CAPSML: Bridging the Gap Between Clinicians, Lived Experience, and AI Systems for Glucose Regulation in Type 1 Diabetes.

cCoP
Clinical Community of Practice
Annual Meeting 2024

CAPSML

Chirath Hettiarachchi

School of Computing, College of Engineering, Computing, and Cybernetics,
Australian National University

April 17, 2024

Australian National University

NCI AUSTRALIA

Presentation Overview

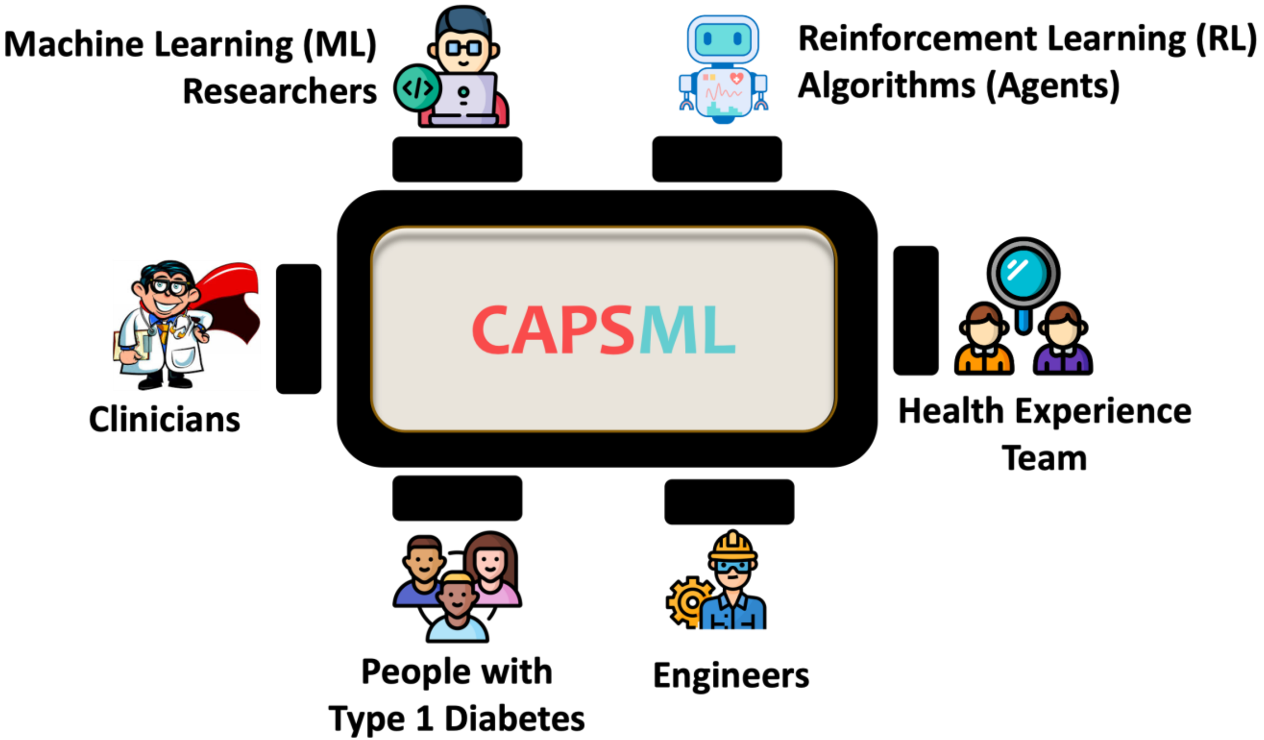

The tool is publicly available as an online tool at “capsml.com”

Agenda

1.1 Background: Glucose Regulation

Sleep

Exercise

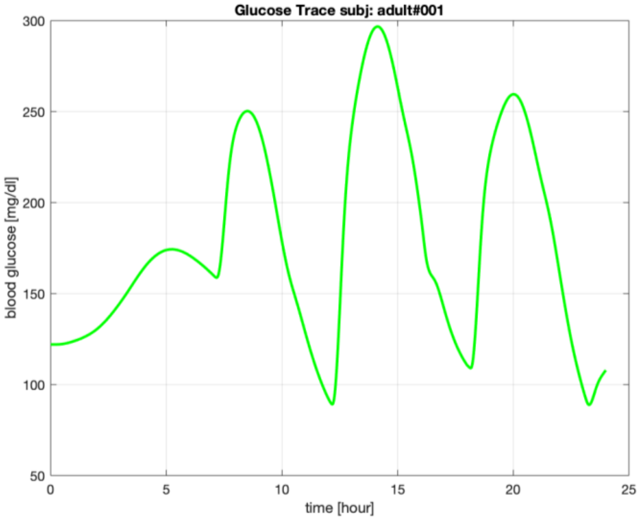

Meals

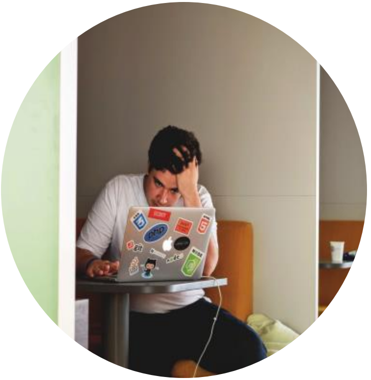

Stress

The fluctuation of blood glucose levels.

1.1 Background: Glucose Regulation

Sleep

Meals

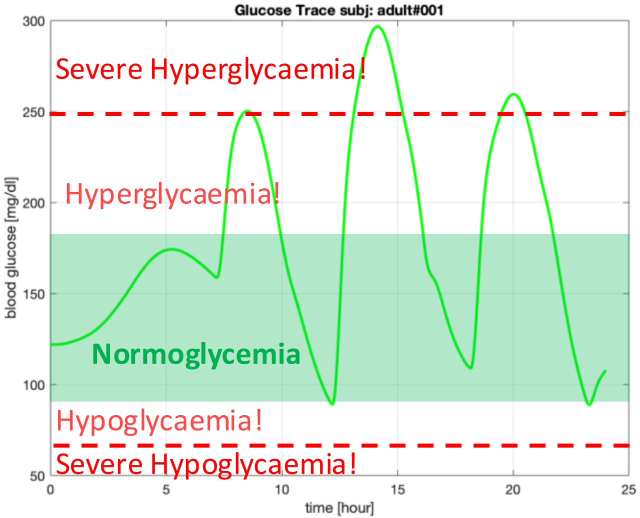

Exercise

Stress

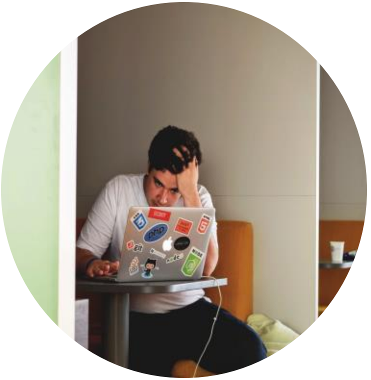

Maintaining glucose homeostasis is vital!

1.2 Background: Artificial Pancreas Systems (APS)

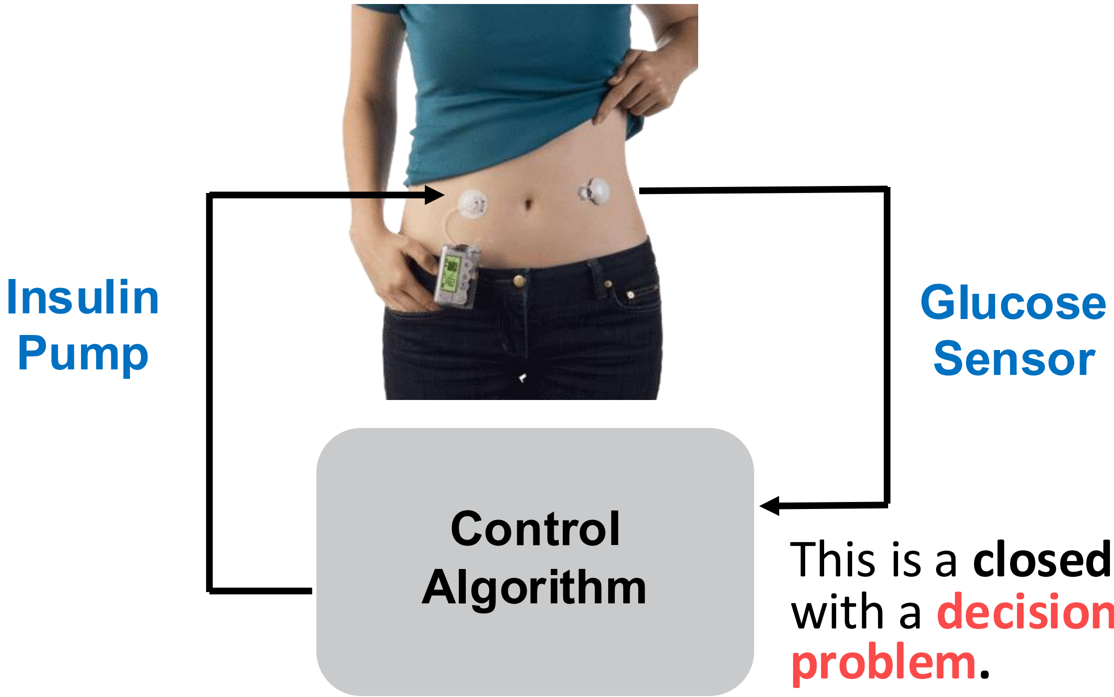

This is a **closed-loop system**, with a **decision-making problem**.

