

Health-driven personalized metabolic models of postprandial glucose responses to mixed meals

Anurag Das

Department of Computer
Science and Engineering
Texas A&M University
College Station, USA
adas@tamu.edu

Ghady Nasrallah

Department of Computer
Science and Engineering
Texas A&M University
College Station, USA
ghadynasrallah@tamu.edu

Sicong Huang

Department of Computer
Science and Engineering
Texas A&M University
College Station, USA
siconghuang@tamu.edu

Bobak J. Mortazavi

Department of Computer
Science and Engineering
Texas A&M University
College Station, USA
bobakm@tamu.edu

Ricardo Gutierrez-Osuna

Department of Computer
Science and Engineering
Texas A&M University
College Station, USA
rgutier@tamu.edu

Abstract—The relationship between the macronutrient composition of a meal and the resulting post-prandial glucose response is complex given the large inter-individual differences in metabolism. We present JointCGMacros, a computational model that learns a joint embedding of meal macronutrients and post-prandial glucose, mediated by demographics, metabolic health, and gut microbiota variables. The model extracts parallel embeddings from (1) postprandial glucose responses to a meal and (2) the meal’s macronutrient composition conditioned on health parameters using a triplet loss. The macronutrient embedding is an interpretable parametric expression that captures how health parameters modulate the effect of individual macronutrients. We evaluated the model on an experimental dataset containing postprandial glucose responses to a variety of mixed meals from subjects with different metabolic health status (healthy, pre-diabetes, type 2 diabetes). JointCGMacros significantly outperforms a model that attempts to predict macronutrients directly from postprandial glucose. These findings may lead to the development of automatic dietary monitoring using off-the-shelf wearable devices.

Index Terms—Continuous glucose monitors, diet monitoring, triplet loss, metabolic syndrome.

I. INTRODUCTION

Sustained levels of high blood glucose can have serious health consequences, increasing the risk of developing diabetes and its long-term consequences (e.g., kidney failure, blindness, limb amputations). Thus, monitoring and controlling diet is critical to preventing and managing diabetes. However, manual recording of food intake is often time-consuming and error-prone [1]. Several technologies have been proposed to assist in automatic diet monitoring (e.g., inertial measurements, microphones, wearable cameras), but these solutions only capture limited information, such as the timing, type, and amount of food, but not their macronutrient composition.

An alternative approach is to analyze the glucose response after consuming a meal, i.e., the post-prandial glucose response (PPGR), which can be measured using off-the-shelf continuous glucose monitors (CGMs). The primary contributor to elevated PPGRs is the amount (and type) of carbohydrates, but PPGRs are also influenced by other, non-glycemic macronutrients in the meal. For example, adding protein and fat to a carbohydrate-rich meal can reduce the initial peak

and slow down recovery to baseline [2], [3]. Thus, PPGRs contain information that may be used to recover, in part, the macronutrient composition of meals. The main challenge with this approach is the large inter-individual difference in PPGRs to identical meals.

To address this challenge, we propose **JointCGMacros**, a machine-learning model that learns a joint embedding of PPGRs and macronutrients that captures their underlying relationship, conditioned on each individual’s health and demographic factors. JointCGMacros consists of two modules: (1) a PPGR encoder that consumes CGM data to produce a low-dimensional embedding of the PPGR to a meal, and (2) a health encoder that consumes health and demographics factors to produce a personalized macronutrient embedding. We train JointCGMacros using a triplet loss [4] that aligns PPGR and macronutrient embedding pairs from the same meal while forcing PPGR and macronutrient embedding pairs from different meals to be away from each other. We evaluate the performance of the proposed joint-learning model against an equivalent model that predicts a generic (i.e., not personalized) embedding of macronutrients.

II. RELATED WORK

Studies over the past decade have explored the use of CGMs to develop personalized nutrition programs that account for inter-individual differences. Zeevi et al. [5] collected CGM data and meal logs from an 800-person cohort, and developed a machine-learning (ML) model that integrated clinical features, dietary habits, physical activity, and gut microbiome profiles to predict glycemic responses to meals. Their approach was validated in a randomized controlled trial, where personalized dietary recommendations significantly reduced postprandial glucose excursions. Tily et al. [6] also highlighted the importance of gut microbiome activity in predicting PPGRs. ML models trained on food composition, anthropometrics, and microbial pathway activity showed that gut microbiome features significantly improved predictions.

