

Macronutrient constraints and priors improve carbohydrate predictions from continuous glucose monitors

Anurag Das
Dept. Computer Science and Engineering
Texas A&M University
College Station, TX, USA
adas@tamu.edu

Wendy Bevier
Sansum Diabetes Research Institute
Santa Barbara, CA, USA
wbevier@sansum.org

Bobak J. Mortazavi
Dept. Computer Science and Engineering
Texas A&M University
College Station, TX, USA
bobakm@tamu.edu

Edmund Do
Dept. Computer Science and Engineering
Texas A&M University
College Station, TX, USA
eds@tamu.edu

Rony Santiago
Office of Risk Services
University of California Santa Cruz
Santa Cruz, CA, USA
rofsanti@ucsc.edu

Ricardo Gutierrez-Osuna
Dept. Computer Science and Engineering
Texas A&M University
College Station, TX, USA
rgutier@tamu.edu

Namino Glanz
Santa Barbara County Education Office
Children & Family Resource Services
Santa Barbara, CA, USA
nglanz@sbceo.org

David Kerr
Diabetes Technology Society
Burlingame, CA
kerr@diabetestech.org

Abstract—We propose an approach to estimate the macronutrients in a meal *automatically* by analyzing the meal’s glucose response using off-the-shelf wearable sensors (continuous glucose monitors). We rely on the fact that the shape of the glucose response to a meal depends on all the macronutrients in the meal, not just its carbohydrates (carbs). However, protein, fat, and fiber tend to affect the glucose response in similar ways, so recovering their individual amounts is numerically ill-conditioned. To address this problem, our approach compresses macronutrients into a latent variable that captures their correlated effects on glucose. Then, we train a machine learning model to predict the latent variable from the glucose response of a meal. Finally, we recover the amount of the original macronutrients by incorporating prior knowledge of how they co-occur in conventional meals. Using experimental data from 45 participants, we show that predicting carbs indirectly (through the latent variable) reduces the prediction error when compared to predicting carbs directly, i.e., without considering the protein and fats in the meal.

Index Terms—continuous glucose monitors, diabetes, diet monitoring, macronutrients

I. INTRODUCTION

Foods¹ rich in carbohydrates (carbs) lead to high levels of blood glucose after a meal, which increases the risk of developing diabetes (T2D). To manage diabetes, patients have to monitor their diet, particularly carbs. However, conventional diet monitoring methods are cumbersome (e.g., manually logging meals in a diary) [1] and error-prone (e.g., 24-hour recalls) [2], [3]. A potential solution is to use continuous glucose monitors (CGMs) to estimate the macronutrients (macros) in a meal, in a manner akin to how fitness trackers

use accelerometers to estimate physical activity. CGMs are wearable devices that measure glucose every 5-15 minutes using a flexible electrode inserted under the skin. To predict macros from CGMs, we rely on the fact that glucose levels after a meal—known as the post-prandial glucose response (PPGRs), are affected not only by carbs, but also by other macros. For example, adding protein, fat and fiber to a carb-rich meal reduces and delays the peak glucose, and also slows baseline recovery. This suggests that the shape of the PPGR can be used to estimate the macros in a meal.

However, predicting the individual amounts of protein and fat is challenging since they have similar effects on PPGRs. Further, the macros in conventional meals tend to be correlated: macros are not generally single-sourced (e.g., when you eat avocado for its 15% in healthy fats, you are also consuming 9% in carbs), and they may be added to improve shelf life (e.g., breads with fat dry out more slowly) or ensure that the meal is palatable.

To address these two issues, we propose an approach that uses (1) an embedding equation to compress macros into a latent variable, and (2) a prior distribution over macros, as illustrated in Fig. 1. The embedding equation generates a latent variable $Z = f(C, P, F, B)$ that captures domain knowledge about how macros affect the PPGR (e.g., carbs increase it whereas protein and fat reduce it). It is this latent variable Z , rather than the individual macros, that we predict from a meal’s PPGR via machine learning. Once the model predicts Z , the original macros can be recovered using the inverse function $f^{-1}(\cdot)$ and a joint prior distribution of macros based on the dietary habits of an individual, a region, or a culture [4] [5].

II. RELATED WORK

Multiple studies have been conducted to understand how different macros affect PPGRs. The main determinant is the

¹This work was supported by the National Science Foundation award IIS 2014475.

¹The abstract of this manuscript has been accepted to the Annual Nutrition Meeting (NUTRITION 2024). The abstract presented contains preliminary results on a subset of the participants in the study. This submission describes all details and includes results on all participants.

