1 September 2025

Enabling Healthcare Operations with Agentic AI - DISHA Overview and Architecture

INFOCOMM MEDIA DEVELOPMENT AUTHORITY

Dr Sucheendra K Palaniappan, VP of Data Science & Engineering of SBX Corporation

Sucheendra K. Palaniappan is Vice President of Data Science and Engineering at SBX Corporation and Chief Technology Officer of SBX Technologies, a subsidiary focused on healthcare innovation. He is also a Senior Scientist at the Systems Biology Institute in Tokyo.

He holds a Ph.D. in Computer Science from the National University of Singapore and has primarily worked at the intersection of biomedicine and healthcare, with a focus on translational data science and applied machine learning.

Enabling Healthcare Operations with Agentic AI : DISHA Overview and Architecture

Setting the direction for the future of
Clinical Decision Intelligence (CD_x)

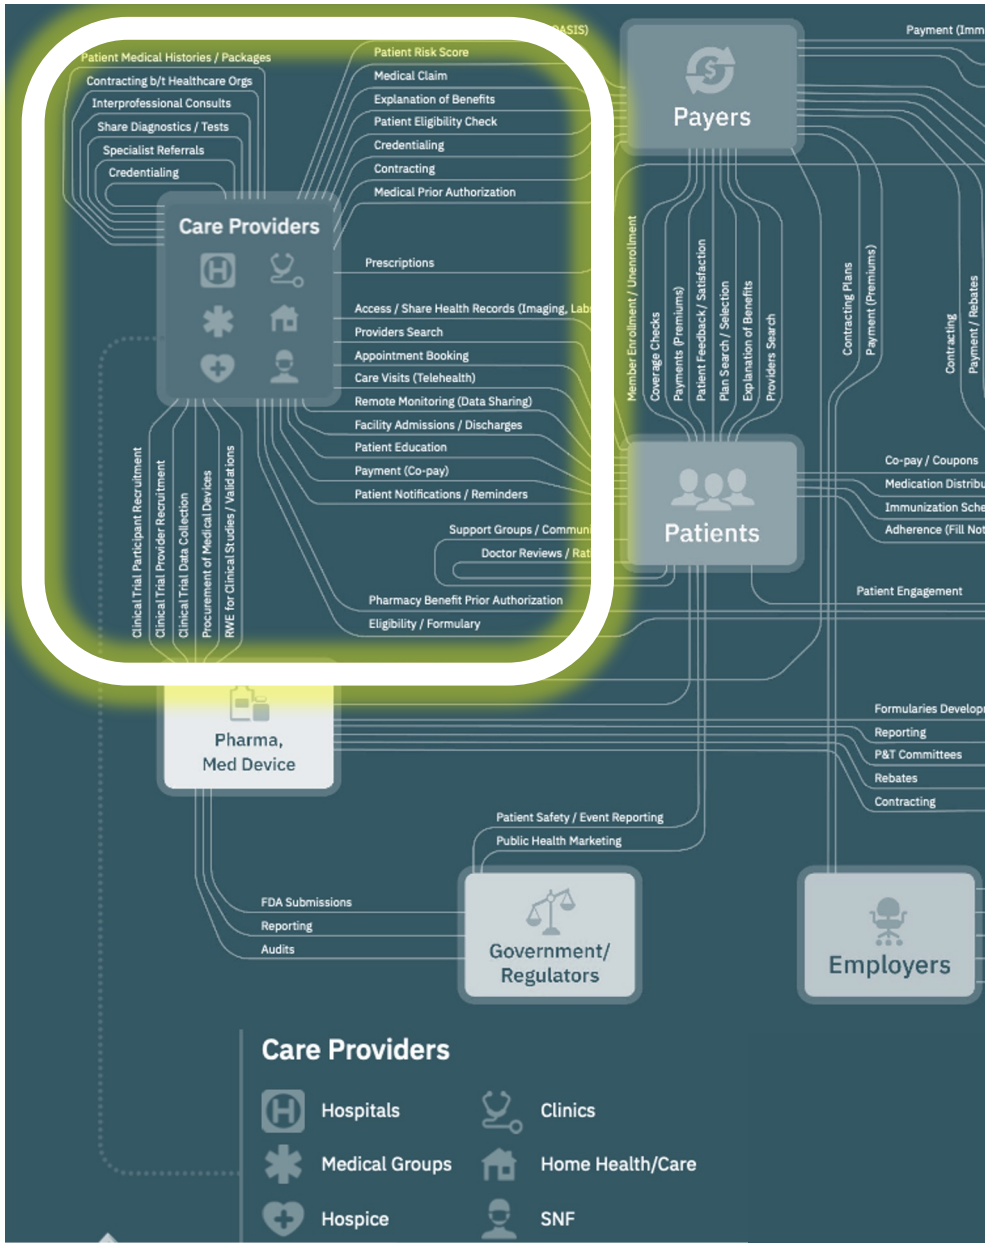

Disha
enabling agency in healthcare

SBX Group
Circa 2000 - Present

SBI The Systems Biology Institute | SBX | SBX Technologies | SBX BioSciences

The Supply Side Bottleneck: Lack of "Agency"

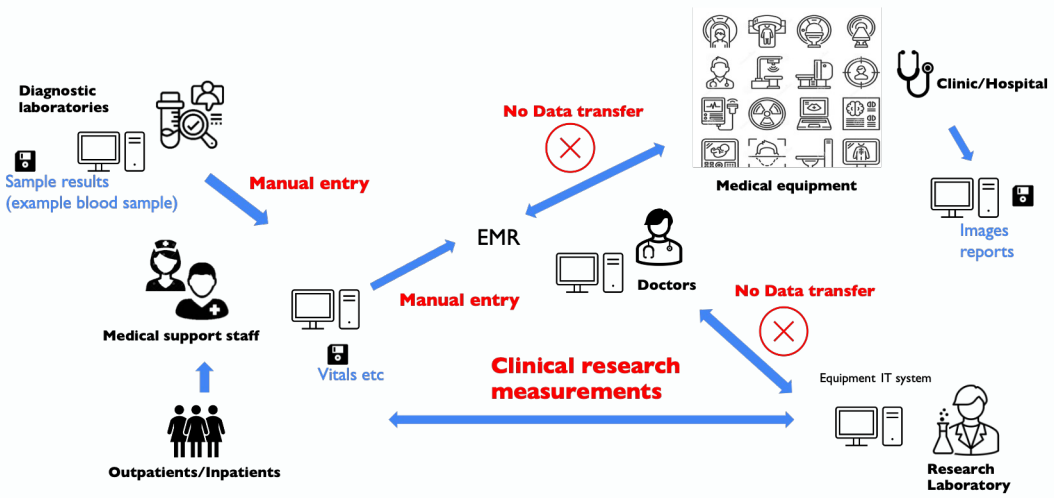

Typical workflow in a small hospital/clinic

- Fragmented workflows and siloed data disrupt coordination between teams and increase administrative burden for clinicians
- Manual processes and legacy IT slow processes, introduce errors, and add to provider workload
- Poor data readiness blocks analytics, AI, and real-time clinical decision support
- Inefficiency erodes margins, limits scalability, and contributes to staff burnout
- Inflexible, non-interoperable systems make modernization costly and frustrate care delivery

Disha
enabling agency in healthcare

Overview of Agentic AI

Agentic ≠ LLMs

LLMs answer on prompt; no tools/state

Reasoning systems can plan/Chain of Thought but still passive/reactive

Agentic systems plan → act (tools/APIs) → observe → adapt with aspects of Memory, policies, goal-directed outcomes

Agents add structure and orchestration for goal-directed behavior

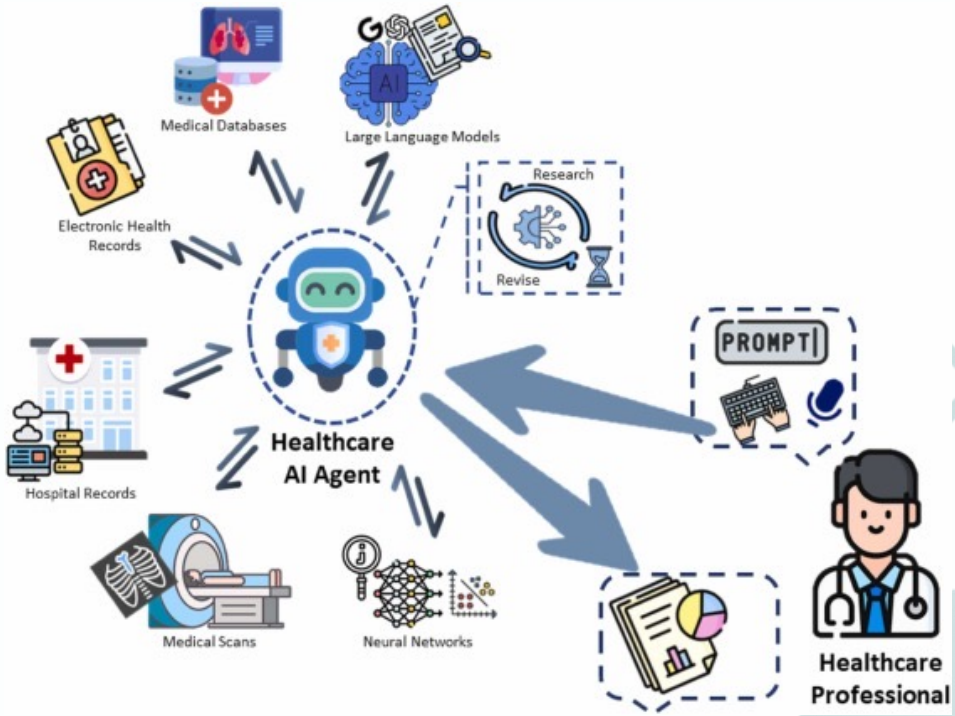

Karunanayake, Nalan. "Next-generation agentic AI for transforming healthcare." Informatics and Health 2.2 (2025): 73-83.