These studies seek to estimate PPGRs from the macronutrient amounts of meals (i.e., a *direct* problem). Its *inverse* counterpart, i.e., predicting macronutrients from PPGRs, is far more challenging. It is a one-to-many problem in which metabolic health parameters (e.g., HbA1c, insulin sensitivity) and demographics (e.g., sex, age) play a major role. Thus,

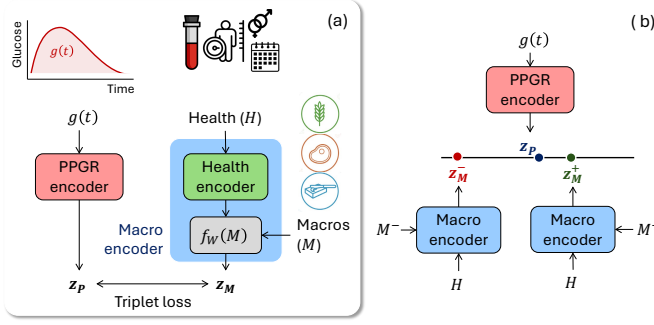


Fig. 1. (a) Block diagram of JointCGMacros. A PPGR $g(t)$ is passed to an encoder network to produce a one-dimensional embedding z_P . Health parameters H are fed to a second encoder network to produce a set of weights that are combined with the meal macronutrients M to produce a macro embedding z_M using a parametric function $f_W(M)$. (b) The two encoders are then jointly trained using triplet loss, which takes a PPGR embedding z_P as an anchor and macro embeddings from the same meal (M^+) and from different meals (M^-) to bring z_P and z_M^+ closer to each other and push z_P and z_M^- apart.

PPGRs from two individuals cannot be compared without controlling for these health/demographic factors. Our proposed model (JointCGMacros) addresses this complex issue by learning a joint embedding that aligns PPGRs with the macronutrient composition of a meal, *conditioned* on metabolic and demographic parameters. The end result is an interpretable and personalized embedding of a meal that accounts for inter-individual differences.

III. METHODS

A. Joint embedding model

The proposed model is illustrated in Figure 1a. It consists of two encoder networks trained jointly in a data-driven manner. The PPGR encoder consumes the 3-hour postprandial glucose response to a meal (as measured by a CGM every 15 min) as a 13-dimensional vector ($g(t)$), and generates a one-dimensional PPGR embedding z_P . In turn, the health encoder consumes a vector of metabolic health and demographics parameters of each individual (H) and generates a personalized set of weights that are combined with the macronutrient amounts of a meal (M) to produce an embedding z_M as:

$$z_M = \frac{C - w_P P - w_F F - w_B B}{1 + w'_C C + w'_P P + w'_F F + w'_B B} \quad (1)$$

where C , P , F , and B are the amount of net carbs, protein, fat and fiber, respectively, in calories¹. Coefficients w_P , w_F , and w_B in the numerator denote how much protein, fat, and fiber subtract from the effect of carbs, whereas coefficients w'_C , w'_P , w'_F , and w'_B in the denominator produce an adjusted caloric content. As such, the embedding z_M can be interpreted as a generalization of the carb-caloric-ratio, the proportion of calories in a meal that come from carbohydrates, a measure of the glycemic load of a meal:

$$CCR = \frac{C}{C + P + F + B} \quad (2)$$

¹Macronutrient amounts in grams are converted into calories as follows: 4 calories/gram for carbs, 4 calories/gram for protein, 9 calories/gram for fat, and 2 calories/gram for fiber

For $(w'_C, w'_P, w'_F, w'_B) \rightarrow 0$, the denominator approaches one, in which case z_M only considers the subtractive effect of non-glycemic macronutrients. In turn, for $w'_C = w'_P = w'_F = w'_B = 1$ and $w_P = w_F = w_B = 0$, z_M reduces to the CCR in equation (1). This parametric embedding has two main advantages. First, it allows the model to learn the optimum combination of macronutrients that is related to PPGRs (either linearly or non-linearly) for each individual. Second, the resulting weights are interpretable (one per macronutrient) and can be traced back to the health parameters of each individual. Both z_M and z_P are one-dimensional embeddings to prioritize interpretability and reduce the risk of overfitting.

