

CoralMD: A Multi-Modal Dashboard for Preventive and Personalized Medicine

Tanner DeGrazia

Pomona College

Claremont, California, USA

trdh2022@mymail.pomona.edu

Abstract

Modern healthcare systems remain largely reactive, addressing disease only after symptoms appear. At the same time, genomic sequencing, wearable sensors, and electronic health records (EHRs) now make it possible to track individual risk continuously. CoralMD is a prototype multimodal dashboard that integrates three data streams, genomic variants, physiological signals from wearables, and EHR data, into a single, interpretable view of patient risk. Instead of building a highly optimized black box model, CoralMD uses a transparent, rule based scoring scheme and narrative visualizations to surface how each stream of data contributes to future risk across four domains (cardio & cerebrovascular, metabolic, neurodegenerative, and cancer). The system is implemented as a Streamlit application with patient and clinician facing views, backed by a small specifically curated synthetic case study dataset that mimics realistic encounters and variant annotations to display the potential of the prototype. This paper describes the motivation, system architecture, data design, and heuristic risk model underlying CoralMD, presents a case study illustrating how the dashboard supports forward reasoning, and reflects on ethical challenges around risk visualization, normalization, and bias. All towards the effort of shifting healthcare to a more individualized and proactive system.

Keywords

personalized medicine, clinical decision support, chronic disease, multimodal learning, machine learning, interpretability, ethics

1 Introduction

Cancer, cardiovascular disease, neurodegenerative disorders, and diabetes are among the leading causes of death in the United States. Yet care delivery for these conditions is still mostly reactive. Clinicians treat disease after it has been diagnosed rather than acting early based on trajectories of risk. The core promise of personalized or “precision” medicine is to reverse this pattern by combining genomics, continuous physiology measures, and clinical history to anticipate disease and intervene early [3, 16].

In practice, this vision has not yet been fully realized or implemented at the point of care. Clinical workflows remain fragmented across institutions and data systems. Genomic reports may live in one portal, Apple Watch or Fitbit data in another, and EHR data in yet another. Even when machine learning models are available, they are often tuned for accuracy rather than interpretability, making it difficult for clinicians to understand why a model recommends an action or if they should trust its outputs in these high stake environments. The result is a gap between increasingly rich data about

individuals appearing all over the place and the tools clinicians actually have to reason with that data lagging behind.

CoralMD is an attempt to explore what a more proactive, multimodal system could look like at the prototype level. The project takes inspiration from “Medicine 3.0”, a phrase popularized by longevity scientist, Peter Attia to describe a shift from reactive, disease centered care to proactive, prevention focused care, all grounded in long term risk trajectories [1]. Rather than striving for a production ready platform right now, CoralMD focuses on three goals:

- (1) **Integrate multiple data streams**, genomics, wearables, and EHR, into a coherent data model and user interface.
- (2) **Provide transparent, explanation first risk estimates** built from simple, inspectable heuristics rather than a large black box model.
- (3) **Investigate ethical and representational questions** around how risk is visualized, compared, and communicated to clinicians and patients.

The resulting system is a Streamlit dashboard with separate patient and practitioner views and a central “Storyline” screen that summarizes how inherited baseline, clinical history, and everyday behavior push risk up or down in four disease domains. Under the hood, a tight multimodal scoring function uses EHR diagnoses and labs, a number of handpicked genetic variants, and wearable summaries (steps, sleep, resting heart rate, heart rate variability) to produce domain level scores normalized to the unit interval.

Because there is no publicly available dataset that jointly contains dense genomics, wearables, and EHR for the same individuals, the current CoralMD prototype uses a synthetic case study dataset designed to showcase realistic patterns and use cases. All while the architecture is designed to scale to real world sources such as MIMIC-IV, ClinVar, and gnomAD (which will be implemented in future work). Thus, the emphasis in this capstone project is on multimodal integration, interpretability, and ethical visualization, not on maximizing predictive performance (yet).

2 Background and Related Work

Personalized medicine has increasingly shifted toward the integration of multimodal health data, including genomic sequences, wearable sensor streams, and EHRs, to create forward looking, individualized risk predictions. The related work that motivated CoralMD falls into four themes: (1) clinical trajectory prediction, (2) automation and human/AI collaboration, (3) decision support models for specific workflows, and (4) challenges of interpretability, fairness, and multimodal integration.