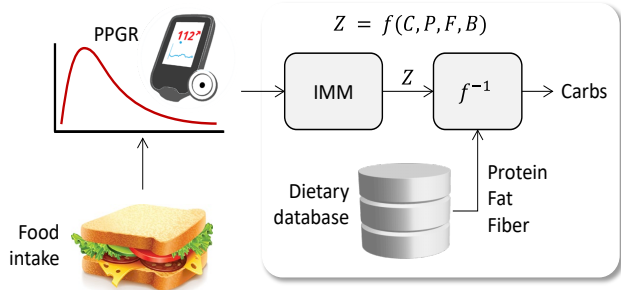


Fig. 1. Overview of the proposed approach. We measure the post-prandial glucose response (PPGR) of a meal with a Continuous Glucose Monitor (CGM), and then pass it to an inverse metabolic model (IMM), which predicts a macronutrient embedding (Z). The embedding is combined with priors from a dietary database to recover the original macros—in this case, carbs.

amount and type of carbs. Carbs are typically compared by their glycemic index (GI) [6], a measure of how they raise glucose levels compared to a reference food (typically glucose) defined as having a GI=100. Most vegetables have low GIs (<55), whereas sugars and starches have high GI (>70). However, the GI does not consider that most meals have other macros as well. Specifically, adding protein, fat or dietary fiber to a meal reduces and/or slows down the PPGR [7], [8], typically due to gastric emptying or insulin secretion ([9]–[11]). To illustrate, [9] conducted a study where participants consumed breakfasts rich in carbs, protein or fat before consuming white bread. The glucose response to the bread was lower after consuming a protein- or fat-rich breakfast than after eating a carb-rich breakfast. Similar studies [10] have found that foods containing protein and fat have lower PPGRs than those only containing glucose.

Recent studies have focused on personalized nutrition by modeling individual differences in food metabolism. As an example, Zeevi et al. [12] used CGMs to track the glucose levels of 800 subjects for one week while they kept detailed records of their diet. Then, the authors developed a direct metabolic model that predicted the PPGR to different meals based on their macros. To account for individual differences, the model could be personalized to each patient’s “phenotype”, e.g., anthropometric features, blood panels and gut microbiota. To validate the model, the authors used an independent group of 100 subjects who consumed two types of meals: those the model predicted would lead to low PPGRs (“good” diet), or the opposite (“bad” diet). Post-prandial glucose excursions when following the “good” diet were significantly lower than those for the “bad” diet. Compared to [12], our work addresses the inverse problem: predicting macros from PPGRs [13]. In a prior study, we had participants consume breakfast shakes with various macros, and then rest for 8 hours. To predict macros from PPGRs, we built a machine-learning (ML) model using eXtreme Gradient Boosting (XGBoost) [14]. We found that subtracting the baseline of each meal (i.e., the first glucose reading) and then normalizing each patients’ table of PPGRs by its mean/variance (z-score) improved predictions substantially, compared to using un-normalized PPGRs. The best model achieved an NRMSE (normalized root-mean-square error) of 22% for carbs, 50% for protein, and 40% for fat.

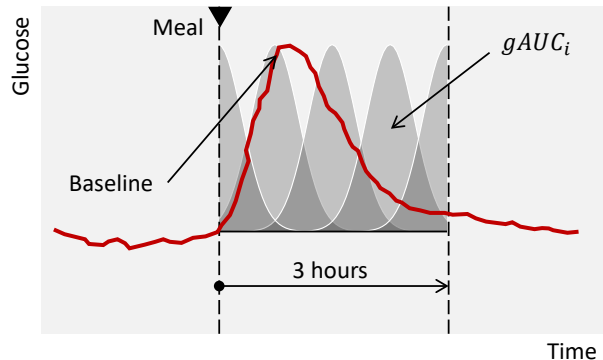


Fig. 2. Extracting temporal features (gAUC) using 5 Gaussian kernels from the post-prandial glucose response (shown in red) using a family of gaussian kernels, (shown in gray).

In a follow-up study [15], we tested if post-prandial responses of other nutritional biomarkers (amino-acids, triglycerides, and insulin), which we had also measured, would further improve predictions. Feeding those additional biomarkers to the model reduced errors significantly for protein (from 50% to 23%) and fat (from 40% to 32%)

III. METHODS

This section describes our procedures to extract features from PPGRs, compress macro amounts into a latent variable (Z), predict Z from PPGRs, and convert Z back to macros.