Disha
enabling agency in healthcare

Core components of an Agentic Stack

- Planner & policy engine
- Tool executor (connectors to systems – Model context protocol)
- Memory/state & caches
- Guardrails (allow-lists, abuse filters)
- Observability (traces, audit), Human-in-loop gates

Standards for Communication

•**MCP (Model Context Protocol):** consistent agent ↔ tool communication

•**A2A (Agent-to-Agent, Google):** discovery & messaging across agents/systems

Disha
enabling agency in healthcare

Some Agentic Frameworks

Category	Framework / Tool	Highlights / Example Use Case
Orchestration Runtimes	LangGraph	Stateful graph programming for long-running, agentic flows with human-in-loop.
	LangChain Agents (ReAct)	Prompt-based agent model; now evolving into LangGraph for production.
Workflow Engine	n8n (AI Agent Node)	Low-code orchestrator; integrates agents ↔ triggers ↔ APIs seamlessly.
Multi-agent Coordination (A2A)	AutoGen	Facilitates multi-agent dialogue, tool use, and supervisor agents.
	CrewAI	Manages role-based agent teams ("crews") for collaborative workflows.
Guardrail Toolkit	NVIDIA NeMo Guardrails	Policy layer to restrict/outbound calls and unsafe LLM generations.

Healthcare : Agentic ≠ Agency

- Safety & Compliance Risks with Unbounded Autonomy
- PHI compliance (Edge vs Cloud), What is shared? where it is shared? how it is shared? With whom it is shared?
- **Strong** Need for Human Oversight (Human in control)
- Need for plug and play solutions + easy IT integration + minimal setup
- Role scoped solutions

Disha
enabling agency in healthcare

Building a system of intelligent agents

Disha Platform powers hospital operations with a **system of interoperable, task-specific agents**

- Agents automate key tasks and workflows, reducing manual effort and errors
- Seamlessly operate within existing hospital systems-no disruption
- Each agent delivers the 3 Is: Information, Intelligence, and Interface for actionable results

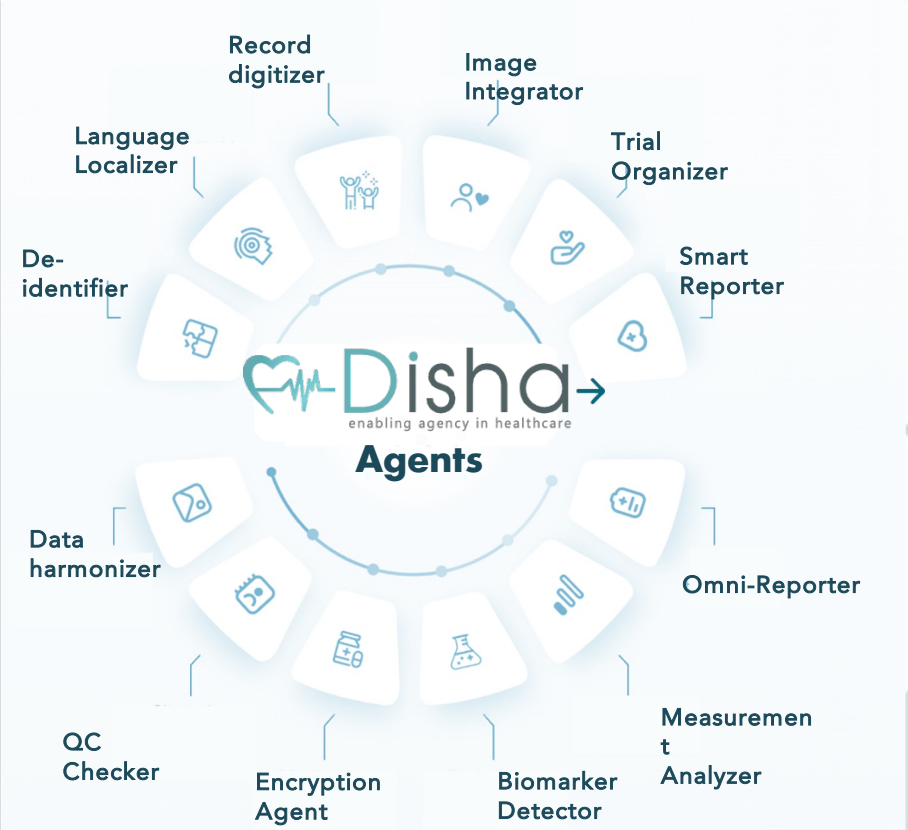

Disha
enabling agency in healthcare

Unlock operational efficiency and smarter care with Disha's intelligent agents.

Disha Box

Flexible, Secure, and Future-Ready Deployment

- Deploy Anywhere
  - On-premises, cloud, or as a portable “Disha Box” (edge device)
  - Local deployment ensures patient data privacy and regulatory compliance-no data leaves your premises
- Plug & Play Integration
  - Secure APIs, custom UIs, or hybrid modes
  - Designed to be interoperable with diverse devices, and legacy systems
- Minimal IT Burden
  - Pre-configured Device, Minimizes need for specialist IT staff and reduces operational complexity

Disha Box*

(comes with agents preinstalled, plug and play)

- 2048-core Ampere GPU, 64GB LPDDR5 RAM, 4TB NVMe SSD, with HDMI, USB, Gigabit Ethernet, and CSI camera ports.
- DISHA Box runs compact LLMs and multimodal models tuned for hospital use—up to 12B parameters—locally on-site.
- Plug-and-play setup with built-in privacy; designed for steady, reliable performance.

*under development

Disha
enabling agency in healthcare

Disha Box

Flexible, Secure, and Future-Ready Deployment

- Scalable & Modular
  - Add/remove agents as needs change; supports multi-site and multi-department rollouts
- Attach physical sensors seamlessly
  - Camera, Microphone, Speaker etc
- Role based Provisioning
  - Each box can be provisioned for a specific role (i.e bed, Nurse, technician etc)
- Future-Ready
  - Built for new protocols, instantly compatible with LLMs and AI tools

Disha Box*

(comes with agents preinstalled, plug and play)

- 2048-core Ampere GPU, 64GB LPDDR5 RAM, 4TB NVMe SSD, with HDMI, USB, Gigabit Ethernet, and CSI camera ports.
- DISHA Box runs compact LLMs and multimodal models tuned for hospital use—up to 12B parameters—locally on-site.
- Plug-and-play setup with built-in privacy; designed for steady, reliable performance.

*under development

Disha
enabling agency in healthcare

Disha box setup

Host System

The host system can be any laptop or desktop where the Disha Box is connected using a USB Type-A and/or Type-C connector cable.

User can access the agents using browser only. The box will availble by the domain name https://disha.local

