

A SPARSE CODING APPROACH TO AUTOMATIC DIET MONITORING WITH CONTINUOUS GLUCOSE MONITORS

Anurag Das¹, Seyedhooman Sajjadi¹, Bobak Mortazavi¹, Theodora Chaspari¹, Projna Paromita¹, Laura Ruebush², Nicolaas Deutz² and Ricardo Gutierrez-Osuna¹

¹Department of Computer Science and Engineering, Texas A&M University

²Center for Translational Research in Aging and Longevity, Texas A&M University
{adas, bobakm, chaspari, hooman130, projna.paromita, ltruebush, nep.deutz, rgutier}@tamu.edu

ABSTRACT

Measuring dietary intake is a major challenge in the management of chronic diseases. Current methods rely on self-report measures, which are cumbersome to obtain and often unreliable. This article presents an approach to estimate dietary intake automatically by analyzing the post-prandial glucose response (PPGR) of a meal, as measured with continuous glucose monitors. In particular, we propose a sparse-coding technique that can be used to estimate the amounts of macronutrients (carbohydrates, protein, fat) in a meal from the meal's PPGR. We use Lasso regularization to represent the PPGR of a new meal as a sparse combination of PPGRs in a dictionary, then combine the sparse weights with the macronutrient amounts in the dictionary's meals to estimate the macronutrients in the new meal. We evaluate the approach on a dataset containing nine standardized meals and their corresponding PPGRs, consumed by fifteen participants. The proposed technique consistently outperforms two baseline systems based on ridge regression and nearest-neighbors, in terms of correlation and normalized root mean square error of the predictions.

Index Terms— Continuous glucose monitors, macronutrient prediction, sparse coding.

1. INTRODUCTION

Diet plays a major role in the development of chronic diseases, such as type 2 diabetes, obesity and heart disease [1]. Thus, monitoring diet is an essential component of many clinical interventions. Unfortunately, conventional methods for tracking dietary intake rely on self-report tools, such as food diaries and 24-hour recall, which are problematic. Food diaries require manual input, which is tedious and often leads to low adherence rates. Further, 24-hour records suffer from memory recall, which can lead to severe over and under-reporting [2]. To address this issue, several technologies have been explored to facilitate diet monitoring. For example, wearable sensors have been used to detect eating behaviors [3], and computer vision has been used to estimate nutritional content from food photographs [4, 5]. These methods can reduce burden to the user, but at present they are inaccurate [6]. As such, accurately measuring dietary intake remains a major challenge in dietary research [2].

This article examines a potential solution to this problem based on continuous glucose monitors (CGMs) and machine learning. A

This work was supported by funding from the NSF (Award # 2014475) and Engineering Research Center for Precise Advanced Technologies and Health Systems for Underserved Populations (PATHS-UP) (Award # 1648451), and a grant from the Kleberg Foundation.

CGM consists of a small electrode inserted in the skin, which measures glucose in the interstitial fluid. The use of CGMs has traditionally been limited to type 1 diabetes or poorly-controlled type 2 diabetes, but CGMs are at an inflection point due to cost reductions and improvements in accuracy [7], so their use in other applications, such as obesity and pre-diabetes, is promising. The mechanism by which CGMs may be used to monitor diet is based on the fact that the evolution of blood glucose after a meal, also known as the post-prandial glucose response (PPGR), depends on the macronutrients in the meal (e.g., carbohydrates, protein, fat, fiber). The major determinant of PPGRs is the amount of carbohydrates, but adding protein, fat, or fiber to a meal generally yields smaller spikes and lengthier responses [8, 9]. This suggest that the shape of the PPGR can be used to recover the macronutrient composition of the meal.

To test this hypothesis, we propose a sparse coding technique to predict macronutrients by analyzing PPGRs. The technique uses a dictionary of meals, each meal defined by its macronutrients and the corresponding PPGR. Given the PPGR of a new meal, the technique uses the PPGR dictionary to generate a sparse code via Lasso regularization [10], and then combines the sparse weights with the known macronutrients of the meals in the library. We validate the approach on a database consisting of PPGRs for fifteen individuals who consumed nine different meals of known macronutrients. The proposed approach outperforms two baseline techniques (ridge regression, nearest neighbors) in both subject-dependent and subject-independent scenarios.