A. Feature extraction from PPGRs

Following prior work [13], we compress a PPGR into a small number of Gaussian-windowed area-under-the-curve (gAUC) features at various time delays, as illustrated in Fig. 2. First, we align the start of the PPGR with the time of meal intake by resampling the PPGR using linear interpolation. Then, we subtract the baseline (glucose at time $t = 0$), as it correlates with an individual’s metabolic health (e.g., HbA1c). Next, we extract the gAUCs and z-score them, as in [13].

B. Macronutrient encoding

Our embedding is adapted from a model that has been proposed to estimate the GI of mixed meals [8]. The original model separates glycemic macros (carbs) from non-glycemic macros (protein, fat, and fiber), and defines the GI as the ratio of glycemic macros to all macros. To make model coefficients easier to interpret, we use a linear model to compute the latent variable Z :

$$Z = 1 + \frac{C - B}{\max(C - B)} - \alpha_P \frac{P}{\max P} - \alpha_F \frac{F}{\max F} \quad (1)$$

where C, P, F, B are the amount of carbohydrates, protein, fat, and fiber expressed in calories to account for their relative caloric effect², $(C - B)$ are *net* carbs (i.e., carbs minus fiber), and α_P , and α_F describe how much protein and fat (P, F) counteract the effect of net carbs. Finally, we normalize each macro by its maximum amount across all meals, and add an intercept term to minimize scaling effects.

²To compute calories, we multiply C, P, F, B in grams (Table 1) by 4, 4, 9, 2.

C. Inverse metabolic model (IMM)

The IMM uses XGBoost to predict Z from gAUCs. We chose XGBoost, as opposed to other regression models, as it has achieved state-of-the-art results on multiple ML challenges, e.g., text classification, customer behavior prediction, product categorization. XGBoost outperforms other gradient boosting models and runs much faster and scales more easily to new examples [14]. Moreover, our previous study [13] confirmed that XGBoost is a good choice for our problem, when coupled with baseline subtraction and z-scoring. We use a leave-one-subject-out procedure to optimize hyperparameters: number of trees, depth of trees, and regularization term³. Namely, we split data into training (41 subjects), validation (1 subject), and test (1 subject). Using grid search, we train a model on 41 subjects, optimize hyper-parameters on the validation set, and report the final results on the test set. We repeat this process exhaustively for all possible splits.

TABLE I

MACRONUTRIENT COMPOSITION OF BREAKFASTS. THE NUMBERS IN RED INDICATE HIGH AMOUNTS OF MACRONUTRIENTS

Study day	Carb (g)	Protein (g)	Fat (g)	Fiber (g)
1	66	22	11	00
2	66	66	11	00
3	66	22	42	00
4	73	66	42	07
5	24	22	11	00
6	66	22	11	00
7	66	66	11	00
8	66	22	42	00
9	24	22	11	00
10	66	22	42	07

D. Inverting the latent variable back to macros

Once Z has been estimated from gAUCs, we can recover the original macros by inverting eq. (1). As an example, net carbs can be obtained by rearranging eq. (1) as

$$C - B = \max_{C-B} \times (Z - 1 + \alpha_P \frac{P}{\max_P} + \alpha_F \frac{F}{\max_F}) \quad (2)$$

Note that eq. (2) needs an estimate of the amount of protein and fat in the meal. This information can be obtained from a variety of sources, such as an individual’s dietary habits, the menu of a restaurant, or the culinary traditions of a region or a culture [4], [5]. When predictions can be reduced to a few meal choices, it may be advantageous to feed each meal’s (P, F) values to eq. 2, and then compute confidence intervals from the multiple predictions for net carbs.

E. Evaluation

Following prior work [13], we use the normalized root mean squared error (NRMSE) to measure model performance:

$$NRMSE = \sqrt{\frac{1}{N} \sum_{i=0}^N \left(\frac{\hat{y} - y}{y} \right)^2} \quad (3)$$

where \hat{y} and y are the predicted and ground truth amounts of the macros respectively, and N is the number of samples in the test set. Because the NRMSE is a percentage, it makes it easier to compare the prediction errors for different amounts of food and different macros.

³The search space of each hyperparameter is as follows: number of trees: 20, 30, depth of trees: 2, 3, regularization: 1, 2




Fig. 3. Prediction error (NRMSE) for carbohydrates using different embedding coefficients for protein and fat.