† Existing APS are **hybrid-closed loop systems**, requires manual inputs **(meal announcement & carbohydrate estimation)**.
First commercial systems introduced in USA (2016); Australia (2019).

1.3 Background: Reinforcement Learning (RL)

Artificial Intelligence (AI)
└→ Machine Learning (ML)
  └→ Reinforcement Learning (RL)

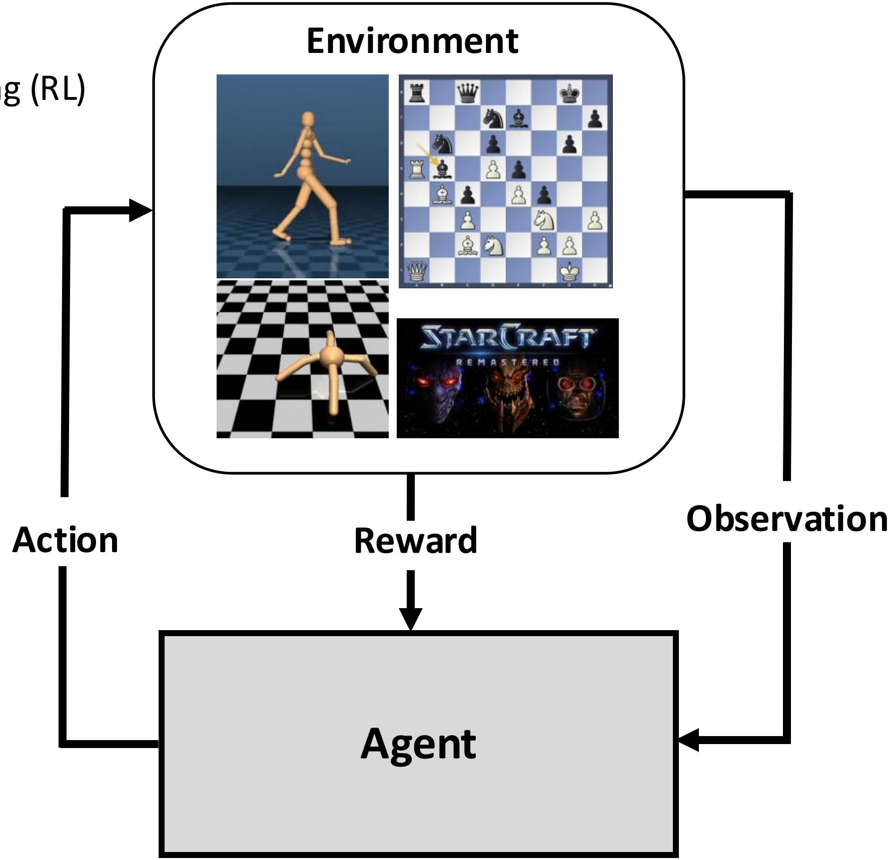

1.4 Background: RL-based APS – Towards full automation.

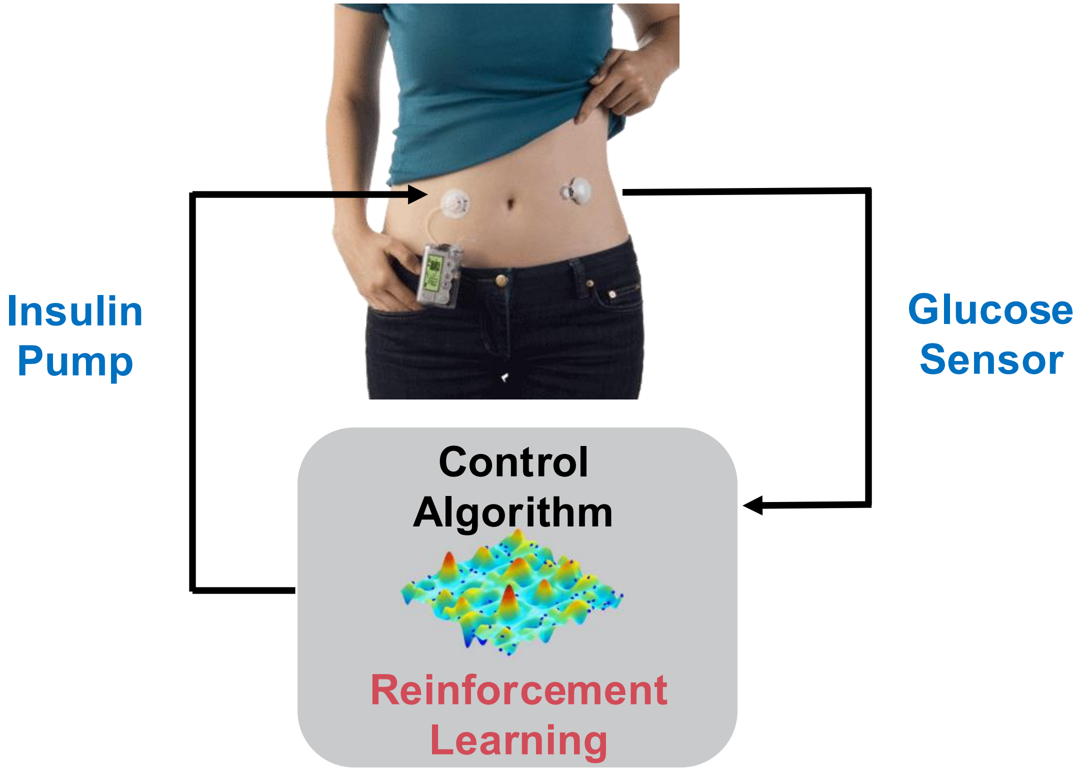

1.4 Our Aim: RL-based APS – Towards full automation.

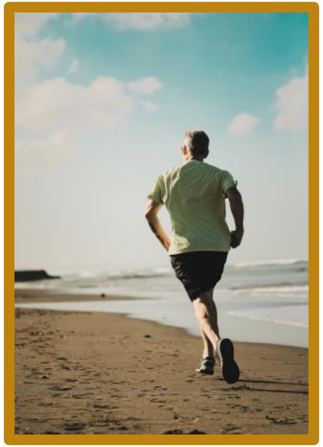

How much **Insulin** do I need?

† Meal announcement (typically 20 minutes in advance) affects the quality of life & carbohydrate estimation is challenging and mistakes lead to poor glucose regulation.

02. Introduction: CAPSML

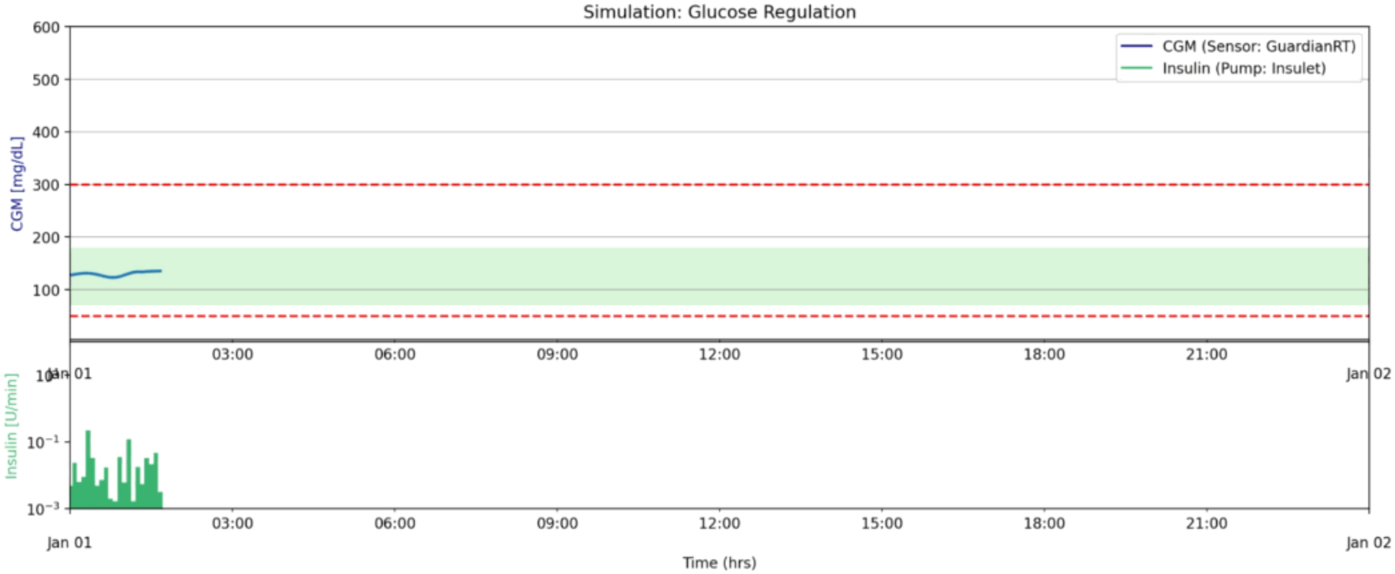

The treatment strategies learnt by the RL algorithms are complicated and hard to understand.

02. Introduction: CAPSML

Table. Clinical performance comparison of RL-based APS and clinical treatment.