B. Triplet loss

We use a triplet loss to learn the joint embedding space by aligning PPGR embeddings with their corresponding macronutrient embeddings. As illustrated in Figure 1b, each training example consists of a triplet (z_P, z_M^+, z_M^-) , where z_P is the PPGR embedding (anchor), and z_M^+ and z_M^- are the macronutrient embeddings from the same meal (positive) and a different meal (negative), respectively. By changing only the macronutrient in each triplet, we force the network to distinguish between positive and negative macros. Thus, the loss function encourages the PPGR embedding to be closer to meals that induce similar responses and far away from those that induce different responses:

$$L_{\text{triplet}} = \sum_{i=0}^N \max(\|z_P^{(i)} - z_{M^+}^{(i)}\|_2^2 - \|z_P^{(i)} - z_{M^-}^{(i)}\|_2^2 + \alpha, 0) \quad (3)$$

where N is the number of triplets², and α is the margin separating the positive examples from the negative examples. We use the L2 norm as the distance metric between PPGR and macro embeddings.

C. Network architecture

The PPGR encoder consists of two hidden layers with 64 nodes and 32 nodes, each followed by BatchNorm [7] and ReLU activation [8], and dropout set to 0.1. Health parameters H are passed to an encoder with two hidden layers, the first one of size 64 followed by BatchNorm and ReLU activation, and the second one of size 7, i.e., the number of learnable weights in eq. (1). To force the weights to be positive, we pass them through a softplus activation and compute z_M using eq. (1). The network is trained using a learning rate of 5×10^{-3} with a batch size of 64. The triplet margin (α) is set to 1.0.

D. Predicting macro embedding

Once trained, the model can be used to predict the macro embedding (z_M) of a meal from the corresponding PPGR, as follows. First, the PPGR of a test meal is passed to the PPGR encoder to obtain the corresponding embedding (z_P). Then, an individual's health parameters H are passed to the second encoder and combined with a macronutrient vector M to produce the embedding z_M . Since M is unknown for a

²The dataset used for this study –see section III-F– contains six unique meal combinations. This results in five distinct triplets per CGM, each sharing the same anchor and positive, but differing in the negative.

TABLE I
MACRONUTRIENT COMPOSITION OF BREAKFAST MEALS, CODED AS
HAVING HIGH (H) OR LOW (L) AMOUNTS OF C , P , F AND B .

Day	Meal	Carb (g)	Prot (g)	Fat (g)	FiBer (g)	CCR
4	HHHH	66	66	42	07	0.287
5	LLLL1	24	22	11	00	0.345
9	LLLL2	24	22	11	00	0.345
10	HLHH	66	22	42	07	0.355
3	HLHL1	66	22	42	00	0.362
8	HLHL2	66	22	42	00	0.362
2	HHLL1	66	66	11	00	0.424
7	HHLL2	66	66	11	00	0.424
1	HLLL1	66	22	11	00	0.591
6	HLLL2	66	22	11	00	0.591

test sample, we generate multiple embeddings ($z_{M,1}, z_{M,2}, \dots, z_{M,n}$) from a list of potential meals (M_1, M_2, \dots, M_n). The meal M_i whose macronutrient embedding $z_{M,i}$ is closest to z_P is treated as the prediction.

E. Validation

We compare JointCGMacros against a baseline model that predicts the CCR in eq. (2) from its corresponding PPGR. The baseline model is a feedforward network with the same number of parameters as the PPGR encoder in JointCGMacros. It consists of two hidden layers, each with 64 units and each followed by BatchNorm, Dropout, and ReLU activation. A final fully-connected (FC) layer takes the output of the hidden layer and produces the CCR. *Because the baseline model does not use health parameters H , this comparison allows us to determine if the personalized macro embedding z_M can reduce the effect of inter-individual differences.*

We train both models using 10-fold cross-validation. For each training fold, we randomly select 80% of the data for training and 20% for validation. We use validation data for early stopping (10 epochs) to prevent overfitting. We report two performance measures on the test fold: the Pearson correlation and the normalized root mean squared error (NRMSE) between ground truth and prediction. Because the NRMSE is the percentage of error relative to ground truth, it allows us to compare errors between target variables of different scales.