2.1 AI and Data Driven Clinical Trajectory Prediction

A growing body of work models disease trajectories rather than static outcomes, emphasizing quality of life and early intervention. Hou et al. propose a hybrid CNN–GRU model for real time sepsis prediction that updates estimates every two hours to distinguish between fast decline and slow recovery phenotypes [8]. This phenotype centered approach illustrates how real time physiological signals can enable actionable early interventions, an idea central to CoralMD’s focus on trajectory aware risk models.

Du et al. extend this idea with a graph based representation of intensive care unit (ICU) patient state, embedding relationships between organs, interventions, and outcomes to support “what if” reasoning [4]. Trevena et al. similarly construct a digital twin framework that maps patient pathways into a directed acyclic graph stored in Neo4j, enabling exploratory simulation of physiological responses [17]. Together, these works highlight the importance of modeling how health evolves over time, motivating CoralMD’s storyline and domain based view of risk.

2.2 Automation, Clinical Workflow Support, and Human/AI Collaboration

A second theme focuses on the use of AI and robotics to automate clinical tasks while maintaining clinician control. Huang et al. develop an imitation learning based robotic ultrasound system capable of following clinical imaging protocols autonomously [9]. Bernardes et al. create a robotic needle trajectory correction mechanism that compensates for soft tissue deflection in MRI guided procedures, achieving millimeter level precision without reinsertion [2]. These systems demonstrate an emerging paradigm of closed loop, data responsive systems that assist rather than replace human expertise.

Sanz-Peña et al. present a fully 3D printed ankle exoskeleton that delivers individualized torque assistance based on patient specific gait patterns [15]. Their work underscores the value of adaptable, user specific assistive technology, an insight that CoralMD adopts by tailoring risk explanations to each individual’s genomic, clinical, and behavioral context.

2.3 Clinical Decision Support and Risk Focused Modeling

Other work focuses on predictive modeling for targeted clinical decision support. Zhu et al. develop a transition flow model to predict 7 and 30 day revisit risk for elderly diabetes patients with fall related injuries, helping identify intervention points [18]. Eskandari and Lee apply Markov Decision Processes to optimize postoperative care plans after joint replacement, integrating performance and patient reported outcomes to determine cost effective rehabilitation schedules [6]. These models demonstrate how structured clinical data can reveal decision relevant patterns, mirroring CoralMD’s emphasis on actionable contributions rather than opaque scores.

Li et al. examine system level decision making through an optimization model that helps hospitals allocate resources between in person and evisit services [13]. Their results show how digital

health tools can improve efficiency and accessibility, contextualizing how a system like CoralMD would need to fit into real clinical workflows.

2.4 Ethical, Interpretability, and Generalization Challenges

Across these studies, several challenges recur. High performing models are often opaque, limiting clinician trust. Datasets are frequently institution specific or demographically skewed, raising concerns about fairness and generalization. And very few systems attempt true multimodal integration. Most focus on either physiological signals, robotic control, or clinical outcomes in isolation.

CoralMD responds to these gaps in three ways. First, instead of a deep neural network, it uses a deliberately simple, rule based risk model whose logic can be inspected line by line. Second, it foregrounds interpretability and uncertainty in the UI, using narrative explanations and tables rather than single traffic light scores. Third, it proposes an architecture able to ingest genomic, wearable, and EHR data in parallel, even though the current prototype uses a synthetic case study rather than full scale clinical data.

3 System Design and Implementation

This section describes the current CoralMD implementation: the overall architecture, the dashboard views, and the underlying multimodal risk model.

3.1 Architecture Overview

CoralMD is implemented as a Streamlit app with multiple pages. Figure 1 summarizes the data and component flow.