Method	Failure (%)	Severe Hypo.(%) TBR-Level 2	Hypo. (%) TBR-Level 1	Normo. (%) TIR	Hyper. (%) TBR-Level 1	Severe Hyper.(%) TBR-Level 2	RI	LBGI	HBGI
BBI	0.39	0.00† 0.00-0.00* 0.06(0.40)‡	0.00 0.00-0.00 0.64(2.01)	70.83 61.46-79.17 71.02(11.29)	22.57 17.36-31.94 24.44(10.57)	2.08 0.00-4.86 3.85(5.72)	7.68 6.07-9.38 8.35(4.05)	0.81 0.28-1.68 1.33(1.60)	6.72 4.97-7.94 7.02(2.79)
BBHE	0.35	0.00 0.00-0.00 0.05(0.36)	0.00 0.00-0.00 0.40(1.42)	69.79 60.42-78.47 69.78(11.29)	23.61 18.06-32.29 25.44(10.57)	3.12 0.00-5.21 4.33(5.64)	7.62 5.78-8.69 8.00(3.74)	0.41 0.11-0.93 0.88(1.37)	6.92 5.11-8.11 7.11(2.67)
PPO	2.79	0.00 0.00-0.00 0.32(1.51)	0.00 0.00-1.04 1.19(2.77)	69.44 62.15-76.04 69.12(10.53)	20.83 16.32-25.35 20.93(7.11)	7.99 2.43-12.85 8.44(6.88)	9.56 6.89-12.04 9.79(3.66)	0.89 0.28-2.17 1.64(2.11)	8.05 5.86-10.10 8.14(3.05)
G2P2C	1.62	0.00 0.00-0.00 0.22(1.26)	0.00 0.00-1.04 1.11(2.50)	72.57 66.32-79.86 72.69(9.53)	18.75 15.28-23.26 19.37(6.46)	6.60 0.69-10.76 6.61(5.53)	9.00 6.21-11.04 8.94(3.18)	1.04 0.43-2.14 1.58(1.74)	7.42 5.18-9.35 7.36(2.60)

Statistical performance metrics only provide limited information into the RL-based strategies and doesn't provide any insights towards the operation of the algorithm.

02. Introduction: CAPSML

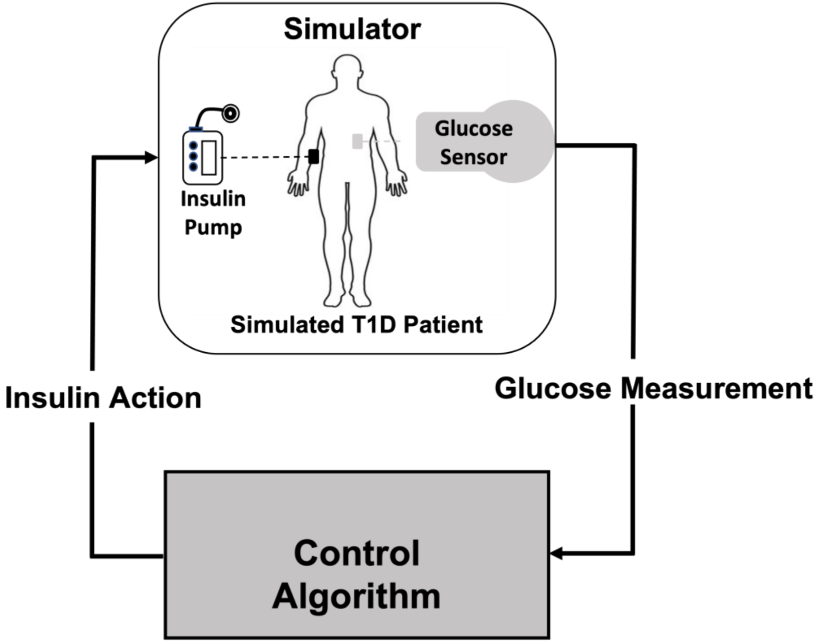

Existing simulation tools require technical knowledge and inaccessible to the general public.

2.1 Quick Start

The tool is publicly available online at: "capsml.com" under the MIT license.
A demonstration of the system is available at: https://youtu.be/JO5MkPCuqCw.

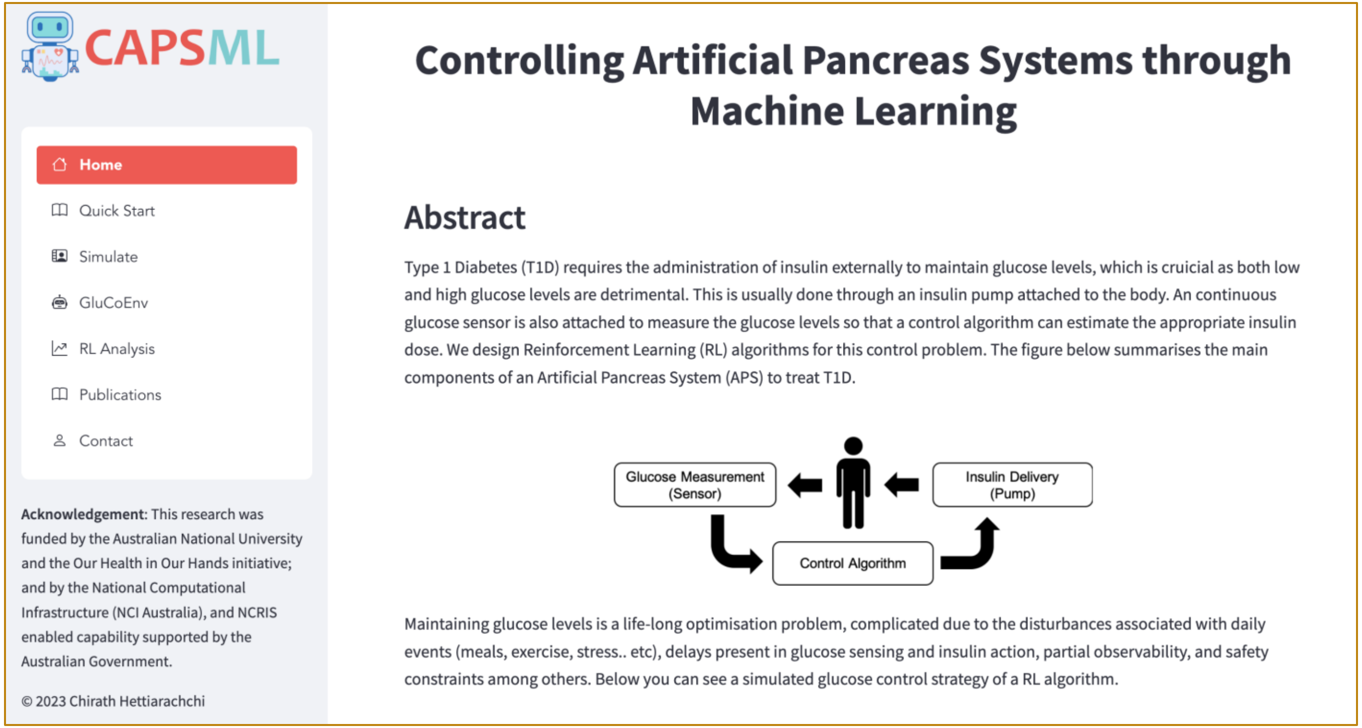

2.1 Quick Start

The Simulator

- The simulator used in CAPSML is based on the open-source **Simglucose Simulator**[1], which uses the **FDA-approved UVA/PADOVA 2008**[2] model.

- **Used for pre-clinical trials.**

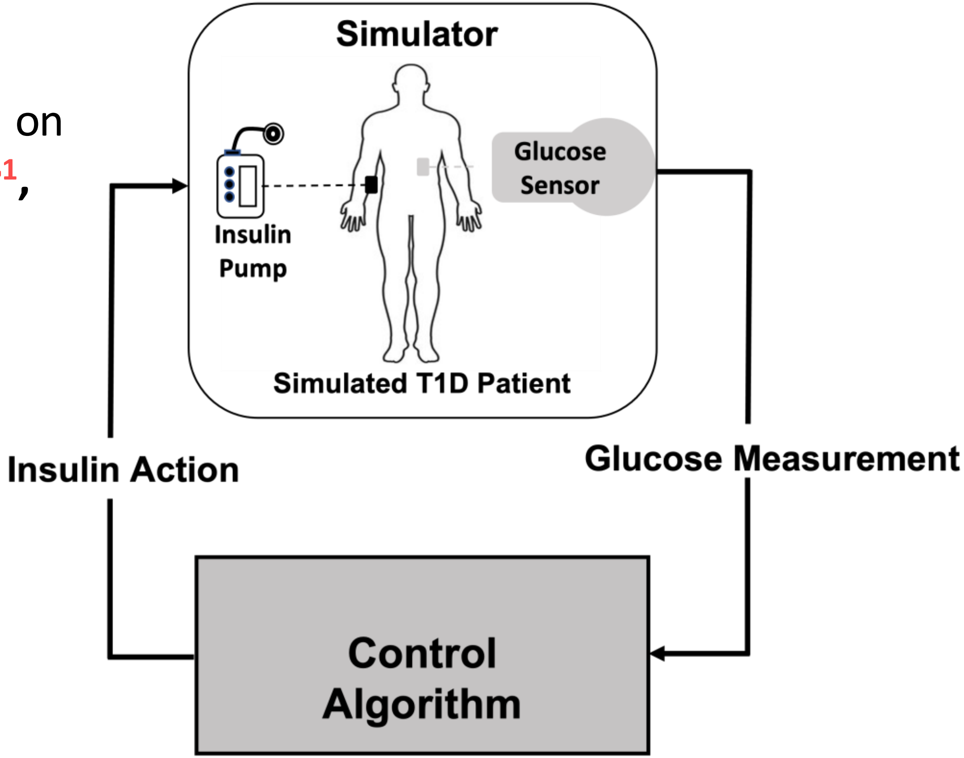

[1] Xie, Jinyu. "Simglucose v0. 2.1 (2018)." *URL https://github. com/jxx123/simglucose* (2018).
[2] Kovatchev, Boris P., et al. "In silico preclinical trials: a proof of concept in closed-loop control of type 1 diabetes." (2009): 44-55.