2. RELATED WORK

Hundreds of studies have been conducted over the past 50 years to understand the effect of various macronutrients on PPGRs. The main determinant of PPGRs is the amount and type of carbohydrates. Carbohydrates are generally described by their glycemic index (GI) [11], which measures the potential for a carbohydrate to raise glucose levels compared to a reference food (typically glucose, defined as having a GI=100). However, the GI is controversial because it does not account for the influence of other macronutrients that are typically present in mixed meals.

Recent work has focused on individual differences in food metabolism. As an example, Zeevi et al. [12] tracked 800 participants for a week using CGMs, and found high inter-personal variability in the glucose response to identical meals. To address this issue, the authors developed a machine-learning model that could predict the PPGR of a meal for each participant by accounting for individual factors (e.g., anthropometric variables, gut microbiota). Also recently, Tily et al. [13] tracked 550 adults for two weeks with CGMs, while they consumed a set of standardized meals

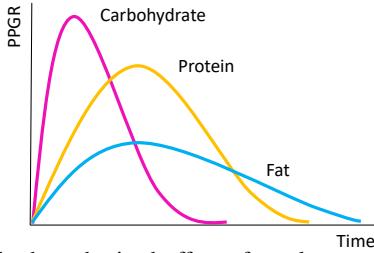


Fig. 1. The hypothesized effect of meal macronutrients on post-prandial glucose

designed to cover a broad range of carbohydrates, proteins, fats and fiber. Then, they built a multilevel mixed-effects regression model to predict PPGRs. This allowed the authors to quantify the influence of meal composition, anthropometric, gut microbiome and lifestyle variables on PPGRs.

Much of the signal-processing work for glucose analytics has focused on time-series forecasting. This is an important problem, since the ability to predict future glucose levels based on past measurements can lead to a number of therapeutic solutions, e.g., warning patients of future hypoglycemic events, delivering insulin automatically, i.e., the artificial pancreas [14]. A vast pool of techniques have been used for this purpose, including autoregressive models, neural networks, random forests, Kalman filters and support vector regression, to mention a few [15]. These models may be trained to predict the actual glucose level at a future time or to generate hypoglycemia alerts [16].

3. METHODS

Figure 1 illustrates the prototypical PPGR to a meal. Depending on the amount of carbohydrates, blood glucose starts to rise 15-30 minutes after the meal, reaches a peak within the first 1-2 hours, and returns to pre-prandial levels within 3-4 hours [17]. Illustrated in the figure is also the hypothesized effect of adding protein and fat to a meal.

Let us denote by $x_{ms}(t)$ the post-prandial glucose level of subject s time t units after consuming meal m , and by $z_{ms}(i)$ the amount of the i -th macronutrient in the meal, where $i \in \{C, P, F\}$. The PPGR can then be characterized by a vector $X_{ms} = \{x_{ms}(1), \dots, x_{ms}(T)\}$, where T is the length of the recording. Consider a dictionary X containing the PPGRs of S subjects after consuming M different meals, each meal with its corresponding macronutrient stored in Z . Assume that a new subject s' has consumed meal m' , which has resulted in a PPGR defined by the vector $X_{m's'}$. Our approach starts by representing this new meal as a linear combination of the PPGRs in the dataset X :

$$X_{m's'} = \sum_s^S \sum_m^M w_{ms} X_{ms} \quad (1)$$

where w_{ms} are the weights learnt by the model. Further assume that the distance between meals in PPGR space is correlated with the distance between them in macronutrient space; that is, for any two meals i and j , we assume $\|X_i - X_j\|^2 \sim \|Z_i - Z_j\|^2$. From this, it follows that the macronutrient composition of the test meal $Z_{m's'}$ can be approximated as a linear combination of the macronutrients in the training meals, using the linear weights in equation (1).