- **Data ingestion layer.** Three CSV files (`geneFive.csv`, `ehrFive.csv`, `wearFive.csv`) provide synthetic genomic, EHR, and wearable data for a small panel of five cases. Each file is loaded into a pandas DataFrame and passed to specific processing functions.
- **Risk modeling layer.** The module `ml_models.py` implements a toy multimodal risk function, `compute_multimodal_risk`, which aggregates scores from EHR diagnoses and labs, genomic variants, and wearable summaries into four domain level risk scores and an overall score.
- **Visualization layer.** The module `visualizations.py` uses Plotly to render time series plots, bar charts of domain scores, contribution tables, and simple narrative text blocks. These components are shared across pages so that the same underlying risk object can be explored from different angles.
- **App pages.** The top level file `app.py` coordinates navigation between Streamlit pages: Patient Home, Practitioner Home, EHR Explorer, Wearables Explorer, Genomics Explorer, and Storyline View.

3.2 Dashboard Views

Patient Home. The Patient Home page provides a simple entry point for a patient (or instructor) to select one of the synthetic patients and view a concise summary of their data. The page surfaces demographic information, recent wearable statistics (average steps, sleep, resting heart rate, HRV), and high level risk scores. In the




Figure 1: High level architecture of the CoralMD prototype. Synthetic genomic, EHR, and wearable CSV files feed into a rule based multimodal risk model and Streamlit visualization layers.

current prototype, patients cannot upload their own data, but the UI is designed to support that extension.

Practitioner Home. The Practitioner Home page is oriented around summarizing multiple patients at once. For this course project, it uses a local SQLite database to create risk scores and notes across sessions, and then presents a table of patients with filters and links into deeper views. For the selected patient, it shows tabbed summaries for EHR, genomics, physiology, and storyline notes, mirroring the information hierarchy of the Storyline View.

EHR Explorer. The EHR Explorer focuses on diagnoses and labs. For a given synthetic patient, it displays a list of encounters (e.g., annual physical, overweight consultation, borderline hypertension), along with lab values such as LDL cholesterol, systolic and diastolic blood pressure, and HbA1c. Time series plots show how key labs have evolved across encounters, highlighting patterns like rising LDL over time.

Wearables Explorer. The Wearables Explorer shows daily steps, sleep hours, resting heart rate, and HRV over a two week window (Figure 2). Summary metrics emphasize the average number of steps per day (around 4,700 for the case study patient), average sleep (about 5.8 hours), and elevated resting heart rate. These views help motivate why wearables are a valuable complement to EHR and genomics for ongoing risk monitoring.

Genomics Explorer. The Genomics Explorer exposes a small table of handpicked variants for each synthetic patient, including rs7903146 (TCF7L2), rs9939609 (FTO), rs2228671 (LDLR), rs2819742 (APOE), rs17782313 (MC4R), and rs5186 (AGTR1). For each variant, the table lists the gene, a qualitative effect label (pathogenic, likely pathogenic, benign), a pathogenicity score between 0 and 1, and the associated condition (e.g., type 2 diabetes susceptibility, obesity risk, elevated LDL cholesterol, Alzheimer's disease susceptibility, hypertension risk).




Figure 2: Wearables Explorer view showing daily steps, sleep, resting heart rate, and HRV for the synthetic case study patient.



Figure 3: Storyline View summarizing how genomics, EHR, and wearables contribute to domain level risk for the synthetic case study patient.

Storyline View. The Storyline View is the central interpretability feature of CoralMD. It aggregates scores from all three modalities and presents them as:

- a stacked bar chart showing how each data stream (EHR, genomics, wearables) pushes each disease domain up or down;
- a factor level table listing which diagnoses, labs, variants, and wearable features contributed to risk;
- a narrative explanation organized into three phases, inherited baseline, clinical snapshot, and everyday behavior, with changes in domain risk shown as delta values; and
- a small “what if sandbox” that lets the user toggle hypothetical changes (e.g., bringing LDL < 130 mg/dL or increasing steps above 7,000/day) and see how domain scores would change under the toy model.

4 Data

CoralMD is designed with real multimodal datasets in mind, but the current implementation uses a small synthetic dataset for a case study patient. This section explains both the design goals and the concrete data used in the prototype.

4.1 Design Goals and Real World Sources

In a full deployment, CoralMD would draw on three types of data:

- **Genomic data** from genotype arrays or whole genome sequencing, annotated with resources such as the GRCh38 reference, 1000 Genomes, ClinVar, and gnomAD [11, 12].
- **Wearable data** from devices like Apple Watch or Fitbit, including heart rate, step counts, sleep, and heart rate variability [7].
- **EHR data** from hospital information systems, such as the MIMIC-IV database of identified ICU patients [10], which provides demographics, diagnoses, encounters, and laboratory measurements.