F. Dataset description

We used the publicly available CGMacros dataset [9] on PhysioNet. CGMacros contains PPGRs, meal macronutrients, and health/demographics parameters for 45 subjects (ages: 18-69; BMI: 21-46 kg/m²; 15 healthy, 16 with pre-diabetes, and 14 with type 2 diabetes). Participants wore an Abbott FreeStyle Libre Pro CGM (15-min sampling period) and a Dexcom G6 Pro CGM (5-min) on the upper arm and abdomen, respectively. For 10 days, participants consumed breakfast, lunch, and dinner meals with different amounts of carbs, protein, fat, and fiber following the average American diet [10]. Breakfast compositions are shown in Table I, sorted by CCR. Some of the breakfasts were consumed twice during the 10-day study. Lunches were ordered from a fast-casual chain restaurant (Chipotle Mexican Grill Inc.). For dinners, participants ate foods of their own choice and recorded the contents on MyFitnessPal. To minimize interferences in glucose responses

TABLE II
DEMOGRAPHIC, METABOLIC HEALTH, AND GUT MICROBIOME
PARAMETERS USED IN THE MODEL

Demographics (2):
- Age, gender
Health (10):
- Body Mass Index (BMI), Fasting Glucose (FG), Glycated Hemoglobin (HbA1c), Insulin (Ins), Homeostatic Model Assessment of Insulin Resistance (HOMA-IR)
- High-Density Lipoprotein (HDL), Cholesterol (Chot), Cholesterol HDL (chHDL), non-HDL, Very Low-Density Lipoprotein (VLDL),
Gut microbiome (22):
- <u>Overall</u> : Gut Microbiome Health (MicroHealth), Inflammatory Activity (Inflam), Digestive Efficiency (DigEff), Metabolic Fitness (MetFit), Gas Production (Gas), Gut Lining Health (GutLin), Microbiome-Induced Stress (Stress), Protein Fermentation (ProtFerm), Active Microbial Diversity (Div)
- <u>Pathways</u> for Flagellar Assembly (Flag), Chemotaxis/Virulence (ChemVir), Lipopolysaccharide Biosynthesis (LipoSacc), Salt Stress (Salt)
- <u>Metabolism Pathways</u> for Bile Acid (Bile) and Oxalate (Oxal)
- <u>Production Pathways</u> for Ammonia (Amm), Butyrate (But), Sulfide Gas (Sulf), Putrescine, (Putr), Trimethylamine (TMA), Methane Gas (Mth) and Uric Acid (Uric)

from prior meals, participants were instructed to eat lunch at least 3 hours after breakfast, with only water or coffee (without sugar) in between, and dinner at least 3 hours after lunch. They also took meal photographs before and after eating, from which we extracted the meal timestamps and the proportion of the meal they consumed. We report results from the breakfast meals using PPGRs from the Abbott FreeStyle CGM. We used 34 variables available on the CGMMacros dataset as health/demographics parameters -see Table II.

IV. RESULTS

A. Comparison against baseline

In a first step, we compare the performance of the baseline model against a JointCGMacros model that uses all 34 health parameters in Table II. Results are summarized in Figure 2a. JointCGMacros strongly outperformed (lower RMSE, higher correlation) the baseline model, indicating that our proposed model is able to account for inter-individual differences in postprandial glucose responses. Note that, for the baseline model, the correlation coefficient can become negative and the NRMSE can exceed an error of 100%, indicating that PPGRs are not predictive of the macronutrient composition of meals *unless individual health/demographics parameters are considered*.