$$Z_{m's'} = \sum_s^S \sum_m^M w_{ms} Z_{ms} \quad (2)$$

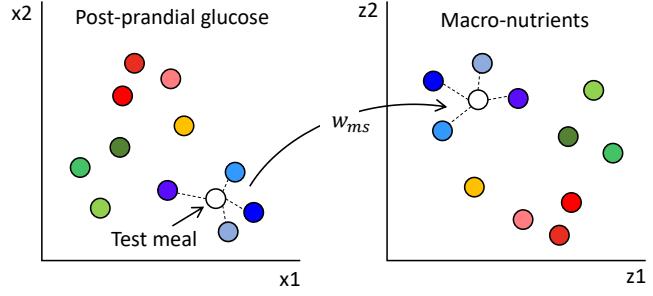


Fig. 2. Illustration of the sparse-coding approach

Table 1. Macronutrient composition of nine meals

Meal	Carbs (g)	Protein (g)	Fat (ml)
C1P1F1	52.25	15	13
C1P2F2	52.25	30	26
C2P2F1	94.75	30	13
C2P1F2	94.75	15	26
C2P2F2	94.75	30	26
C2P2F3	94.75	30	52
C2P3F2	94.75	60	26
C3P2F2	179.75	30	26
C3P3F3	179.75	60	52

The approach preserves local neighbourhoods as illustrated in Figure 2. The plot on the left shows the meals that are close to the test meals in the PPGR space. The same meals are still close to the test meal in the macronutrient space on the right, even though the position of the meals changes. Similar techniques, known as exemplar-based, have been used with success in the voice-conversion literature [18, 19], where the feature vectors X_{ms} and Z_{ms} represent time-aligned vectors of acoustic features (e.g., MFCCs) from a source and a target speaker, respectively. Work in exemplar-based methods has shown that improved reconstruction can be achieved by forcing most of the weights to become zero. Accordingly, we enforce a sparse non-negativity constraint on the weights w_{ms} by solving the minimization problem:

$$w_{ms} = \operatorname{argmin} \|X_{m's'} - w_{ms} X_{ms}\|^2 + \lambda w_{ms} s.t. w_{ms} \geq 0 \quad (3)$$

where λ prevents overfitting by penalizing the L1 norm of the weights, w_{ms} . We use Lasso regularization to learn these weights. It is possible that one of the weights learnt using Lasso becomes very large which may cause the predicted macronutrients to become unbounded. We therefore bound the predicted macronutrients using a normalization step that forces the weights to add up to one before making a prediction with equation (2).

$$\hat{w}_{ms} = \frac{w_{ms}}{\sum w_{ms}} \quad (4)$$

3.1. Baseline models and performance measures

We evaluate the sparse-coding approach against two baseline techniques: (1) ridge regression (RR), as a representative of regularization methods, and (2) a Nearest-Neighbor classifier operating in a Linear Discriminant Analysis subspace (LDA-kNN), as a representative of distance-based classifiers. We do not consider more complex techniques due to the small size of the dataset. The RR baseline

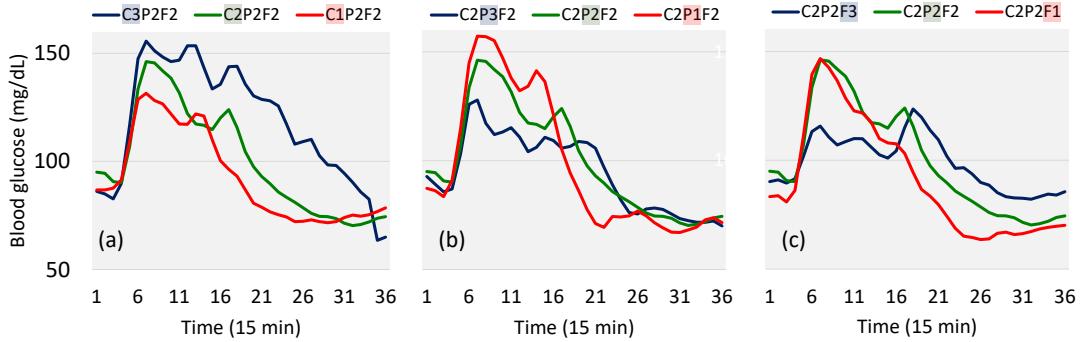


Fig. 3. (a) Glucose response at increasing levels of C, with P and F at fixed levels. (b) Glucose response at increasing levels of P, with C and F at fixed levels. (c) Glucose response at increasing levels of F, with C and P at fixed levels.