2.1 Quick Start

Alice Bob

Chloe David

Midnight Breakfast Lunch Dinner

8am 12pm 7pm

40g 80g 100g

Control Algorithm

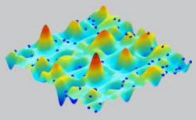

Run Simulation

Step 1
Select *in-silico* subject

Step 2
Setup meal protocol

Step 3
Select control algorithm

Step 4
Run simulation

2.1 Quick Start

Step 1. Select T1D *in-silico* Subject

Alice **Bob** **Chloe** **David**

- Select from 10 Adolescents & 10 Adults.
- Subjects have different body Weights, Ages, Insulin to Carbohydrate Ratios (ICR), & Insulin Sensitivity Factors (ISF).

Table. Parameters of *in-silico* subjects.

	Subject Name	Body Weight (Kg)	Age	Total Daily Insulin (TDI)	Insulin to Carbohydrate Ratio (ICR)	Insulin Sensitivity Factor (ISF)
0	adolescent0	68.7060	18	36.7339	12	15.0360
1	adolescent1	51.0460	19	62.0310	5	13.1751
2	adolescent2	44.7910	15	24.2428	23	33.5260
3	adolescent3	49.5640	17	35.2529	14	21.8225
4	adolescent4	47.0740	16	34.0047	12	20.9084
5	adolescent5	45.4080	14	49.5813	7	17.6966
6	adolescent6	37.8980	16	43.6381	8	12.4932
7	adolescent7	41.2180	14	63.3867	4	11.9355
8	adolescent8	43.8850	19	24.0782	21	20.0139
9	adolescent9	47.3780	17	33.1735	14	31.8685
10	adult0	102.3200	61	50.4167	10	8.7731
11	adult1	111.1000	65	57.8688	8	9.2128
12	adult2	81.6310	27	56.4297	9	17.9346
13	adult3	63.0000	66	33.8079	16	42.6534
14	adult4	94.0740	52	68.3159	5	8.2313
15	adult5	66.0970	26	61.3888	10	18.2133
16	adult6	91.2290	35	42.0066	22	26.1531
17	adult7	102.7900	48	42.7788	13	12.2506
18	adult8	74.6040	68	67.2115	5	7.6432
19	adult9	73.8590	68	64.4485	5	10.6926

2.1 Quick Start

Step 2. Select Meal Protocol - Set up meals based on the time and carbohydrate content or portion size of different meals/food items.

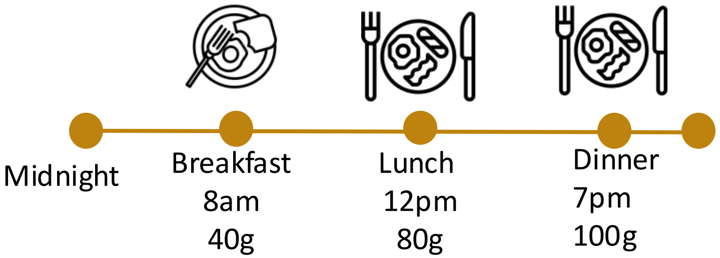

(A)

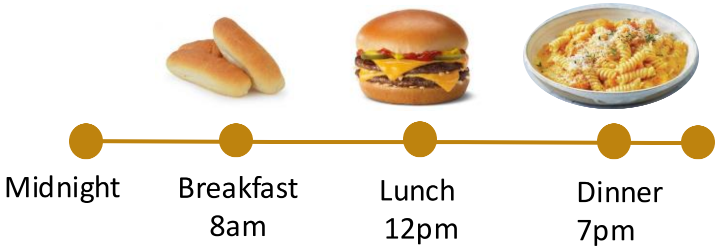

(B)

†In-between snacks can also be incorporated[1], currently not implemented on CAPSML.

[1] Hettiarachchi, Chirath, et al. "A Reinforcement Learning Based System for Blood Glucose Control without Carbohydrate Estimation in Type 1 Diabetes: In Silico Validation." 2022.

2.1 Quick Start

Step 3. Select Control Algorithm

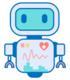

(A) Clinical Treatment (Basal - Bolus)	(B) Reinforcement Learning Strategies
Insulin Pump Therapy: Fixed basal & bolus setting • **Fixed basal insulin delivery** based on personalised Total Daily Insulin (TDI). • **Meal Bolus** based on the Carbohydrate to Insulin Ratio (CIR). • **Correction Bolus** based on the Insulin Sensitivity Factor (ISF).	A closed-loop system that monitors historical glucose trends to estimate insulin requirements.
Meal Announcement – 20 minutes in advance.	**No Meal Announcements.**
Manual Meal Carbohydrate (CHO) Estimation.	**No Meal Carbohydrate (CHO) Estimation.**
E.g., **BBI** – Ideal CHO estimates, **BBHE** – Human error in CHO estimates	E.g., **G2P2C** (novel algorithm), **PPO** (state-of-the-art RL algorithm adapted for glucose control)

[1] Hettiarachchi, C et al. "G2P2C–A modular Reinforcement Learning Algorithm for Glucose Control by Glucose Prediction and Planning in Type 1 Diabetes", *BSPC 2024.*

2.1 Quick Start

Step 4. Run Simulation

Adolescent0

Meals

Glucose

Insulin

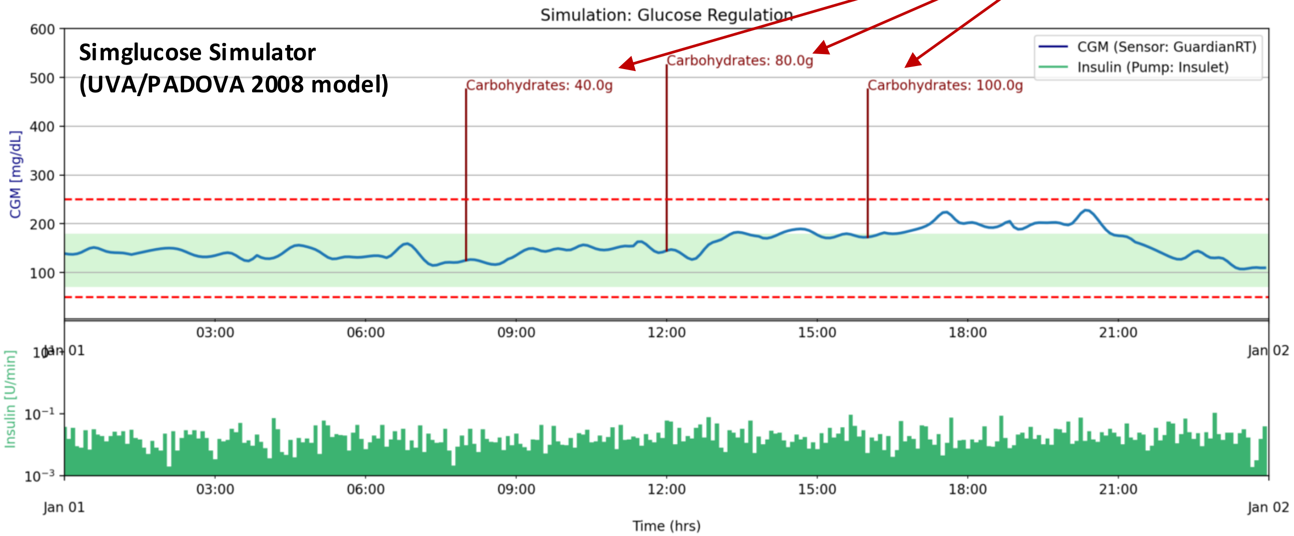

Control Algorithm (e.g., *our G2P2C algorithm*[1,2])

[1] Hettiarachchi, C et al. "G2P2C–A Modular Reinforcement Learning Algorithm for Glucose Control by Glucose Prediction and Planning in Type 1 Diabetes." *BSPC 2024.*
[2] Hettiarachchi, C, et al. "A Reinforcement Learning Based System for Blood Glucose Control without Carbohydrate Estimation in Type 1 Diabetes: In Silico Validation." *EMBC*. IEEE, 2022.

