Poster: GlucOS: A secure, safe and extensible system for automated insulin delivery

Hari Venugopalan
hvenugopalan@ucdavis.edu
UC Davis

Shyreas Madhav Ambattur Vijayanand
smvijay@ucdavis.edu
UC Davis

Samuel T. King
kingst@ucdavis.edu
UC Davis

***Abstract*—Type 1 Diabetes (T1D) is a metabolic disorder where an individual's pancreas stops producing insulin. To compensate, they inject synthetic insulin. The dynamics of insulin and the impact of various intrinsic and extraneous factors on glucose complicate the management of T1D [19]. While research has shown that modern machine learning (ML) based algorithms [19], [8], [18], [16] are well suited for managing T1D, existing automated insulin delivery systems [10], [13], [4], [2], [1] do not support ML.**

Incorrect predictions from ML, either from a malicious model [6], [9] or from the blind spots of a benign model, can kill people when used to control an insulin pump. Automated insulin delivery systems also have to contend with other practical challenges in the real world such as skin infections [12] and pump failures [7]. These risks along with the lack of explainability of complex ML discourages people from adopting ML in practice to manage T1D [3].

Thus, current system builders and users opt for explainable and deterministic algorithms that offer modest control in managing T1D over the most accurate predictions from state-of-the-art ML. However, people living with T1D miss out on a viable opportunity to improve their long-term health and reduce their T1D management burden by not using using ML.

In this paper, we take on the challenge of supporting powerful predictive ML algorithms for managing T1D. We build the first automated insulin delivery system that can adopt any algorithm to control an insulin pump securely and safely. Our contribution is in our novel system, called GlucOS, that we design and implement from scratch, and in our novel security mechanisms that handle both security and safety when automatically dosing insulin. GlucOSis not tied to any specific algorithm and can safely support any algorithm, including ML-based [20], [15], physiological-based [5], [17], control theoretic [11], [14], and heuristic-based [13], [10].

Our key insight for security is that over a long enough period of time, all correct algorithms will dose the same amount of insulin. For example, on a given day, a person will need to dose a fixed amount of insulin to absorb all of the glucose from the food they eat. Since synthetic insulin action is slow, predicting future metabolic states to inject insulin early is critical for long-term health. GlucOS's support for modern ML-based algorithms provides users with this predictive power to dose insulin. For security, GlucOSpairs a predictive ML model with a conservative and safe model to provide the ML with enough flexibility to control the insulin pump proactively while staying within the bounds dictated by the safe model.

Rather than throw away the boring old traditional algorithms that current systems use and replacing them with fancy ML, we repurpose these simple algorithms to serve as the foundation for our security logic. By grounding our security logic in easy to understand and effective algorithms, we inherit the explainability and determinism that come with them.

We report our experiences running GlucOSon one individual for 2.5 months to manage their T1D. Running our software on a real human forces us to design a practical and real system. Our evaluation using virtual humans in a simulator show that our security and safety mechanisms generalize beyond the specific individual.

Our novel contributions include:

- **The clean-slate design and implementation of a flexible automated insulin delivery system that supports any predictive algorithm to manage T1D safely.**
- **Security mechanisms grounded by safe physiological models that protect individuals from egregious errors and malicious predictions.**
- **A case study describing our experiences from a real-world deployment, and results from simulation to show that our techniques generalize.**