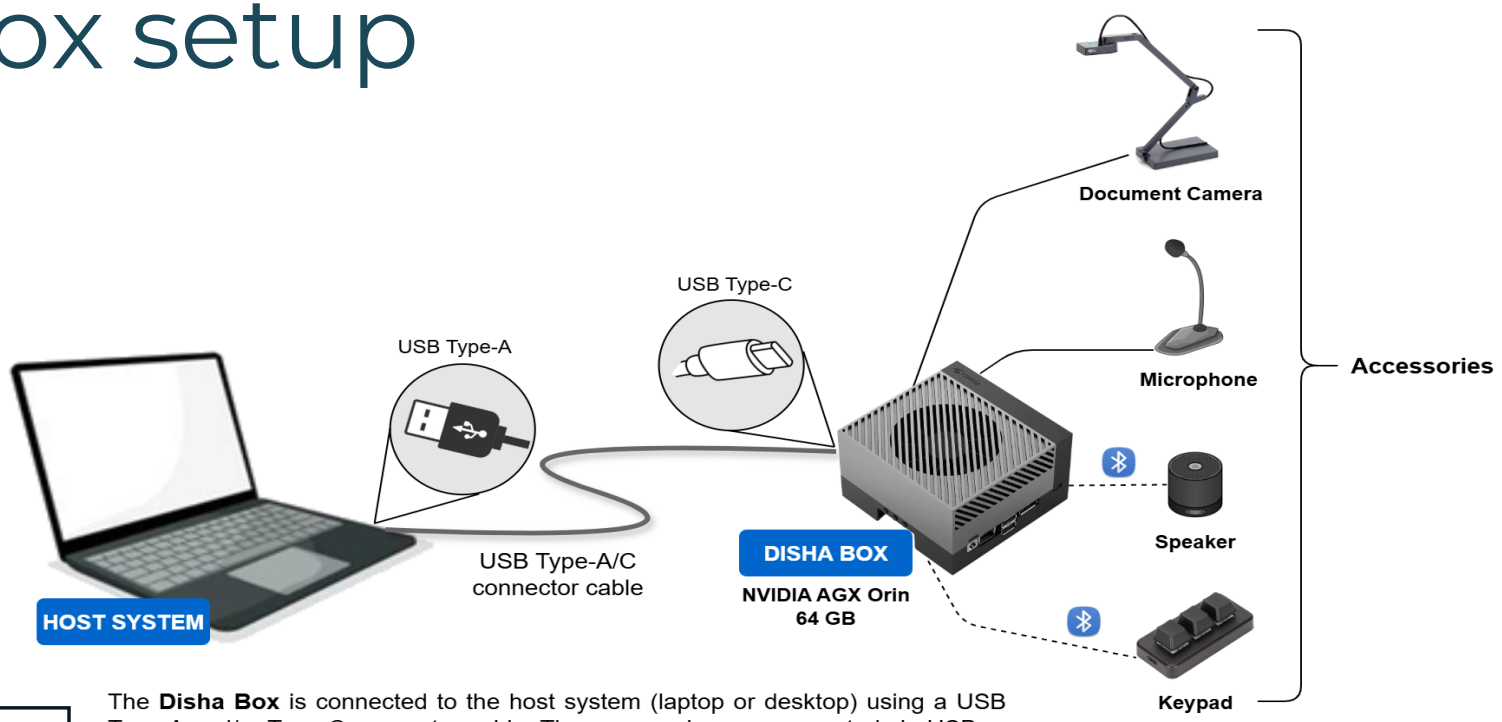

NVIDIA AGX Orin 64GB Specification

GPU: NVIDIA Ampere architecture with 2048 NVIDIA CUDA® cores and 64 Tensor Cores

CPU: 12-core Arm Cortex-A78AE v8.2 64-bit CPU

Memory: 64GB LPDDR5 (256-bit interface) with 204.8 GB/s bandwidth

Storage: 64GB eMMC 5.1 + 2TB SSD

The **Disha Box** is connected to the host system (laptop or desktop) using a USB Type-A and/or Type-C connector cable. The accessories are connected via USB or bluetooth.

- NO INTERNET CONNECTION REQUIRED
- POWERED BY LARGE LANGUAGE MODELS (LLM)
- PLUG AND PLAY SUPPORT
- AGENT-DRIVEN SERVICES
- MEETS ALL HIPPA AND MEDICAL GUIDELINES

enabling agency in healthcare

Model(s) Summary

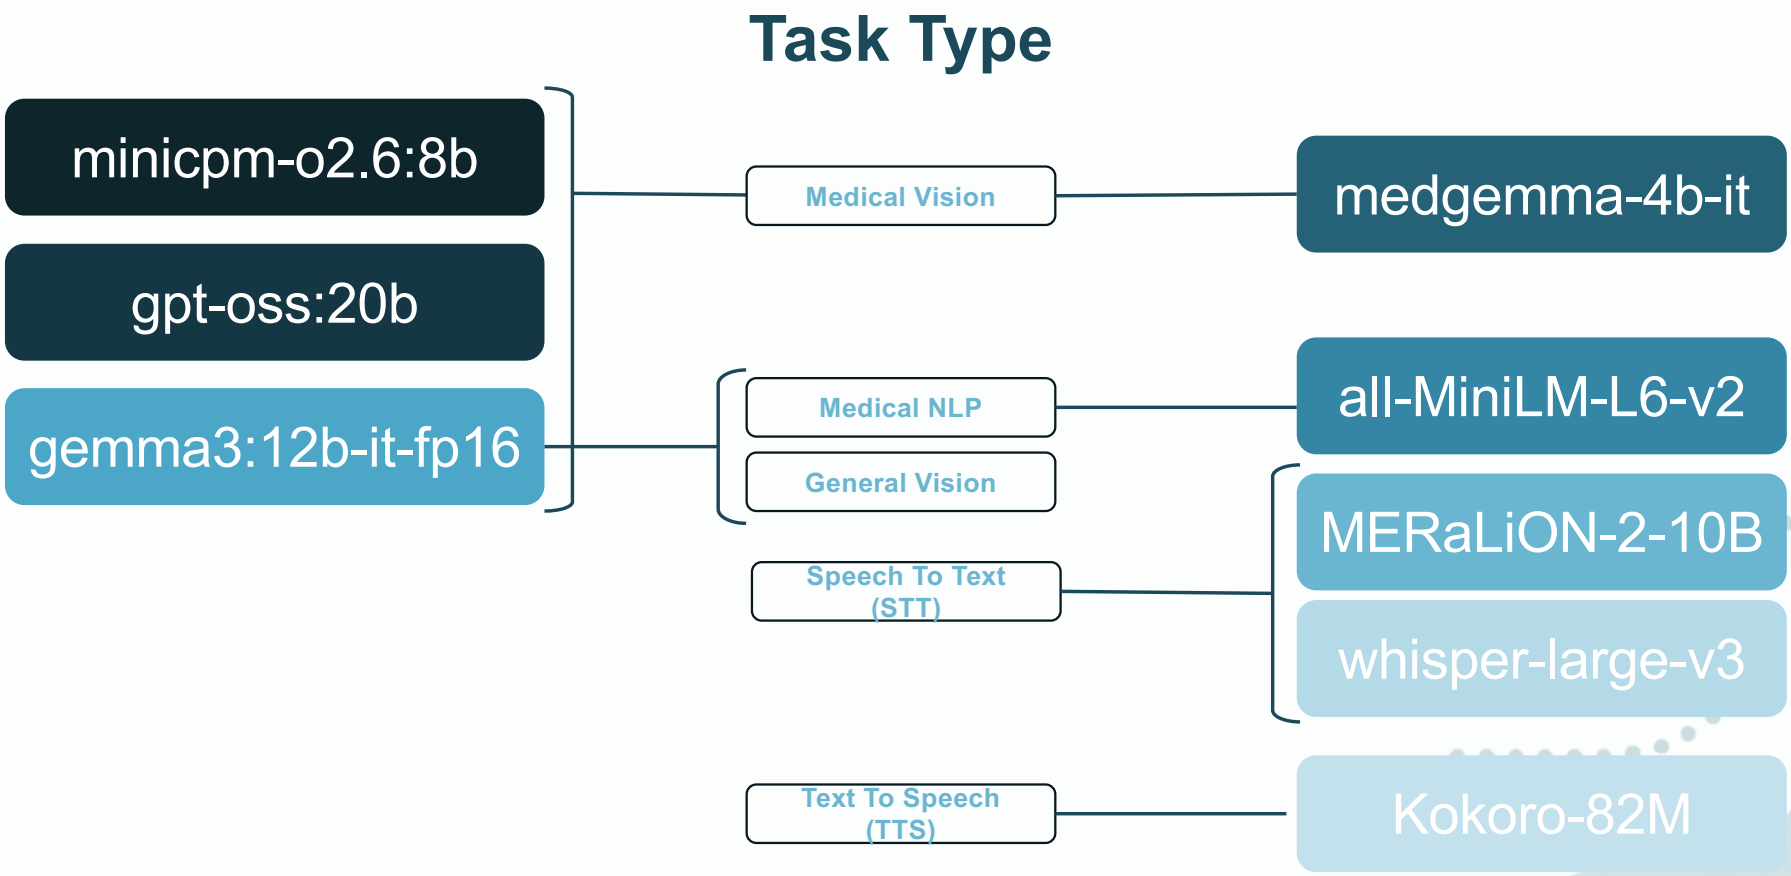

Disha
enabling agency in healthcare

Agent Model Mapping

Agent	Primary Model	Secondary Model	Task Type
Prescription Reader	google/medgemma-4b-it	gemma3:12b-it-fp16	Medical Vision
Discharge Summary	gpt-oss:20b	openbmb/minicpm-o2.6:8b	Medical Vision
Speech To Text (MERaLiON-2-10B)	MERaLiON/MERaLiON-2-10B		STT
ICDMapper	gemma3:12b-it-fp16	all-MiniLM-L6-v2	Medical NLP
Notes Reader	gemma3:12b-it-fp16		General Vision
Radiology Explainer	google/medgemma-4b-it	gemma3:12b-it-fp16	Medical Vision
Measurement Analyzer	google/medgemma-4b-it	gemma3:12b-it-fp16	Medical Vision
Text To Speech (TTS)	hexgrad/Kokoro-82M		TTS
Speech To Text (STT)	openai/whisper-large-v3		STT
Report of Reports (RoR)	gpt-oss:20b	google/medgemma-4b-it openbmb/minicpm-o2.6:8b	Medical Vision

