Introduction to AI and Digital Health Interventions

Hassan Ghasemzadeh, Shovito Soumma, Saman Khamesian

College of Health Solutions
Embedded Machine Intelligence Lab
Arizona State University

ABRC Workshop on AI/ML for Digital Health Interventions, November 7, 2025

What is AI?

- Artificial Intelligence (AI) is the field of computer science focused on making machines think and act intelligently—like humans.

- That doesn't mean they have emotions or consciousness, but they can solve problems, make decisions, understand language, and even learn from experience.

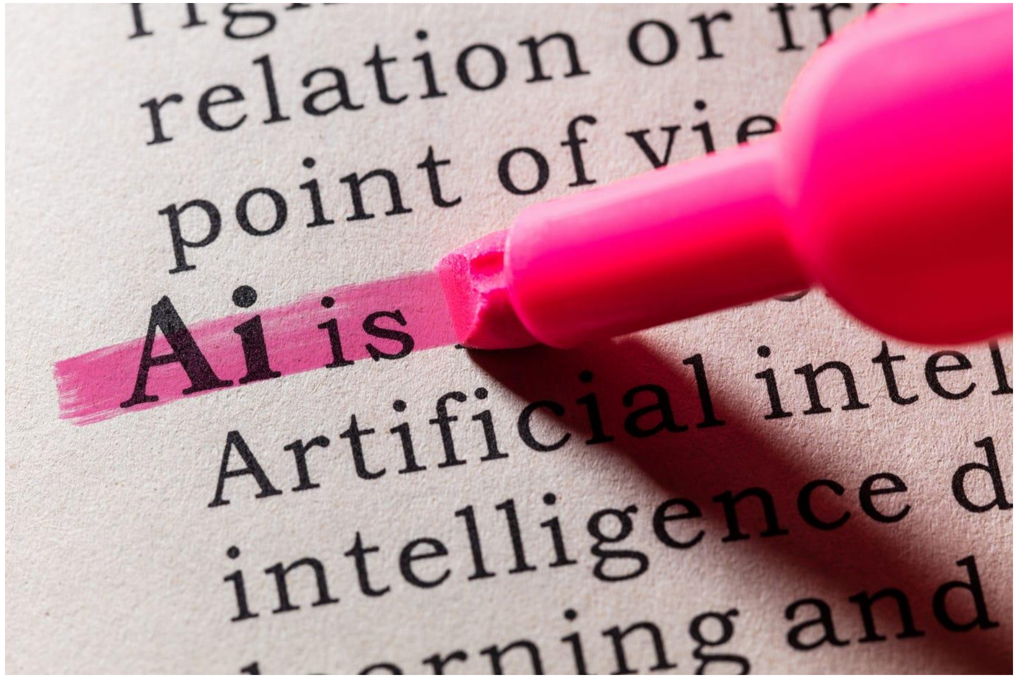

What is AI?

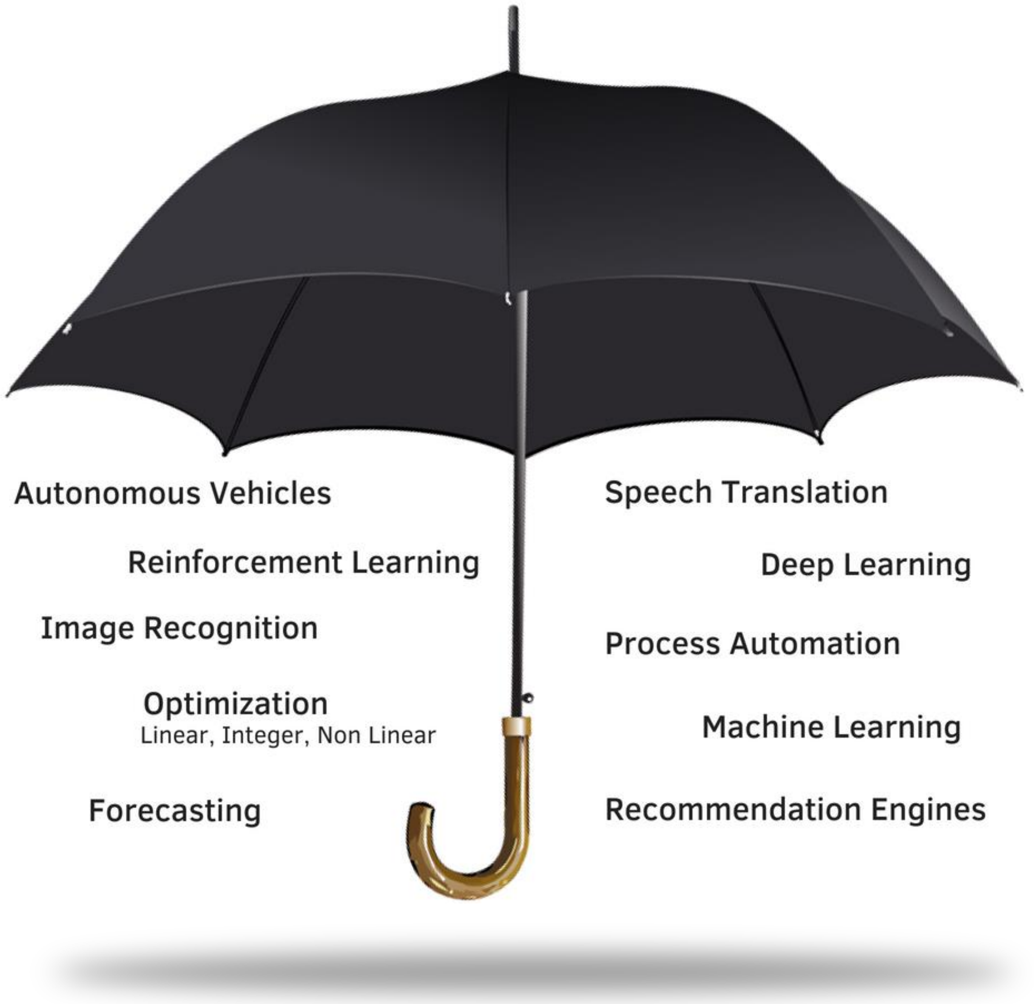

- Think of AI as the **big umbrella**.
- Under this umbrella, we have many techniques like machine learning, robotics, expert systems, and more.

Easy Analogy

- Imagine teaching a dog tricks. You want the dog to learn how to fetch, sit, or roll over.

- If a computer learns how to do these kinds of things — recognize patterns, respond to voice, or plan actions — we say it's artificially intelligent.

Examples of AI

- **Medical diagnosis systems**
- **Virtual assistants** like Siri or Alexa
- **Autonomous vehicles** like Tesla and Waymo
- **Recommendation Systems:** Netflix or Spotify suggesting movies or music based on your past choices.
- **AI in healthcare:** A system that looks at retina images to detect early signs of diabetic retinopathy.
- **Fraud Detection:** AI detects unusual patterns in bank transactions (e.g., sudden spending in another country).

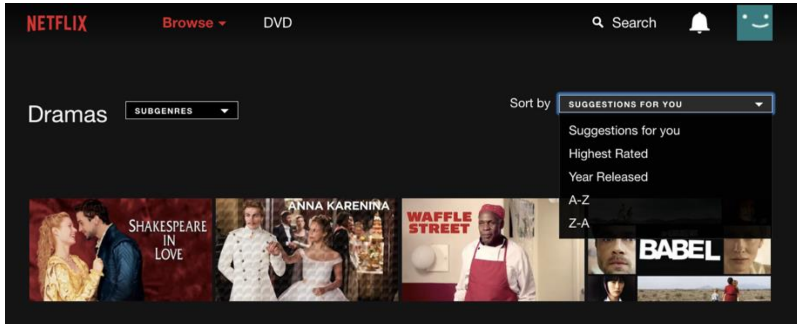

What is ML?

"Learning from Experience"

- Machine Learning is a subset of AI where the machine learns from data rather than being explicitly programmed.

- Instead of telling the computer exactly what to do, we give it examples, and it figures out the rules on its own.

What is ML?

- Let’s say you want to teach a machine to recognize apples and oranges.
- In traditional programming, you'd write code: "if color is red and size is small, then apple".
- In ML, you'd show the computer hundreds of images of apples and oranges, and it learns the patterns on its own.

Supervised Learning

- You give the computer both inputs and correct outputs.
- Example: Diagnosing whether a tumor is benign or malignant, based on labeled patient data.
- Analogy: Like a teacher showing students the correct answers.