References

[1] Insulet omnipod 5. https://www.omnipod.com/.
[2] Tandem control iq. https://www.tandemdiabetes.com/products/automated-insulin-delivery/control-iq.
[3] Openapi: Frequently asked questions, 2024. https://openaps.org/frequently-asked-questions/.
[4] Luz E. Castellanos, Courtney A. Balliro, Jordan S. Sherwood, Rabab Jafri, Mallory A. Hillard, Evelyn Greaux, Rajendranath Selagamsetty, Hui Zheng, Firas H. El-Khatib, Edward R. Damiano, and Steven J. Russell. Performance of the Insulin-Only iLet Bionic Pancreas and the Bihormonal iLet Using Dasiglucagon in Adults With Type 1 Diabetes in a Home-Use Setting. *Diabetes Care*, 44(6):e118–e120, 06 2021.
[5] Claudio Cobelli, Eric Renard, and Boris Kovatchev. Artificial Pancreas: Past, Present, Future. *Diabetes*, 60(11):2672–2682, 10 2011.
[6] Mohammed Elnawawy, Mohammadreza Hallajiyan, Gargi Mitra, Shahrear Iqbal, and Karthik Pattabiraman. Systematically assessing the security risks of ai/ml-enabled connected healthcare systems. *arXiv preprint arXiv:2401.17136*, 2024.
[7] Lutz Heinemann, G Alexander Fleming, John R Petrie, Reinhard W Holl, Richard M Bergenstal, and Anne L Peters. Insulin pump risks and benefits: a clinical appraisal of pump safety standards, adverse event reporting, and research needs: a joint statement of the european association for the study of diabetes and the american diabetes association diabetes technology working group. *Diabetes care*, 38(4):716–722, 2015.
[8] Peter G. Jacobs, Pau Herrero, Andrea Facchinetti, Josep Vehi, Boris Kovatchev, Marc D. Breton, Ali Cinar, Konstantina S. Nikita, Francis J. Doyle, Jorge Bondia, Tadej Battelino, Jessica R. Castle, Konstantia Zarkogianni, Rahul Narayan, and Clara Mosquera-Lopez. Artificial intelligence and machine learning for improving glycemic control in diabetes: Best practices, pitfalls, and opportunities. *IEEE Reviews in Biomedical Engineering*, 17:19–41, 2024.
[9] Tamar Levy-Loboda, Eitam Sheetrit, Idit F. Liberty, Alon Haim, and Nir Nissim. Personalized insulin dose manipulation attack and its detection using interval-based temporal patterns and machine learning algorithms. *Journal of Biomedical Informatics*, 132:104129, 2022.
[10] Loop. An automated insulin delivery app for ios, built on loopkit. https://github.com/LoopKit/Loop.

[11] Katrin Lunze, Tarunraj Singh, Marian Walter, Mathias D Brendel, and Steffen Leonhardt. Blood glucose control algorithms for type 1 diabetic patients: A methodological review. *Biomedical signal processing and control*, 8(2):107–119, 2013.

[12] Robert S Mecklenburg and Terin S Guinn. Complications of insulin pump therapy: the effect of insulin preparation. *Diabetes Care*, 8(4):367–370, 1985.

[13] OpenAPS. The open artificial pancreas system project. https://github.com/openaps.

[14] Griselda Quiroz. The evolution of control algorithms in artificial pancreas: A historical perspective. *Annual Reviews in Control*, 48:222–232, 2019.

[15] N. S. Tyler, C. M. Mosquera-Lopez, L. M. Wilson, and et al. An artificial intelligence decision support system for the management of type 1 diabetes. *Nat Metab*, 2:612–619, 2020.

[16] Josep Vehí, Iván Contreras, Silvia Oviedo, Lyvia Biagi, and Arthur Bertachi. Prediction and prevention of hypoglycaemic events in type-1 diabetic patients using machine learning. *Health Informatics Journal*, 26(1):703–718, 2020. PMID: 31195880.

[17] Roberto Visentin, Michele Schiavon, Rita Basu, Ananda Basu, Chiara Dalla Man, and Claudio Cobelli. Chapter 6 - physiological models for artificial pancreas development. In Ricardo S. Sánchez-Peña and Daniel R. Cherñavvsky, editors, *The Artificial Pancreas*, pages 123–152. Academic Press, 2019.

[18] Meng Zhang, Kevin B. Flores, and Hien T. Tran. Deep learning and regression approaches to forecasting blood glucose levels for type 1 diabetes. *Biomedical Signal Processing and Control*, 69:102923, 2021.

[19] T. Zhu, C. Uduku, K. Li, and et al. Enhancing self-management in type 1 diabetes with wearables and deep learning. *npj Digital Medicine*, 5:78, 2022.

[20] Taiyu Zhu, Kezhi Li, Pau Herrero, and Pantelis Georgiou. Basal glucose control in type 1 diabetes using deep reinforcement learning: An in silico validation. *IEEE Journal of Biomedical and Health Informatics*, 25(4):1223–1232, 2021.