Disha
enabling agency in healthcare

Tech Stack

Frontend

React.JS
v19.0 RC

Middleware

Next.JS
v15.0

Backend

Python
v3.10

DB / Object Store / Secret Store

MongoDB
v8.0.4

MinIO
2024-12-18T13:15:44Z

Vault
v1.20.0

Deployment

Docker
Docker Compose

Disha
enabling agency in healthcare

The Disha Agency

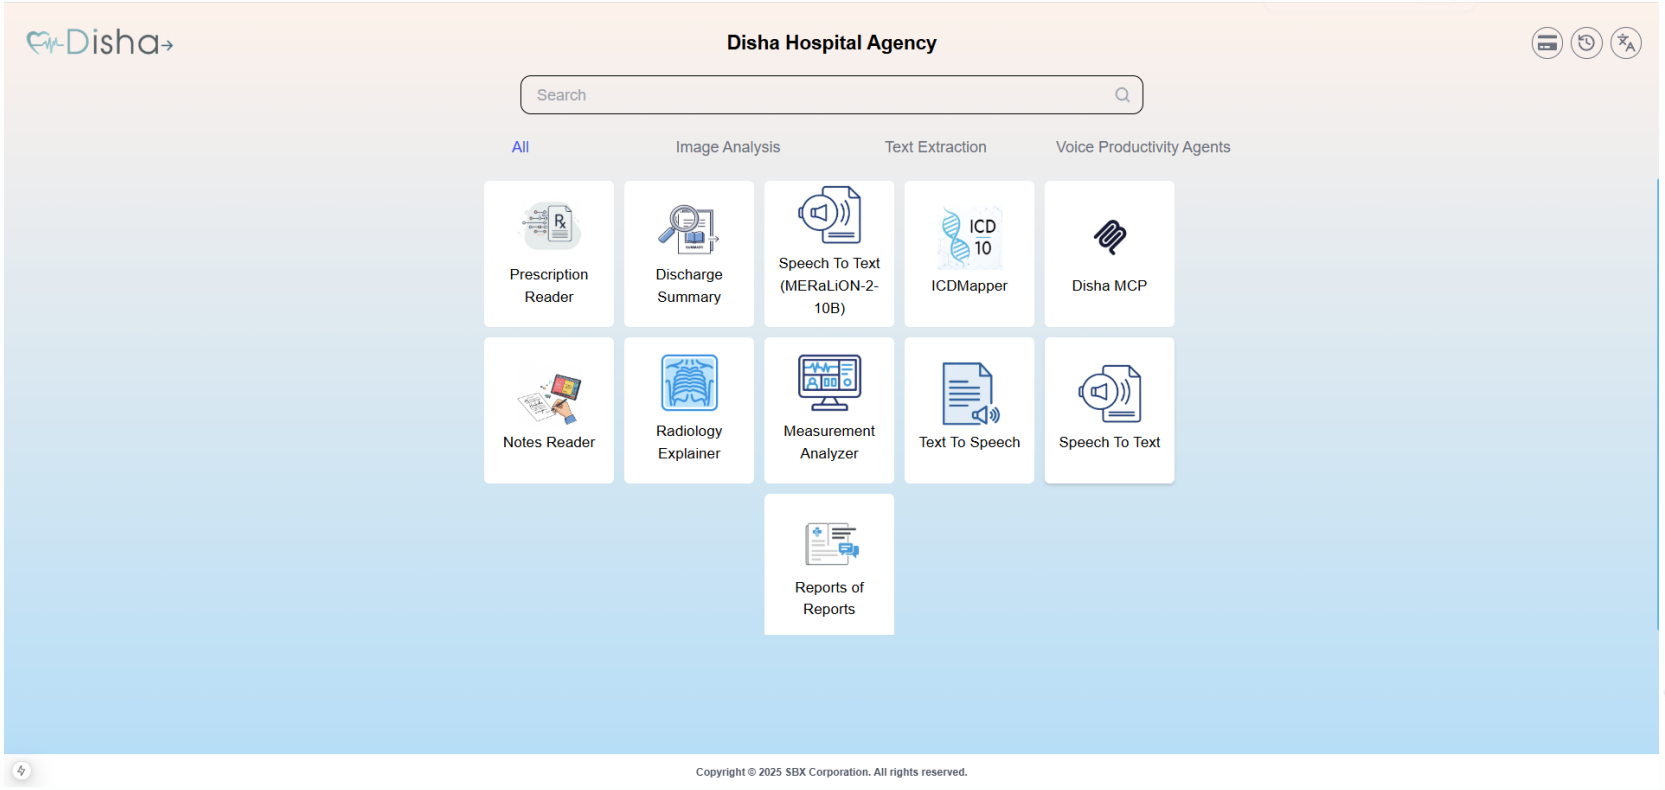

Disha
enabling agency in healthcare

Agency Showcase: ROI

Reporting & acti**o**nable **I**nsights

- Reports (both in text narrative, tabular quantities and images) are a critical component of hospital operations and clinical practice
- Deliberate Integration: Data and reports need to operate with existing HMS
- Ability to generate actionable insights in a trustable, explainable and interpretable manner is central to healthcare

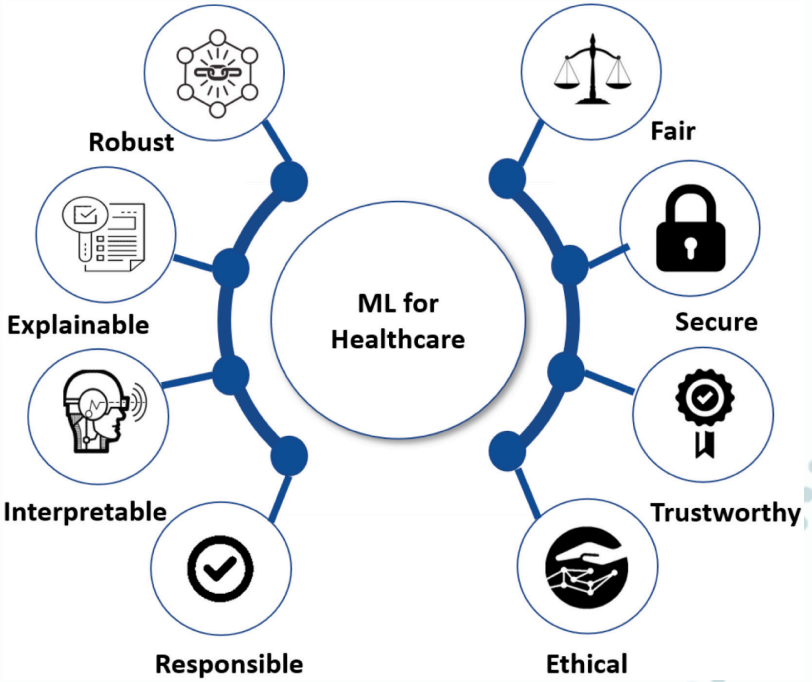

Rasheed, Khansa, et al. "Explainable, trustworthy, and ethical machine learning for healthcare: A survey." Computers in Biology and Medicine (2022): 106043.