Supervised Learning

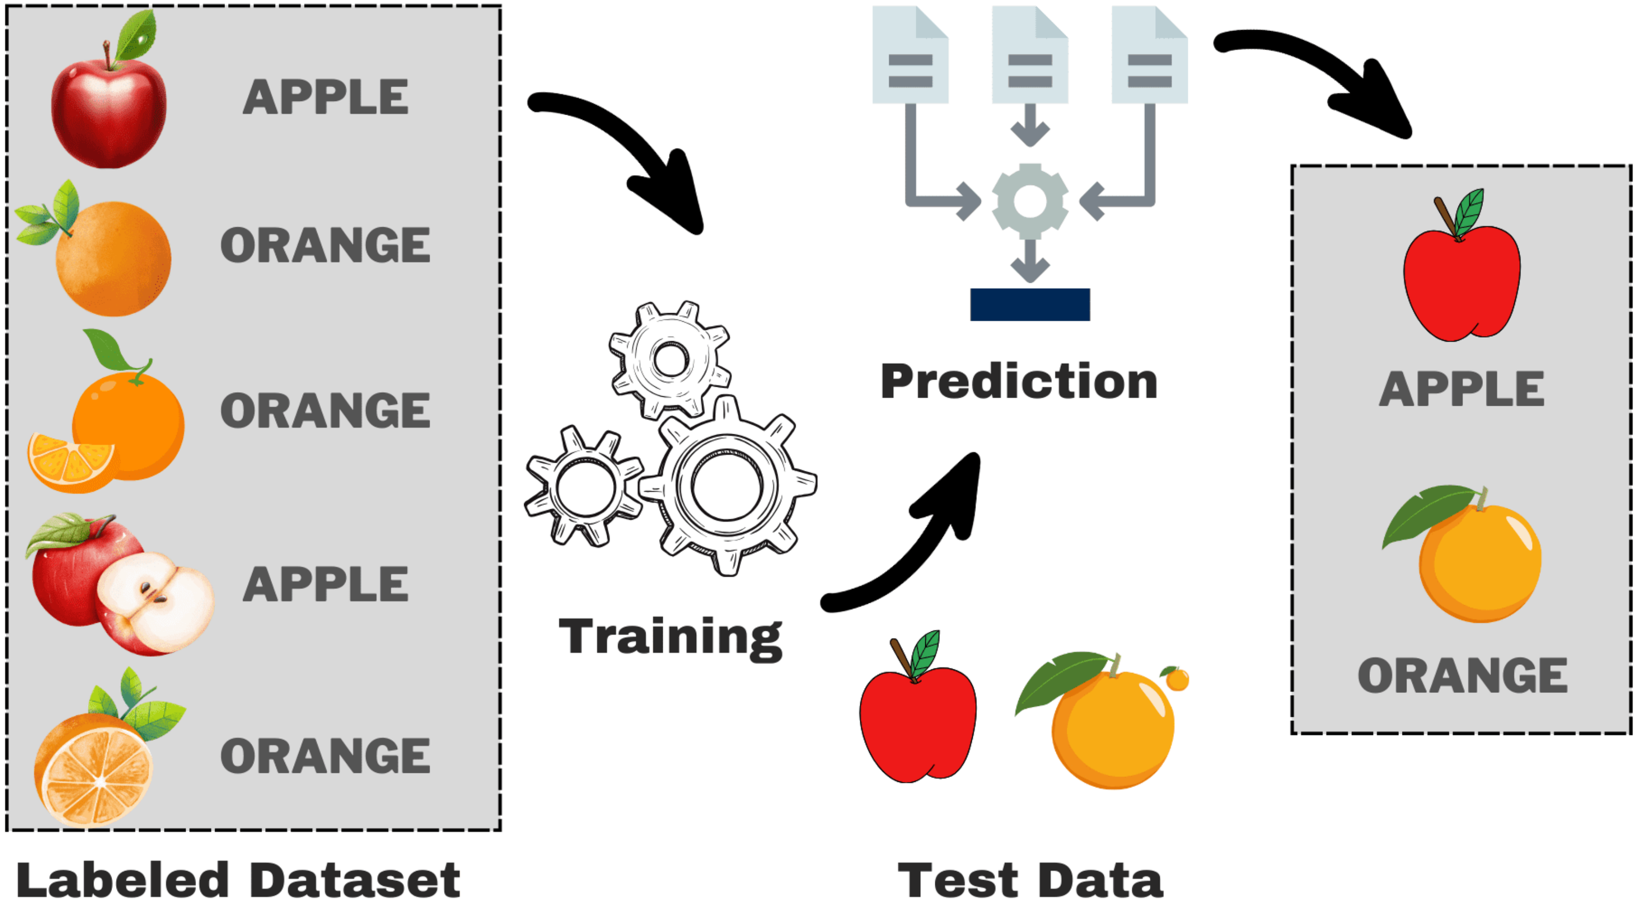

Predictive AI

- Medical imaging
- Biomedical diagnostic
- Activity recognition
- Glucose forecasting
- Stress monitoring
- Behavior modeling
- Fall detection
- Dietary assessment

Now What?

To know the cliff is ahead is ***knowledge***

To build the bridge around it is ***agency***

Glucose control

- Keeping blood glucose level within range
- Glucose control helps to lose weight, optimize mental health, suppress food carvings, improve sleep, prevent diabetes
- Poor glucose control increases risks of cardiovascular disease, kidney disease, stroke, cancer, reduced eyesight

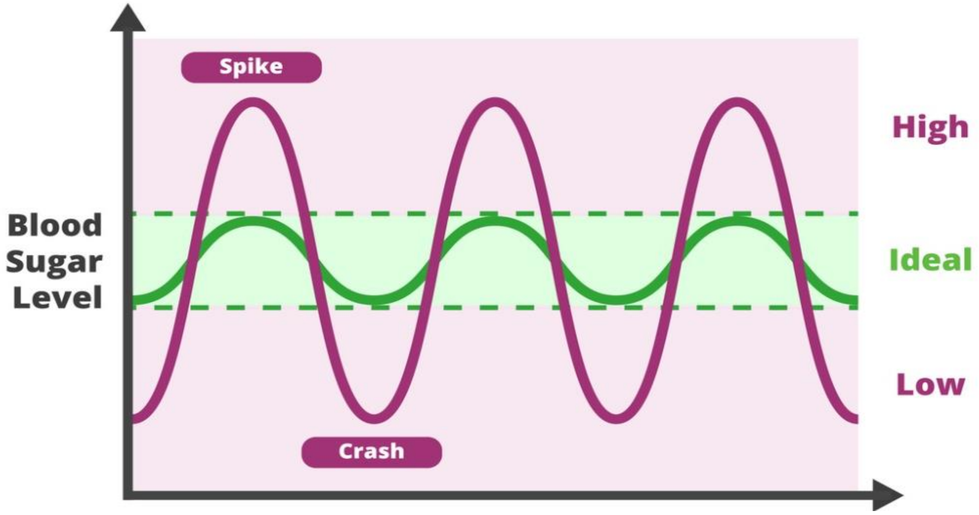

Credit: https://www.precisionnutrition.com/

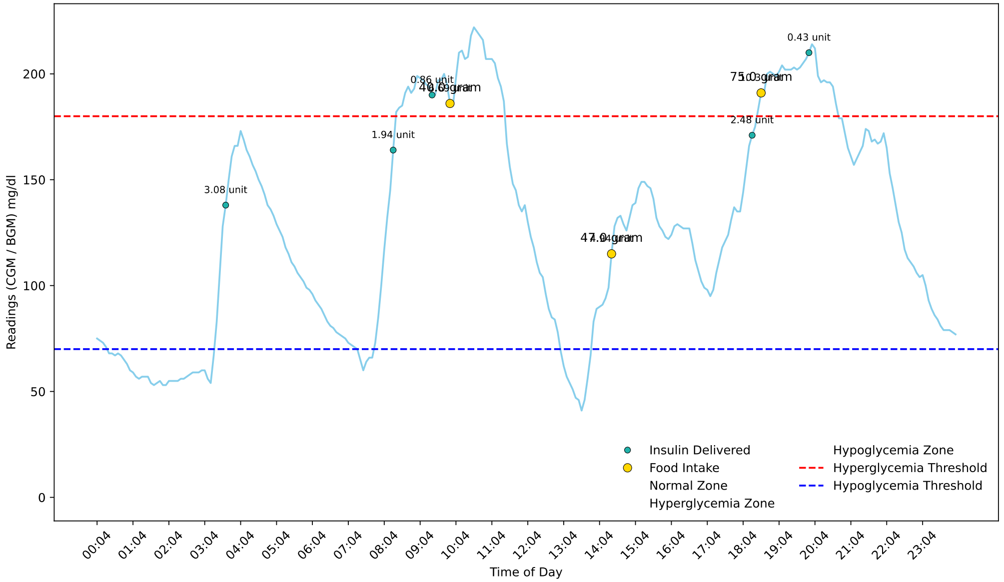

- Jarvis et al. 2023. Continuous glucose monitoring in a healthy population. *Metabolism*.
- Arefeen et al. 2022. Forewarning Postprandial Hyperglycemia with Interpretations using Machine Learning. IEEE BSN.