consists of three separate models, one for each of the three macronutrients. In contrast, LDA-kNN is a joint model for the three macronutrients. Namely, we use LDA to find an 8-dim projection (i.e., the number of classes minus one) of the PPGRs that maximizes separability among the 9 meals; then, we use the kNN rule to classify a new meal in the LDA subspace. Note that, in contrast with the two baselines, the sparse-coding technique is unsupervised since it does not use class information to learn the weights.

We evaluate the models using two measures: (1) the Pearson correlation between ground truth and predicted amounts of macronutrients¹, and (2) the normalized root mean squared error, $NRMSE = \sqrt{\frac{1}{N} \sum (\hat{y} - y)^2 / y^2}$ where \hat{y} and y are the predicted and ground truth amount of macronutrient, respectively, and N is the number of samples.

4. EXPERIMENTAL SETUP

4.1. Dataset

To test the proposed model, we recruited 15 healthy subjects ages 60–85 years and Body Mass Index(BMI) of 25–35. Each subject participated in 9 study days in which they consumed a predefined meal in a randomized design. Subjects were asked to be overnight fasted prior to the meal intake on each study day so that the first blood glucose reading would be their fasting glucose level. After taking a baseline blood sample the morning of a study visit, each subject consumed a predefined meal. Subjects remained in a sedentary state and were not allowed to consume any other food for the next 8 hours. This study was approved by the Texas A&M Institutional Review Board (IRB #2017-0886). Table 1 shows the composition of the nine meals. Each meal had a known amount of the same carbohydrates (C: Maltodextrine), protein (P: Whey Protein), and fat (F: Sunflower Oil), which we denote as $CxPxFx$, where x represents the amount of each macronutrient in the meal (1: low; 2: medium; 3: high). For the duration of the study, participants wore a CGM (Abbott Freestyle Libre Pro), which measured interstitial glucose every 15 minutes. We perform our analysis using the first 32 PPGR readings (8 hours) from the time the meal was consumed. To account for differences in fasting glucose levels across participants, we subtracted the baseline glucose of each PPGR prior to performing the sparse reconstruction.

Figure 3(a) shows the average response across subjects as we increase the amount of carbohydrates (C1, C2, C3) while maintaining

the other two macronutrients at a fixed level (P2, F2). The PPGR becomes more pronounced at higher levels of carbohydrates, both in terms of the maximum value and the overall area under the curve. Figure 3(b) shows the average response across subjects as we increase the amount of protein (P1, P2, P3) while maintaining the other two macronutrients fixed (C2, F2). As we increase the amount of protein, the glucose response becomes more moderate, with lower maximum levels and slower return to the baseline. Figure 3(c) shows a similar effect as we increase the amount of fat (F1, F2, F3) while maintaining the other two macronutrients fixed (C2, P2). These results provide support to our overall strategy, as they show that the shape of the glucose response depends on the constituents of the meal.

4.2. Hyperparameter selection

We used a two-step cross-validation procedure to optimize model hyperparameters: a regularization term λ for the sparse coding method and RR, and the number of neighbors k for LDA-kNN. Namely, given a test subject and a validation subject (different from the test subject), we train a model on the remaining set of 13 subjects for all possible hyperparameter values². Then, we evaluate the model on the validation subject and record the NRMSE. Keeping the test subject fixed, we repeat this procedure until all combinations of validation and training subjects are covered. Finally, we evaluate the test subject using the hyperparameter value that yields the lowest average NRMSE across all validation subjects. In this way, the test set is never used to optimize the models' hyperparameters.

