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

Chirath Hettiarachchi

Research Fellow,
School of Computing, College of Systems & Society
Australian National University

March 24, 2025

Acknowledgement of Country

Overview

Content

Reinforcement Learning (RL)

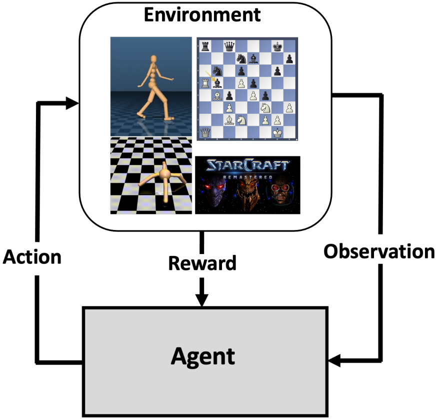

The "Law of Effect" in psychology introduced learning by trial and error, which described the effect of reinforcing events based on the tendency to select actions [Thorndike, 1911].

The Glucoregulatory System & Type 1 Diabetes

Sleep

Meals

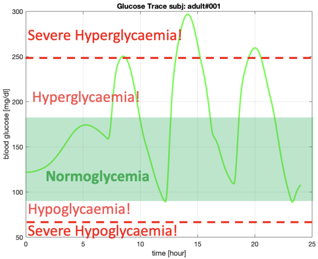

Exercise

Stress

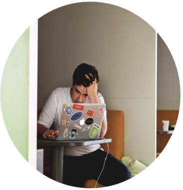

Maintaining glucose homeostasis is vital!

Clinical objective is to improve Time in Normoglycemic Range (TIR),
while avoiding hypoglycemic & hyperglycemic risk.

Artificial Pancreas Systems (APS)

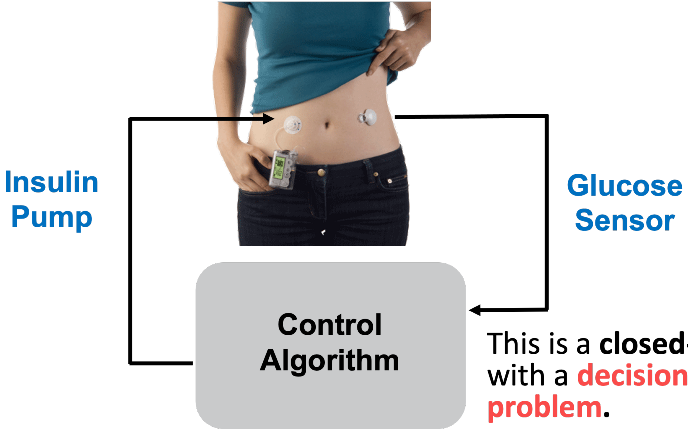

This is a **closed-loop system**, with a **decision-making problem**.

APS are **high-risk medical devices**. Existing commercial APS are **hybrid systems** (**manual decision-making required**), and designed using classical control algorithms (PID & MPC).

Content

Improve Performance & Eliminate Cognitive Burden

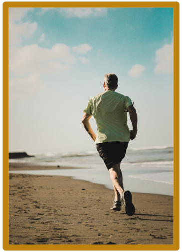

How much **Insulin** do I need?

Brew-Sam, et al. "Experiences of young people and their caregivers of using technology to manage type 1 diabetes mellitus: Systematic literature review and narrative synthesis", JMIR Diabetes, 2021.

Hettiarachchi, et al. "Integrating multiple inputs into an artificial pancreas system: Narrative literature review", JMIR Diabetes, 2022.

Brew-Sam, et al. "Toward Diabetes Device Development That Is Mindful to the Needs of Young People Living With Type 1 Diabetes: A Data-and Theory-Driven Qualitative Study", JMIR Diabetes, 2023.

Existing Clinical Treatment (Basal-Bolus)

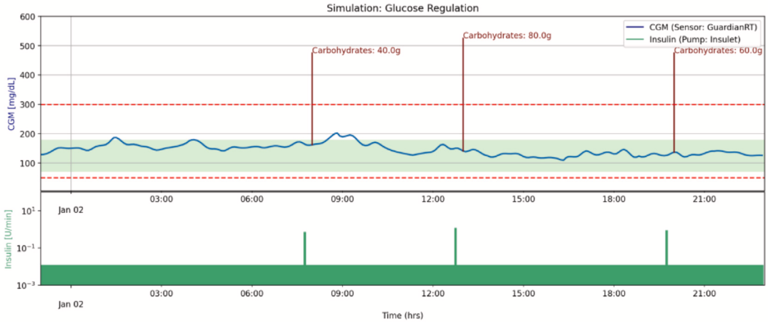

Figure: An ideal basal-bolus insulin treatment strategy without human error.