*Patient experiencing **3 hypoglycemic** events leading to **260 minutes** of hypoglycemia in one day plus **2 hyperglycemic** events resulting in **320 minutes** of hyperglycemia in 24 hours.*

Glucose forecasting

- Design neural network models to predict glucose response

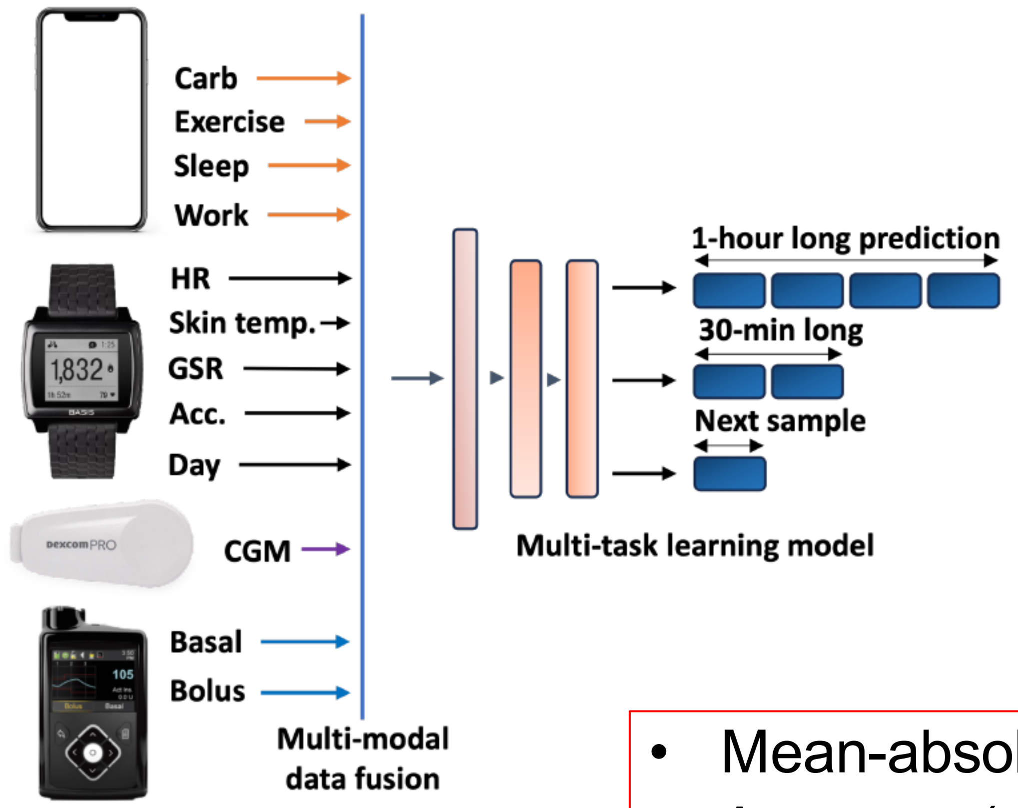

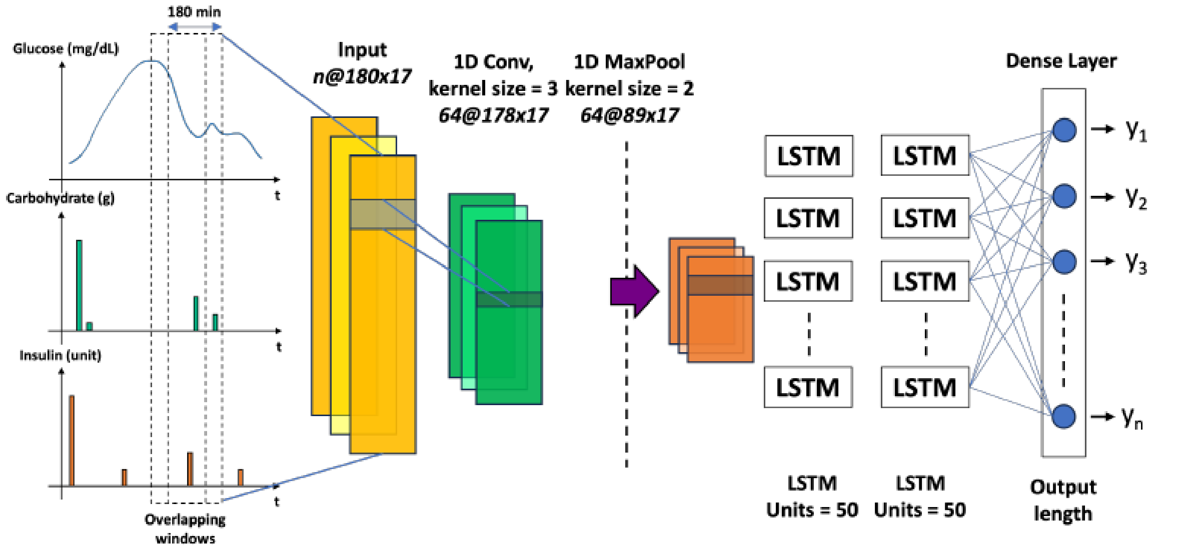

- Stacked **CNN-LSTM** architecture for glucose prediction
- Input to the model is 3-hour long multimodal data

- Mean-absolute-error (60-minute horizon): 16.5 mg/dL
- Accuracy (predicting hyperglycemia): 89%

Arefeen, A. and Ghasemzadeh, H., 2023, October. Glysim: Modeling and simulating glycemic response for behavioral lifestyle interventions. In *2023 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)* (pp. 1-5). IEEE.

Multimodal glucose forecasting

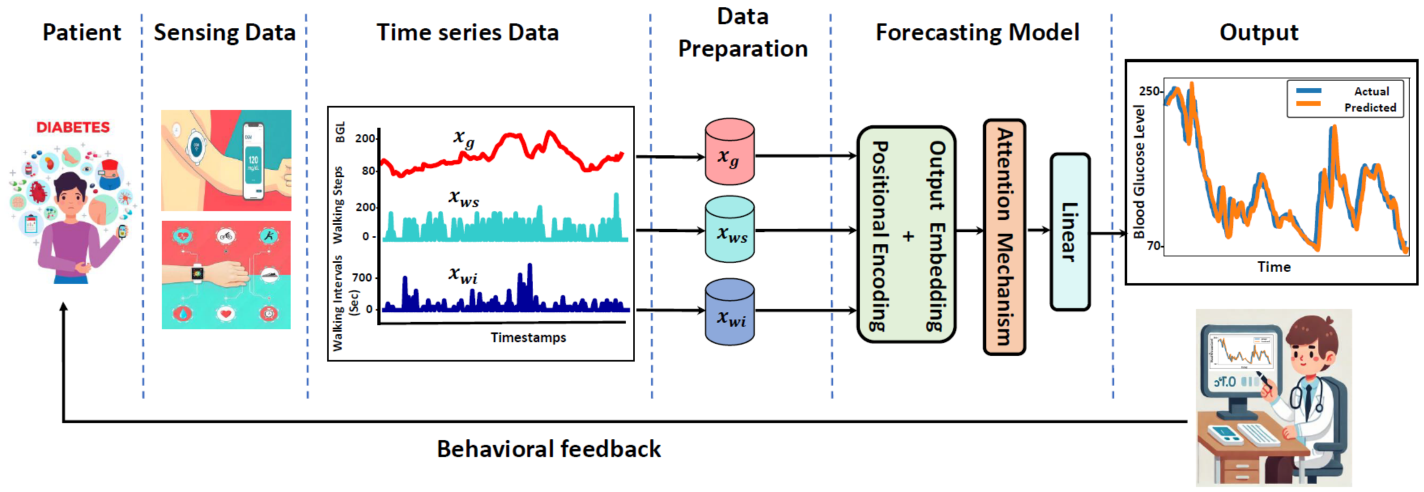

Cohort	RMSE		MAE	
	Baseline	AttenGluco	Baseline	AttenGluco
Healthy	18.04	**16.05**	13.02	**11.12**
Pre-T2DM	19.95	**18.27**	15.12	**13.65**
Oral	25.01	**22.56**	17.9	**15.74**
Insulin	29.9	**27.18**	22.28	**19.93**

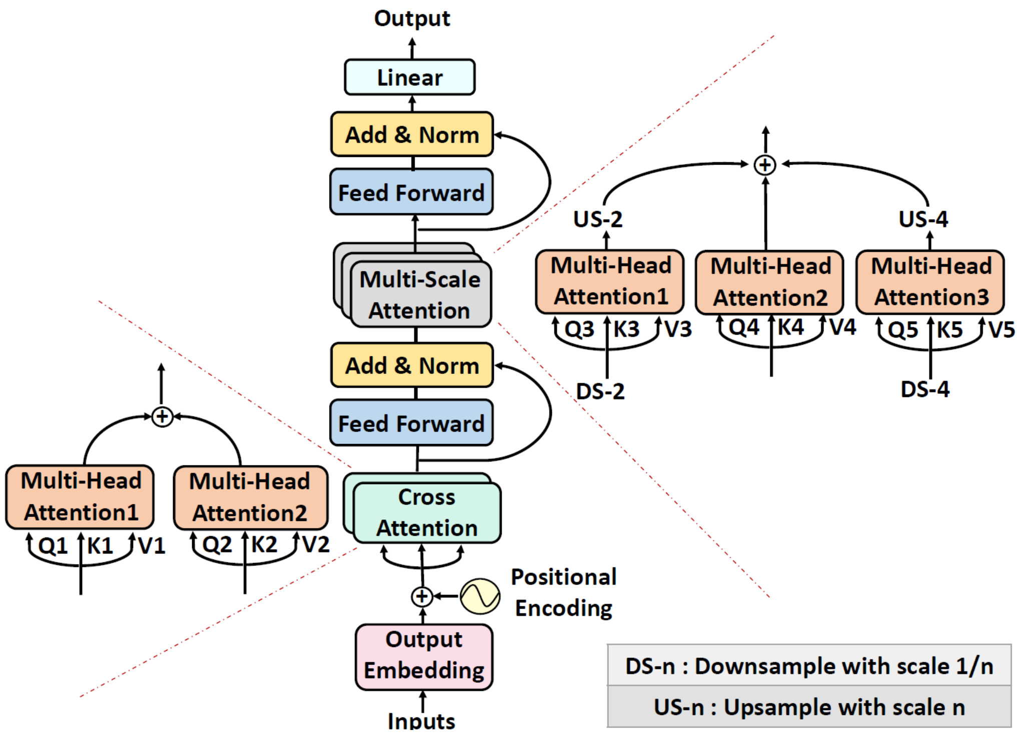

Beyond Prediction

- Prevention
  - How to prevent adverse event
  - Reason about negative outcomes
  - Provide explanations
  - Generate intervention options
- Empower patients and clinicians
  - More efficient clinical decision making
  - Scalability

