ASU College of Health Solutions
Arizona State University

EMBEDDED MACHINE INTELLIGENCE LAB

Case Study: Reinforcement Learning in Metabolic Health

Saman Khamesian, Hassan Ghasemzadeh

College of Health Solutions
Embedded Machine Intelligence Lab
Arizona State University

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

From Walking to Type 1 Diabetes: A Real RL Challenge

- So far, we've looked at a simple example — using RL to encourage daily walking.
- It's a familiar, low-stakes problem that helps us understand how to define (or design) the environment, state, action, and reward.
- But in real-world healthcare, problems are often far more **complex**.
- One such example is the **management of Type 1 Diabetes**.
- In this context, decisions about **insulin** and **meals** influence blood glucose levels in ways that are **delayed**, **nonlinear**, and **potentially dangerous** if handled incorrectly.

Type 1 Diabetes: Problem Definition

- People with T1D don’t produce insulin naturally, so they need to carefully decide when to eat and how much insulin to take.

- Making the wrong decision can lead to dangerous situations:
  - **Hypoglycemia:** When blood glucose drops below 70 mg/dL → dizziness, confusion, or even coma.
  - **Hyperglycemia:** When it goes above 180 mg/dL → long-term damage to the heart, kidneys, eyes, and nerves.

Type 1 Diabetes: Problem Definition

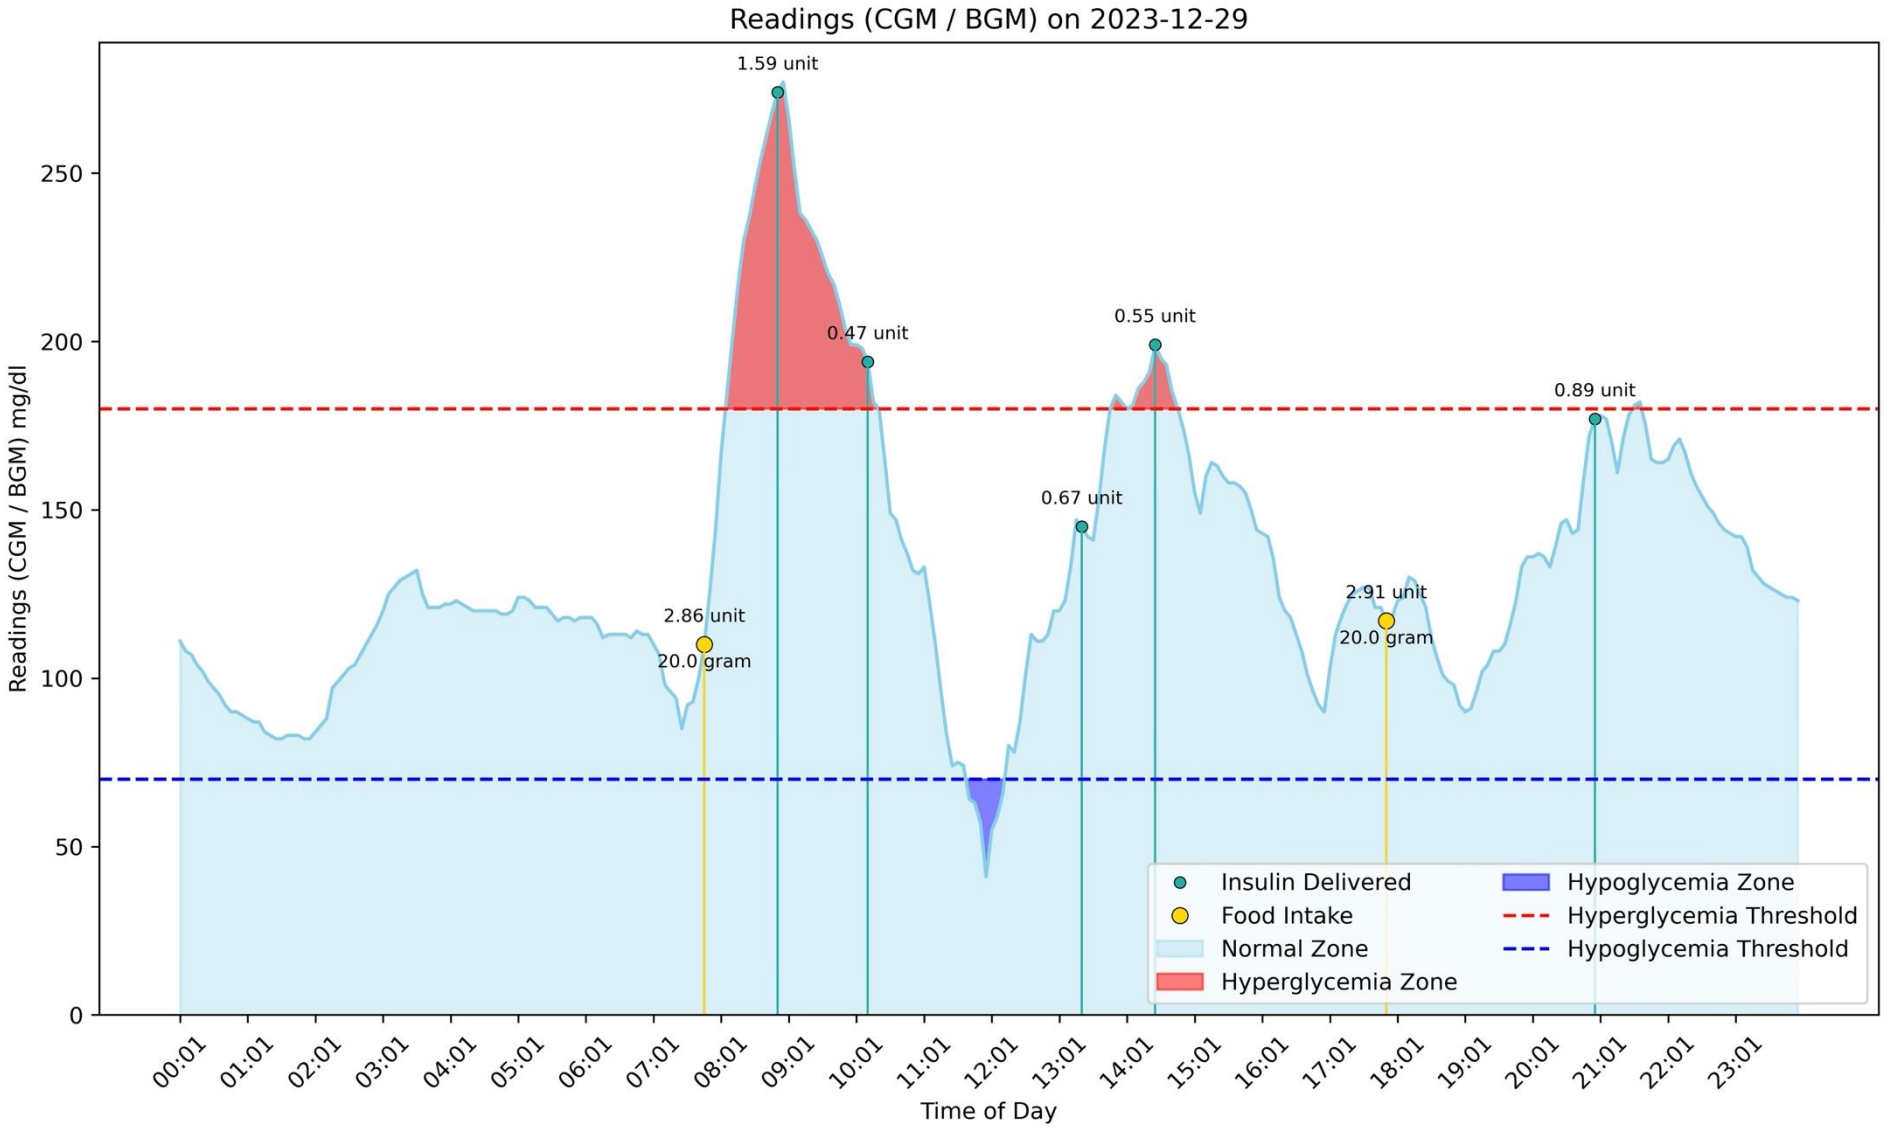

Type 1 Diabetes: Problem Definition

- The goal is to **keep blood glucose in a healthy range (70-180 mg/dL)** — we call this the **Time in Range (TIR)** metric.