5. RESULTS

In a first step, we evaluate the sparse-coding model using the leave-one-subject-out (LOSO) procedure outlined above, where we train the models on 14 subjects and test on the held-out subject. We repeat this procedure until all subjects are tested. Because the models are trained on data from multiple subjects, they can be thought of being *subject-independent*. Results for the sparsity parameter λ are illustrated in Figure 4. The lowest error is achieved for values of $\lambda = 5, 15$ and 10 , for C, P and F, respectively, which confirms that there is an advantage to imposing a sparsity constraint on the representation. Note that, for $\lambda = 0$, the model becomes equivalent to least-squares regression.

¹We calculate the correlation by collecting the predicted and ground truth macronutrients for all test subjects and meals and then calculate a single value (i.e., a pooled correlation)

²The hyperparameter ranges are (1) sparse reconstruction: $\lambda = \{0.005, 0.5, 1, 1.5, 2, 5, 10, 15, 25, 50, 100\}$, (2) ridge regression: $\lambda = \{0.05, 0.5, 1, 1.5, 2, 5, 10, 15, 25, 50, 100\}$ (3) nearest neighbors: $k = \{1, 3, 5, 7, 9\}$.

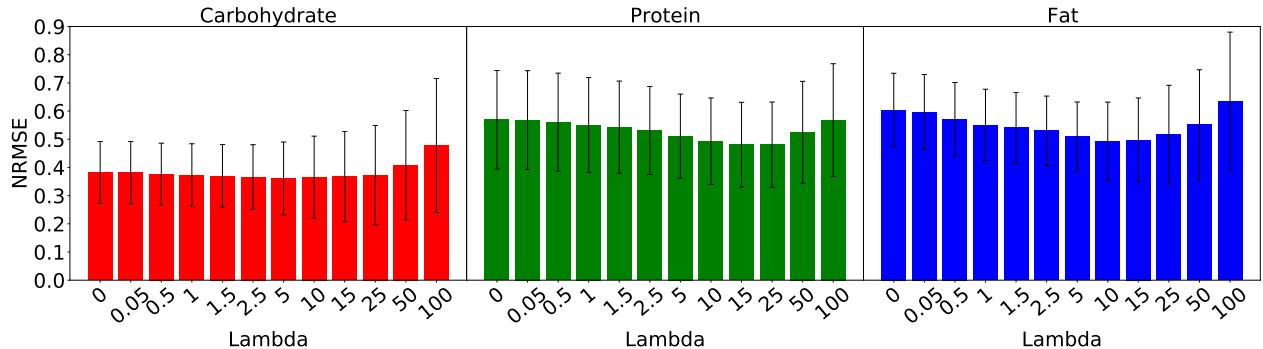


Fig. 4. NRMSE for the sparse reconstruction model, as a function of the regularization parameter λ

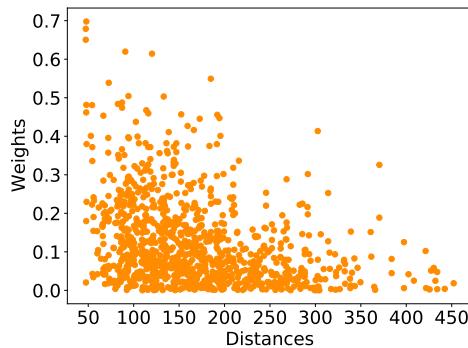


Fig. 5. Distribution of non-zero weights as a function of the distance between a test meal and those in the dictionary.

Further, we examined the distribution of sparse weights as a function of the distance $\|X_i - X_j\|$ between a PPGR in the test set and those in the dictionary. Results are shown in Figure 5. Meals with similar PPGRs (i.e., smaller distances) tend to receive higher weights. In other words, as illustrated in Figure 2, the sparse representation is *locally linear*.

Next, we compared the performance of the approach against the two baselines. Results are summarized in Table 2. Sparse coding consistently outperforms the two baselines, for all three macronutrients, achieving higher correlation and lower NRMSE. Further, predictions for carbohydrates are more accurate than for fats and proteins. This is an intuitive result since the largest determinant of post-prandial glucose is the amount of carbohydrates in the meal.