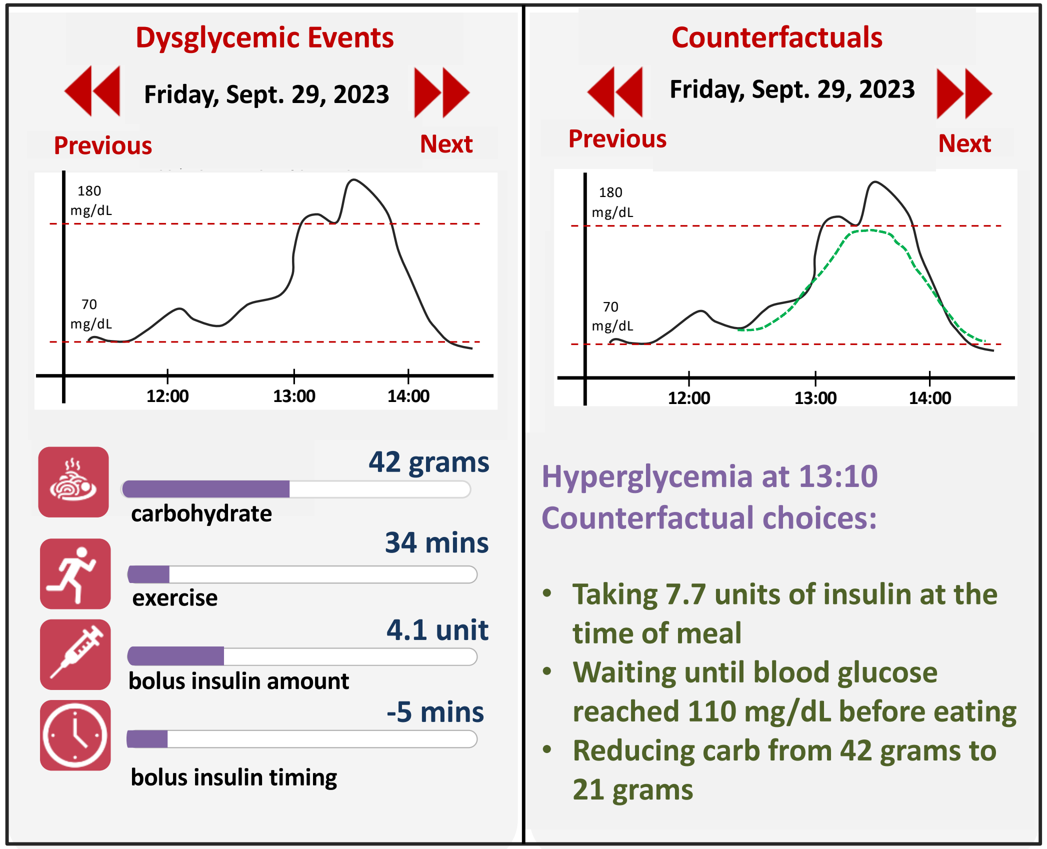

Counterfactual explanations

- Toy example: loan approval system
- AI model makes recommendations (approve/decline)
- Counterfactuals are data points that change the prediction outcome

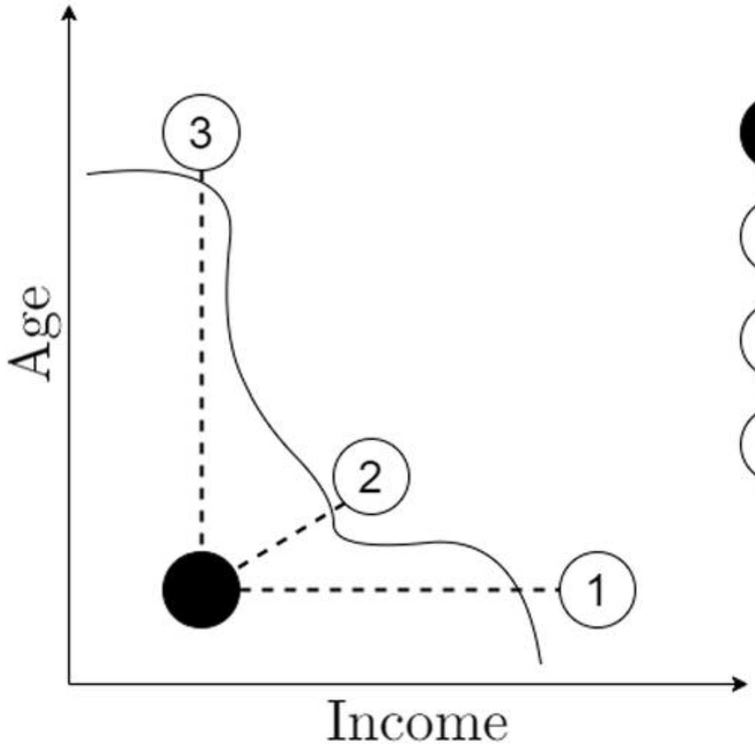

	Income	Age
●	$ 32,000	39
1	$ 40,000	39
2	$ 36,000	44
3	$ 32,000	140

Brughmans et al. 2023. Nice: an algorithm for nearest instance counterfactual explanations. *Data Mining and Knowledge Discovery*.

Counterfactual explanations

Counterfactual explanation (CF) is a type of XAI method that helps understand model decisions by showing how **small**, meaningful changes to the input could have led to a **different outcome**.

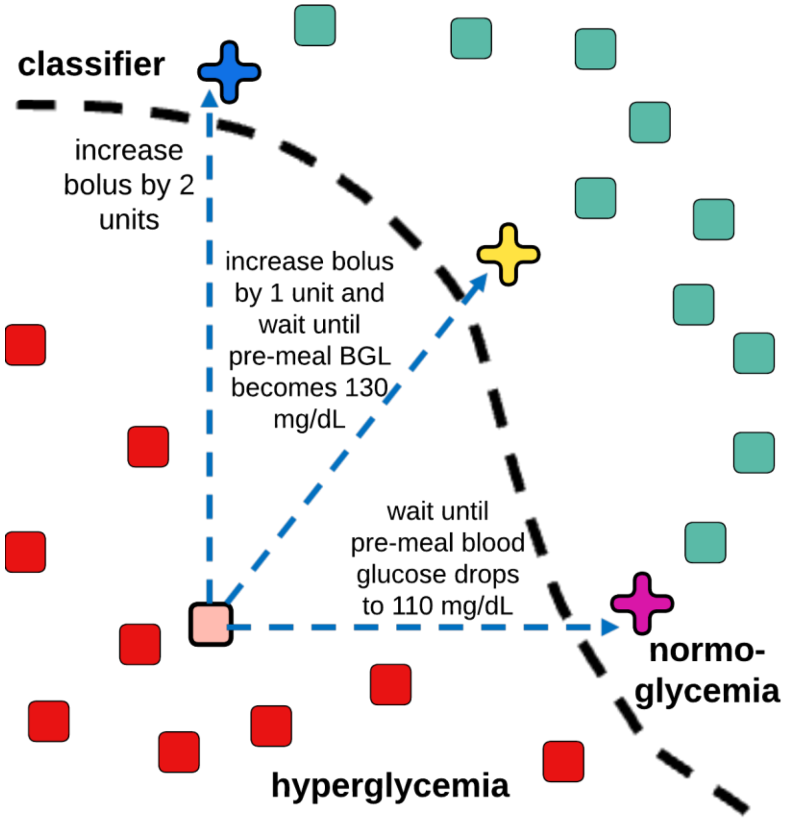

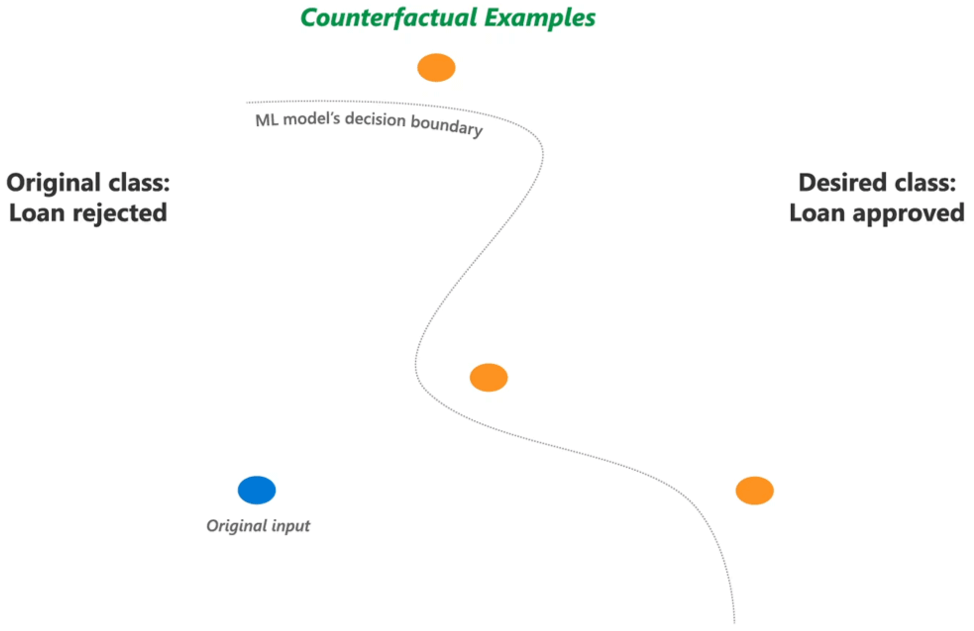

GlyRAG: augmented prediction with LLMs

- Fuse textual physiological context (from LLMs) with numerical CGM sequences.
- Use retrieval-augmented generation (RAG) to leverage similar historical glucose patterns during prediction.
- Improve forecasting accuracy and clinical reliability (especially for long horizons).