Because no single dataset combining all three modalities is publicly available for the same individuals, the current version of CoralMD uses synthetic data shaped by these sources rather than directly using them.

4.2 Synthetic Case Study Dataset

Three small CSV files encode multimodal data for a single patient and a few additional placeholder rows.

Wearable data. The file `wearFive.csv` contains 14 days of daily summaries with columns: `date`, `steps`, `sleep_hours`, `resting_hr`, and `hrv`. For the main case study patient, the average number of steps per day is approximately 4,737, average sleep is 5.8 hours, resting heart rate is around 80 bpm, and HRV averages in the low 30s. These values place the patient below common activity targets, with short sleep and slightly elevated resting heart rate.

EHR data. The file `ehrFive.csv` encodes four encounters for subject 9, including an annual physical, a visit labeled “Overweight”, and two visits related to elevated blood pressure and borderline hypertension. Each row includes `subject_id`, `age`, `gender`, `encounter_id`, `encounter_date`, `diagnosis`, `lab_name`, and `lab_value`. Lab values include LDL cholesterol (ranging from 125 to 155 mg/dL), HbA1c (5.4–5.8%), BMI (28.5), and systolic/diastolic blood pressure. These features are chosen to capture common cardiometabolic risk factors in a concise format.

Genomic data. The file `geneFive.csv` lists six variants. `TCF7L2` (rs7903146), `FTO` (rs9939609), `MC4R` (rs17782313), and `LDLR` (rs2228671) are labeled as pathogenic or likely pathogenic for type 2 diabetes susceptibility, obesity risk, or elevated LDL cholesterol. An `APOE` variant is labeled benign with low pathogenicity, representing baseline Alzheimer’s susceptibility, and an `AGTR1` variant contributes to hypertension risk. Each row provides `variant_id`, `gene`, `effect` (pathogenic vs. likely pathogenic vs. benign), a `pathogenicity_score` between 0 and 1, and a textual condition.

4.3 Data Flow in the Prototype

Within the Streamlit app, these three tables flow through the following steps:

- (1) Data are loaded into pandas DataFrames and normalized (column names lowercased, dates parsed).
- (2) EHR rows are grouped by patient and encounter; wearable rows are sorted by time and restricted to the most recent 30 days.

- (3) The risk model functions in `ml_models.py` compute domain level scores from each modality and then aggregate them.
- (4) Visualization functions in `visualizations.py` use the resulting scores and raw data to render plots, tables, and narratives on the dashboard pages.

Although the dataset is small, using realistic column names and value ranges makes it straightforward to swap in real data in the future without changing the modeling code.

5 Multimodal Risk Model and Analysis

Because CoralMD is an early stage prototype, the “analysis” in this write up focuses on the design and behavior of the toy multimodal risk model, along with a qualitative case study, rather than on large scale statistical measures.

5.1 Toy Heuristic Multimodal Risk Model

The function `compute_multimodal_risk` in `ml_models.py` creates scores from three helper functions:

- `_score_metabolic_from_ehr(df_ehr)`, which uses diagnoses and labs,
- `_score_from_genomics(df_gen)`, which uses variant level annotations, and
- `_score_from_wearables(df_wear)`, which uses recent steps, sleep, resting heart rate, and HRV.

Each helper function returns contributions for four domains: cardio & cerebrovascular, metabolic, neurodegenerative, and cancer.

EHR based scoring. The EHR scoring function normalizes column names and then:

- searches diagnosis strings for keywords like “overweight”, “obesity”, “prediabetes”, and “type 2 diabetes”, adding increasing weights to the metabolic domain;
- extracts lab values for LDL cholesterol and HbA1c and adds additional metabolic risk if LDL is ≥ 130 mg/dL or HbA1c is in the prediabetic (5.7–6.4%) or diabetic ($\geq 6.5\%$) range.

This scheme encodes a simplified version of clinical reasoning: multiple moderately abnormal findings can accumulate into substantial metabolic risk.