1. How can we increase their time-in-range?
2. How can we help patients prevent hyperglycemia and hypoglycemia?
3. What is the best timing and dosage of bolus insulin and meals for each patient?

Type 1 Diabetes: Questions

- The idea is to build an RL framework that helps patients stay in range, by suggesting low-frequency, high-impact actions — like when to eat, or how much insulin to inject.

- On the other hand, we need a human body (an individual with T1D) to take these actions and observe the glucose response.

- But it's DANGEROUS, Impractical, and Costly!

Type 1 Diabetes: Road to RL Framework

- Human body is our environment.

- Now, what if we replace the human body with a simulator that can model the glucose response for us? It's like a prediction model that estimates glucose levels for us.

- We simplify the environment this way, given the real-world challenges and limitations of input data.

Type 1 Diabetes: Road to RL Framework

- For example, if the RL system suggests injecting 2 units of insulin or consuming 50 grams of carbohydrates, the simulator **predicts** how these actions will affect blood glucose levels over time.

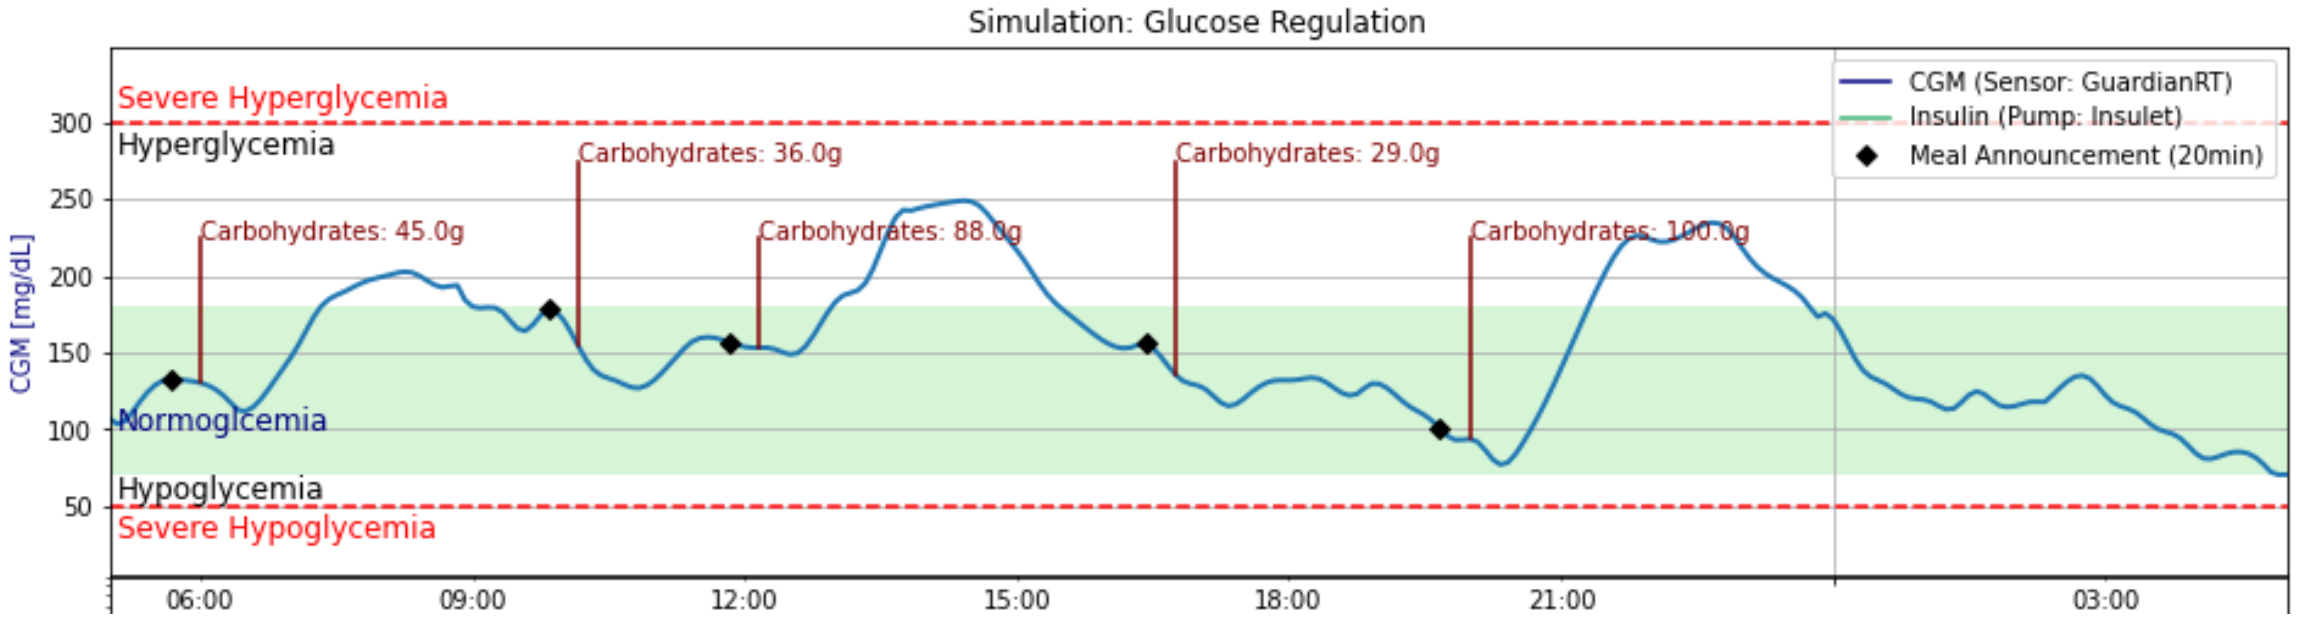

Type 1 Diabetes: Road to RL Framework

Send Action

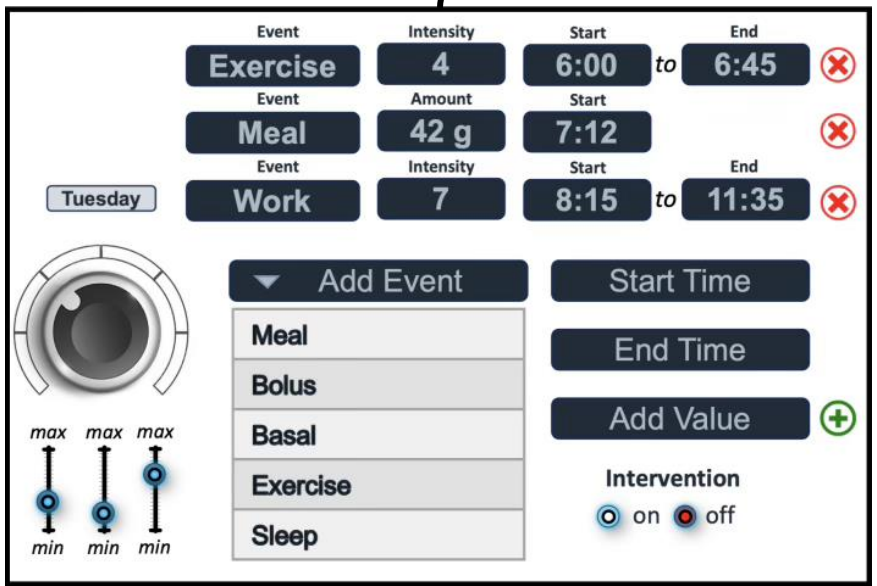

RL Framework
(Generates Recommended Actions)

The simulator acts as a controlled environment that provides feedback to the RL framework.