GlucOS: A secure, safe and extensible system for automated insulin delivery

Hari Venugopalan, Shreyas Madhav Ambattur Vijayanand and Samuel T King
University of California, Davis

UC DAVIS UNIVERSITY OF CALIFORNIA

IEEE S&P

Problem Statement

- ML based algorithms are well suited for managing Type 1 Diabetes (T1D).
  - Calculating insulin doses depends on various intrinsic and extraneous factors.
  - ML can predict insulin dynamics.
  - Millions of papers published on using ML to manage T1D.
- Automated insulin delivery systems do not use ML.
  - Incorrect ML predictions can kill individuals.
  - Looser control and long-term health complications.
- GlucOS to the rescue:
  - **Extensible:** GlucOS supports any algorithm for automated insulin delivery, ML or otherwise.
  - **Safe and Secure:** GlucOS repurposes traditional algorithms to provide the foundation for security.
  - Allows ML to be proactive with insulin delivery while staying within the bounds of traditional algorithms for safety and security.

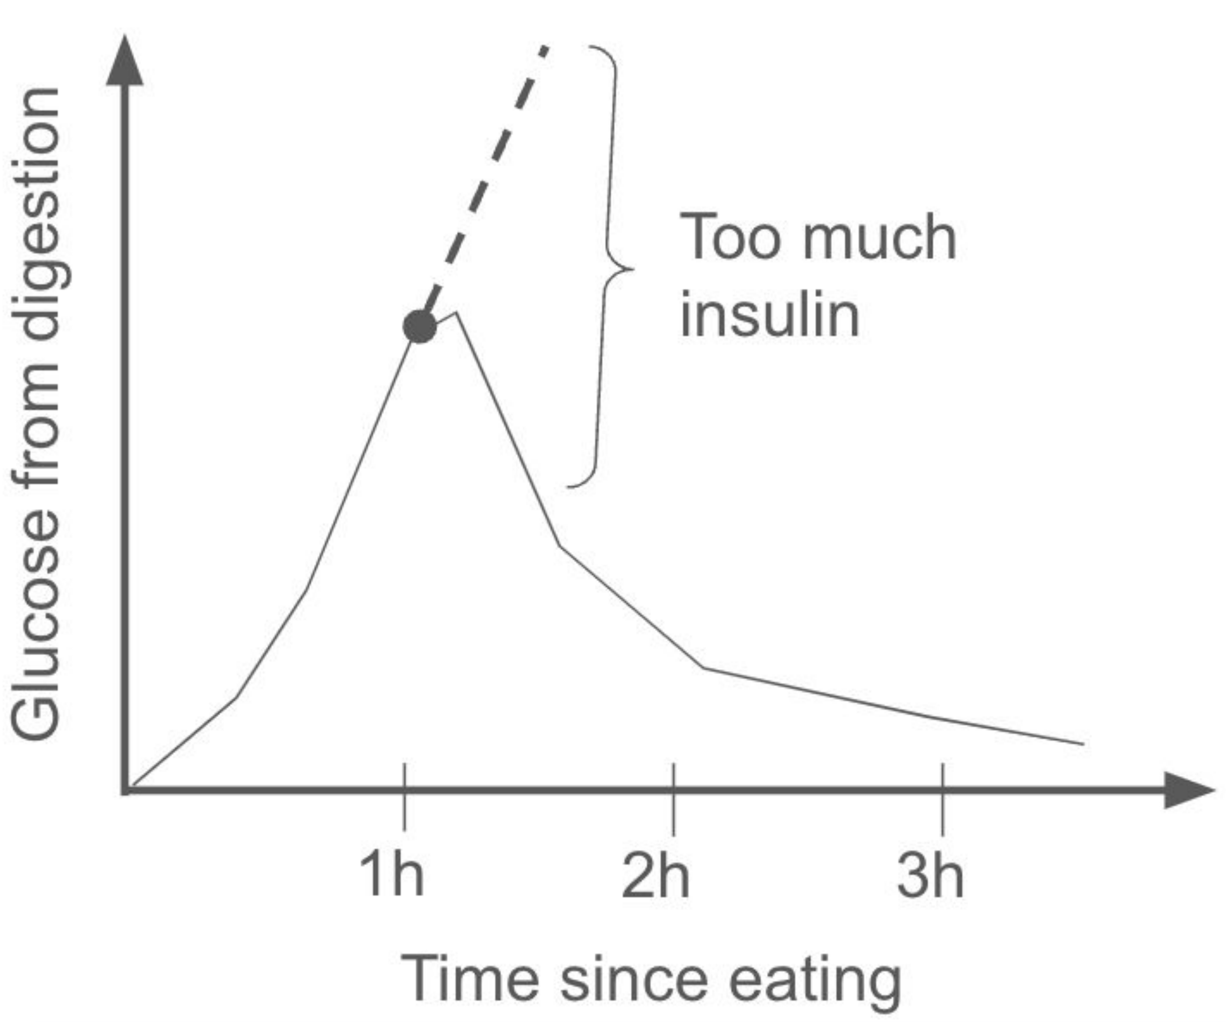

System Design

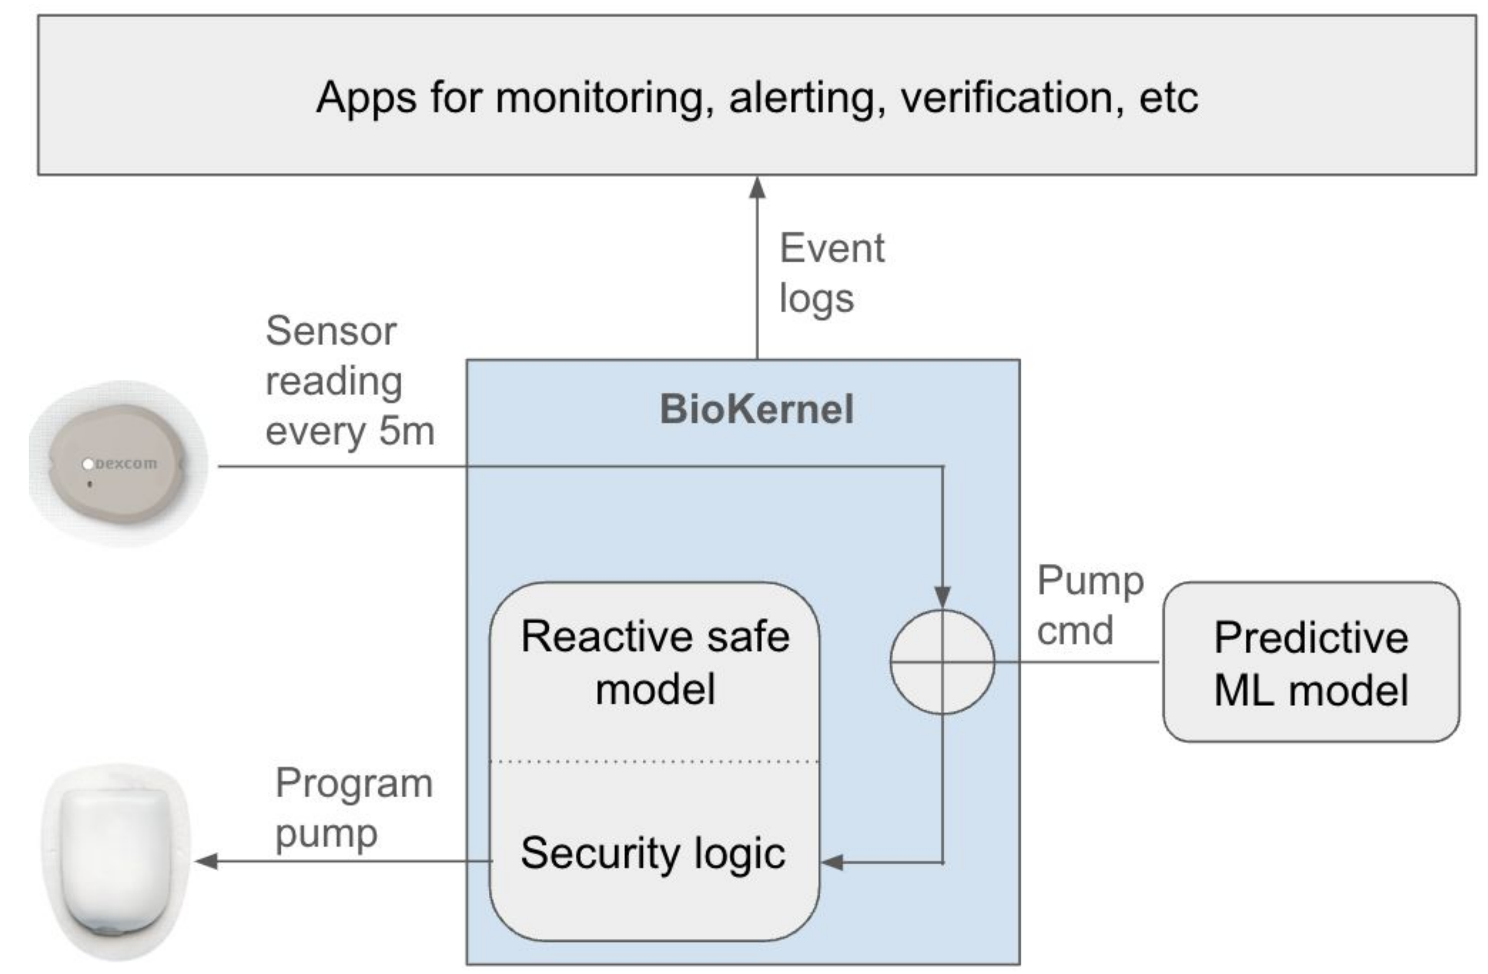

Results

Malicious algorithm dosing 10x the required insulin **without** GlucOS (instant death)

