Name:

Mihika Bobbarjung

Email Address:

mihikabob10@gmail.com

Research Topic:

Physics-Informed Machine Learning for Meal-Driven Blood Glucose Prediction

Abstract: Diabetes affects over 537 million people worldwide, yet current glucose management tools remain expensive and prone to failure (Cheng et al., 2024). This study proposes a hybrid model for meal-driven blood glucose prediction, combining a calibrated physiological ordinary differential equation (ODE) with a Long Short-Term Memory (LSTM) neural network. Using the CGMacros dataset, which pairs continuous glucose monitor readings with macronutrient intake logs from 45 participants, two models will be trained on dietary features alone, deliberately excluding glucose history to prevent shortcut learning. A pure LSTM baseline will be compared against a hybrid model that additionally receives a Bergman Minimal Model-derived physiological prediction as a structured input. This framework tests whether physiological inputs provide measurable improvements in post-meal glucose prediction beyond data-driven approaches, with potential applications in CGM failure scenarios and real-time dietary decision support.

Introduction: For many individuals with diabetes, continuous blood glucose monitoring is the difference between safety and a life-threatening episode (Zhang et al., 2023). Continuous Glucose Monitoring (CGM) has become a key tool in diabetes management, providing real-time glucose data (Farahmand et al., 2025a; Zhang et al., 2023). However, CGM is expensive, prone to sensor failure, and requires uninterrupted data to be useful (Zhang et al., 2023; Kovur et al., 2025). Most existing predictive models rely on past glucose readings as inputs. Because glucose changes slowly, these models learn to predict near-persistence rather than learning how food actually drives glucose up or down, a critical failure mode for dietary intervention (Alredaini et al., 2025). Notably, prior work has shown that naive persistence baselines, predicting that glucose will remain close to its most recent value, can achieve performance comparable to more complex models over short prediction horizons, showing that apparent accuracy may reflect autocorrelation rather than true physiological understanding (Rodríguez-Rodríguez et al., 2019). This addresses the known issue that short-horizon glucose prediction models can achieve deceptively strong performance by exploiting temporal autocorrelation rather than learning meal-driven physiological responses. This study addresses the following question: Does incorporating a Bergman Minimal Model-derived physiological input reduce post-meal glucose prediction error compared to a purely data-driven LSTM when glucose history is withheld from both models? Glucose history is intentionally excluded as an experimental control to prevent models from exploiting temporal autocorrelation, therefore isolating whether they can learn causal relationships between meal composition and glucose response rather than relying on near-persistence signals. This is urgent: a meal-driven predictor that does not depend on continuous sensor data could serve millions of patients during CGM gaps and guide real-time dietary decisions without requiring hardware.

Literature Review: Recurrent neural networks, particularly LSTMs and GRUs, became the most widely used approach for CGM prediction because they can learn patterns across time, with LSTMs specifically solving the training instability that limited earlier networks (Bian et al., 2024; Lim et al., 2025). Transformer-based models have shown stronger results on longer prediction windows by using attention mechanisms to weigh multiple data streams at once (Farahmand et al., 2025a; Farahmand et al., 2025b). Combining both, Bian et al. (2024) showed that a Transformer-LSTM hybrid outperforms either model alone. Recent work explores federated and retrieval-based systems for personalization. Generalizability across patient groups has been studied by Rancati et al. (2024) and Haleem et al. (2025), while Yu et al. (2022) showed that transfer learning accelerates adaptation to new patients. Despite this progress, nearly all models receive past glucose values as inputs, achieving apparent accuracy by exploiting

persistence rather than learning meal dynamics, a fundamental confound that undermines their clinical value for dietary guidance (Rodríguez-Rodríguez et al., 2019; Soumma & Ghasemzadeh, 2026). Existing approaches prioritize predictive accuracy without explicitly enforcing physiological structure, obscuring whether models learn causal relationships or exploit glucose regularities. Ghimire et al. (2026) uses a hybrid neural network physiological model to predict glucose, but doesn't use a physiological equation as a direct model input. Limited studies have incorporated a Bergman Minimal Model-derived prediction as an explicit structured input to a neural network while simultaneously withholding glucose history, leaving the independent contribution of physiological input to meal-driven prediction largely unquantified.

Methodology: Dataset & Preprocessing: This study will use CGMacros, a dataset pairing CGM readings (recorded every 5 minutes by a FreeStyle Libre sensor) with self-reported meal logs from 45 participants (Gutierrez-Osuna et al., 2025). Features include glucose (mg/dL), carbohydrates, protein, fat, and fiber (grams). All features will be scaled using training data only, and the dataset will be split 80/20 by time to prevent data leakage. Glucose history is excluded to isolate meal-response learning and prevent reliance on autoregressive signals. **Feature Engineering:** Neither model will receive past glucose values. Inputs include: (1) macronutrient absorption curves modeled as exponential decay signals (carbohydrates: 60-min; protein: 120-min; fat: 180-min); (2) delayed meal signals at 15, 30, 60, and 120 minutes; (3) rolling 60-minute meal totals and a fiber-adjusted carbohydrate load score; and (4) time-of-day signals encoded as sine and cosine to capture phenomena like the dawn effect (O'Neal et al., 2023). **Physiological Model:** A Bergman Minimal Model-derived baseline will estimate glucose response to meals. Three parameters, glucose rise rate, decay speed, and resting level, will be fit per participant via numerical optimization on training data. This is a simplified physiological model of glucose–insulin behavior. It does not fully simulate the body, but instead provides a helpful biologically informed input that guides the machine learning model (Bergman, 2021). **Model Architecture & Evaluation:** A standard stacked LSTM with dropout and dense output layers will serve as the baseline, trained on meal features only. The hybrid model receives identical inputs plus the physiological prediction, allowing it to learn when to use or override it. Both models will be evaluated on RMSE of standardized glucose values (standardized per participant), reinforced by mean absolute error (MAE) to improve interpretability and a focused analysis of postprandial error (0–120 minutes) to assess clinical relevance during rapid glucose change.