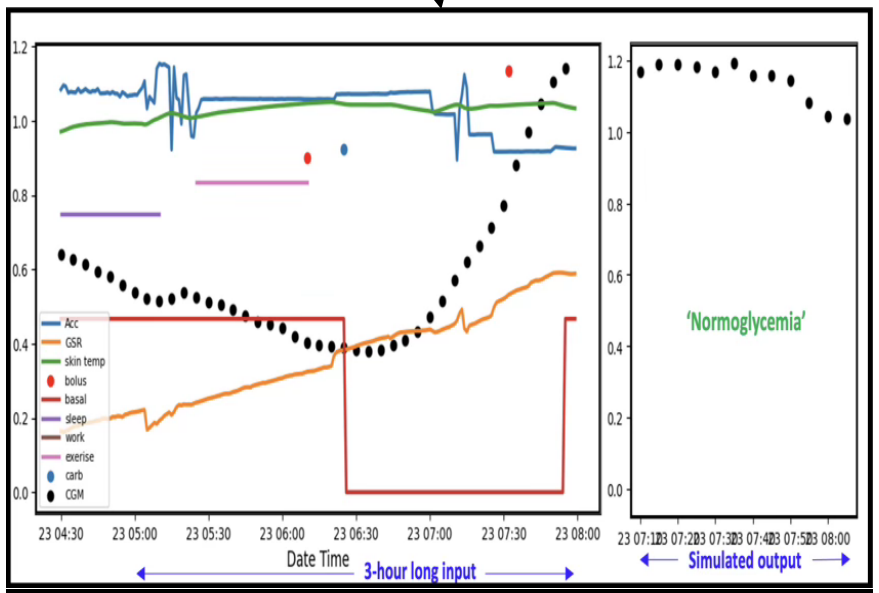

Simulator
(Generates Simulated BG level)

Simulated BGL

Type 1 Diabetes: Road to RL Framework

- The RL agent acts as the brain of the system.
- It processes real-world signals like:
  - CGM data,
  - insulin delivery history,
  - activity levels, and
  - nutrition values to model the patient's state.
- to understand the current state of the patient and decide what to do.

Type 1 Diabetes: Road to RL Framework

What the agent observes (input features)

At each step, the agent sees the past 6 hours (72 timesteps) of 7 signals:

$$s_t = \left[H,\, Z,\, G,\, C,\, I^{\text{bolus}},\, \Delta^{\text{meal}},\, \Delta^{\text{inject}}\right]$$

- H: Hour of day (0–23)
- Z: Sleep indicator (0 or 1)
- G: CGM glucose values
- C: Carb intake
- I^{bolus}: Insulin doses
- Δ^{meal}, Δ^{inject}: Time since last meal/injection

Context-aware input helps personalize and stabilize decision-making.

Type 1 Diabetes: Road to RL Framework

What the agent can do (outputs)

Each action is a 6-dimensional vector:

$$a_t = [\text{do nothing},\ \text{eat},\ \text{inject},\ c_t,\ i_t,\ k_t]$$

- `do nothing`, `eat`, `inject`: The agent chooses **one action**
- $c_t \in [5, 50]$: How many carbs to eat (if choosing to eat)
- $i_t \in [2, 15]$: How many insulin units to inject (if choosing to inject)
- $k_t \in \{0, \ldots, 11\}$: **When** to take the action in the next hour (in 5-minute steps)

🎯 Helps make **realistic and timely decisions** based on the situation.

Type 1 Diabetes: Turning Actions into Human Advice

- Even if our RL system learns the best action, we're still left with a problem:
- Imagine a patient receives: "Eat 7g of carbs in 25 minutes"
  - ❓ Will they understand why?
  - ❓ What should they eat?
  - ❓ Is it okay if they're about to go to sleep?
- We need personalized explanations and context.
- That's why we need **personalized explanations**, not just numbers.

Type 1 Diabetes: Turning Actions into Human Advice

- This is where LLMs — like ChatGPT — can play a key role.
- LLMs can interpret these numeric actions:
  - Understand the user's history
  - Add context (e.g., time of day, recent meals, sleep)
  - Transform "Eat 7g carbs"

  → "Have a small apple slice now to avoid low blood sugar."
- They become health advisors, not just calculators.