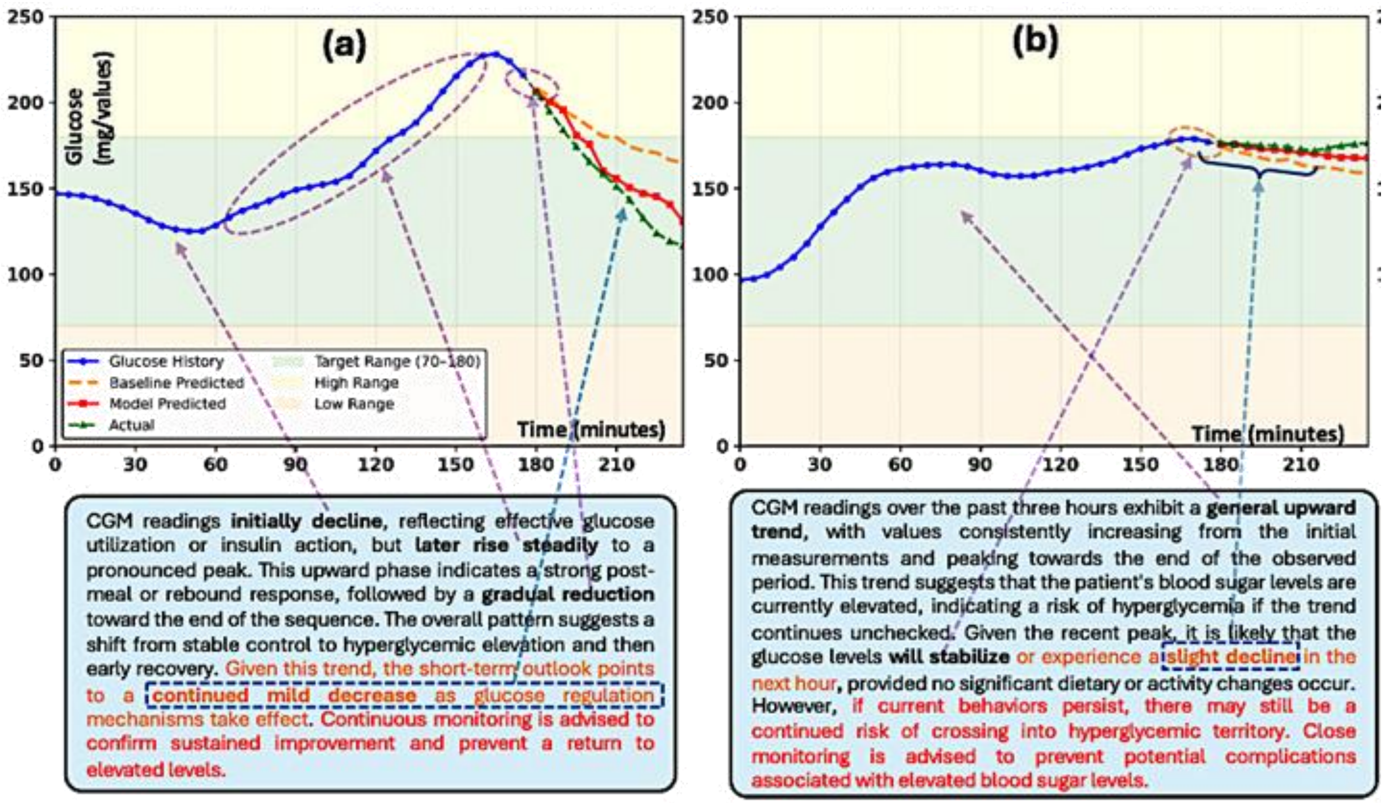

GlyRAG: results

- Context Generation: LLM analyzes past glucose sequences and generates textual summaries describing: Trend (rising/falling/stable), variability, rate of change, and hypo/hyper risk.
- Multimodal Transformer Encoder (MMT): Fuses CGM time-series and text embeddings using cross-attention.
- Cross-translational Loss: to align modalities (text ↔ glucose).
- Retrieval-Augmented Forecasting: Retrieves top-K similar embeddings from historical data to refine predictions.

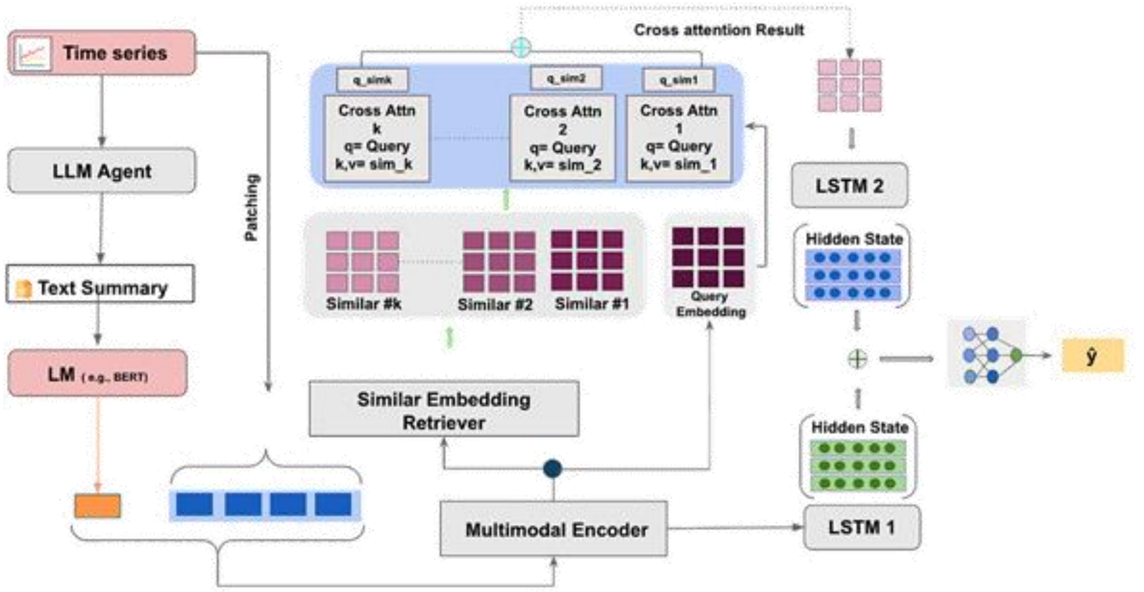

Ohio Dataset, PH = 60 minutes

Method	Sensitivity		Clarke Error Grid Regions (%)				
	Hyper	Hypo	A↑	B↓	C↓	D↓	E↓
GLIMMER [14]	86	42	85.4	13.26	0.13	1.10	0.02
CNN-LSTM	78	16	74.31	23.12	**0.11**	2.43	0.03
TimesFM [7]	97	81	75.95	22.36	0.29	1.33	0.07
GlyRAG	**97**	**92**	**85.53**	**13.59**	0.15	**0.83**	**0.025**

