recorded. If dry, the function sets the heart rate target 110 for atrial fibrillation patients, 70 for everyone else. If heart rate is above that target and blood pressure is acceptable, the decision is UPTITRATE. If the patient also has COPD, the output notes that only bisoprolol or metoprolol are appropriate. If the heart rate is between 50 and 60, the decision is HOLD. If the heart rate is below 50, the decision is REDUCE.

Step 6 handles SGLT2 and MRA independently in the same step. For SGLT2, Type 1 diabetes is checked first and if present, the output is CONTRAINDICATED. If eGFR is below 20, it is HOLD. Otherwise, it is MAINTAIN/START 10mg. For MRA, all three potassium below 5.0, eGFR at or above 30, and creatinine at or below 2.5 must be satisfied. If all pass the output is MAINTAIN/ADD. If potassium is above 5.5, it is REDUCE. Otherwise, HOLD.

Step 7 checks trajectory by looking at heart rate and blood pressure together in the same row. If heart rate is above 100 and systolic BP is below 100 simultaneously, the trajectory is WORSENING that specific combination is a hemodynamic stress pattern that is clinically significant for decompensating heart failure. If heart rate is below 80 and SBP is above 110, the trajectory is IMPROVING. Everything else is STABLE.

The function returns all of these in a dictionary. The code then applies this function to every row in the dataset using pandas to apply, which processes all 2.97 million rows and returns a results data frame with one decision row per timestamp. The results are saved to a CSV file.

When we ran the code on the first 100 patients the system produced clear output distributions a meaningful number of EMERGENCY flags on rows where values

were genuinely critical, WET classifications driving diuretic increases, RAAS being held on rows where blood pressure was low, beta blockers being skipped wherever fluid state was not dry, and WORSENING trajectory flags appearing on rows where high heart rate and low blood pressure were occurring together. The outputs made clinical sense across all seven steps.

7. System Flowchart

The flowchart maps the complete path that every row in the dataset takes when the logic engine runs on it.

It starts at the top with a START node that feeds into the input reading block. This block shows all the values being pulled into the function heart rate, systolic BP,

SpO2, respiratory rate, from the sensor columns, creatinine, potassium, eGFR from the lab columns, drug doses from the medication columns, and the comorbidity flags.
Everything the system needs is read in before any decision is made.

From the input block, the flow moves immediately to the emergency gate check. This is the first decision point, and it sits above everything else. If any of the six emergency conditions are true, the flow branches out to an emergency exit block on the side showing that all medication steps are set to HOLD; the alert reason is recorded, and the function returns. The main flow never continues past this point if an emergency is detected. The visual design here uses red to make it clear this is a stop condition and not just a hold.

If the emergency check passes, the flow moves down to the fluid classification block. This is where the respiratory rate determines whether the patient is WET, BORDERLINE, or DRY. The flowchart shows three parallel cards side by side, one for each fluid state, each listing what that state means for the downstream decisions. WET shows that diuretics should increase, and beta blockers should be skipped. BORDERLINE shows that both should be held, and monitoring should continue. DRY shows that beta blocker titration is now on the table, and the system can proceed through the remaining steps.

From the fluid classification, the flow moves through each drug decision step in order. Diuretics, then RAAS, then beta blockers, then SGLT2 and MRA together. Each step is a separate block with its own colour and its own rules listed inside. The colour coding is consistent throughout red for the emergency step, teal for classification, blue for diuretic, green for RAAS, purple for beta blocker, orange for SGLT2 and MRA, and mint green for trajectory.

After all, five drug steps, the trajectory check runs. This block looks at heart rate and blood pressure together and produces a WORSENING, IMPROVING, or STABLE assessment for the row.

From the trajectory, the flow moves to the output collection block. This shows all seven decision outputs being gathered emergency flag, fluid state, diuretic action, RAAS action, beta blocker action, SGLT2 and MRA actions, and trajectory. These all become columns in the output row.