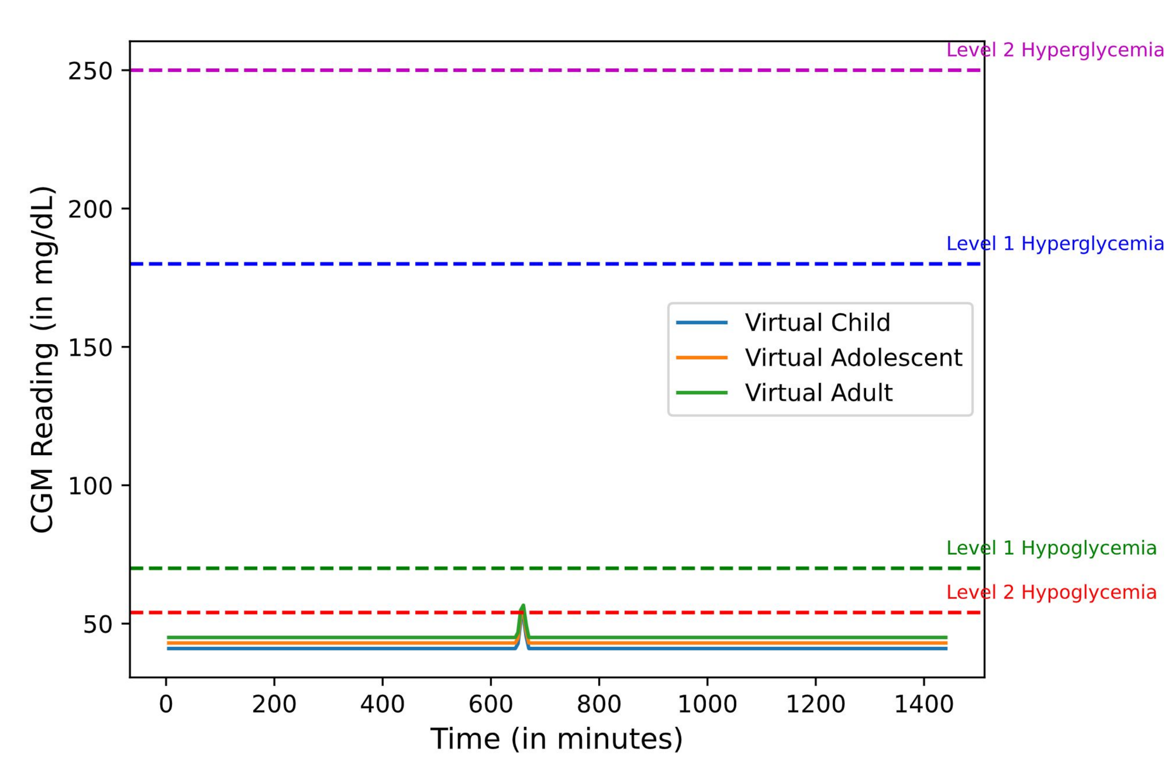

Malicious algorithm dosing 10x the required insulin **with** GlucOS

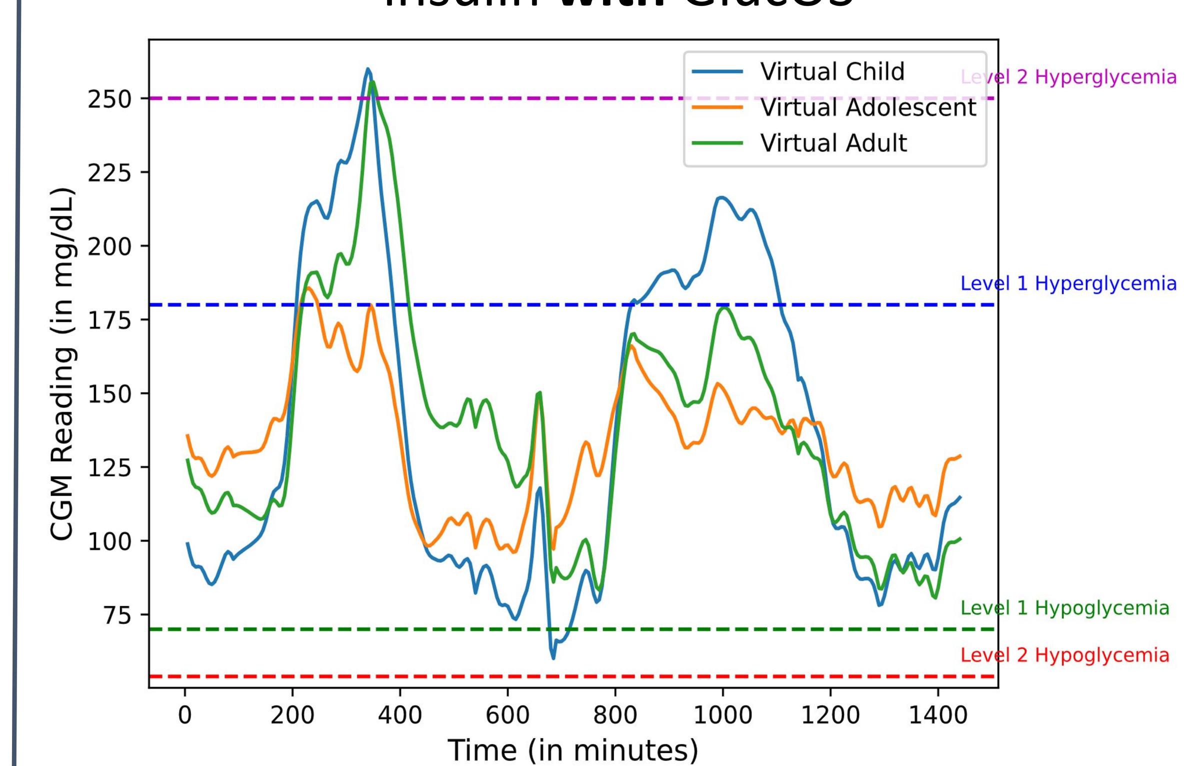

Malicious algorithm dosing 0.1x the required insulin **without** GlucOS (health complications)