End-to-end system integration

- Generate, explain, simulate
- Predict, explain, prevent

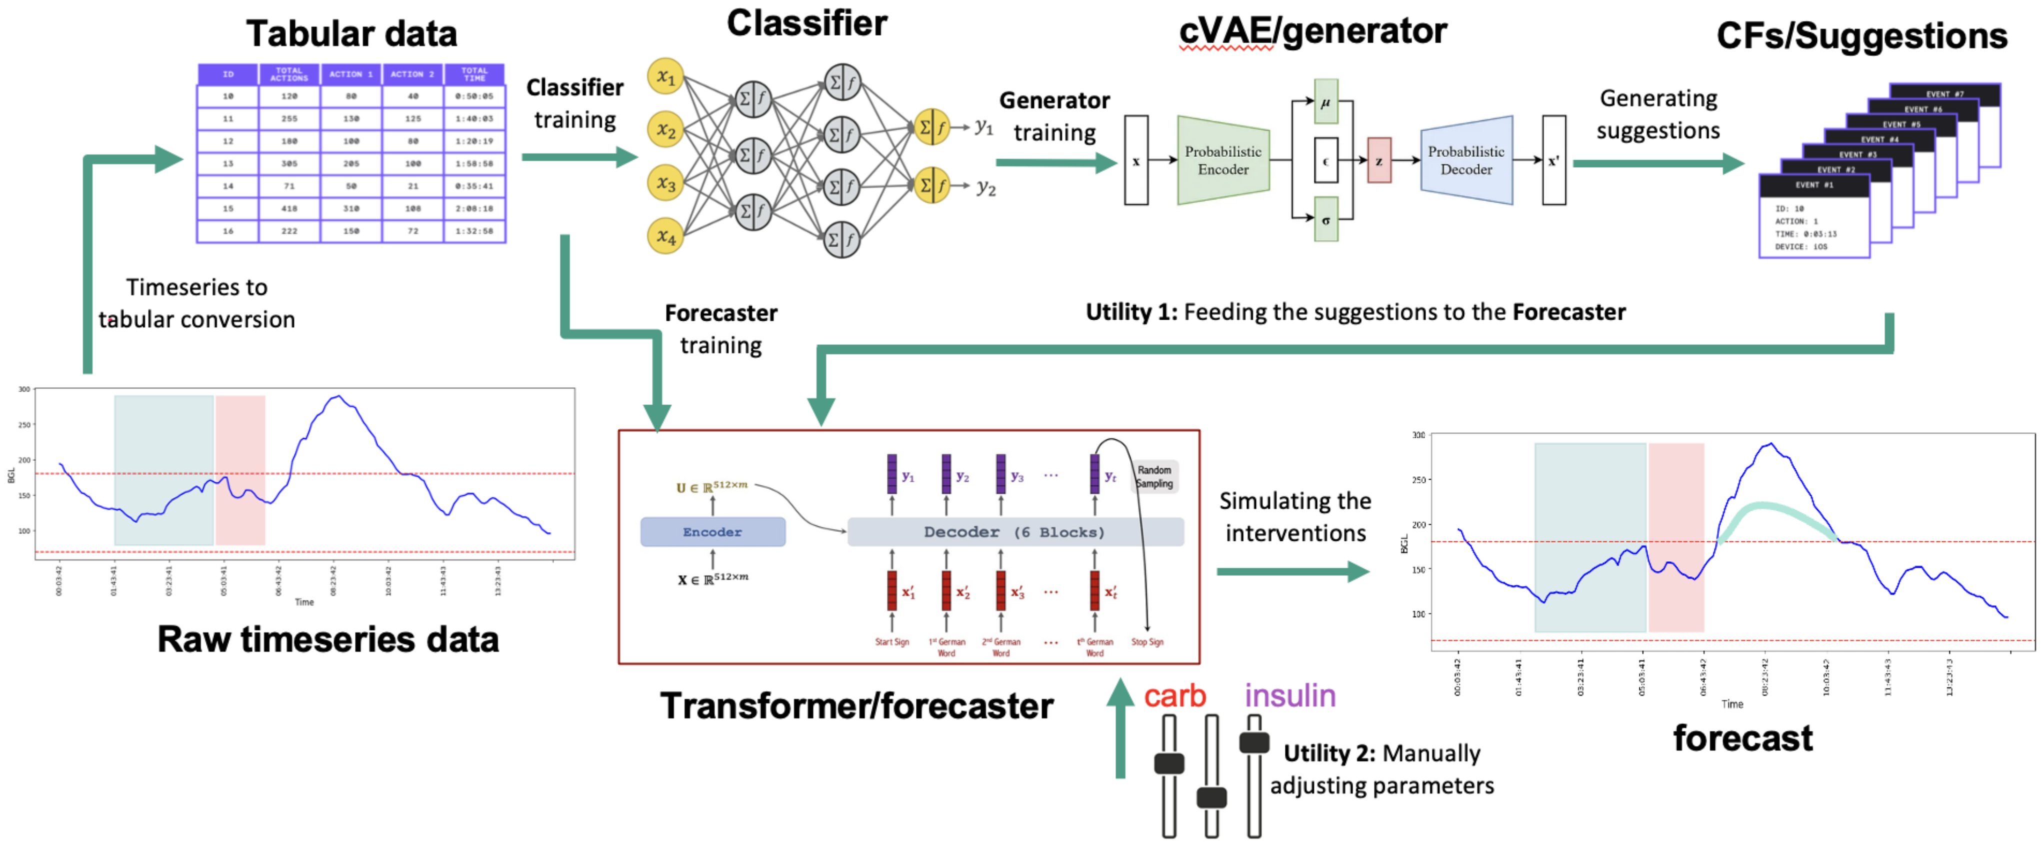

Clinical decision support

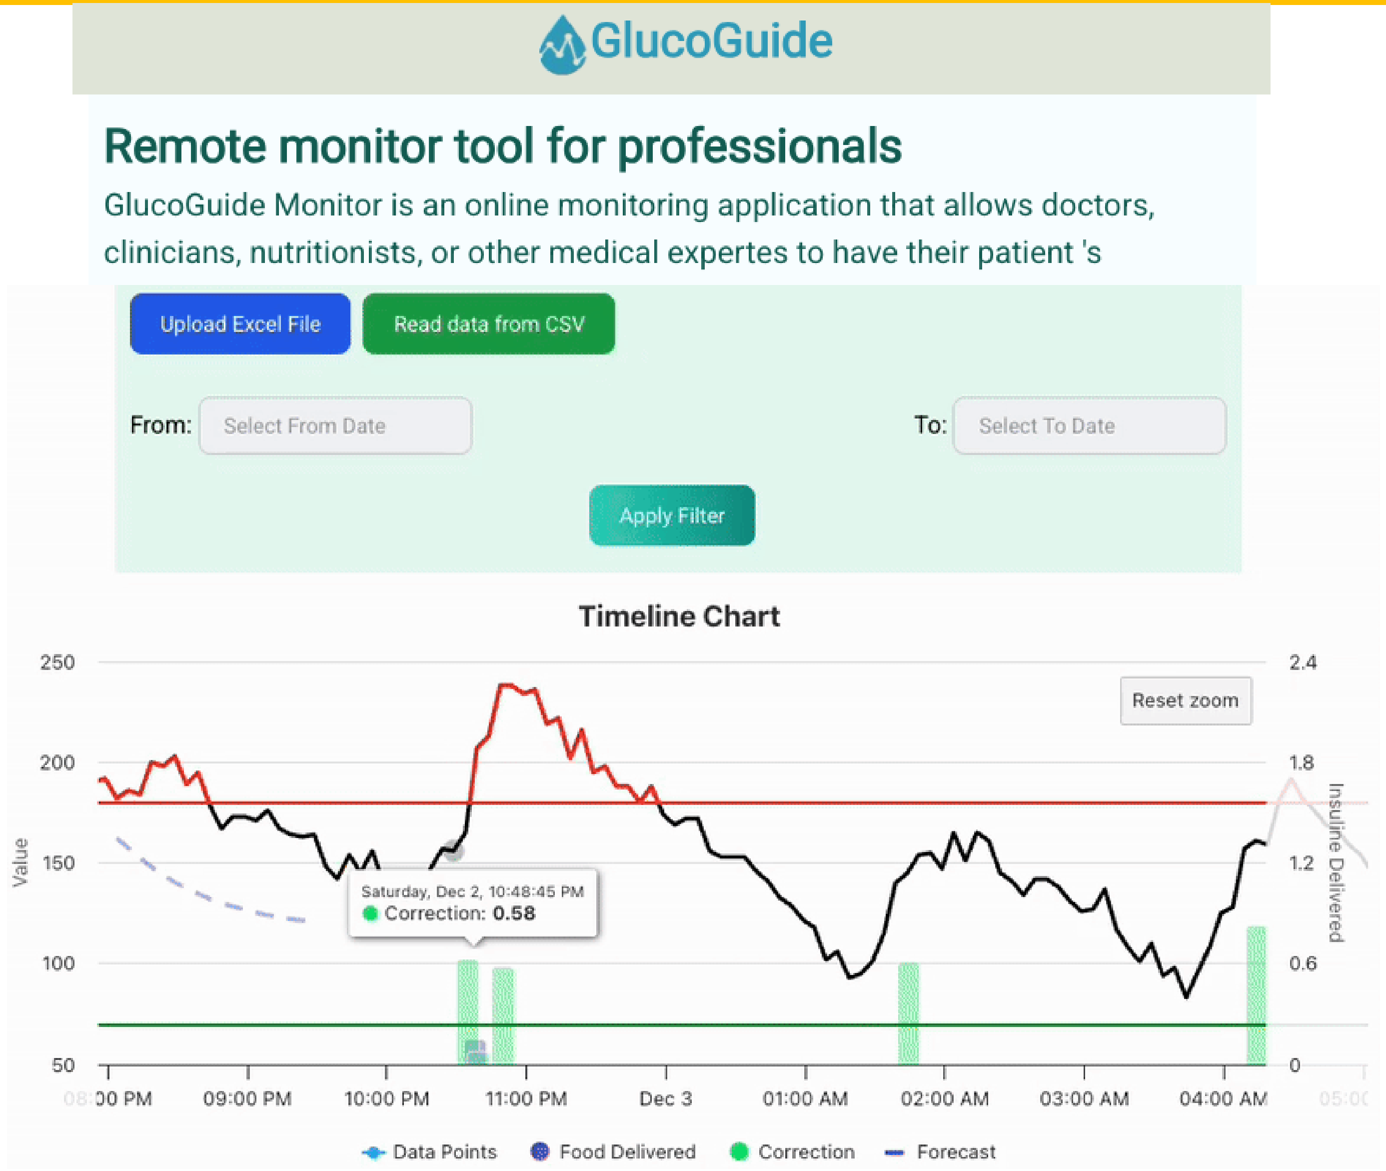

LIVE DEMO

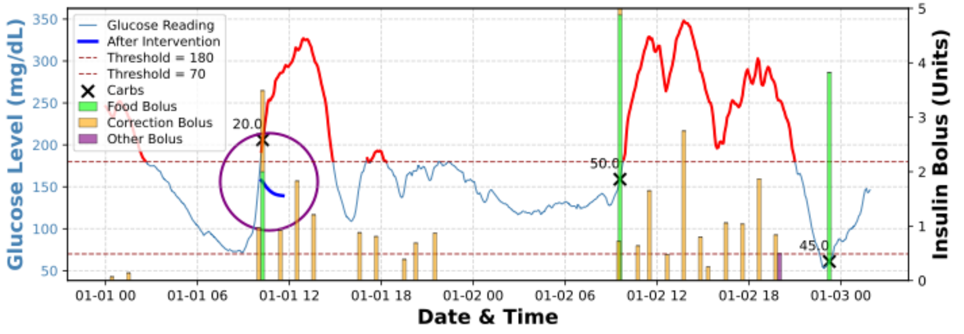

Practicum: LLM in XAI

- Go to this website: https://tinyurl.com/GlyRAG

- Install the Libraries

- Load the data

- Prompt LLM
  - *for context generation*
  - *CGM forecasting*

Reinforcement Learning

- A type of ML where an agent (learner) **interacts with an environment**, tries different **actions**, and **learns** by receiving **rewards** or **penalties**.

- Learning by trial and error, just like a human learning to ride a bike or play a game.

- No teacher gives the correct answer—the agent **learns from experience**.

Reinforcement Learning

- **Key Concepts:**
  - **Agent:** The decision-maker (e.g., a game bot, a smart thermostat).
  - **Environment:** The world the agent lives in (e.g., a game level, a house).
  - **State (S):** The current situation (e.g., current room temperature).
  - **Action (A):** What the agent can do (e.g., turn the heater on).
  - **Reward (R):** The feedback (e.g., positive if the room stays comfortable).

Reinforcement Learning

- **Key Concepts:**
  - The agent **observes** the current state.