The flow ends at a LOG and REPEAT node which represents the system writing the results for this row and moving to the next one. In a real deployment, this would also represent waiting for the next monitoring cycle before running again.

The key design feature of the flowchart is that the emergency gate sits entirely above everything else and has its own exit path. Nothing no drug decision, no trajectory check, no output happens if an emergency gate fires. Safety exits the flow completely rather than being one consideration among many. And fluid state acts as the central junction that the diuretic path and the beta blocker path both pass through, making it visually clear why it is the most important classification in the whole system.

8. Results and Final Output

Once the logic engine ran the first 100 patients from the dataset, the results came back in a way that made clinical sense, which was a good sign.

On the emergency gates side, a meaningful number of rows came back flagged as EMERGENCY. When we looked at what triggered those flags, it was mostly low SpO2 readings, critically high potassium, and very low blood pressure. These were real ICU patients with real acute conditions, so seeing those flags fire on those rows confirmed the gates were working the way they should.

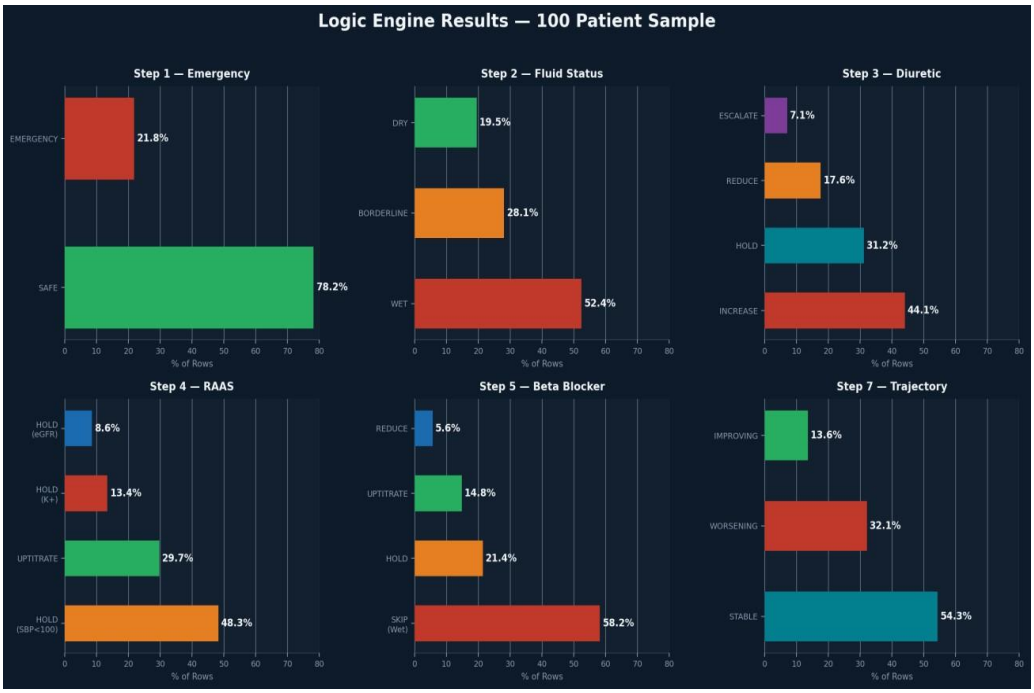

The fluid classification split out across all three states. A significant portion of rows came back as WET which lines up with what you would expect from an ICU heart failure population these are sick patients who are often admitted precisely because of fluid overload. A smaller proportion came back DRY, and the rest sat in the BORDERLINE range.

For diuretics the most common output on wet rows was INCREASE which again is exactly what you would expect fluid overloaded patients need more diuresis. On dry rows with rising creatinine, the REDUCE flag came through correctly picking up over-diuresis scenarios.

RAAS showed a high proportion of HOLD outputs and when we looked at why it was mostly SBP below 100. That makes sense for an ICU population where hypotension is common, and it confirms the blood pressure gate is doing its job of stopping unsafe RAAS up titration.