Requires manual user input on meal announcements & carbohydrate estimation, which leads to errors & sub-optimal glucose regulation.

Content

Method: Problem Formulation & Experimental Setup

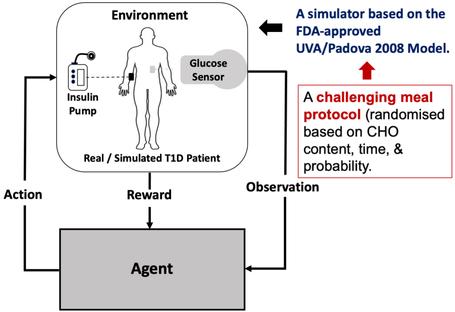

A Partially Observable Markov Decision Process (POMDP). Use the open-source Simglucose simulator, based on the FDA-accepted UVA/PADOVA 2008 model for pre-clinical trials.

Hettiarachchi, et al. "A Reinforcement Learning Based System for Blood Glucose Control without Carbohydrate Estimation In Type 1 Diabetes: In Silico Validation", EMBC2022.

Hettiarachchi, et al. "Non-linear continuous action spaces for reinforcement learning in type 1 diabetes", AJCAI2022.

PPO (2022)

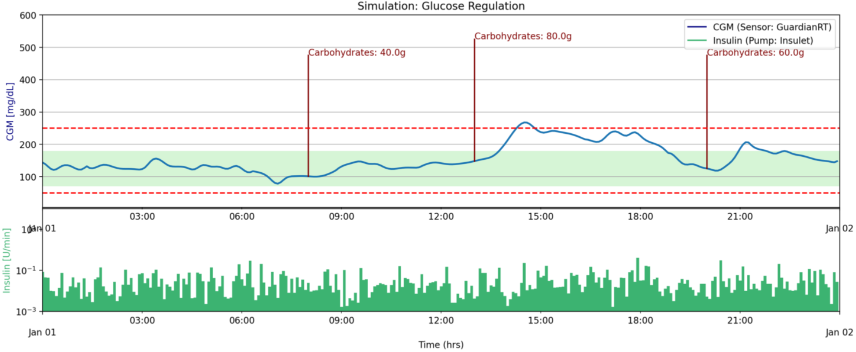

No meal announcements.
No meal carbohydrate (CHO) estimation.

Hettiarachchi, et al. "A Reinforcement Learning Based System for Blood Glucose Control without Carbohydrate Estimation In Type 1 Diabetes: In Silico Validation", In 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) 2022 Jul 11 (pp. 950-956). IEEE.

Content

CAPSML: Co-creation & synergies when people and AI work together

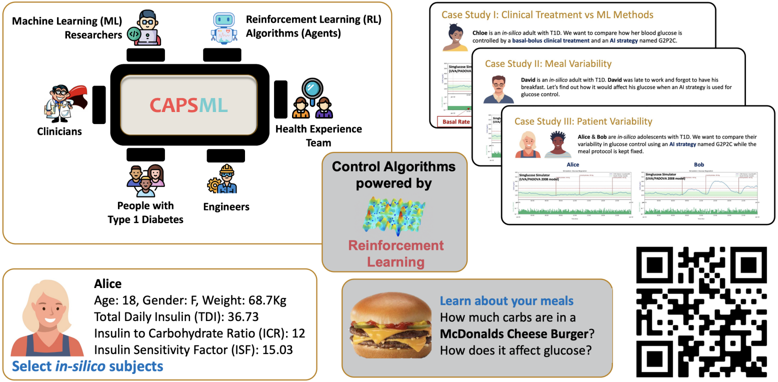

Desborough, et al. "A Framework for Involving Coproduction Partners in Research About Young People with Type 1 Diabetes", Health Expectations, 2022.

Hettiarachchi, et al. "CAPSML: Bridging the Gap Between Clinicians, Lived Experience Experts, and Artificial Intelligence Systems for Glucose Regulation in Type 1 Diabetes", 34th Medical Informatics Europe (MIE), 2023.