  - It chooses an **action** (based on a policy).
  - The **environment** responds with a **new state** and a **reward**.
  - The agent **updates** its policy to do better next time.
  - This loop keeps going until the agent **learns the best strategy!**

Reinforcement Learning

- **Key Concepts:**

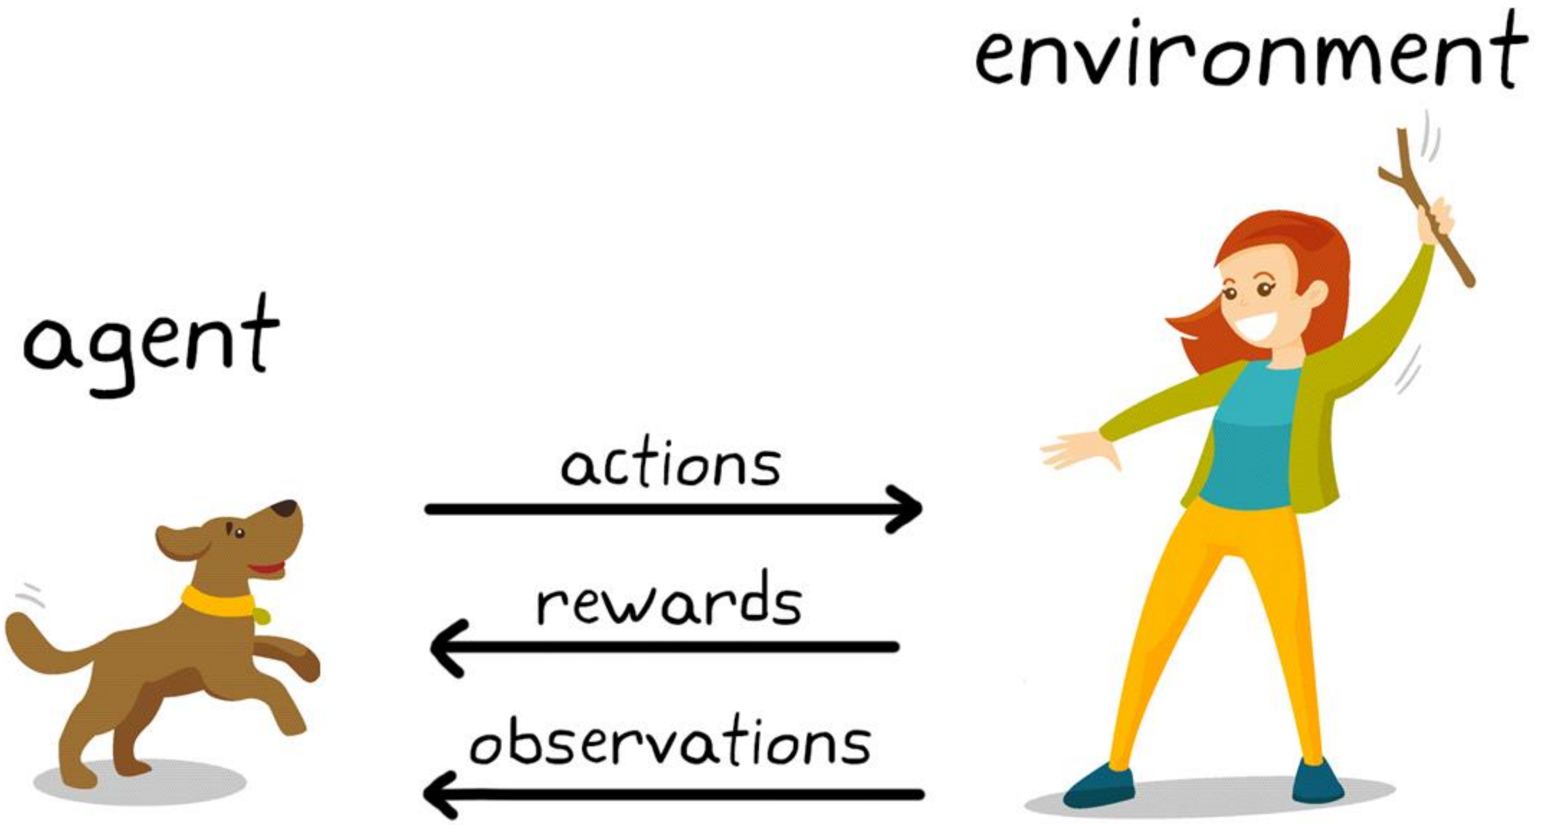

Practicum

- Go to this website: ***https://tinyurl.com/bdhvv9sp***

- Now that we understand how reinforcement learning works in theory — let's make it **visual and interactive**.

- We'll use a tool called ***The Beginner's RL Playground***, which lets us build a tiny world and watch an agent learn in real time.

Beginner’s RL Playground

- **Beginner’s RL Playground**

  - This grid represents the **environment**.
  - You can place rewards 💎, penalties ☠️, and walls 🚧.
  - The agent 🤖 starts somewhere and tries to learn how to **reach the goal** while **avoiding danger**.
  - We can control how hard or easy the world is.
  - We can move things and see **how the agent adapts**.

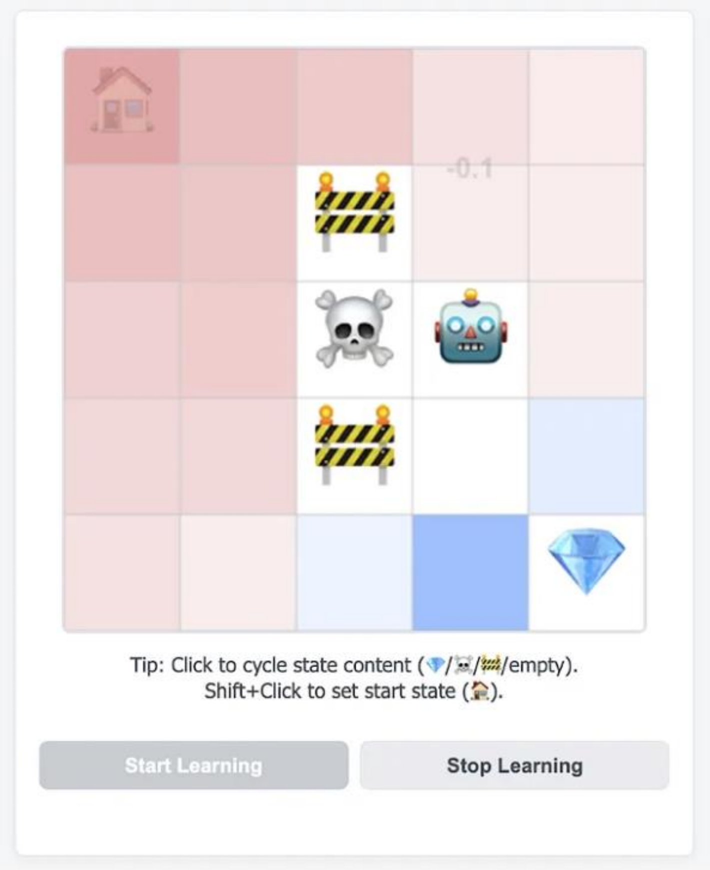

Reinforcement Learning

- **Environment Settings (Right Panel):**
  - Change Step Penalty – discourage long paths.
  - Modify Gem Reward and Skull Penalty.
  - Set Speed – slow down to observe learning.
  - Try moving the gem or hazards during learning!