B. Model interpretation

Figure 2b shows a scatterplot of the predicted macro and PPGR embeddings for the JointCGMacros model, color-coded by the ground-truth CCR in eq. (2). The embeddings towards the bottom left (dark blue) have the lowest CCR, which represents meals with high amounts of carbs, protein, fat, and fiber (e.g., meal HHHH in Table I). Since these meals have the highest amounts of non-glycemic macronutrients, they lead to the lowest PPGRs. Meals with the highest CCR (dark red) appear towards the top right. These meals have a high amount of carbs and low amounts of other macronutrients (HLLL in Table I), which leads to the highest PPGRs.

Table III shows the distribution of the weights in eq. (1) learned by the model. In the numerator, fiber ($w_B = 3.95$) has the strongest effect, followed by protein ($w_P = 1.22$) and

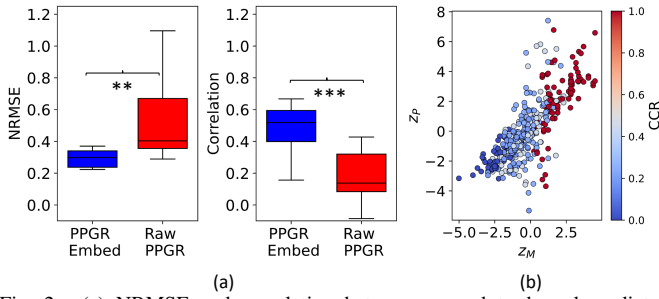


Fig. 2. (a) NRMSE and correlation between ground truth and predicted macronutrients for the baseline model (red) and JointCGMacros (blue) with 34 health parameters (**: $0.001 < p < 0.01$, ***: $p < 0.001$). (b) Scatterplot of macro z_M vs. PPGR z_P embeddings, color coded by the CCR in eq. (2).

TABLE III

MACRONUTRIENT WEIGHTS LEARNED BY THE JOINT MODEL ($\mu \pm \sigma$)

	C	P	F	B
Numerator		1.22(0.34)	0.33(0.21)	3.95(2.04)
Denominator	0.33(0.14)	0.04(0.08)	0.01(0.01)	0.01(0.01)

fat ($w_F = 0.34$), suggesting healthy alternatives to control elevated post-prandial glucose response (i.e., add fiber and protein, but not fat). In the denominator, net carbs (w'_C) have the strongest contribution ($w'_C = 0.33$), whereas the effect of the three non-glycemic macronutrients is negligible ($w'_P = 0.03$, $w'_F = 0.01$, $w'_B = 0.01$). Ignoring the smallest weights, these results indicate that, *when factoring health parameters*, post-prandial glucose responses are proportional to the expression $(C - 3.95B - 1.22P)/(1 + 0.33C)$.

C. Feature importance

In a final analysis, we performed a bidirectional search to identify the most critical health parameters. Bidirectional search alternates between forward subset selection (FSS) from the empty set and backward selection (BSS) from the full set, with the constraints that features added by FSS cannot be removed by BSS, and vice versa. As a result, both searches converge in the middle. Results are shown in Figure 3. The two most important features are HOMA-IR, a clinical measure of insulin resistance and beta-cell function, and fasting glucose, a key parameter for the diagnosis of T2D. The optimal feature set (lowest NRMSE) contains 17 features, including a mix of demographics (age, gender), metabolic health (lipid profiles, HbA1c), and several pathways from gut microbiome analytics. This optimal set achieves a lower NRMSE than the full feature set (rightmost boxplot, in brown), a difference that is statistically significant (0.247 vs. 0.291, $p = 0.03$).

V. DISCUSSION

Prior research has shown that the relationship between post-prandial glucose and meal macronutrients is largely driven by inter-individual differences. Our results show that this three-way relationship can be unveiled by learning a joint embedding of PPGRs and macronutrient composition conditioned on a few demographic, metabolic health, and gut microbiota parameters. We find that the key contributors to elevated postprandial glucose are net carbs (as expected) with a linear correction for fiber and protein. It is the expression