CAPSML Interface *(www.capsml.com)*

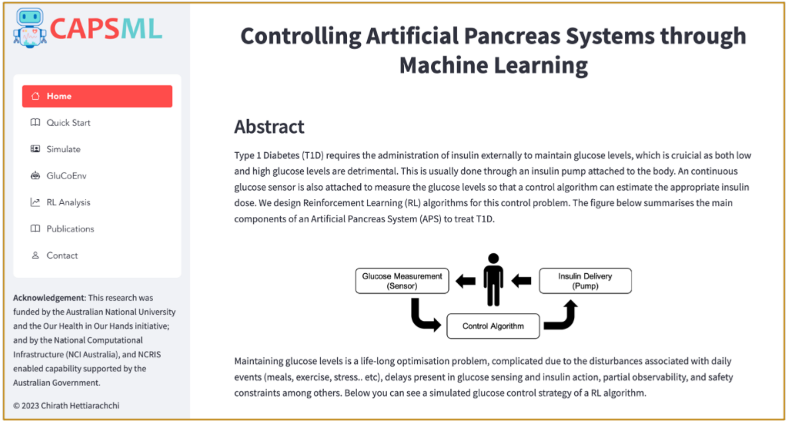

Publicly available online and released under the MIT license.

System demonstration: `https://youtu.be/J05MkPCuqCw`.

Running a simulation...

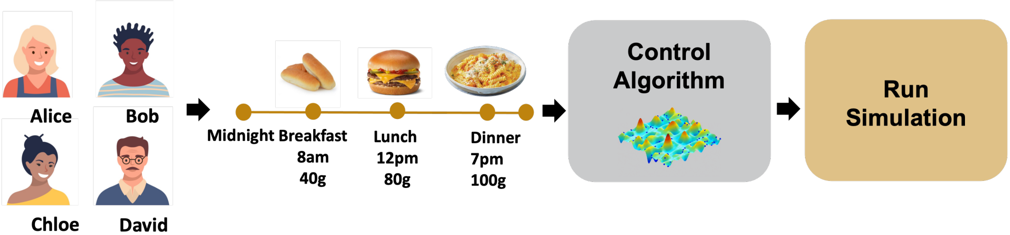

Step 1	**Step 2**	**Step 3**	**Step 4**
Select *in-silico* subject	**Setup meal protocol**	**Select control algorithm**	**Run simulation**

Four simple steps, on a desktop/mobile device without any specialized hardware.
Previous simulation tools require technical knowledge & inaccessible to the public.

Sample simulation output

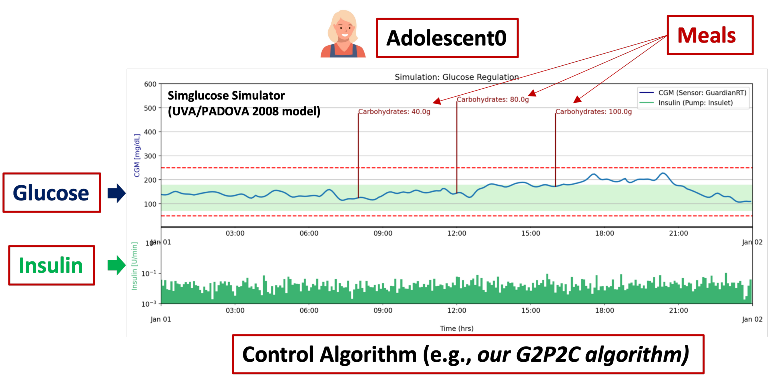

Table. Clinical performance metrics.

	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

Simple visualizations & specialized clinical metrics / features for power users.