Beta blockers were skipped on most rows because most rows were either WET or BORDERLINE and the dry-before-you-try rule kicked in. On the DRY rows where heart rate was above targeting the UPTITRATE output came through correctly.

SGLT2 came back as MAINTAIN/START on the majority of rows where no contraindication was present. The CONTRAINDICATED flag appeared correctly on rows where the Type 1 diabetes comorbidity flag was set.

MRA came back as MAINTAIN/ADD on most rows where potassium and kidney function were within a safe range. On rows where potassium was pushing above 5.5, the REDUCE flag fired and on rows above 6.0 the STOP output appeared.

Trajectory showed WORSENING on rows where high heart rate and low blood pressure were happening at the same time that hemodynamic stress pattern the logic checks for. IMPROVING appeared in rows where the heart rate was low and blood pressure was good. Most rows came back to STABLE, which is the expected distribution across a mixed ICU population.

The total alert count across the 100-patient sample confirmed that a real and clinically appropriate number of rows were flagging serious conditions not so many that the system would be crying wolf constantly and not so few that it would be missing things.

The final output was saved as a CSV file with one row per patient timestamp and one column per decision step1 through step7 plus the alert flag and alert reason. That structure means the results can be reviewed row by row, filtered by decision type, or aggregated across patients to look at population-level patterns.

9. Challenges and Limitations

A few things came up during the project that are worth being honest about.

The biggest one is the fluid status proxy. Our logic was designed around thoracic impedance data from a wearable patch that gives a direct read on how much fluid is in the lungs. MIMIC-IV does not have that. What it does have is a respiratory rate which we used as a proxy elevated respiratory rate is a sign of fluid overload and correlates reasonably well with the clinical picture. But it is not the same thing. A patient could have a high respiratory rate for reasons other than fluid infection, anxiety, and pain. So, the fluid classification in our current implementation is an approximation, and the results need to be interpreted with that in mind.

The second limitation is that SGLT2 coverage in the dataset is thin. Dapagliflozin only got its HFrEF approval in 2020 and MIMIC-IV ends in 2022, so there are only two years of data where SGLT2 prescribing would show up. Most admissions in the dataset predate that approval. So, the SGLT2 decision outputs in our results

are largely based on eligibility checking rather than actual dose management which limits what we can say about that drug class from this data.

The third one is that we are working with retrospective ICU data. Real medication titration for heart failure happens mostly in outpatient settings over weeks and months after a hospital discharge. The ICU setting in MIMIC-IV is a different clinical context patients are sicker, more monitored, and managed by teams rather than outpatient cardiologists. The logic translates but the environment it was designed for, and the environment the data comes from is not perfectly matched.

The fourth thing is that our trajectory check in the current code is simplified. The full logic design called for tracking trends across three consecutive readings per patient. In the current implementation, we check heart rate and blood pressure within a single row because building a proper rolling window across the 2.97 million row dataset requires more engineering work. That is the next thing to fix.

10. Next Steps

A few things are lined up for the next phase of the project.

The most immediate one is fixing the trajectory logic to properly track trends across consecutive readings rather than just looking at a single row. That means building a rolling window calculation that groups readings by patient and admission and checks whether the last three are all moving in the wrong direction.

After that we want to start comparing what the logic engine recommended against what happened in the data. For each patient row, we can look at what our system said to do and then look at what the real clinical team did and measure the agreement. Where there is disagreement, those are the interesting cases either the system got it wrong, and we need to fix a rule, or the real clinical team did something suboptimal, and our system was right.

The prediction model conversation we had with the client is also moving forward. The plan is to build a model on top of the existing dataset that predicts things like deterioration risk or readmission likelihood using the features already in the 48column dataset. That would sit alongside the rule-based logic rather than replacing it the GOFAI engine handles the medication decisions, and the ML model handles the forecasting.

We also need to get the impedance patch data situation sorted out. Either we find a dataset that has actual thoracic impedance readings, we can combine with MIMICIV, or we spend more time validating that respiratory rate is a good enough proxy and document the limitations clearly. That decision affects how confident we can be in the fluid classification of outputs.