LIMS: Laboratory Information Management System
HMS: Hospital Management Systems

Disha
enabling agency in healthcare

ROI Agents showcase

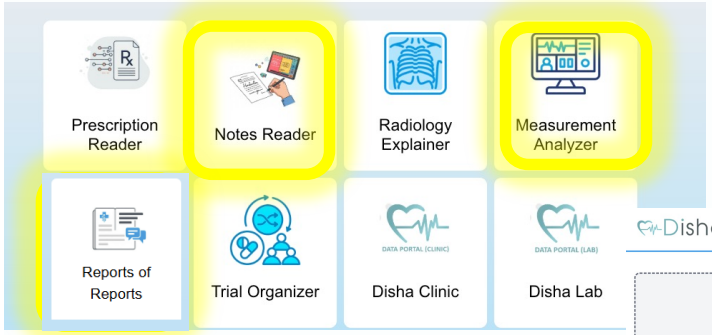

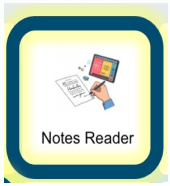

Reading, structuring and reporting of clinical notes

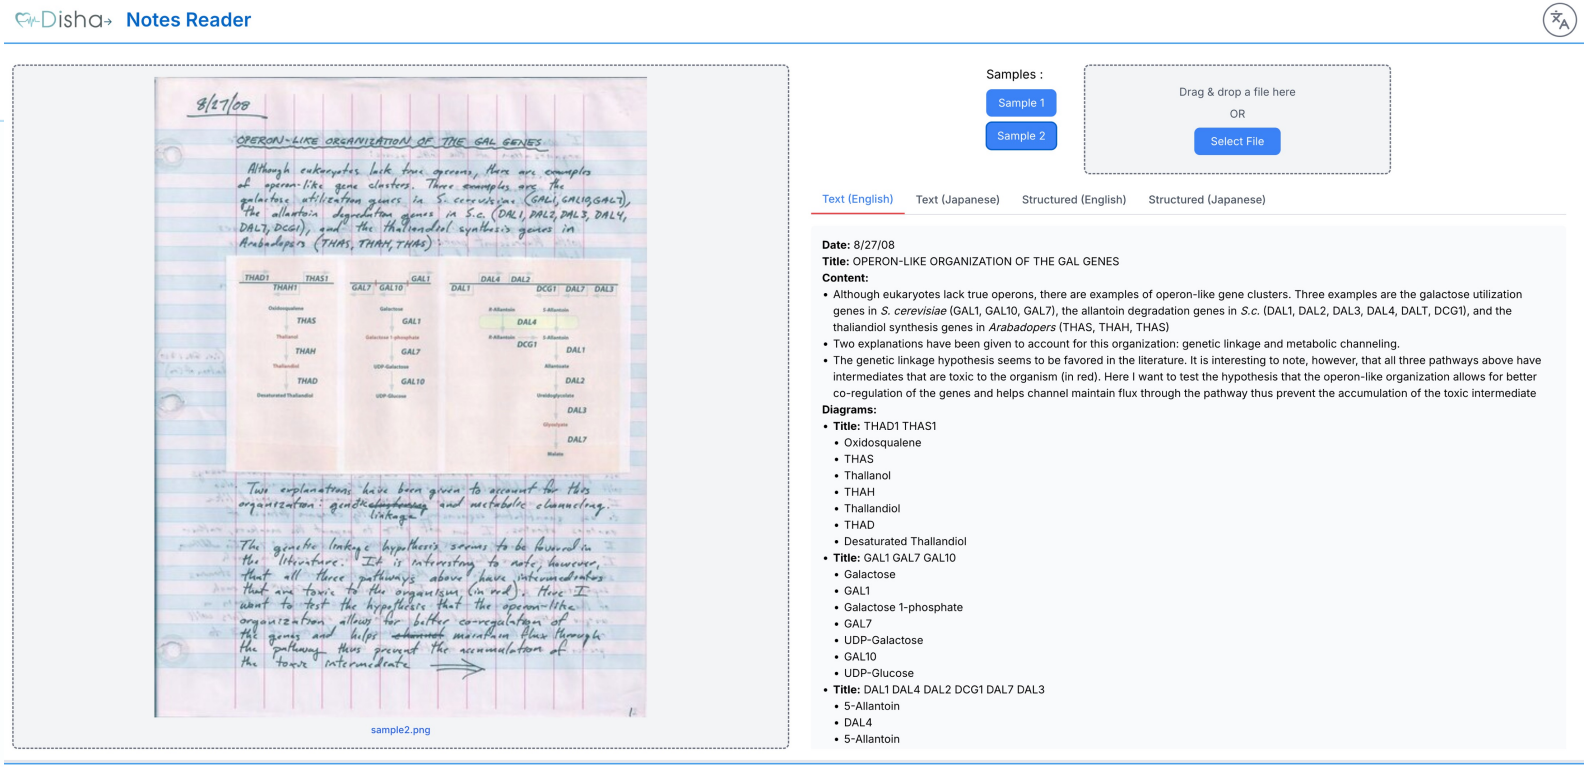

Disha
enabling agency in healthcare

ROI Agents showcase

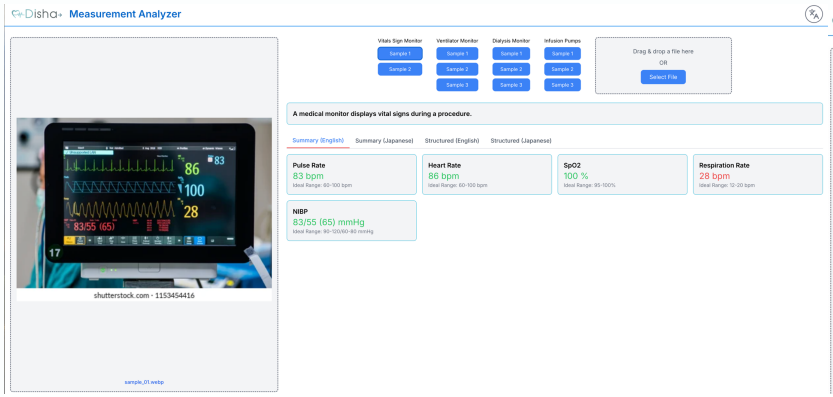

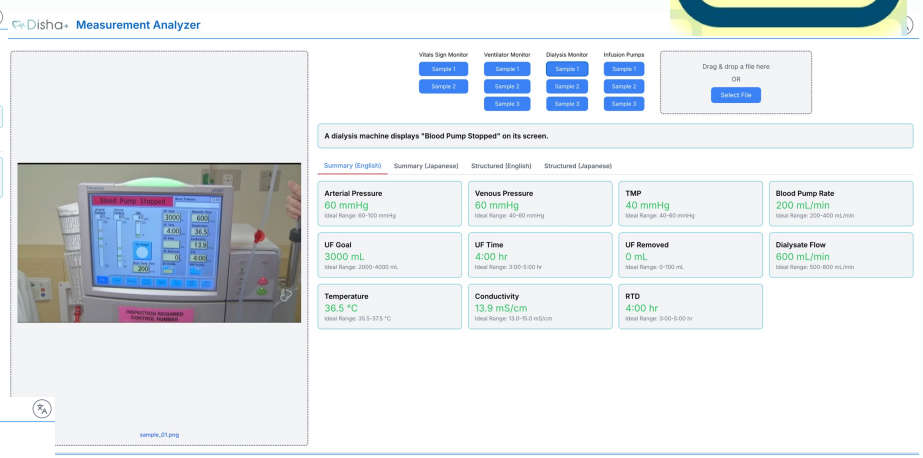

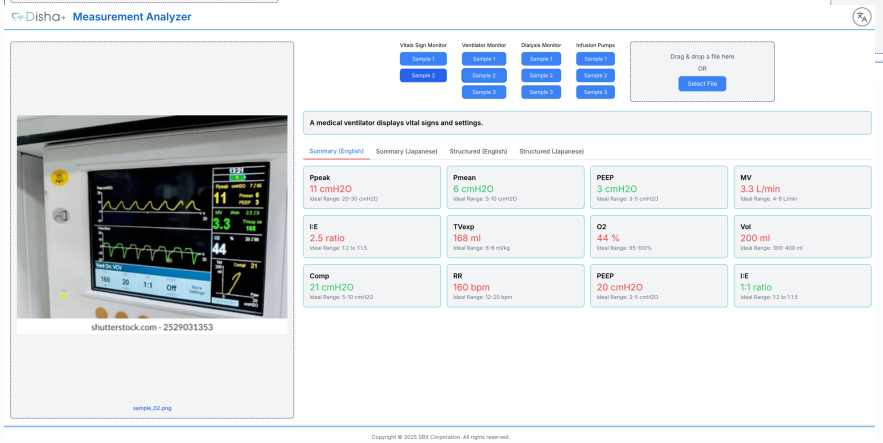

Reading,structuring & Reporting of measurement from diverse devices

Disha
enabling agency in healthcare

ROI Agents showcase

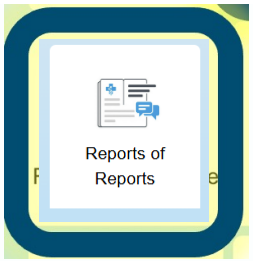

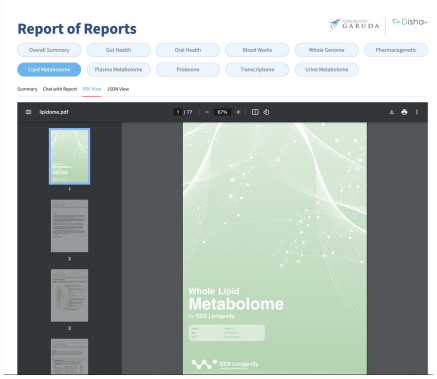

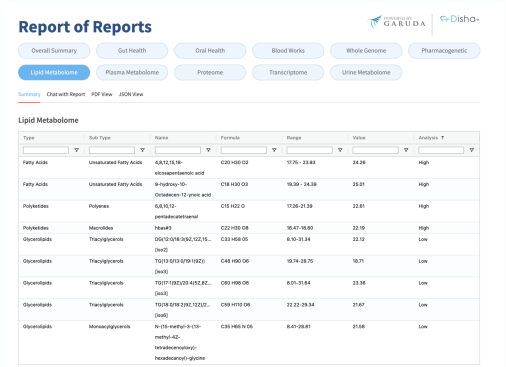

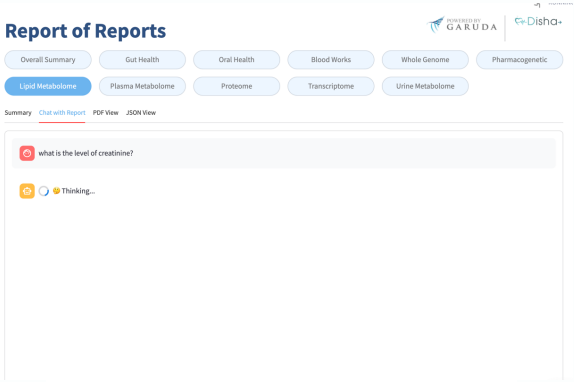

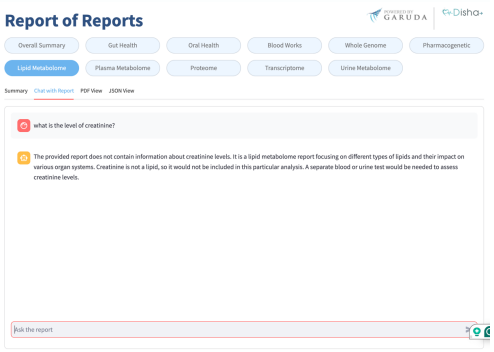

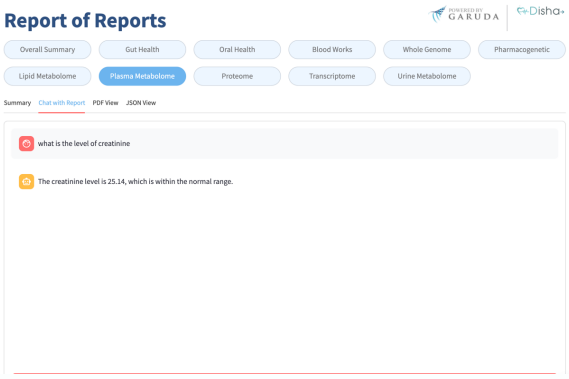

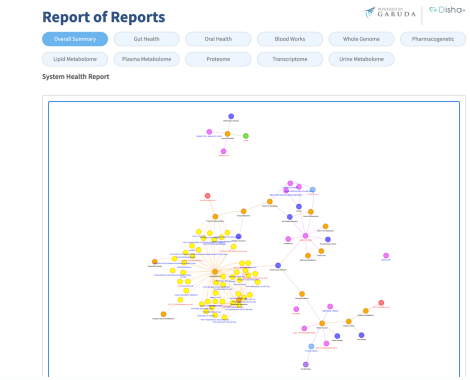

Disha
enabling agency in healthcare

Agency Stories

Discharge Summary Agent

Prescription Reader Agent

Speech To Text (MERaLiON) | ICDMapper

Report of Reports (RoR) | Text To Speech (TTS)

Measurement Analyzer | EHR (XLS)