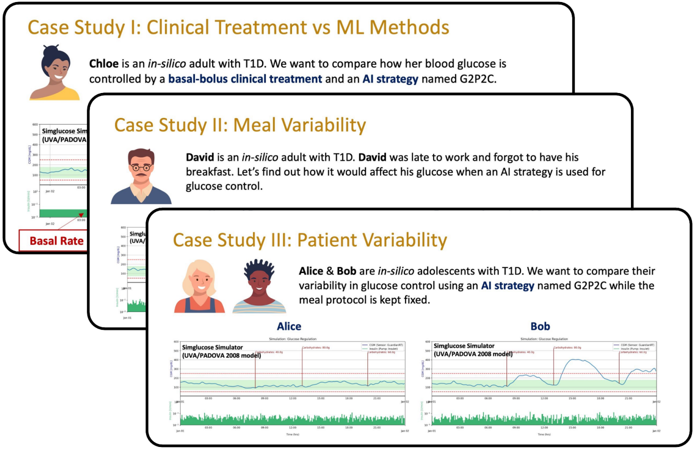
Case Study I: Clinical Treatment vs ML Methods
Chloe is an *in-silico* adult with T1D. We want to compare how her blood glucose is controlled by a **basal-bolus clinical treatment** and an **AI strategy** named G2P2C.
Basal Rate
Case Study II: Meal Variability
David is an *in-silico* adult with T1D. **David** was late to work and forgot to have his breakfast. Let's find out how it would affect his glucose when an AI strategy is used for glucose control.
Case Study III: Patient Variability
Alice & **Bob** are *in-silico* adolescents with T1D. We want to compare their variability in glucose control using an **AI strategy** named G2P2C while the meal protocol is kept fixed.
Alice
Bob
Simglucose Simulator
(UVA/PADOVA 2008 model)
Simglucose Simulator
(UVA/PADOVA 2008 model)

Sample Case Study I: Clinical Treatment vs RL Methods

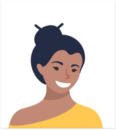

Chloe is an *in-silico* adult with T1D. We want to compare how her blood glucose is controlled by a **basal-bolus clinical treatment** and an **RL strategy** named G2P2C.

Basal-bolus Treatment

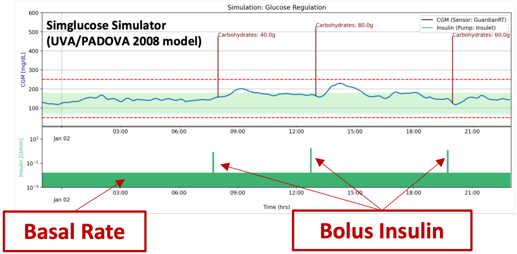

- Requires manual user input on meal announcements & carbohydrate estimation.
- Easy to understand.

RL Strategy

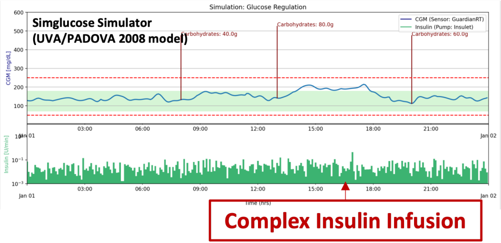

- The RL treatment strategy is more complex.
- Highlights the importance of collaboration & tools to improve explainability.

Sample Case Study II: Patient Variability

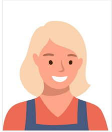

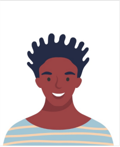

Alice & **Bob** are *in-silico* adolescents with T1D. We want to compare their variability in glucose control using an **RL strategy** named G2P2C while the meal protocol is kept fixed.

Alice

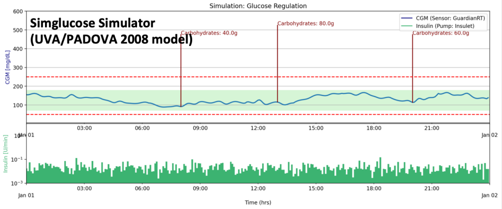

Bob

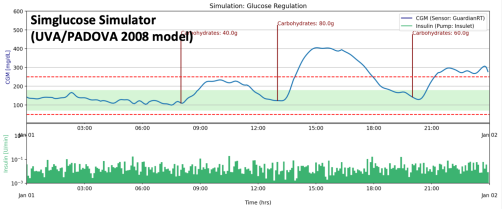

- Controlling blood glucose in Alice is easier compared to Bob.
- The RL strategies need to be carefully validated across all cohorts and subjects. CAPSML provides the capability to run simulations and visualise.

Sample 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

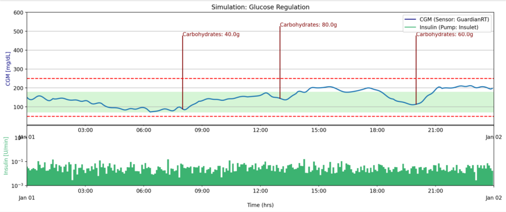

Forgot to have breakfast

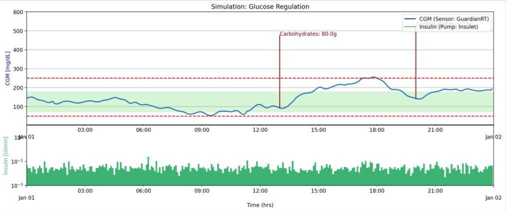

- Simulating real-life events to evaluate potential failures of RL strategies is important.