Genomics based scoring. The genomics scoring function looks at the variant table and:

- defines a high impact mask as any variant with pathogenicity score ≥ 0.8 or an effect label containing “pathogenic”;
- increments the metabolic score if any high impact variant is associated with diabetes or obesity;
- increments the cardio & cerebro score if any high impact variant is associated with coronary disease or LDL cholesterol;
- increments the neurodegenerative score for Alzheimer’s related variants; and
- applies additional gene level, meaning changes occur if certain genes (`APOE`, `LDLR`, `PCSK9`, `APOB`, `TCF7L2`, `FTO`, `MC4R`) are present.

Wearable based scoring. The wearables scoring function first restricts to the most recent 30 days (if date information is available)

and then computes mean daily steps, sleep, resting heart rate, and HRV. It applies simple thresholds such as:

- steps < 4000 \Rightarrow increase metabolic and cardio & cerebro risk;
- steps between 4000 and 7000 \Rightarrow modest increases;
- average sleep < 6 hours \Rightarrow increase metabolic and neurodegenerative risk;
- resting heart rate > 80 bpm or > 90 bpm \Rightarrow increase cardio & cerebro risk; and
- HRV < 25 ms \Rightarrow small increases in cardio & cerebro and metabolic risk.

Aggregation and normalization. The domain scores from EHR, genomics, and wearables are summed and then divided by hand tuned maximum values (e.g., 12 for cardio & cerebro, 10 for metabolic) to produce normalized scores between 0 and 1. The overall risk score is the mean of the four domain scores.

5.2 Case Study: Synthetic Patient - Johnny Cash

For the main synthetic patient (nicknamed Mr. Cash) with the data described in Section 4. The model produces the following normalized risk scores for each respective chronic disease category:

- cardio & cerebro: ≈ 0.71 ,
- metabolic: 1.00 (saturating the toy scale),
- neurodegenerative: ≈ 0.69 , and
- cancer: 0.00.

The overall risk score is therefore around 0.60 on a 0–1 scale. In the Storyline View, these scores are presented as the starting point for a narrative explanation.

Narrative breakdown. The narrative explanation divides contributions into three phases:

Inherited baseline. Genomic variants in TCF7L2, FTO, MC4R, and LDLR raise baseline metabolic risk and modestly elevate cardio & cerebro and neurodegenerative domains.

Clinical snapshot. EHR diagnoses of overweight and borderline hypertension, combined with rising LDL and a slightly elevated BMI, further increase metabolic and cardio & cerebro risk.

Everyday behavior. Wearable data indicate below target activity (around 4,700 steps/day), short sleep (5.8 hours), and resting heart rate around 80 bpm, pushing metabolic risk to the top of the toy scale and nudging cardio & cerebro and neurodegenerative domains upward.

Importantly, cancer risk remains at baseline, as neither the variants nor the EHR or wearable features in this toy dataset directly relate to cancer.

What if sandbox. The what if sandbox in the Storyline View allows the user to toggle hypothetical improvements, such as increasing steps to at least 7,000/day or lowering LDL below 130 mg/dL. Under the toy model, these changes reduce the metabolic and cardio & cerebro scores, giving a qualitative sense of which levers (physical activity, lipids, weight) are most influential.

5.3 Limitations of the Current Analysis

The current analysis has clear limitations:

- The multimodal risk model is rule based and hand tuned, not trained from large datasets.
- The dataset is small and synthetic, so no claims are made about real world predictive performance.
- The normalization constants and thresholds are illustrative and would require clinical validation.
- There is no formal user study assessing whether clinicians find the explanations helpful.

Nevertheless, the prototype fulfills its primary goal for this course: it demonstrates how heterogeneous data streams can be combined into a single, interpretable risk narrative and surfaces key ethical questions about how risk is visualized and communicated.

6 Ethical Considerations

Designing CoralMD raised ethical questions not only about which data to use but also about how to represent risk and comparison.

6.1 Visualizing Risk Without Over Persuading

Mahony and Hulme's analysis of the IPCC's "burning embers" diagram [14] shows how a seemingly neutral risk visualization can become a powerful rhetorical device. By choosing particular color scales and thresholds, the diagram helped shape public and policy debates about climate danger. Their work underscores that risk visualizations are never purely descriptive; they encode aesthetic and moral decisions.