Measurement Analyzer | Text To Speech (TTS)

LangGraph workflows

Disha
enabling agency in healthcare

Agentic AI for Clinical Decision Intelligence: The TGWA Case Study

Setting the direction for the future of Clinical Decision Intelligence (CD_x)

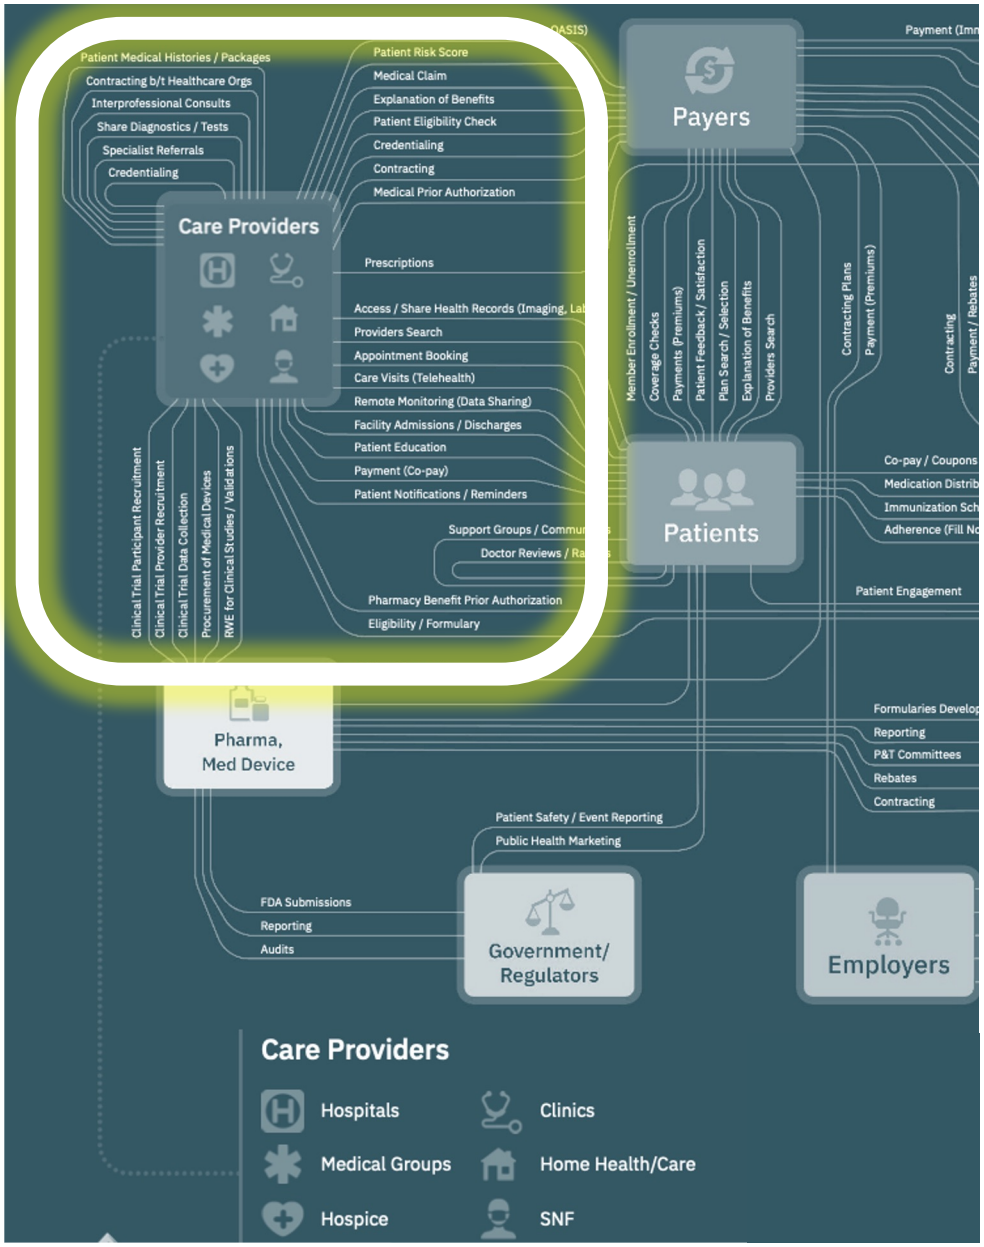

Disha
enabling agency in healthcare

SBX Group
Circa 2000 - Present

SBI The Systems Biology Institute | SBX | SBX Technologies | SBX BioSciences

Agents in Action – AI Clinic

Developed and deployed an AI Clinic with Disha agents to operationalize an Aging and Wellness Clinic in Tokyo, Japan

1,500 Patients

4 Clinical Trials

Aging and Wellness focused Clinic in Tokyo

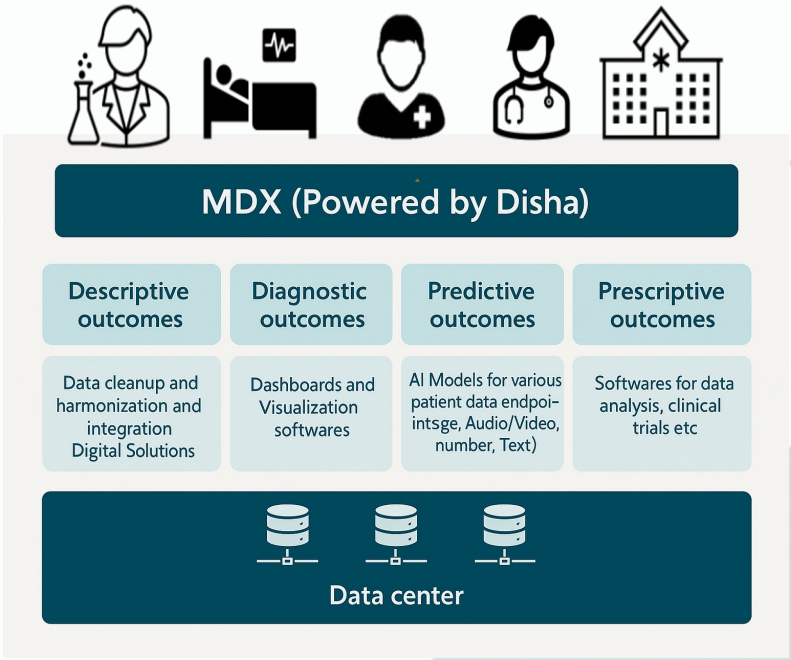

Disha
enabling agency in healthcare

Before that. State of AI

https://www.idgconnect.com/article/3583467/gartner-accelerating-ai-deployments-paths-of-least-resistance.html

Last decade

-Focused on building POC AI models in research.

-Putting models to production and operationalizing was an afterthought

Today

- >50% of POCs make it to production.

- Average turnaround cycle is 9-10 months

Near Future

- Large number of organizations (hospitals) will shift focus from piloting to operationalizing AI

*"IT leaders responsible for AI are discovering **'AI pilot paradox'**, where launching pilots is deceptively easy but deploying them into production is notoriously challenging."*

Gartner

Disha
enabling agency in healthcare

How much of real world AI “systems” is the AI “model”?

Google

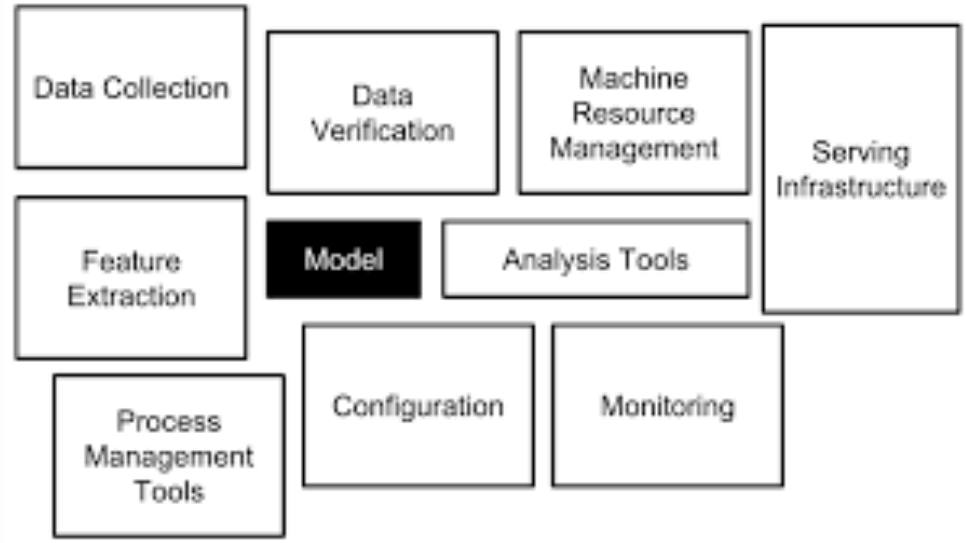

Phase 1
(Research/experimental)
Can we use AI to solve this problem?

Phase 2
How do we implement these models into IT systems at scale?

Disha

https://developers.google.com/machine-learning/crash-course/production-ml-systems

We are now at the point where software engineering was a (few) decades ago!