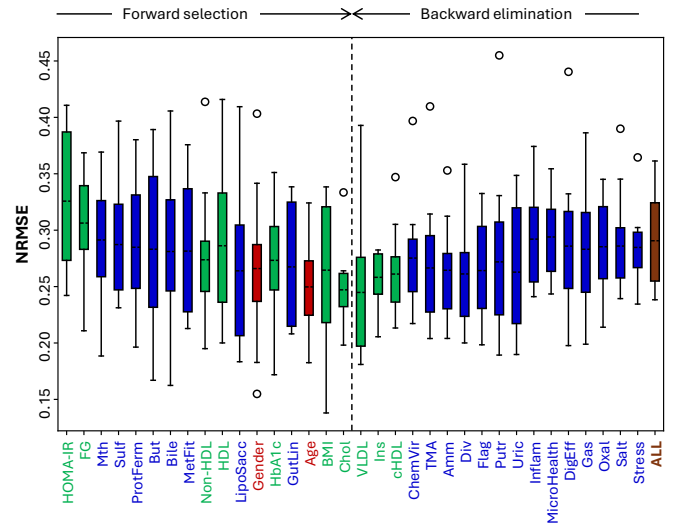


Fig. 3. Bidirectional search to identify the most important demographic (red), metabolic health (green), and gut microbiome (blue) variables.

$(C - 3.95B - 1.22P)/(1 + 0.33C)$, and not the conventional carb-caloric-ratio, that can be recovered from PPGRs, but only when accounting for an individual's health variables. We used the parametric expression in eq. (1) to improve the interpretability of our model, but further improvements in macronutrient prediction from PPGRs may be achieved by relaxing this constraint, e.g., concatenating macronutrient and health vectors $[M, H]$ and feeding them to a deep-learning model, but this would likely require access to a larger dataset than CGMacros, the only of its nature that is publicly available.

REFERENCES

- [1] B. J. Mortazavi and R. Gutierrez-Osuna, "A review of digital innovations for diet monitoring and precision nutrition," *J Diabetes Sci Technol*, vol. 17, no. 1, pp. 217–223, 2023.
- [2] E. Moghaddam, J. A. Vogt, and T. M. Wolever, "The effects of fat and protein on glycemic responses in nondiabetic humans vary with waist circumference, fasting plasma insulin, and dietary fiber intake," *J Nutr*, vol. 136, no. 10, pp. 2506–2511, 2006.
- [3] A. Rytz, D. Adeline, K.-A. Lê, D. Tan, L. Lamothe, O. Roger, and K. Macé, "Predicting glycemic index and glycemic load from macronutrients to accelerate development of foods and beverages with lower glucose responses," *Nutrients*, vol. 11, no. 5, p. 1172, 2019.
- [4] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," *arXiv preprint arXiv:1703.07737*, 2017.
- [5] D. Zeevi, T. Korem, N. Zmora, D. Israeli, D. Rothschild, A. Weinberger, O. Ben-Yacov, D. Lador, T. Avnit-Sagi, M. Lotan-Pompan *et al.*, "Personalized nutrition by prediction of glycemic responses," *Cell*, vol. 163, no. 5, pp. 1079–1094, 2015.
- [6] H. Tily, A. Perlina, E. Patridge, S. Gline, M. Genkin, V. Gopu, H. Lindau, A. Sjuue, N. Klitgord, and M. Vuyisich, "Gut microbiome activity contributes to individual variation in glycemic response in adults," *bioRxiv*, p. 641019, 2019.
- [7] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. ICML*, 2015, pp. 448–456.
- [8] Y. Li and Y. Yuan, "Convergence analysis of two-layer neural networks with relu activation," in *Proc. NeurIPS*, vol. 30, 2017.
- [9] R. Gutierrez-Osuna, D. Kerr, B. Mortazavi, and A. Das, "Cgmacros: a scientific dataset for personalized nutrition and diet monitoring," 2025. [Online]. Available: <https://physionet.org/content/cgmacros>
- [10] E. Cohen, M. Cragg, J. deFonseka, A. Hite, M. Rosenberg, and B. Zhou, "Statistical review of us macronutrient consumption data, 1965–2011: Americans have been following dietary guidelines, coincident with the rise in obesity," *Nutrition*, vol. 31, no. 5, pp. 727–732, 2015.