In CoralMD, this insight motivated several design choices. The dashboard deliberately avoids traffic light style color codes or single "risk scores" that might implicitly tell clinicians what to do. Instead, domain scores are shown on a muted scale with clear uncertainty and are accompanied by textual explanations and factor level tables. The goal is to support interpretation, not to enforce consensus or dictate action.

6.2 Comparison, Normalization, and "Normal" Bodies

Dumit and de Laet's "Curves to Bodies" [5] traces how growth charts, calorie tables, and biomedical curves do more than measure individuals: they construct norms and define which bodies count as healthy or deviant. Presenting a single reference curve can hide important variations across populations and lead to stigmatizing interpretations.

CoralMD confronts this problem by emphasizing within person trajectories and contextualized explanations rather than rigid comparisons to a single "normal" population. For example, the Storyline View focuses on how a patient's own risk changes under different scenarios, rather than ranking them against a fixed percentile. Future versions of the system could incorporate explicit fairness checks across demographic groups and report uncertainty intervals for risk scores to avoid overconfident claims.

6.3 The Stakes of Precision and Transparency

Both Mahony & Hulme and Dumit & de Laet highlight how misleading or oversimplified visualizations can have real world consequences. In personalized medicine, those consequences may involve treatment decisions or long term patient anxiety. For this reason,

CoralMD treats each data comparison and visualization choice as a moral decision. The toy risk model is transparent by design. Clinicians (and patients who care/are able to) can inspect the code to see how scores are computed. Future work would need to extend this transparency to more sophisticated models, for example by using after the fact explanation methods or inherently interpretable architectures.

6.4 Predictive Power as Social Power

Finally, any system that predicts disease risk wield social power. It can influence which patients receive attention, how resources are allocated, and how individuals see their own health trajectories. Following the warnings from Mahony & Hulme and Dumit & de Laet, CoralMD positions itself not as an oracle but as a collaborator. Its design aims to empower human reasoning through clearer representations of data, explicit acknowledgment of uncertainty, and resistance to oversimplification.

7 Conclusion and Future Work

CoralMD is a prototype exploration of what a multimodal, interpretable, preventive medicine dashboard could look like. By combining synthetic genomic variants, EHR encounters and labs, and wearable summaries into a single Storyline View, the system demonstrates how heterogeneous data can be woven into a coherent narrative that highlights modifiable levers of risk.

Although the current risk model is intentionally simple and the dataset is small, the project contributes:

- a concrete architecture for integrating genomics, wearables, and EHR data;
- a heuristic multimodal scoring function that decomposes risk by domain and modality; and
- an ethics informed visualization design that foregrounds interpretability and resists oversimplified traffic light risk scores.

Future work would follow along three directions. First, integrating real multimodal datasets would allow training and validating more sophisticated models while preserving interpretability. Second, deploying CoralMD in a cloud environment (potentially in AWS) with secure data handling would make it accessible beyond the local development environment. Third, conducting user studies with clinicians and patients could empirically evaluate whether the Storyline View and what if sandbox help users reason more effectively about risk and prevention.

Ultimately, the goal is not just to build another dashboard, but to clarify what ethically responsible, data driven personalized medicine might look like in practice, and how data science tools can support, rather than replace, human judgment in the clinic, all gearing towards an effort of overhauling the attitude of our healthcare system.