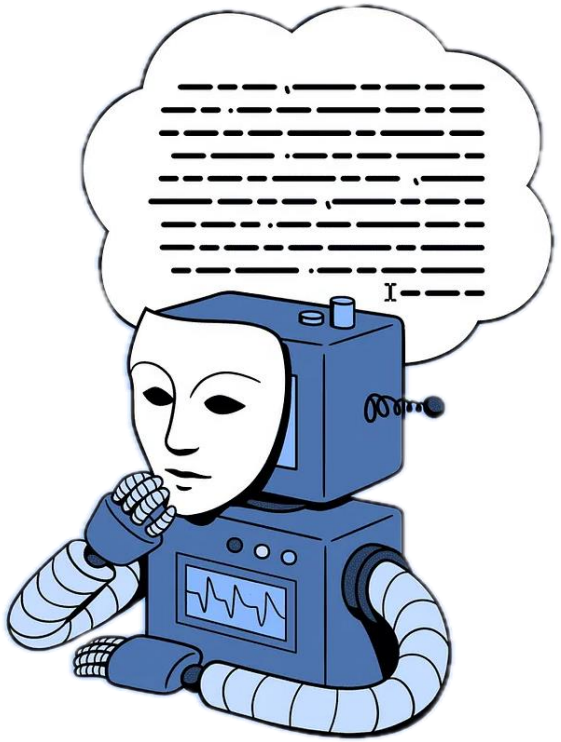

Large Language Model?
A deep learning model trained on vast text data to understand and generate human-like language. It can summarize, translate, answer questions, and create text.

Type 1 Diabetes: Turning Actions into Human Advice

Raw Action	LLM-Based Recommendation
`eat = 10g`, `in 20 mins`	*Have a small apple slice in 20 minutes to prevent a glucose dip.*
`inject = 3u`, `in 5 mins`	*Take 3 units of insulin in 5 minutes to cover your recent meal.*
`do nothing`	*No action needed — your glucose is stable.*
`eat = 15g`, `in 10 mins`, near bedtime	*Consider a small snack like crackers before bed to avoid nighttime lows.*
`inject = 5u`, after large meal	*Take 5 units of insulin now to help manage your post-meal spike.*

Type 1 Diabetes: Prompt Example

Description: You are a diabetes coach helping a patient manage their blood glucose. Based on the patient's current context and the AI agent's recommended action, generate a short, clear, and supportive suggestion.

Patient background:
- Age: 27
- Diagnosis: Type 1 Diabetes
- Currently using CGM and insulin pens
- Usually eats 3 main meals and 1 snack daily
- Recently had hypoglycemia during sleep
- Prefers simple, familiar foods like fruits, toast, crackers

Current situation:
- Time: 8:15 PM
- Glucose: 85 mg/dL and trending slightly downward
- Last meal: 3.5 hours ago
- Patient is getting ready for bed

AI-recommended action:
- Action type: eat
- Carb amount: 10g
- Timing: in 15 minutes

Response format:
Generate a one-sentence recommendation that explains **what** to do and **why**, using clear, supportive language. Avoid medical jargon.

Example output:
“Have a small apple slice in 15 minutes to avoid a low during the night.”

Scan to access to the prompt!

https://shorturl.at/7LT7y

Type 1 Diabetes: Practice

- You can try this prompt as an input for different language models:
  - ChatGPT,
  - Google Gemini,
  - Meta Llama.
- Modify the input and observe how the output changes.
- Then, try another example from a different domain (e.g., diet).

College of Health Solutions
Arizona State University

EMBEDDED MACHINE INTELLIGENCE LAB

Hassan Ghasemzadeh
Director
Arizona State University

Saman Khamesian
Graduate Research Associate
Arizona State University

Embedded Machine Learning Lab (EMIL)

Hassan.Ghasemzadeh@asu.edu
ghasemzadeh.com

College of Health Solutions
Arizona State University

EMBEDDED MACHINE INTELLIGENCE LAB

Thank You!