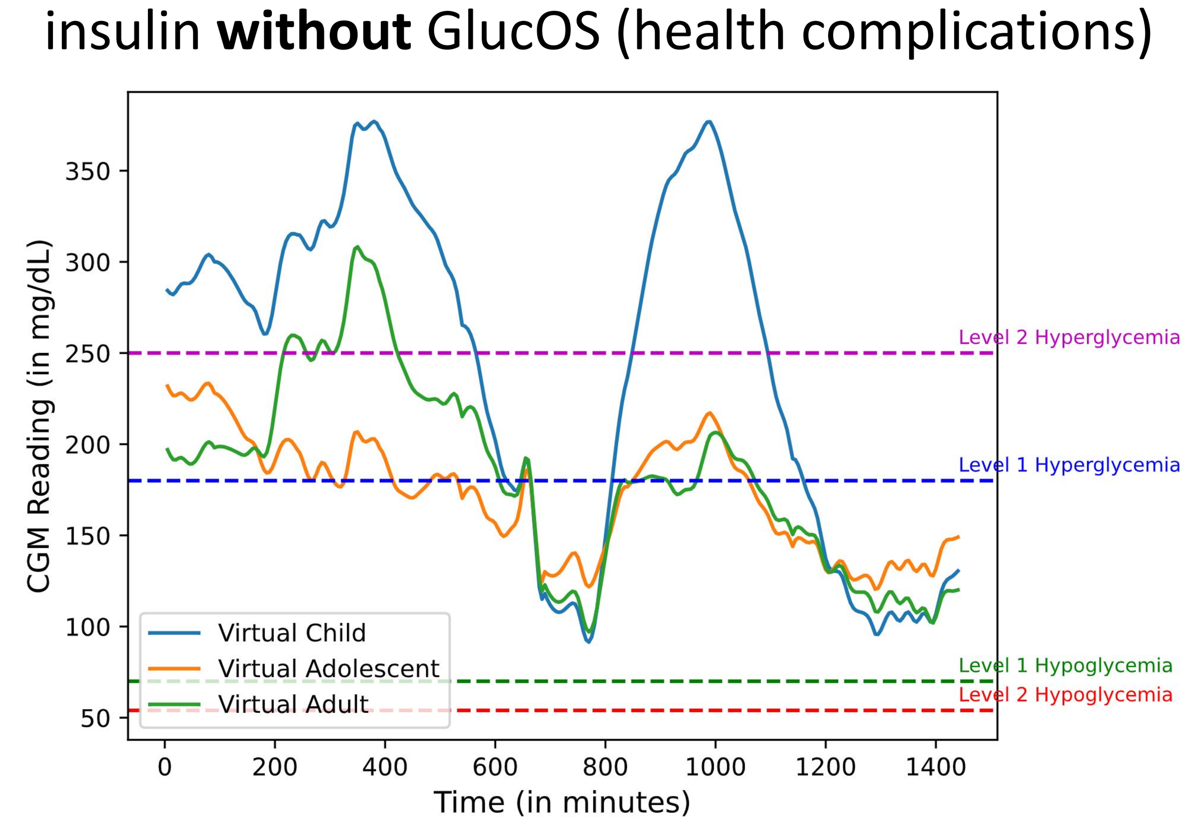

Malicious algorithm dosing 0.1x the required insulin **with** GlucOS

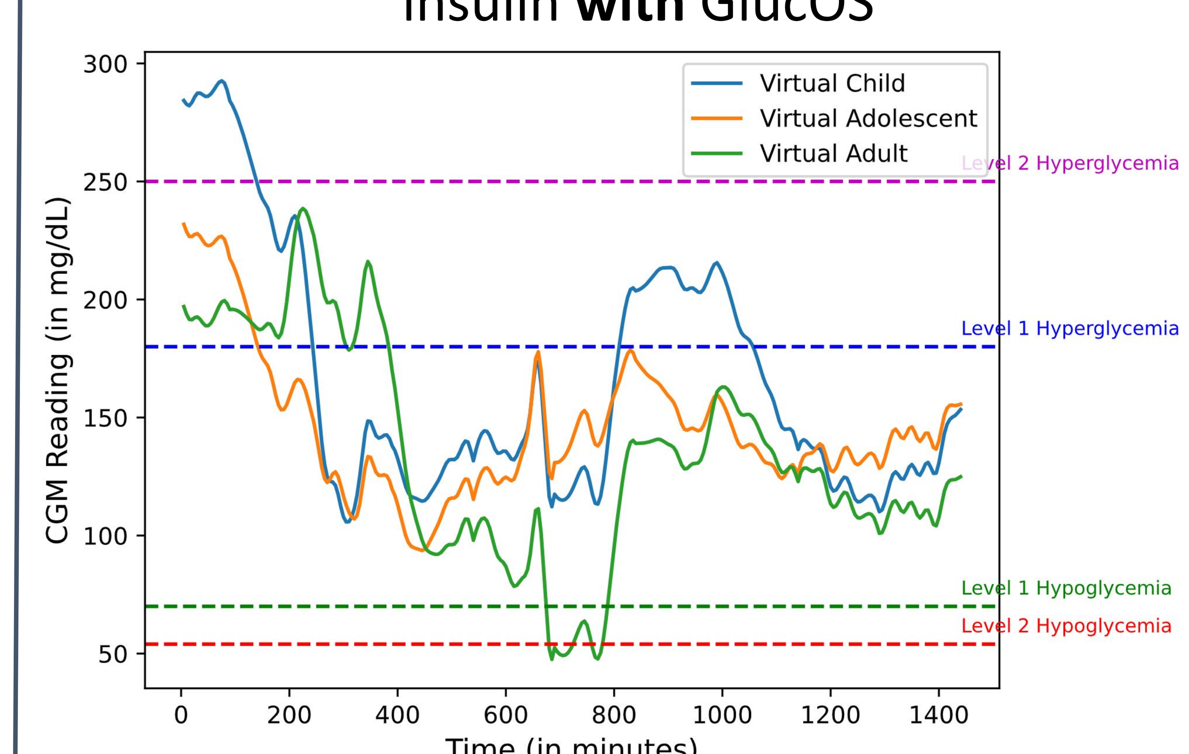

Real-world Deployment

Glucose Range (mg/dl)	Range label	Reactive Model	Predictive ML Model
<54	L2 Hypo	0%	0%
54-69	L1 Hypo	0.48%	0%
70-180	In range	91.9%	97.22%
181-250	L1 Hyper	7.62%	2.78%
>250	L2 Hyper	0%	0%