2.1 Quick Start

Step 4. Run Simulation

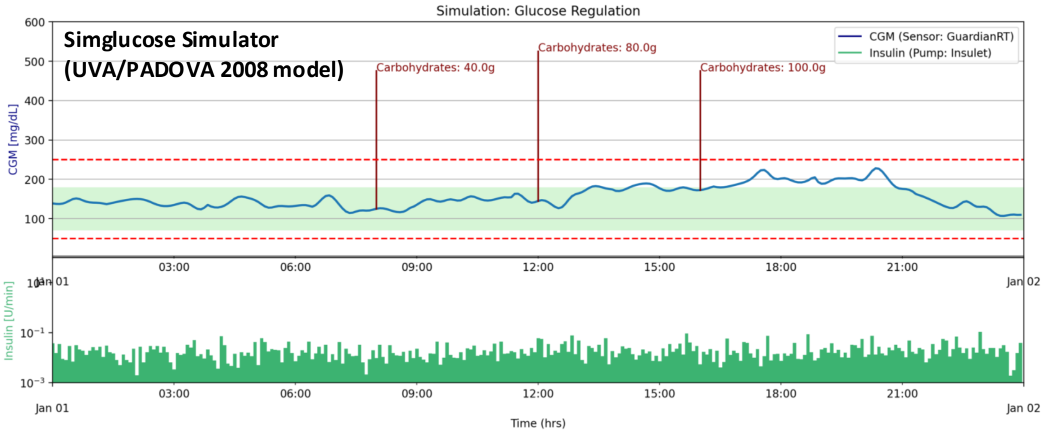

Table. Clinical performance metrics.

	Metric	Value
0	Time Above Range (TAR) - Level2 (>250 mg/dL)	0.0%
1	Time Above Range (TAR) - Level 1 (180 - 250 mg/dL)	14.93%
2	Time In Range (TIR) (70 - 180 mg/dL)	85.07%
3	Time Below Range (TIR) - Level 1 (54 - 70 mg/dL)	0.0%
4	Time Below Range (TIR) - Level 2 (<54 mg/dL)	0.0%
5	Risk Index (RI)	4.21
6	Low Blood Glucose Index (LBGI)	0.44
7	High Blood Glucose Index (HBGI)	3.77

2.2 Running Simulations…

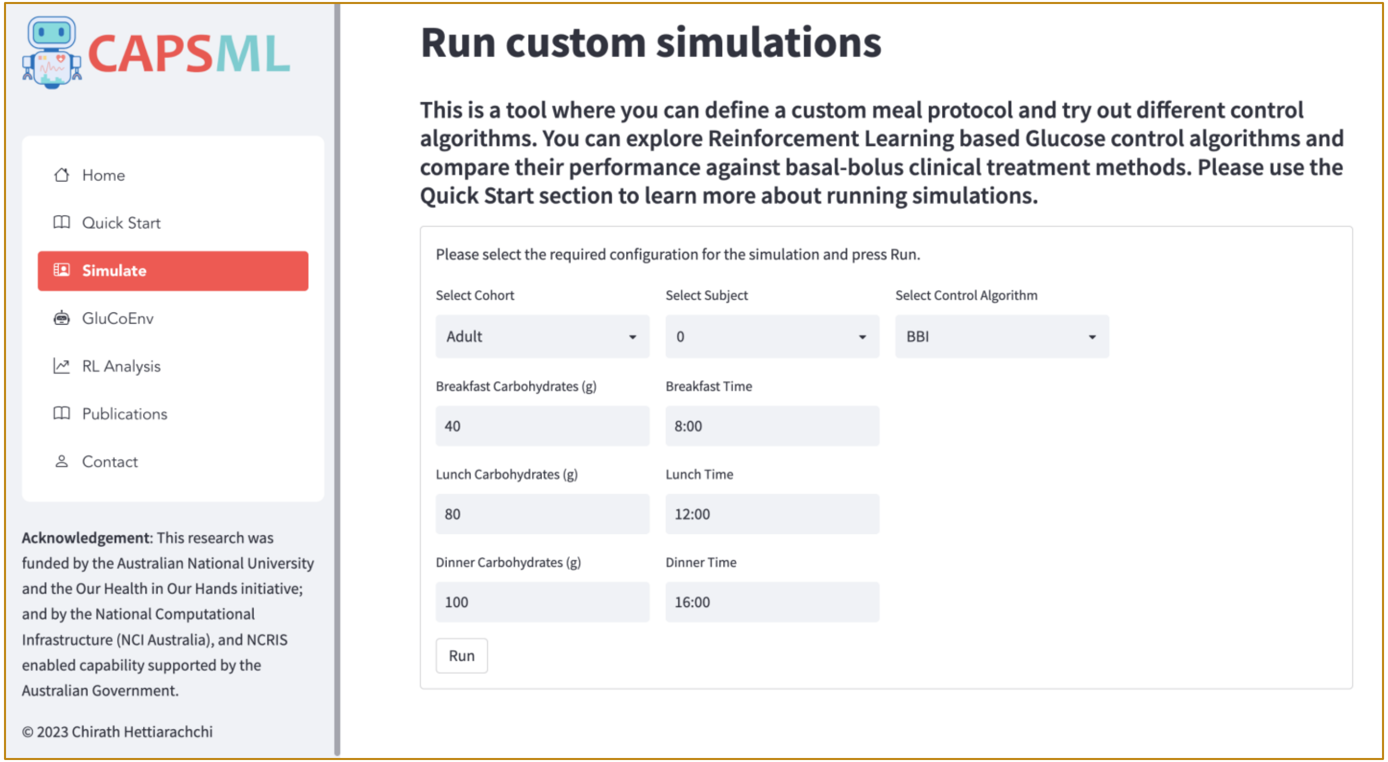

(move to live demo…)

03 Diabetes Research:
Developing RL-based Treatment Strategies

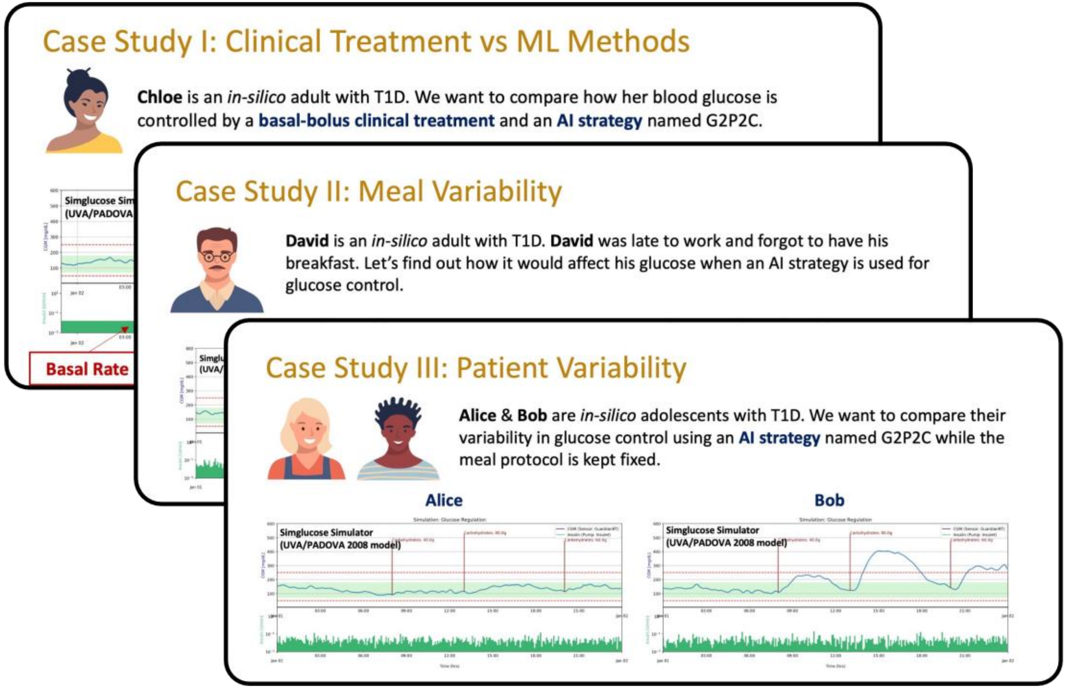

03 (A). Case Study I: Clinical Treatment vs RL Methods

Chloe is an *in-silico* adult with T1D. We want to compare how her blood glucose is controlled by a **basal-bolus clinical treatment** and an **RL strategy** named G2P2C.

Basal-bolus Treatment

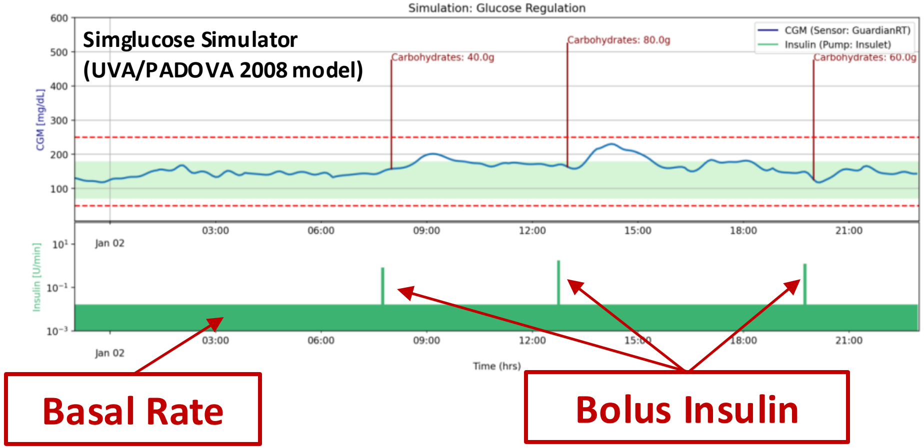

- Requires manual user input on meal announcements & carbohydrate estimation.
- Easy to understand.