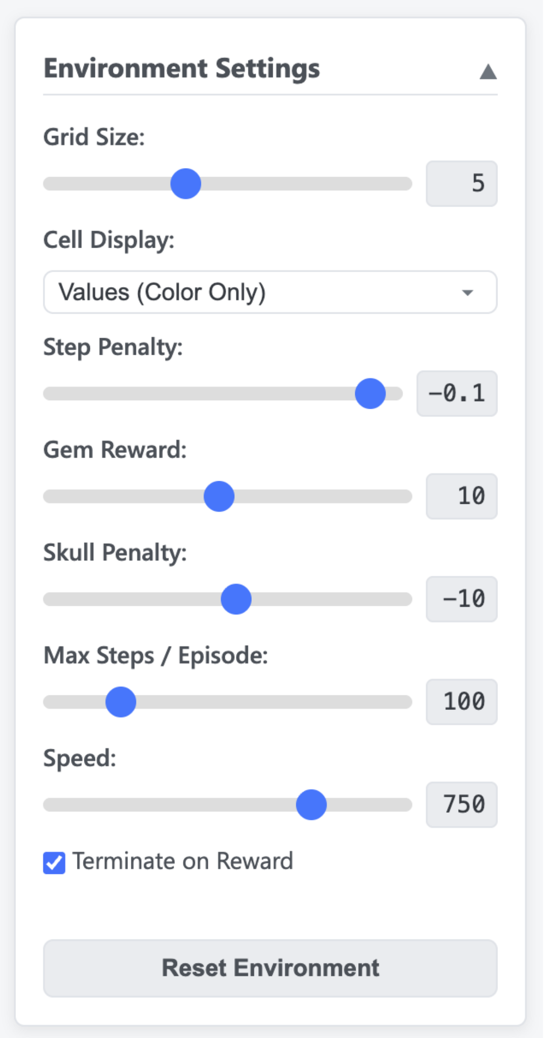

Reinforcement Learning

- **Algorithm Settings (Left Panel):**

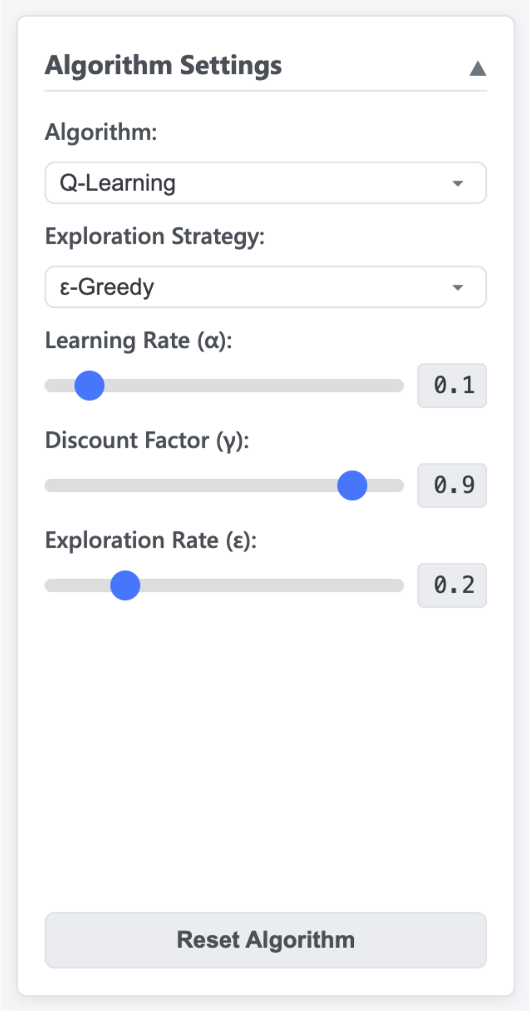

- **Choose between different learning algorithms**
  - Q-Learning, SARSA, Actor-Critic, etc.
- **Adjust Learning Rate (α)**
  - how quickly the agent updates.
- **Adjust Discount Factor (γ)**
  - how much it values future rewards.
- **Adjust Exploration Rate (ε)**
  - how often it tries new paths.

Reinforcement Learning

- **Practicum**

https://tinyurl.com/bdhvv9sp

5 Minutes

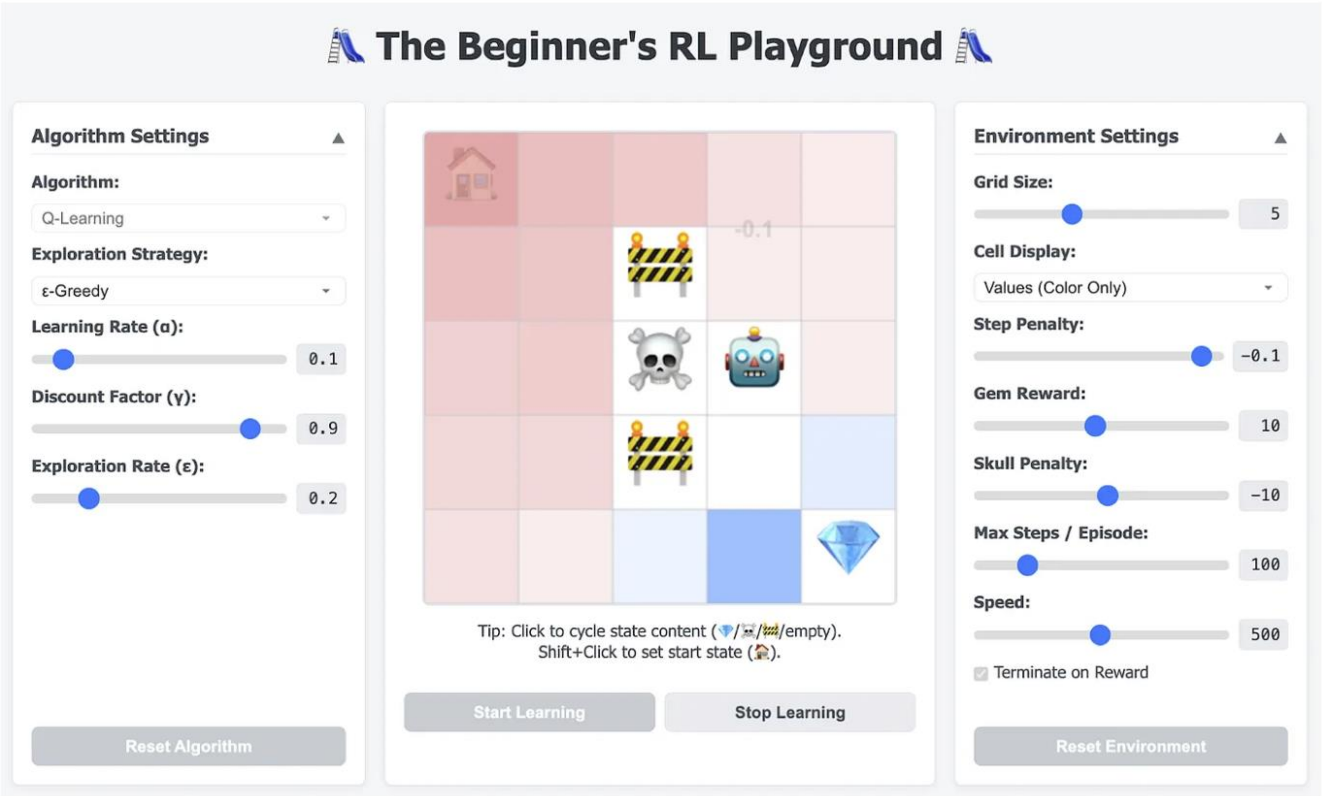

Scan to Access!

Source: *https://awjuliani.github.io/web-rl-playground/*

Practicum

- **Task 1:** Set up a small grid (e.g., 3x3), place the robot far from the gem, and run learning. Then try again with grid size = 7x7 — and place the robot in the same relative corner.

- **Task 2:** In a 5x5 grid, run learning with one layout (e.g., clear path to gem). Then add walls (🚧) and re-run without changing any algorithm settings.

College of Health Solutions
Arizona State University

Embedded Machine Intelligence Lab

Thank You!

Hassan Ghasemzadeh
Director
Arizona State University

Saman Khamesian
Graduate Research Associate
Arizona State University

Shovito Barua Soumma
Graduate Research Associate
Arizona State University

Embedded Machine Learning Lab (EMIL)
Hassan.Ghasemzadeh@asu.edu
ghasemzadeh.com