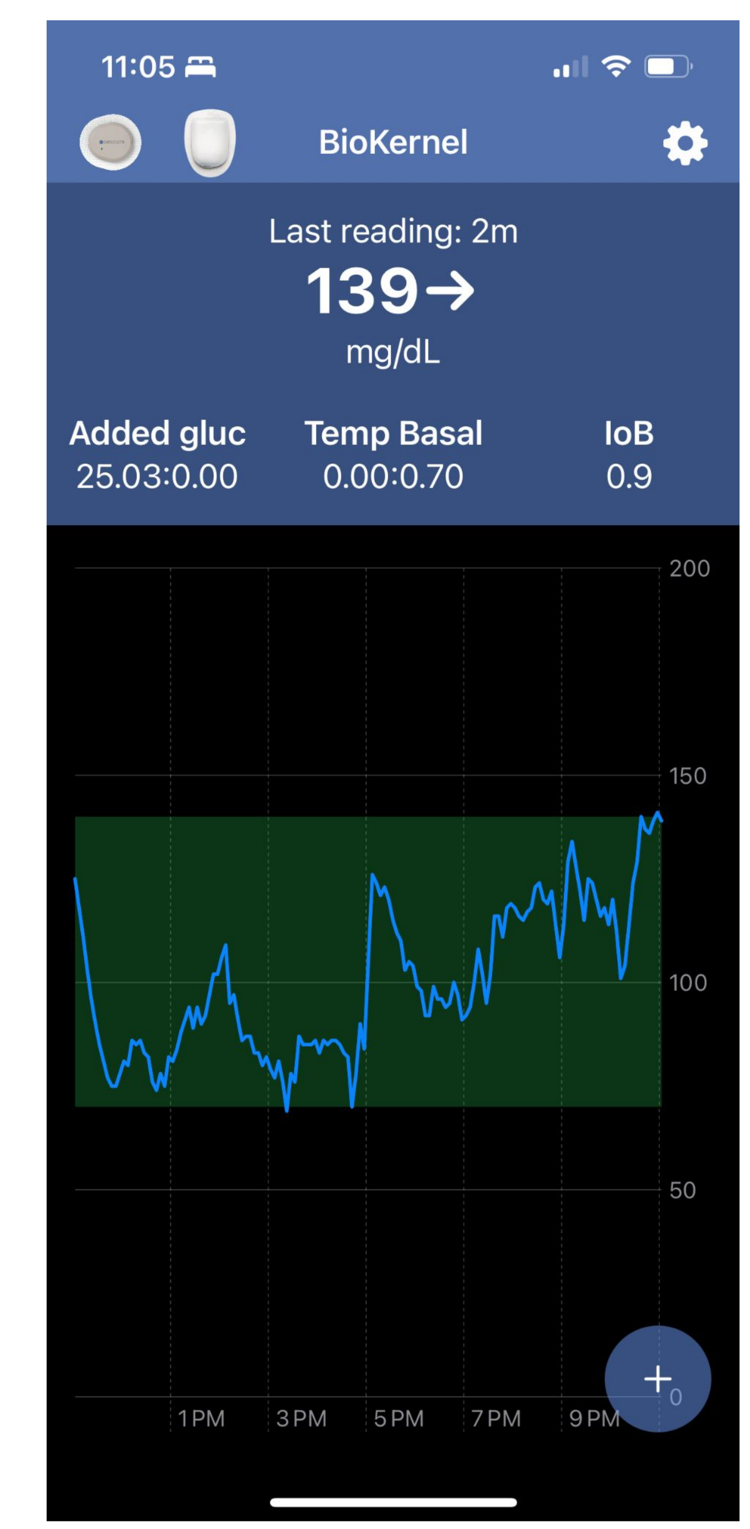

Future Work

- **Authentication:** Repurpose insulin delivery systems to provide seamless authentication for individuals with T1D.
- **Formal Verification:** Prove safety from implementational vulnerabilities/bugs.
- **HCI:** Discover effective alerting methods for individuals who prefer manual injections.