In a final analysis, we evaluate the sparse representation model when only meals within each subject are used. This analysis is motivated by the fact that there exist large individual differences in food metabolism; see section 2.1. Thus, it is possible that forcing the model to only use data from each subject could lead to higher accuracy. To answer this question, we use a leave-one-meal-out (LOMO) procedure, as follows. For each participant, we build a model using 8 of the meals as training data and the remaining meal as the test sample. We repeat this process until all 9 meals are tested for each participant and all participants are covered. As such, this can be thought of as a *subject-dependent model*. Results are shown in Table 3. Both models perform comparably in the case of carbohydrates, but performance degrades markedly for the subject-dependent model in the case of protein and fat. This result suggests that, despite large inter-individual differences in food metabolism, the sparse representation model can take advantage of having training data from multiple par-

Table 2. Pooled correlation and NRMSE using LOSO

Method	Pooled Correlation			Average NRMSE		
	C	P	F	C	P	F
Proposed	0.49***	0.28**	0.39***	0.37	0.49	0.48
RR	0.39***	0.12	0.24**	0.45	0.71	0.60
LDA-kNN	0.36***	0.05	0.28**	0.48	0.64	0.67

***: $p < 0.001$, **: $0.001 < p < 0.05$, *: $0.05 < p < 0.1$

Table 3. Pooled correlation and NRMSE of the predicted macronutrients using LOSO and LOMO cross-validation procedures

Method	Pooled Correlation			Average NRMSE		
	C	P	F	C	P	F
LOSO	0.49***	0.28*	0.39***	0.37	0.49	0.48
LOMO	0.50***	0.07	0.24**	0.41	0.70	0.64

***: $p < 0.001$, **: $0.001 < p < 0.05$, *: $0.05 < p < 0.1$

icipants.

6. DISCUSSION

We have proposed a sparse-coding approach³ to predict meal macronutrients based on PPGRs. First, we represent the PPGR of a new meal as a sparse combination of other meals in a dictionary. Then, we “transfer” the sparse weights onto the macronutrients of the selected meals. On a subject-independent task, we find that the sparse method clearly outperforms two supervised techniques, despite the fact that the sparse weights are selected without using class label information. We also find that the sparse method predicts carbohydrates more accurately than protein and fat, and that it performs better on a subject-independent task (where it can leverage meals from multiple participants) than on a subject-dependent task (where it is constrained to only use meals from the test subject). The improved prediction suggests that higher accuracy may be obtained by having access to a larger dataset containing PPGRs from a large number of participants.

Our approach implicitly assumes that the distance between pairs of meals in the PPGR space is correlated with their distance in the macronutrient space. This suggests a potential improvement where the model would first learn to transform the PPGR space so that it captures pair-wise distances in the macronutrient space, as in multi-dimensional scaling. Improved performance may also be achieved by augmenting PPGRs with physiological variables that impact metabolism, such as body composition. In this way, the model could learn to focus on meals from participants with similar metabolism as the test subject.

³Our implementation of the sparse coding approach is available at: <https://github.com/dasanurag/CGM-Sparse/tree/master>