Conclusion: This study proposes a rigorously controlled test of whether physiological inputs improve meal-driven glucose prediction, a question that has not been directly addressed under this controlled experimental design. The dataset size is limited but sufficient for proof-of-concept; the framework is explicitly designed for extension to larger datasets, and more powerful architectures. All computations will use open-source tools (Python, TensorFlow, scikit-learn, scipy) on standard hardware, ensuring reproducibility. More broadly, this work tests the hypothesis that incorporating domain knowledge into machine learning models can yield improvements that purely data-driven scaling cannot achieve, especially where data contains strong autocorrelation. If the hybrid outperforms the baseline, this study provides evidence that embedding known physiology into machine learning produces higher predictive accuracy, not just a restatement of glucose inertia, with direct implications for CGM-free dietary management in diabetes care. This has implications beyond glucose prediction, suggesting a larger role for hybrid modeling in fields where observational data alone may hide underlying mechanisms.

References

Alredaini, R., Abulkhair, M., & Almisbahi, H. (2025). Interpretable glucose forecasting for type 2 diabetes across traditional, deep, and large language models. Scientific Reports, 16, 2421. https://doi.org/10.1038/s41598-025-32373-4

Bergman R. N. (2021). Origins and History of the Minimal Model of Glucose Regulation. Frontiers in endocrinology, 11, 583016. https://doi.org/10.3389/fendo.2020.583016

Bian, Q., As'arry, A., Cong, X., Rezali, K. A. B. M., & Raja Ahmad, R. M. K. B. (2024). A hybrid Transformer-LSTM model apply to glucose prediction. *PLoS One, 19*(9), e0310084.

Cheng, M., Diao, X., Zhou, Z., Cui, Y., Liu, W., & Cheng, S. (2024). Toward short-term glucose prediction solely based on cgm time series. arXiv preprint arXiv:2404.11924.

Farahmand, E., Azghan, R. R., Chatrudi, N. T., Kim, E., Gudur, G. K., Thomaz, E., & Ghasemzadeh, H. (2025a). AttenGluco: Multimodal Transformer-based blood glucose forecasting on AI-READI dataset. In *2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 1–7). IEEE.

Farahmand, E., Azghan, R. R., Chatrudi, N. T., Ansu-Baidoo, V. Y., Kim, E., Gudur, G. K., & Ghasemzadeh, H. (2025b). GluMind: Multimodal parallel attention and knowledge retention for robust cross-population blood glucose forecasting. *arXiv preprint arXiv:2509.18457.*

Ghimire, S., Celik, T., Gerdes, M., & Omlin, C. W. (2026). Personalized Blood Glucose Prediction Using Physiology- Informed Machine Learning. Machine Learning and Knowledge Extraction, 8(4), 96. https://doi.org/10.3390/make8040096

Gutierrez-Osuna, R., Kerr, D., Mortazavi, B., & Das, A. (2025). CGMacros: a scientific dataset for personalized nutrition and diet monitoring (version 1.0.0). PhysioNet. RRID:SCR_007345. https://doi.org/10.13026/3z8q-x658

Haleem, M. S., Katsarou, D., Georga, E. I., et al. (2025). A multimodal deep learning architecture for predicting interstitial glucose for effective type 2 diabetes management. *Scientific Reports, 15*, 27625.

Kovur, P., Kovur, K. M., Rayat, D. Y., & Wishart, D. S. (2025). POC sensor systems and artificial intelligence — where we are now and where we are going? *Biosensors, 15*(9), 589.

Lim, M. H., Chae, H., Yoon, J., & Shin, I. (2025). A deep learning framework for virtual continuous glucose monitoring and glucose prediction based on life-log data. *Scientific Reports, 15*(1), 16290.

O'Neal, T. B., & Luther, E. E. (2023, May 16). Dawn phenomenon. In StatPearls. StatPearls Publishing. https://www.ncbi.nlm.nih.gov/books/NBK430893/

Rancati, S., Bosoni, P., Schiaffini, R., Deodati, A., Mongini, P. A., Sacchi, L., Toffanin, C., & Bellazzi, R. (2024). Exploration of foundational models for blood glucose forecasting in Type-1 diabetes pediatric patients. *Diabetology, 5*(6), 584–599.

Rodríguez-Rodríguez, I., Chatzigiannakis, I., Rodríguez, J.V., Maranghi, M., Gentili, M., & Zamora-Izquierdo, M.A. (2019). Utility of Big Data in Predicting Short-Term Blood Glucose Levels in Type 1 Diabetes Mellitus Through Machine Learning Techniques. Sensors (Basel, Switzerland), 19.

Soumma, S. B., & Ghasemzadeh, H. (2026). GlyRAG: Context-aware retrieval-augmented framework for blood glucose forecasting. *arXiv preprint arXiv:2601.05353.*

Yu, X., Yang, T., Lu, J., et al. (2022). Deep transfer learning: A novel glucose prediction framework for new subjects with type 2 diabetes. *Complex & Intelligent Systems, 8*, 1875–1887.

Zhang, L., Yang, L., & Zhou, Z. (2023). Data-based modeling for hypoglycemia prediction: Importance, trends, and implications for clinical practice. *Frontiers in Public Health, 11*, 1044059.