IV. EXPERIMENTAL SETUP

In this study, subjects are monitored in free-living conditions for ten days (approved Advarra IRB Protocol Pro00049227). For each study day, subjects are provided two types of prescribed meals with known macros: (1) breakfast shakes, and (2) lunches from a fast-casual restaurant chain (Chipotle Mexican Grill, Inc.). Subjects have free choice of dinners and are not given any dietary restriction other than to avoid food intake for three hours after each prescribed meal. For the ten days, subjects wear an Abbott Freestyle Libre Pro CGM on their upper arm, and a Dexcom G6 Pro CGM in their abdomen.

The data consists of forty five subjects (age: 18-69 years; BMI: 20-46). Based on fasting HbA1c measured at recruitment, 15 subjects can be considered healthy, 16 as having pre-diabetes, and 14 as having type 2 diabetes. We evaluate the proposed model on breakfast shakes using glucose measurements from the Libre CGM. We removed data from two subjects due to sensor failure.

Shown in Table I, the breakfast shakes vary between low (L) and high (H) amounts of carbs, protein, fat, and fiber, according to the average American diet [4]. Some breakfasts are repeated to quantify intra-individual variability to identical meals. Subjects submit before/after photographs of each meal via WhatsApp, from which we extract meal times.

V. RESULTS

We developed XGBoost models for each combination of parameters α_P and α_F in the range (0,1) in increments of 0.1. For each model, we generated ten predictions of net carbs by combining the predicted Z with the protein and fat contents of each of the ten breakfast shakes –see eq. (2), and then averaged the ten predictions. To ensure that NRMSEs are comparable for different α_P and α_F , we report errors on net carbs, the critical macro to monitor in diabetes and pre-diabetes.

Results are summarized in Fig. 3. Several observations can be made. First, NRMSE changes radially from the bottom-left corner, corroborating that the effects of proteins and fat on PPGRs are somewhat interchangeable. Second, the highest NRMSE is obtained at the end of each axis. At the bottom-right ($\alpha_P = 1, \alpha_F = 0$), the model assumes that fat has no effect on PPGRs, and that the effect of carbs can be

cancelled out by adding the same amount of protein (in calories). Likewise, at the top-left ($\alpha_P = 0, \alpha_F = 1$), the model assumes that protein has no effect on PPGRs, and that fat cancels the effect of carbs (calorie-by-calorie). We do not report performance on models with $\alpha_P + \alpha_F > 1$, since Z can become negative for some meals, which is physiologically questionable. Most importantly, the results in Fig. 3 show that predicting Z for combinations around $\alpha_P = \alpha_F = 0.1$ lead to the lowest NRMSE (0.40) for net carbs, whereas predicting net carbs directly ($\alpha_P = \alpha_F = 0$) has an NRMSE of 0.51. A two-tailed t-test reveals that these differences are statistically significant ($p < 0.001$).

In a final step, we compare several approaches to use prior information to recover the net carbs in a meal from eq. (2). Results are shown in Fig. 4. The first bar (blue) shows the NRMSE if net carbs are predicted directly without first learning Z . The second bar (green) is the NRMSE when estimating net carbs for each of the ten meals, and then compute an average (as in Fig. 3). The third bar (red) shows the best-case scenario, when net carbs are predicted using ground truth amounts of protein and fat as the prior, which leads to the lowest NRMSE (0.37). A two-tailed t-test reveals that differences between net carbs predicted using ground truth protein and fats and average estimates of net carbs are statistically significant ($p < 0.01$).

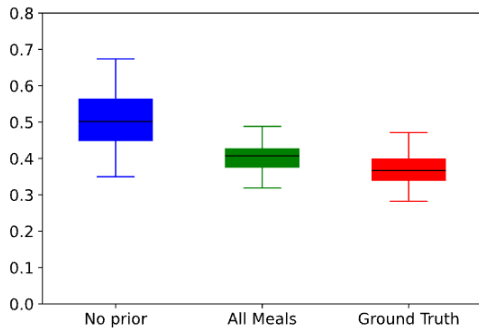


Fig. 4. NRMSE when predicting net carbs directly, as the average prediction for multiple values of protein and fat, and with ground truth amounts of protein and fat.

VI. DISCUSSION

Our results indicate that predicting the amount of net carbs in a meal indirectly –through a latent variable Z that considers the protein and fat in a meal, is superior to predicting net carbs directly. In particular, the lowest least error occurs in the radial neighborhood of $\alpha_P = \alpha_F = 0.1$, which can be viewed as an estimate of the effectiveness of non-glycemic macros (protein, fat) in reducing the post-prandial glucose response of carb-rich meals. Though the best predictions for net carbs are obtained when the ground truth amounts of protein and fat are known, considering the distribution of protein and fat across all meals also improves predictions significantly.