Longer term the goal is to get the system to a point where it can be tested in a real clinical workflow either in a pilot with actual patients or in a simulation environment where clinicians can interact with the outputs and give feedback on whether the recommendations make sense.

11. Conclusion

What we set out to do with this project was build something that addresses a real and well documented problem the gap between what heart failure guidelines say patients should be getting and what they receive in practice. The numbers in the literature are stark. Only around one percent of HFrEF patients is on all four drug classes at the right doses at the same time. A lot of that comes down to the fact that titration is manual, periodic, and dependent on clinic visits that happen weeks or months apart.

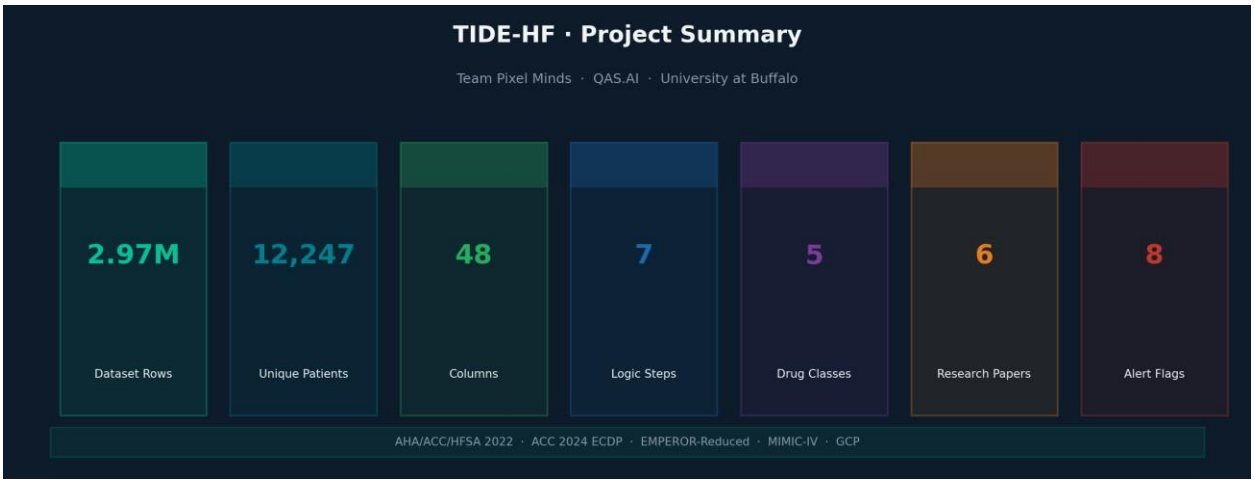

What we build is a system that reads patient data continuously, applies clinical rules that come directly from established guidelines, and produces a medication recommendation for every single timestamp in the dataset. It checks for emergencies first, so it never suggests when the patient is in a dangerous state. It classifies fluid status because that single variable shapes most of what follows. It applies to drug specific safety checks before recommending any dose change. And it flags trends so that slow deterioration does not go unnoticed just because no single reading crossed a critical threshold.

We built the dataset from scratch using five scripts on MIMIC-IV data, ending up with just under three million rows covering over 12,000 unique patients. We wrote the logic engine in Python and ran it on the data. We showed the client the results. The output made clinical sense, and the system behaved the way it was supposed to.

There is still work to do. The trajectory logic needs improvement. The fluid proxy needs validating or replacing. The prediction model layer needs to be built. But the core of what we set out to create a rule-based, explainable, clinically grounded medication decision engine is working and has been tested on real data.

The project has been a genuinely good learning experience. Working with a real clinical dataset, building logic from actual guidelines, dealing with missing data and proxy signals and edge cases all of that pushed us in ways that classroom work does not. And having a real client in QAS.AI who engaged seriously with the work and gave real feedback made a difference to how we approached every decision we made along the way.