Feature Requests: Clinicians

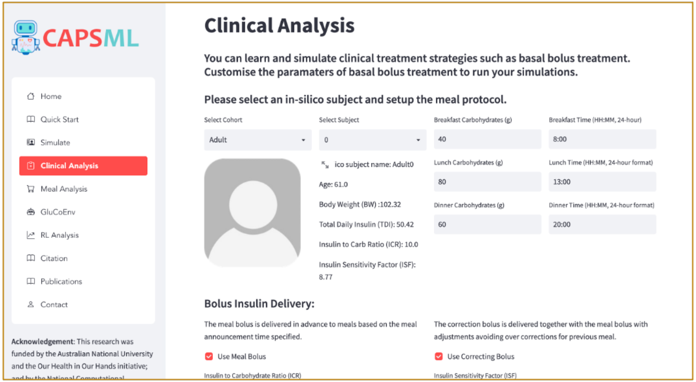

A user friendly, accessible tool for analysing existing clinical treatment strategies.

Current tools require programming knowledge (e.g., MATLAB, Python).

Feature Requests: People with T1D

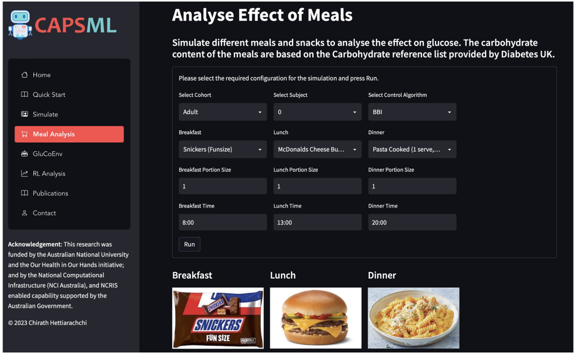

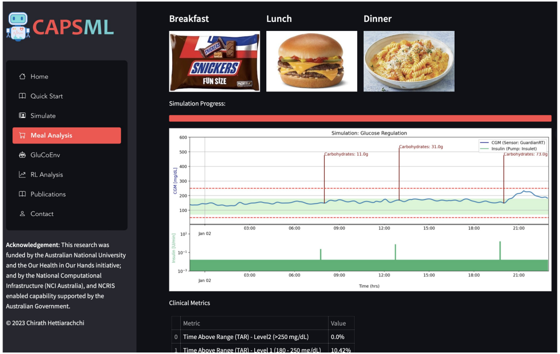

Learn about glucose regulation & closed-loop AI systems; carbohydrate contents of different meals/food and their effect on glucose regulation.

CAPSML Usage & User Demographics

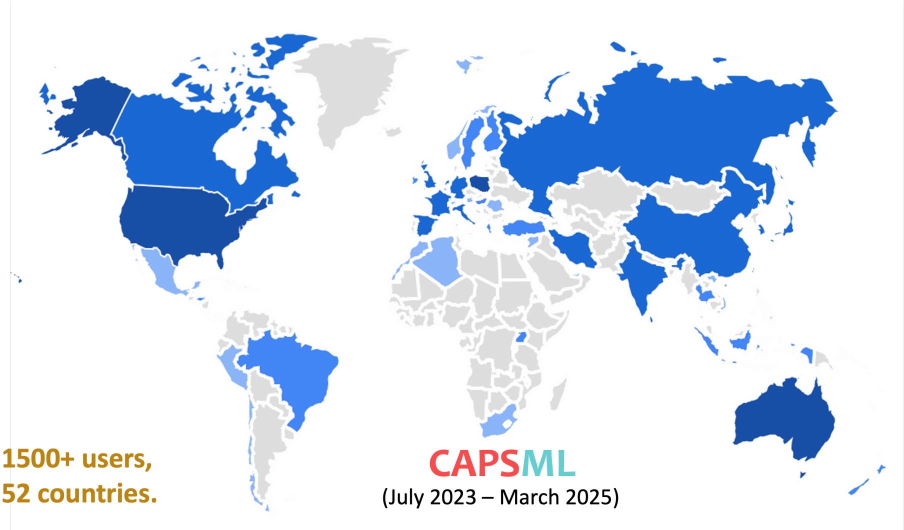

Content

G2P2C (2023)