RL Strategy

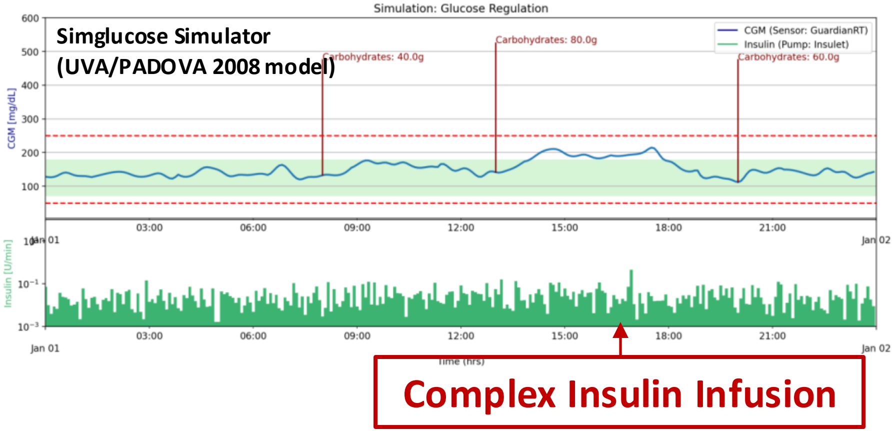

- The RL treatment strategy is more complex.
- Highlights the importance of collaboration & tools to improve explainability.

03 (B). Case Study II: Patient Variability

Alice & **Bob** are *in-silico* adolescents with T1D. We want to compare their variability in glucose control using an **RL strategy** named G2P2C while the meal protocol is kept fixed.

Alice

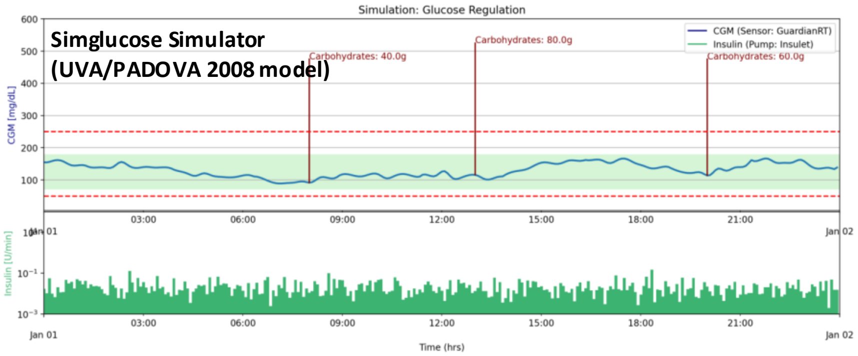

Bob

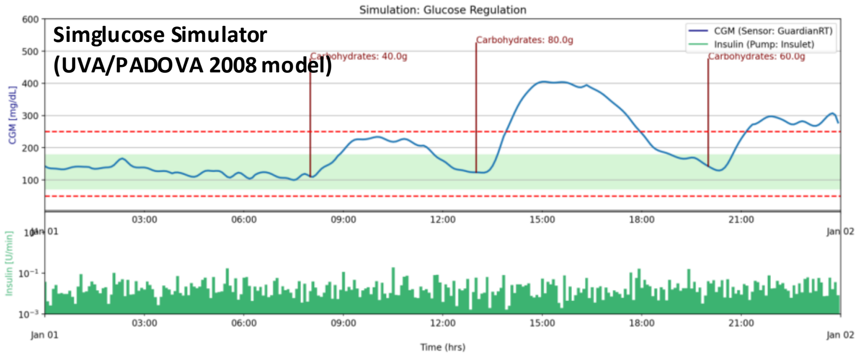

- Controlling blood glucose in Alice is easier compared to Bob.
- The RL strategies need to be carefully validated across all cohorts and subjects. CAPSML provides the capability to run simulations and visualise.

03 (C). Case Study III: Meal Variability

David is an *in-silico* adult with T1D. **David** was late to work and forgot to have his breakfast. Let's find out how it would affect his glucose when an RL strategy is used for glucose control.

Normal Day

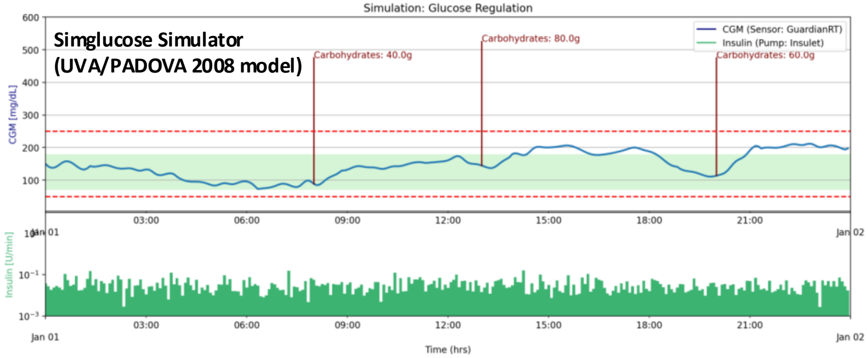

Forgot to have breakfast

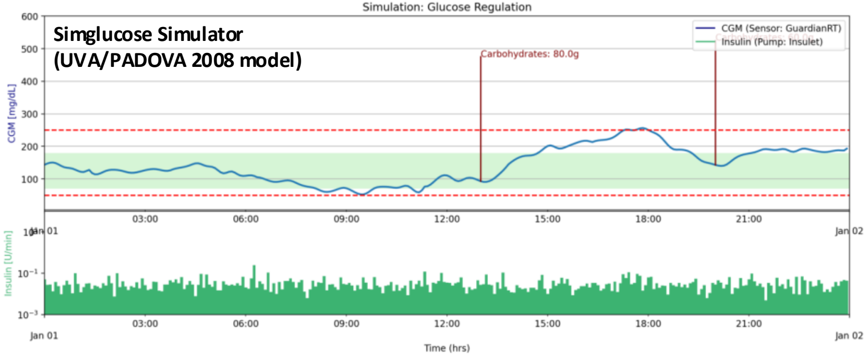

- Simulating real-life events to evaluate potential failures of RL strategies is important.

04 Diabetes Education

- Learn about **glucose regulation and closed-loop systems; carbohydrate content of different meals/food** and their effect on glucose regulation.

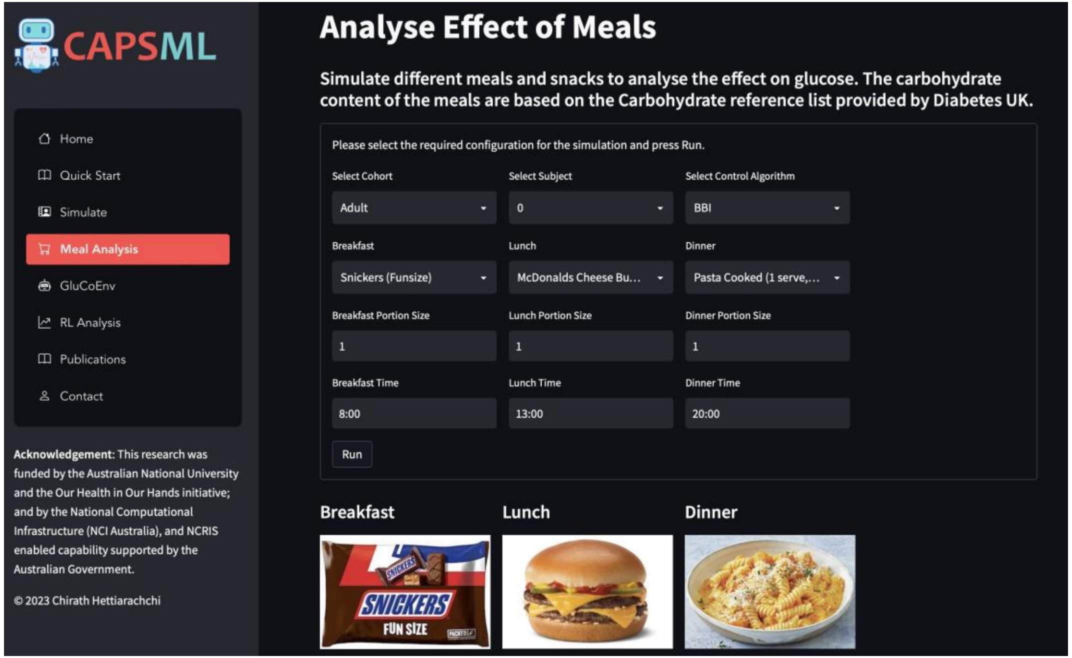

† The carbohydrate reference list by Diabetes UK is integrated in the tool (https://www.diabetes.org.uk/).