References

- [1] Peter Attia and Bill Gifford. 2023. *Outlive: The Science and Art of Longevity*. Harmony.
- [2] Mariana C. Bernardes, Pedro Moreira, Dimitri Lezcano, Lori Foley, Kemal Tuncali, Clare Tempany, Jin Seob Kim, Nobuhiko Hata, Julian Iordachita, and Junichi Tokuda. 2024. In Vivo Feasibility Study: Evaluating Autonomous Data-Driven Robotic Needle Trajectory Correction in MRI-Guided Transperineal Procedures. *IEEE Robotics and Automation Letters* 9, 10 (2024), 8975–8982. doi:10.1109/LRA.2024.3455940
- [3] Francis S. Collins and Harold Varmus. 2015. A New Initiative on Precision Medicine. *New England Journal of Medicine* 372, 9 (2015), 793–795.
- [4] Honglin Du, Muxuan Liang, and Xiang Zhong. 2025. A Data-Driven Graph Model Approach to Critical Care Patient Clinical Trajectory Prediction. *IEEE Robotics and Automation Letters* 10, 11 (2025), 11768–11775. doi:10.1109/LRA.2025.3615534
- [5] Joseph Dumit and Marianne de Laet. 2014. Curves to Bodies. In *Routledge Handbook of Science, Technology, and Society*. Routledge, 71–82.
- [6] Hanieh Eskandari and Hae-Kyung Lee. 2022. A Decision Support System on Postoperative Care Management of Total Joint Replacement Patients Using Clinical Outcome Measurement. *IEEE Robotics and Automation Letters* 7, 4 (2022), 9500–9507. doi:10.1109/LRA.2022.3187618
- [7] Shai G. Hershman et al. 2019. Accuracy of Apple Watch Measurements for Heart Rate and Energy Expenditure. *Journal of Sports Sciences* (2019). Example wearable validation study; update with your preferred source.
- [8] Grace Y. Hou, Andrew C. Hanson, Paul J. Schulthe, Amos Lal, Yue Dong, Ognjen Gajic, and Xiang Zhong. 2025. A Hybrid CNN-GRU Model for Real-Time Prediction of Sepsis Clinical Trajectories in the ICU. *IEEE Robotics and Automation Letters* 10, 9 (2025), 9565–9572. doi:10.1109/LRA.2025.3595040
- [9] Yuhui Huang, Wenzheng Xiao, Chen Wang, Hong Liu, Rui Huang, and Zengqi Sun. 2021. Towards Fully Autonomous Ultrasound Scanning Robot with Imitation Learning Based on Clinical Protocols. *IEEE Robotics and Automation Letters* 6, 2 (2021), 2259–2266. doi:10.1109/LRA.2021.3064283
- [10] Alistair E. W. Johnson et al. 2023. MIMIC-IV, a Freely Accessible Electronic Health Record Dataset. *Scientific Data* 10 (2023), 1–10.
- [11] Konrad J. Karczewski et al. 2020. The Mutational Constraint Spectrum Quantified from Variation in 141,456 Humans. *Nature* 581 (2020), 434–443.
- [12] Melissa J. Landrum et al. 2018. ClinVar: Improving Access to Variant Interpretations and Supporting Evidence. *Nucleic Acids Research* 46, D1 (2018), D1062–D1067.
- [13] Ning Li, Fang Yuan, Ying Li, and Wei Lv. 2024. Public Hospital’s Decision Analysis on Providing E-Visit Services. *IEEE Robotics and Automation Letters* 9, 9 (2024), 7940–7947. doi:10.1109/LRA.2024.3490397
- [14] Martin Mahony and Mike Hulme. 2012. The Colour of Risk: An Exploration of the IPCC’s “Burning Embers” Diagram. *Spontaneous Generations: A Journal for the History and Philosophy of Science* 6, 1 (2012), 75–89.
- [15] Ignacio Sanz-Pena, Hyeongcheol Jeong, and Myunghee Kim. 2023. Personalized Wearable Ankle Robot Using Modular Additive Manufacturing Design. *IEEE Robotics and Automation Letters* 8, 7 (2023), 3734–3741. doi:10.1109/LRA.2023.3290529
- [16] Eric Topol. 2019. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books.
- [17] William Trenava, Amos Lal, Simon Zec, Edin Cubro, Xiang Zhong, Yue Dong, and Ognjen Gajic. 2022. Modeling of Critically Ill Patient Pathways to Support Intensive Care Delivery. *IEEE Robotics and Automation Letters* 7, 3 (2022), 7287–7294. doi:10.1109/LRA.2022.3183253
- [18] Wenjun Zhu, Allie DeLonay, Maureen Smith, Pascale Carayon, and Jingjing Li. 2021. Reducing Fall-Related Revisits for Elderly Diabetes Patients in Emergency Departments: A Transition Flow Model. *IEEE Robotics and Automation Letters* 6, 3 (2021), 5642–5649. doi:10.1109/LRA.2021.3082115