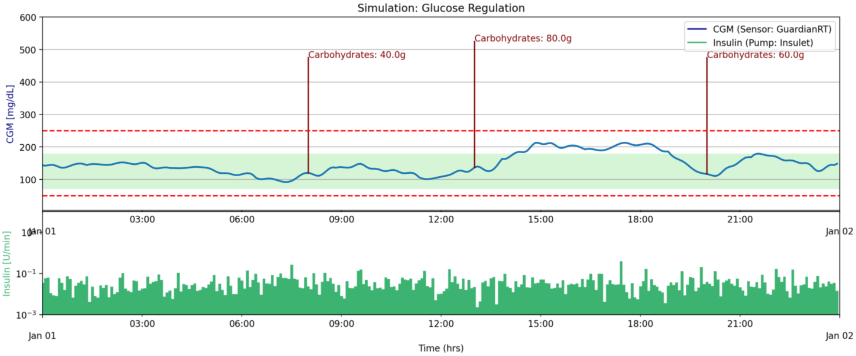

Auxiliary model learning & planning.
Improved safety.

Hettiarachchi, et al. "G2P2C—A modular reinforcement learning algorithm for glucose control by glucose prediction & planning in T1D", Biomedical Signal Processing & Control, 90, p.105839, 2023.

Hettiarachchi, Chirath. "Reinforcement Learning-based Artificial Pancreas Systems to Automate Treatment in Type 1 Diabetes", PhD Thesis, The Australian National University, 2023.

TD3 (2024)

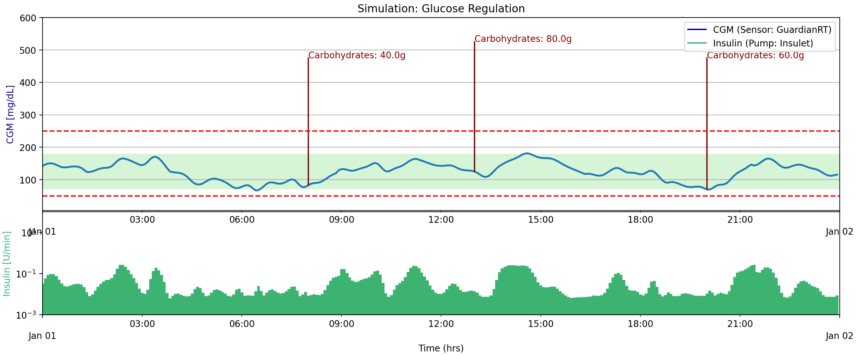

Explored deterministic policies.
Focus on explainability.

Timms, David, et al. "Comparing Deterministic and Stochastic Reinforcement Learning for Glucose Regulation in Type 1 Diabetes", 20th World Congress on Medical and Health Informatics: MedInfo 2025 (Accepted).

Timms, David, et al. "Deterministic Policy Gradient Algorithms for Regulating Glucose in Type 1 Diabetes", (in preparation).

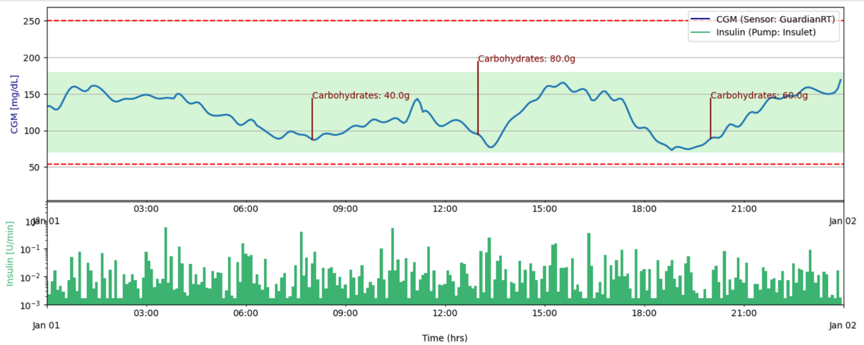

Explored constrained policies.
Improved safety on hard subjects.