Future work will examine alternative embeddings, such as non-linear models that assume non-glycemic macros reduce post-prandial responses in a multiplicative rather than additive manner. Future work will also examine approaches to learn the latent variable in a data-driven fashion (through ML models), as opposed to manually engineering the embeddings.

As an example, auto-encoder (AE) network may be used to reconstruct PPGRs. Auxiliary classifiers may be added to further reduce individual differences, e.g., applying gradient reversal on an auxiliary classifier that attempts to the patient’s HbA1c from the AE bottleneck.

REFERENCES

- [1] L. E. Burke, V. Swigart, M. Warziski Turk, N. Derro, and L. J. Ewing, “Experiences of self-monitoring: successes and struggles during treatment for weight loss,” *Qualitative health research*, vol. 19, no. 6, pp. 815–828, 2009.
- [2] E. K. Choe, S. Abdullah, M. Rabbi, E. Thomaz, D. A. Epstein, F. Cordeiro, M. Kay, G. D. Abowd, T. Choudhury, J. Fogarty *et al.*, “Semi-automated tracking: a balanced approach for self-monitoring applications,” *IEEE Pervasive Computing*, vol. 16, no. 1, pp. 74–84, 2017.
- [3] F. Cordeiro, D. A. Epstein, E. Thomaz, E. Bales, A. K. Jagannathan, G. D. Abowd, and J. Fogarty, “Barriers and negative nudges: Exploring challenges in food journaling,” in *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, 2015, pp. 1159–1162.
- [4] 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.
- [5] B. Karamanos, A. Thanopoulou, F. Angelico, S. Assaad-Khalil, A. Barbato, M. Del Ben, V. Dimitrijevic-Sreckovic, P. Djordjevic, C. Gallotti, N. Katsilambros *et al.*, “Nutritional habits in the mediterranean basin. the macronutrient composition of diet and its relation with the traditional mediterranean diet. multi-centre study of the mediterranean group for the study of diabetes (mgsd),” *European Journal of Clinical Nutrition*, vol. 56, no. 10, pp. 983–991, 2002.
- [6] D. J. Jenkins, T. Wolever, R. H. Taylor, H. Barker, H. Fielden, J. M. Baldwin, A. C. Bowling, H. C. Newman, A. L. Jenkins, and D. V. Goff, “Glycemic index of foods: a physiological basis for carbohydrate exchange,” *The American journal of clinical nutrition*, vol. 34, no. 3, pp. 362–366, 1981.
- [7] T. M. Wolever and C. Bolognesi, “Prediction of glucose and insulin responses of normal subjects after consuming mixed meals varying in energy, protein, fat, carbohydrate and glycemic index,” *The Journal of nutrition*, vol. 126, no. 11, pp. 2807–2812, 1996.
- [8] 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.
- [9] H. Meng, N. R. Matthan, L. M. Ausman, and A. H. Lichtenstein, “Effect of macronutrients and fiber on postprandial glycemic responses and meal glycemic index and glycemic load value determinations,” *The American journal of clinical nutrition*, vol. 105, no. 4, pp. 842–853, 2017.
- [10] 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,” *The Journal of nutrition*, vol. 136, no. 10, pp. 2506–2511, 2006.
- [11] X. Lan-Pidhainy and T. M. Wolever, “The hypoglycemic effect of fat and protein is not attenuated by insulin resistance,” *The American journal of clinical nutrition*, vol. 91, no. 1, pp. 98–105, 2010.
- [12] 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.
- [13] S. Sajjadi, A. Das, R. Gutierrez-Osuna, T. Chaspari, P. Paromita, L. E. Ruebush, N. E. Deutz, and B. J. Mortazavi, “Towards the development of subject-independent inverse metabolic models,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 3970–3974.
- [14] T. Chen, T. He, M. Benesty, V. Khotilovich, Y. Tang, H. Cho, K. Chen, R. Mitchell, I. Cano, T. Zhou *et al.*, “Xgboost: extreme gradient boosting,” *R package version 0.4-2*, vol. 1, no. 4, pp. 1–4, 2015.
- [15] A. Das, B. Mortazavi, S. Sajjadi, T. Chaspari, L. E. Ruebush, N. E. Deutz, G. L. Cote, and R. Gutierrez-Osuna, “Predicting the macronutrient composition of mixed meals from dietary biomarkers in blood,” *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 6, pp. 2726–2736, 2021.