- Devops
  - Software Development + IT Operations is a subarea of software engineering
  - Scalable software products are built by embracing best Devops practices

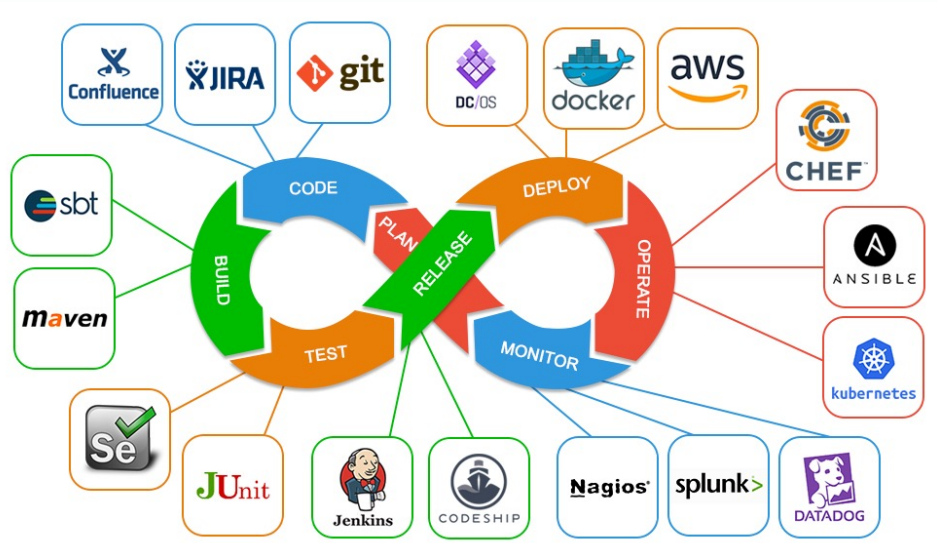

https://medium.com/swlh/how-to-become-an-devops-engineer-in-2020-80b8740d5a52

Embracing similar practices for AI ?

Rise of MLOps

Disha
enabling agency in healthcare

MLOps

- Data + ML + Dev + Ops = MLOps
- Extending Devops for AI-ML

https://www.ml4devs.com/articles/mlops-machine-learning-life-cycle/

- Model versions tracking
- Monitor data versions and transformations
- Automation: data cleanup, transformation, re-training, optimization and scoring
- Monitor model performance (accuracy and robustness)
- Logging and notifications
- Governance: regulatory compliance
- Governance: reporting

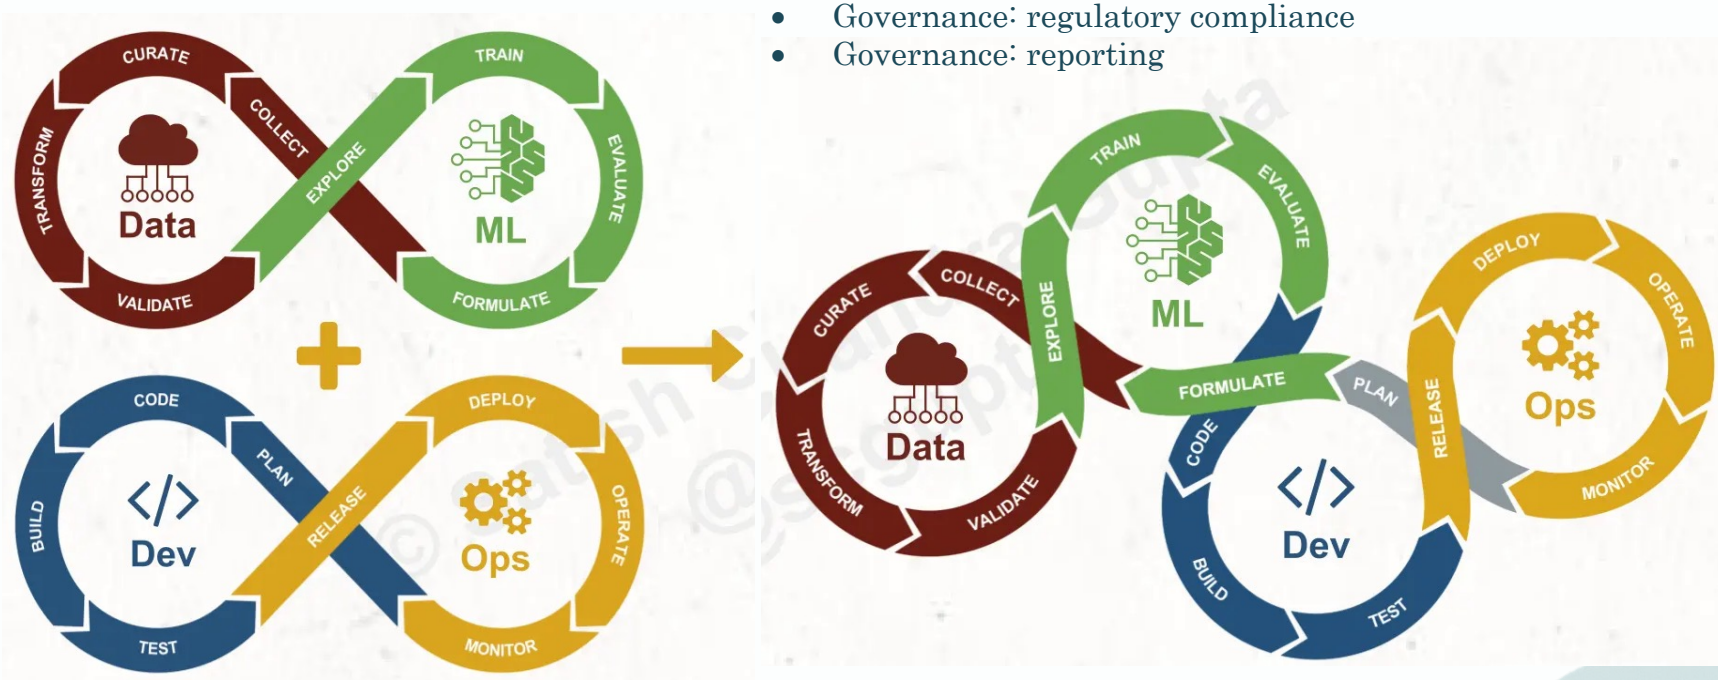

Disha
enabling agency in healthcare

Data and Model

- **Assumptions:** Data Scientist's decisions and assumptions must be explicit.
- **Randomness:** Considering that some machine learning experiments contain pseudo-randomness, this needs to be in some kind of control so it can be reproduced. For example, using "seed".
- **Data:** The same data of the experiment must be available.
- **Settings:** Repeat and reproduce experiments with the same settings from the original.
- **Implementation:** Especially with complex models, different implementations can have different results. This is important to keep in mind when debugging.
- **Environment:** It's crucial to have the same runtime configurations among all data scientists.

Data and model Versioning

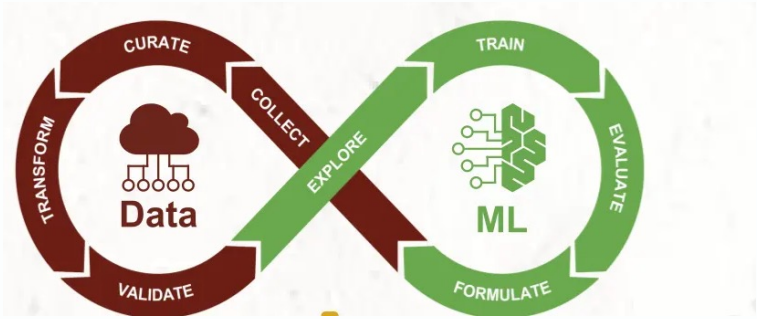

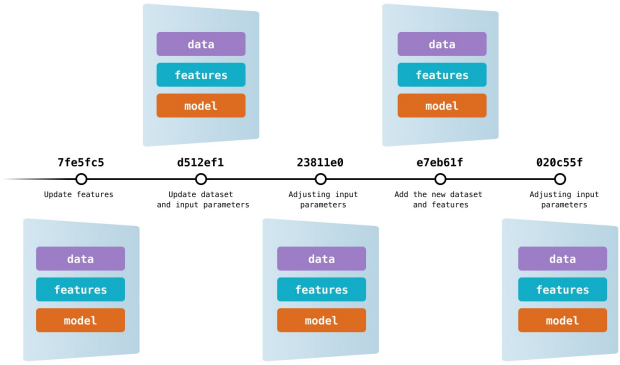

mlflow
DVC
lakeFS

Automation

TensorFlow Extended

ML Stages	Tasks
Data Engineering	Data acquiring, validation, and processing
Model Development	Model training, evaluation, and testing
Continuous Integration	Build and testing
Continuous Delivery	Deployment new implementation of a model as a service
Monitoring	Setting alerts based on pre-defined metrics

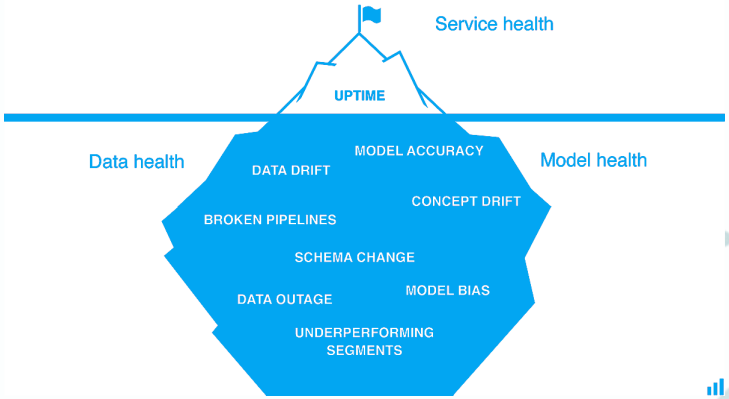

Monitoring

https://mlops-guide.github.io/

Disha
enabling agency in healthcare

AI Clinic Case Study – Overview & Impact

- **Before DISHA, The clinic faced challenges common in healthcare operations**
  - Fragmented data across EMRs and equipment
  - Manual data entry and Limited insights
- **MDx Platform**
  - Centralized data portal that consolidated diverse sources
  - Automated data processing
  - Extensible infrastructure to expand new AI models and applications
- **Key Features of the Platform**
  - In-house compute for seamless processing
  - Full Integration of all equipment & data sources
  - DISHA agents to clean, harmonize, and analyze data
  - MLOps & AIOps in action for continuous optimization
  - AI-ready for future applications