04 Diabetes Education

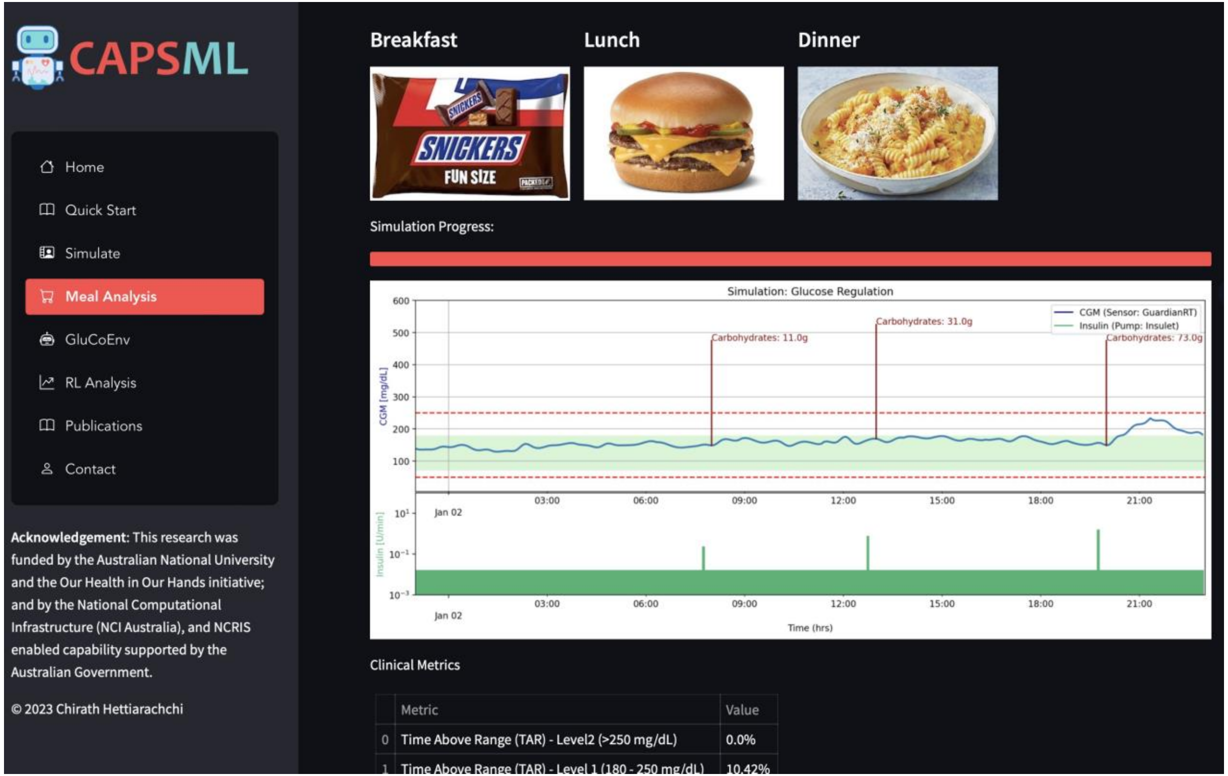

05 Clinical Analysis

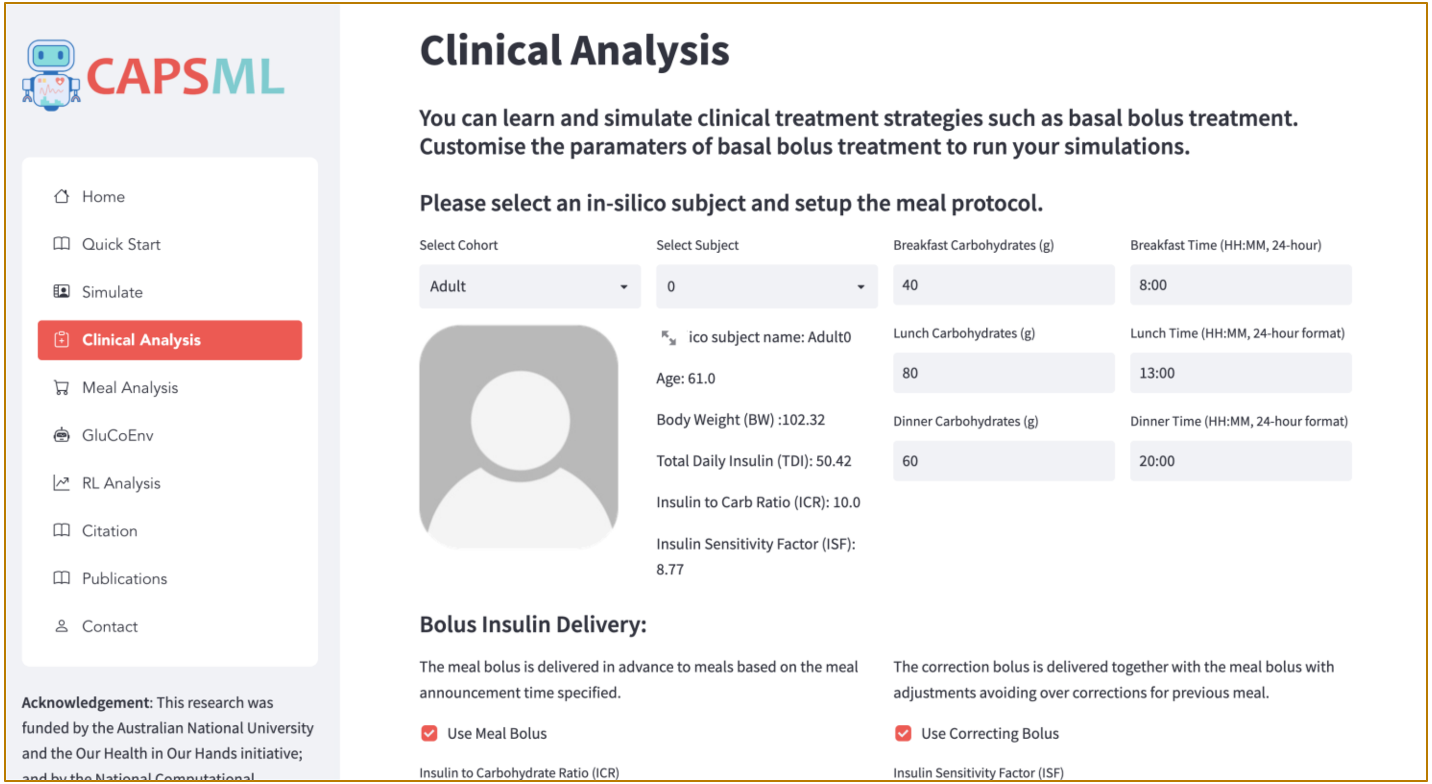

06. Impact & Potential of CAPSML

- CAPSML can be **used on a desktop machine or mobile device without any specialised hardware** to simulate AI systems in glucose control.
- Improve AI systems through **clinicians & health experience expert feedback**.
- Provide additional **customisation capabilities for clinical treatment** strategies.
- A first step towards improving **trust/explainability** of AI systems in glucose control.

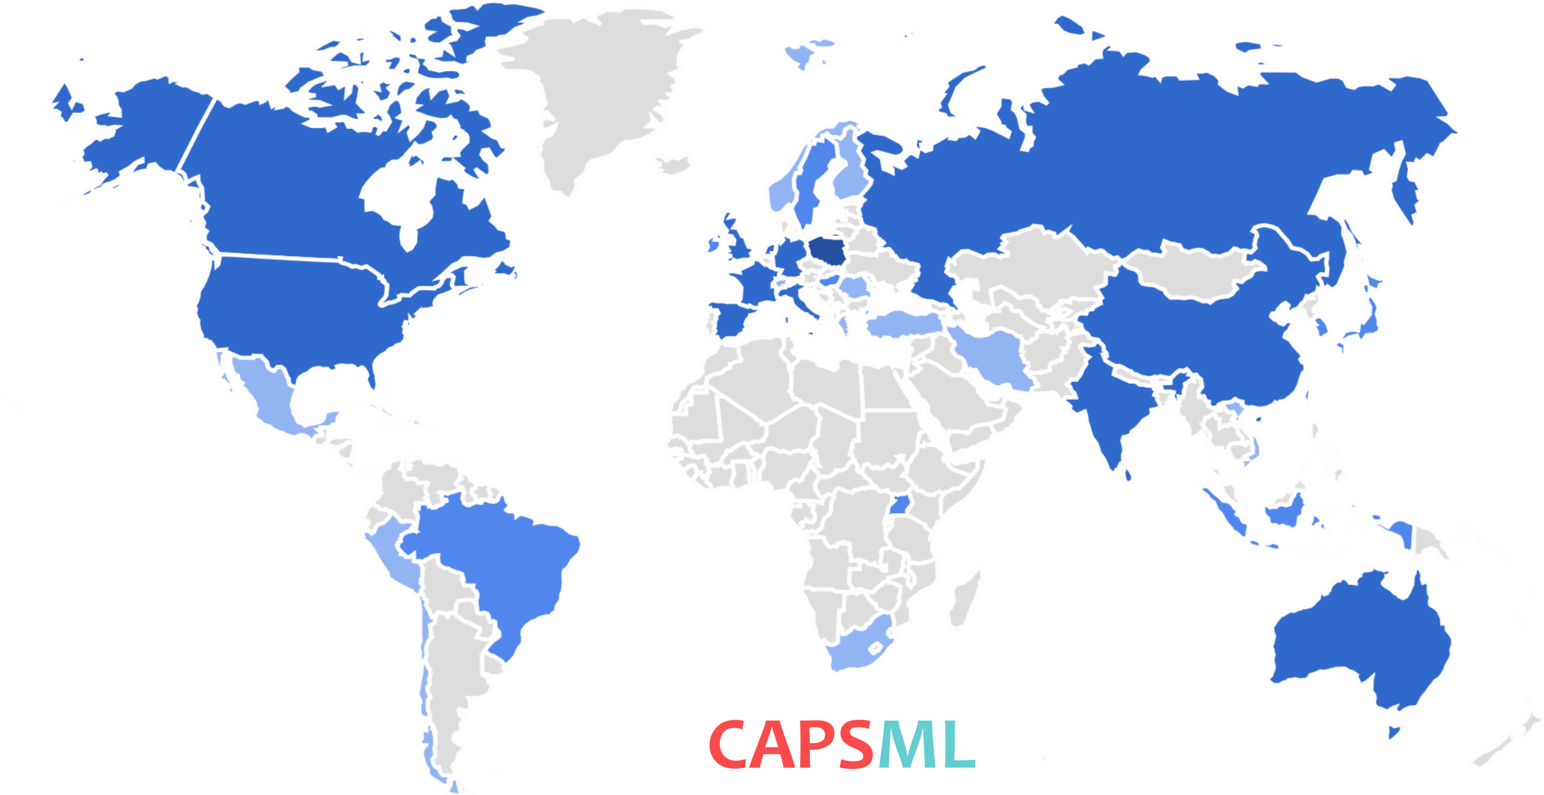

CAPSML

(July 2023 – April 2024)

README MIT license

GluCoEnv

GluCoEnv - Glucose Control Environment is a simulation environment which aims to facilitate the development of Reinforcement Learning based Artificial Pancreas Systems for Glucose Control in Type 1 Diabetes.