External collaboration, University of Moratuwa, Sri Lanka.
Rakshitha, Dilshan; Lorensuhewa, Prasanjith; Anupama, Damika, et al. "Improving Safety in Reinforcement Learning based Artificial Pancreas Systems", (in-preperation).

Other Current Projects

Projects at ANU

1. **Inverse Reinforcement Learning for Glucose Regulation**, *(Honours Project: Jordan Triming, 2024/25).*
2. **Offline Reinforcement Learning for Automated Type 1 Diabetes Systems – Incorporating Clinical and Simulated Data**, *(Honours Project: Samuel Price, 2025).*
3. **A simulation dataset for machine learning applications in type 1 diabetes.** Released Soon...
4. **CAPSML: Co-creating a Deep Analytical Data Science Platform for Therapeutics Research and Education in Type 1 Diabetes.** Open-opportunity HCI Applied Social Science
5. **GluCoEnv: A GPU-based T1D Simulator.** Open-opportunity ML Engineering

Community Projects

1. Explainable RL methods: Integrated gradients, SHAP force plots, and explainability in policy space *(Malay Phadke, Preyes Parab)*.
2. Integrating our current RL algorithm implementations into a mobile phone app.

Open-source Resources: Codebase (GluCoEnv, G2P2C, RL4T1D)

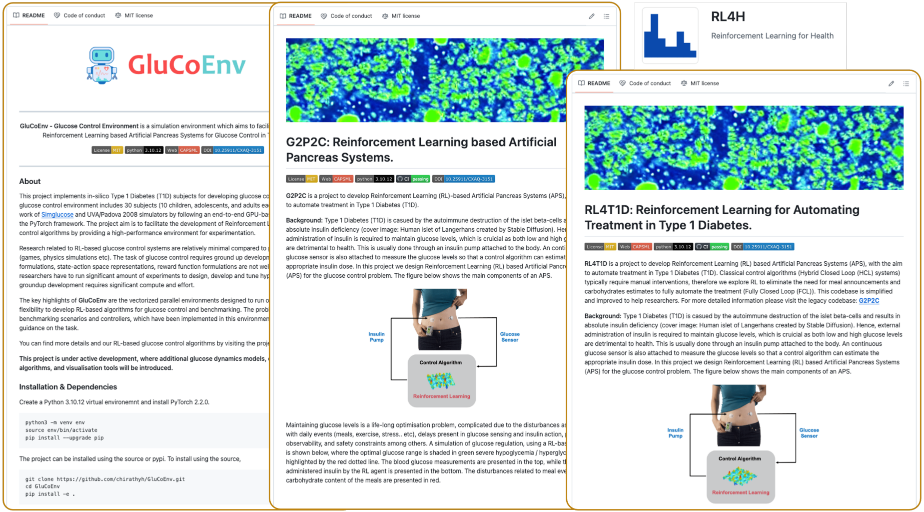

Content

Conclusion

- **Meaningful contributions** emerge through **multi-disciplinary collaborations/research**.
- CAPSML is easily accessible and encourages clinical and lived experience experts to explore AI systems and run custom simulations for learning **(diabetes & AI education)**.
- CAPSML improves research through clinicians & health experience **expert feedback**.
- A first step towards improving **trust/explainability of AI systems** in glucose control.
- CAPSML can be extended to allow researchers to test/compare their systems addressing current limitations in **reproducibility and benchmarking**.
- **The final solution takes time.** Stay patient, be respectful and keep going...

Thank You

Acknowledgement

- This research was funded by the Australian National University (ANU), School of Computing; Our Health in Our Hands (OHIOH) strategic initiative; National Computation Infrastructure (NCI Australia) under the ANU Merit Allocation / ANU Startup Schemes.

Team

- **Collaborators**: OHIOH Team; Dr David O'Neal, Dr Barbora Paldus, & Dr Dale Morrison from the Diabetes Technology Research Group, St Vincent's Hospital, Melbourne; Dr Charith Chitraranjan, University of Moratuwa (UoM).
- **Research Students**: David Timms (Master's Research Project, ANU), Jordan Triming (Honours Project, ANU), Samuel Price (Honours Project, ANU), Dilshan Rakshitha; Prasanjith Lorensuhewa, Damika Anupama (Honours Project, UoM).
- **Supervisors**: Sam Cantrill, Dr Chirath Hettiarachchi, Dr Nicolo Malagutti, Prof. Christoper Nolan, Dr Elena Daskalaki, Prof. Hanna Suominen.