Disha
enabling agency in healthcare

Operations then

- **Diagnostic laboratories**

Sample results
(example blood sample)

Manual entry

Medical support staff

- **Facilitate smooth operation**
- **Provide hassle free service to patients**

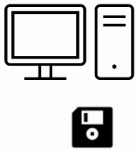

Vitals etc

Manual entry

EMR
(M3 digikar)

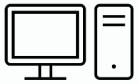

Doctors

Clinic

- **Medical equipment**

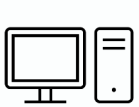

Images
reports

No Data transfer

Equipments IT system

Research
Laboratory

Outpatients

Clinical research measurements

Disha
enabling agency in healthcare

MDx Platform

Improved decision making and outcomes

Descriptive outcomes → Diagnostic outcomes

Predictive outcomes → Prescriptive outcomes

Data cleanup and harmonization and integration Digital Solutions

Dashboards and Visualization softwares

AI Models for various patient data endpoints (Image, Audio/Video, numbers, Text)

Softwares for data analysis, clinical trials etc

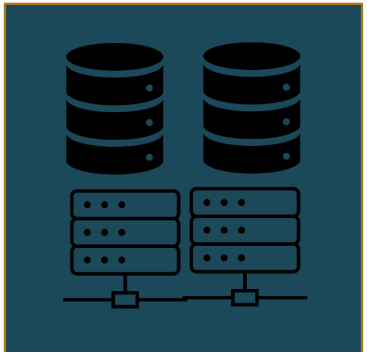

Disha
enabling agency in healthcare

TOKYO GINZA
WELLNESS & AGING CLINIC

Robust, Secure Infrastructure for Scalable AI

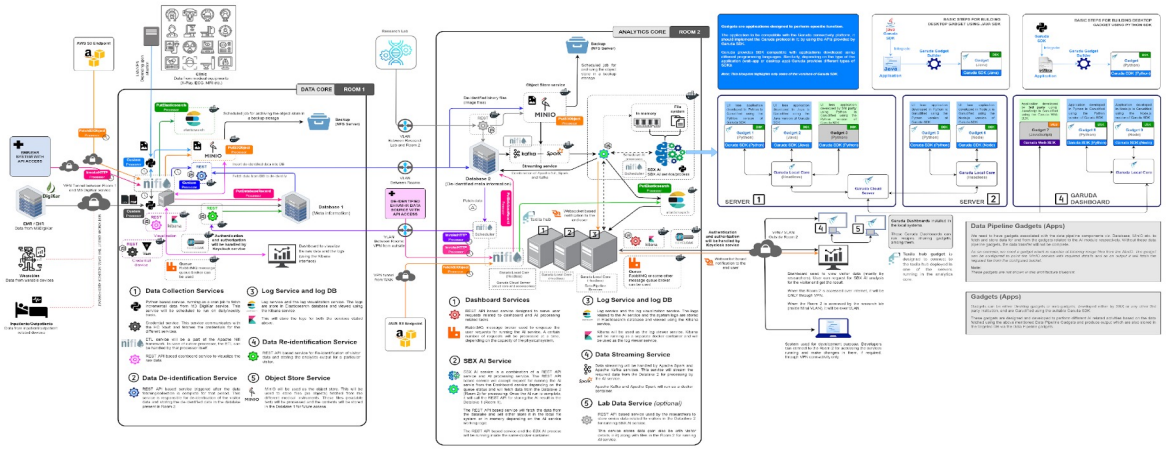

Operationalization (O16n): Technology stack for DevOps & MLOps

Compute and Infrastructure built ground up

- **Intelligent Agents**
- **Operationalization(O16n*)**
- **Data & Security**
- **Compute & Infrastructure**

✓ **Secure, redundant infrastructure ensures compliance, uptime, and data protection**

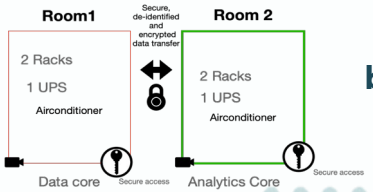

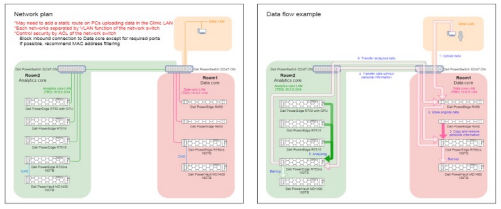

Disha
enabling agency in healthcare

AI Clinic – Data Flow & Analytics in Action

- Visualize predictions, verifications, and effects from clinic data, support doctors in making quick decisions, and provide data-based advice to patients
- Collect and organize data, and enable access and viewing of various data

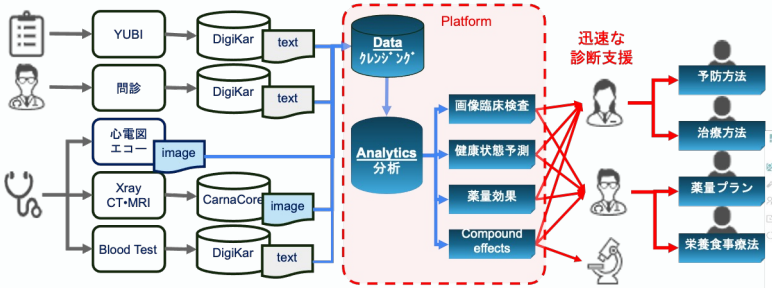

Data Portal
All Clinic and operational data in one place ready for AI

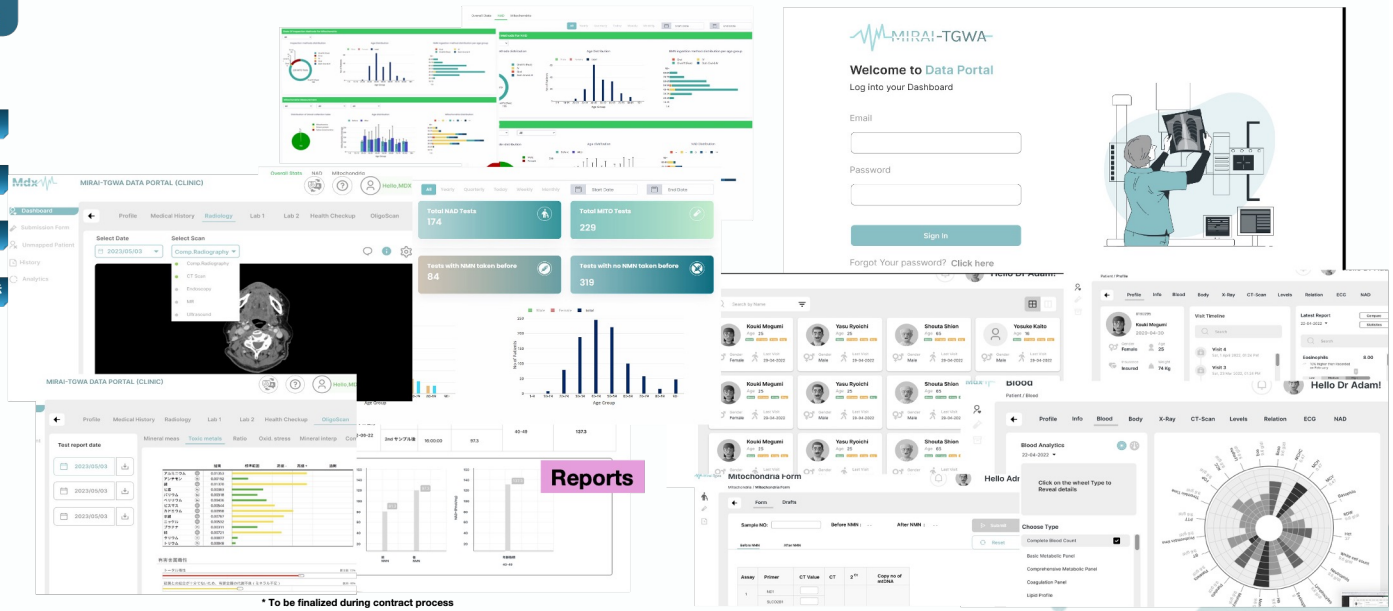

* To be finalized during contract process

Intelligent Agents

Operationalization(O16n*)

Data & Security

Compute & Infrastructure

✓ **Unified data enables real-time analytics, evidence-based advice, and automated reporting-all in one portal**

Disha
enabling agency in healthcare

Disha Pilots

Lighthouse Projects

Reaching out to collaborate with hospitals & clinics to develop, deploy and test

Disha Agents

Current Disha Partners

Hopeful Aging, USA
https://hopefulaging.com/

SRM University and Corporate Hospitals, India

Redkar Hospital, India

The University of Osaka, Faculty of medicine, Japan

Graduate School of Medicine, The University of Tokyo

Keio University Hospital, Japan

Disha
enabling agency in healthcare

THANK YOU