License MIT | python 3.8 | Web CAPSML | Follow @chirathyh

About

This project implements in-silico Type 1 Diabetes (T1D) subjects for developing glucose control algorithms. The glucose control environment includes 30 subjects (10 children, adolescents, and adults each), which extends the work of Simglucose and UVA/Padova 2008 simulators by following an end-to-end GPU-based implmentation using the PyTorch framework. The project aim is to facilitate the development of Reinforcement Learning (RL) based control algorithms by providing a high-performance environment for experimentation.

Research related to RL-based glucose control systems are relatively minimal compared to popular RL tasks (games, physics simulations etc). The task of glucose control requires ground up development where prolem formulations, state-action space representations, reward function formulations are not well established. Hence, researchers have to run significant amount of experiments to design, develop and tune hyperparameters. This groundup development requires significant compute and effort.

The key highlights of **GluCoEnv** are the vectorized parallel environments designed to run on a GPU and the flexibility to develop RL-based algorithms for glucose control and benchmarking. The problem also lacks proper benchmarking scenarios and controllers, which have been implemented in this environment to provide some guidance on the task.

You can find more details and our RL-based glucose control algorithms by visiting the project CAPSML.

This project is under active development, where additional glucose dynamics models, clinical metrics, RL algorithms, and visualisation tools will be introduced.

Installation & Dependencies

Create a Python 3.8.0 virtual environemnt and install PyTorch 1.12.0 with CUDA 11.3.

```
python3 -m venv env
source env/bin/activate
pip install --upgrade pip
```

https://github.com/RL4H/GluCoEnv

README MIT license

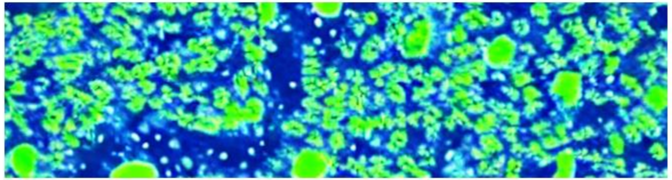

Reinforcement Learning based Artificial Pancreas Systems.

Type 1 Diabetes (T1D) is casued by the autoimmune destruction of the islet beta-cells and results in absolute insulin deficiency (cover image: Human islet of Langerhans created by Stable Diffusion). Hence, external administration of insulin is required to maintain glucose levels, which is cruicial as both low and high glucose levels are detrimental to health. This is usually done through an insulin pump attached to the body. An continuous glucose sensor is also attached to measure the glucose levels so that a control algorithm can estimate the appropriate insulin dose. In this project we design Reinforcement Learning (RL) based Artificial Pancreas Systems (APS) for the glucose control problem. The figure below shows the main components of an APS.

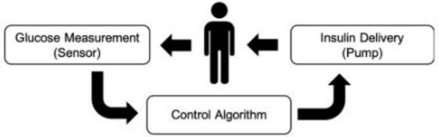

Maintaining glucose levels is a life-long optimisation problem, complicated due to the disturbances associated with daily events (meals, exercise, stress.. etc), delays present in glucose sensing and insulin action, partial observability, and safety constraints among others. A simulation of glucose regulation, using a RL-based strategy is shown below, where the optimal glucose range is shaded in green severe hypoglycemia / hyperglycemia ranges highlighted by the red dotted line. The blood glucose measurements are presented in the top, while the administered insulin by the RL agent is presented in the bottom. The disturbances related to meal events and the carbohydrate content of the meals are presented in red.

https://github.com/RL4H/G2P2C

07. Conclusion & Future Work

- CAPSML welcomes clinical and lived experience experts to experiment with an AI system and run custom simulations to learn about the AI system.

- In future, CAPSML ("**C**ontrolling **A**rtificial **P**ancreas **S**ystems using **M**achine **L**earning") can be extended to allow researchers to test their AI systems for comparison & evaluation.

- It could address the currently limited benchmarking of different AI algorithms and systems in T1D if successfully used at scale.

THANK YOU

Acknowledgement: This research is funded by the Australian National University, School of Computing, Our Health in Our Hands (OHIOH), a strategic initiative of the ANU, the National Computation Infrastructure (NCI Australia) under the ANU Merit Allocation Scheme.

Collaborators: We wish to thank **Dr David O'Neal, Dr Barbora Paldus, and Dr Dale Morrison** from the Diabetes Technology Research Group, St Vincent's Hospital for their valuable insights.

Contact Us





Dr Chirath Hettiarachchi | **Dr Nicolo Malagutti** | **Prof. Christopher Nolan** | **Dr Elena Daskalaki** | **Prof. Hanna Suominen**

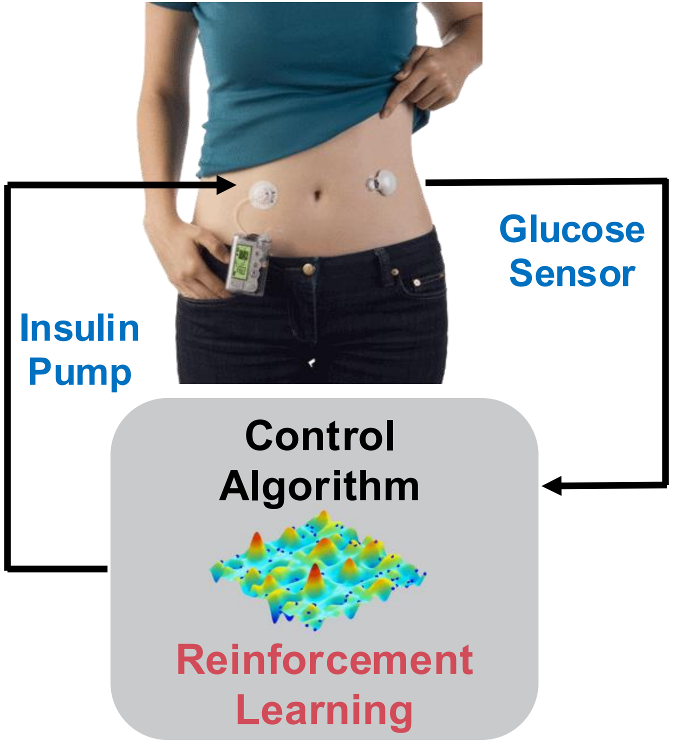
Glucose
Sensor
Insulin
Pump
Control
Algorithm
Reinforcement
Learning

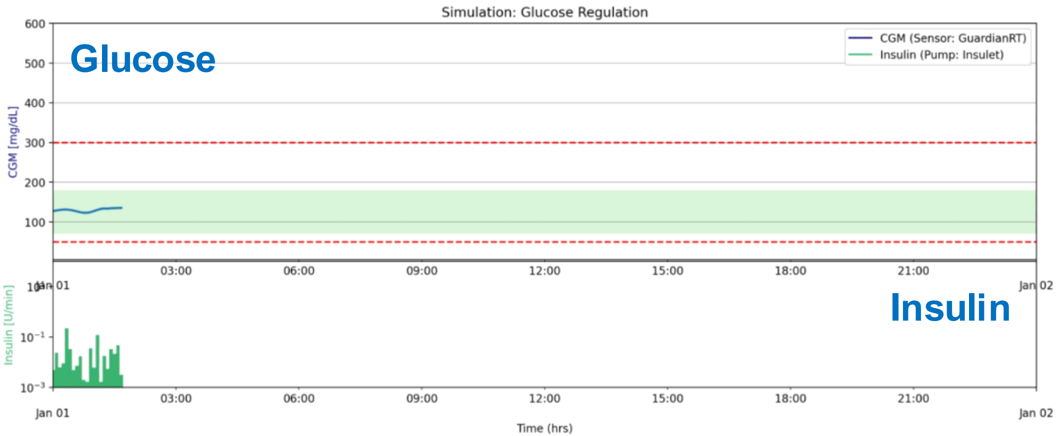
Simulation: Glucose Regulation
Glucose
CGM (Sensor: GuardianRT)
Insulin (Pump: Insulet)
CGM [mg/dL]
600
500
400
300
200
100
Insulin
Insulin [U/min]
03:00
06:00
09:00
12:00
15:00
18:00
21:00
Jan 01
Jan 02
Time (hrs)