7. REFERENCES

- [1] Marian L Neuhouser, “The importance of healthy dietary patterns in chronic disease prevention,” *Nutrition Research*, vol. 70, pp. 3–6, 2019.
- [2] Brooke M Bell, Ridwan Alam, Nabil Alshurafa, Edison Thomaz, Abu S Mondol, Kayla de la Haye, John A Stankovic, John Lach, and Donna Spruijt-Metz, “Automatic, wearable-based, in-field eating detection approaches for public health research: a scoping review,” *NPJ digital medicine*, vol. 3, no. 1, pp. 1–14, 2020.
- [3] Haik Kalantarian, Nabil Alshurafa, Tuan Le, and Majid Sarrafzadeh, “Monitoring eating habits using a piezoelectric sensor-based necklace,” *Computers in biology and medicine*, vol. 58, pp. 46–55, 2015.
- [4] Shaobo Fang, Zeman Shao, Deborah A Kerr, Carol J Boushey, and Fengqing Zhu, “An end-to-end image-based automatic food energy estimation technique based on learned energy distribution images: Protocol and methodology,” *Nutrients*, vol. 11, no. 4, pp. 877, 2019.
- [5] Haik Kalantarian, Nabil Alshurafa, and Majid Sarrafzadeh, “A survey of diet monitoring technology,” *IEEE Pervasive Computing*, vol. 16, no. 1, pp. 57–65, 2017.
- [6] Hamid Hassannejad, Guido Matrella, Paolo Ciampolini, Ilaria De Munari, Monica Mordinini, and Stefano Cagnoni, “Automatic diet monitoring: a review of computer vision and wearable sensor-based methods,” *International journal of food sciences and nutrition*, vol. 68, no. 6, pp. 656–670, 2017.
- [7] David Rodbard, “Continuous glucose monitoring: a review of successes, challenges, and opportunities,” *Diabetes technology & therapeutics*, vol. 18, no. S2, pp. S2–3, 2016.
- [8] Thomas MS Wolever and Claudia 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.
- [9] Andreas Rytz, Dorothée Adeline, Kim-Anne Lê, Denise Tan, Lisa Lamothe, Olivier Roger, and Katherine 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, pp. 1172, 2019.
- [10] Robert Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [11] DJ Jenkins, TM Wolever, Rodney H Taylor, Helen Barker, Hashmein Fielden, Janet M Baldwin, Allen C Bowling, Hillary C Newman, Alexandra L Jenkins, and David 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.
- [12] David Zeevi, Tal Korem, Niv Zmora, David Israeli, Daphna Rothschild, Adina Weinberger, Orly Ben-Yacov, Dar Lador, Tali Avnit-Sagi, Maya Lotan-Pompan, et al., “Personalized nutrition by prediction of glycemic responses,” *Cell*, vol. 163, no. 5, pp. 1079–1094, 2015.
- [13] Hal Tily, Ally Perlina, Eric Patridge, Stephanie Gline, Matvey Genkin, Vishakh Gopu, Haely Lindau, Alisson Sjue, Niels Klitgord, Momchilo Vuyisich, et al., “Gut microbiome activity contributes to individual variation in glycemic response in adults,” *bioRxiv*, p. 641019, 2019.
- [14] Klaus Donsa, Stephan Spat, Peter Beck, Thomas R Pieber, and Andreas Holzinger, “Towards personalization of diabetes therapy using computerized decision support and machine learning: some open problems and challenges,” in *Smart Health*, pp. 237–260. Springer, 2015.
- [15] Silvia Oviedo, Josep Vehí, Remei Calm, and Joaquim Armengol, “A review of personalized blood glucose prediction strategies for t1dm patients,” *International journal for numerical methods in biomedical engineering*, vol. 33, no. 6, pp. e2833, 2017.
- [16] Ashenafi Zebene Woldaregay, Eirik Årsand, Taxiarchis Botassis, David Albers, Lena Mamykina, and Gunnar Hartvigsen, “Data-driven blood glucose pattern classification and anomalies detection: machine-learning applications in type 1 diabetes,” *Journal of medical Internet research*, vol. 21, no. 5, pp. e11030, 2019.
- [17] Guido Freckmann, Sven Hagenlocher, Annette Baumstark, Nina Jendrike, Ralph C Gillen, Katja Rössner, and Cornelia Haug, “Continuous glucose profiles in healthy subjects under everyday life conditions and after different meals,” *Journal of diabetes science and technology*, vol. 1, no. 5, pp. 695–703, 2007.
- [18] Christopher Liberatore, Sandesh Aryal, Zelun Wang, Seth Polley, and Ricardo Gutierrez-Osuna, “Sabr: sparse, anchor-based representation of the speech signal,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [19] Zhizheng Wu, Tuomas Virtanen, Tomi Kinnunen, Eng Siong Chng, and Haizhou Li, “Exemplar-based voice conversion using non-negative spectrogram deconvolution,” in *Eighth ISCA Workshop on Speech Synthesis*, 2013.