

Personalized and Culturally-Tailored Meal Planning Recommendation for Clinical Nutrition

DongHyeon Seo

SEODONGH@USC.EDU

Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292, USA

Susan Kim

SUSAN.KIM2@USC.EDU

Keck School of Medicine, University of Southern California, 1500 San Pablo St, Los Angeles, CA 90033, USA

Keith Burghardt

KEITHAB@ISI.EDU

Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292, USA

Abigail Horn

HORNABIG@USC.EDU

Department of Industrial and Systems Engineering, University of Southern California, 3715 McClintock Ave, Los Angeles, CA 90089, USA

Abstract

Meal prescription interventions are essential for managing patients' dietary needs, yet existing approaches either require manual meal planning or rely on generic apps that lack cultural customization. We introduce HumbleNutri, a meal prescription plan recommender system that generates personalized, culturally tailored meal plans for patients with specific dietary needs under the guidance of a Registered Dietitian Nutritionist (RDN). HumbleNutri begins with a semi-supervised learning step to categorize recipes by meal type and cuisine, enabling culturally informed Medical Nutrition Therapy (MNT) recommendations. The system employs a modular framework that combines collaborative filtering-based recommenders with a bundle optimization model with constraints, suggesting meals that are aligned with patient preferences and MNT guidelines while ensuring that meal combinations satisfy patient-specific nutritional requirements based on their clinical profiles. Meals are organized into daily bundles (breakfast, lunch, dinner) and sequenced into weekly plans that support practical preparation and adherence to MNT targets. HumbleNutri translates clinical diet guidelines into culturally relevant meal plans, offering an equitable platform to deliver precision nutrition with an open-source toolkit and web application.

Keywords: Precision Nutrition, Expert-Guided Recommender Systems, Linear Programming, Semi-Supervised Learning

Data and Code Availability This work uses a publicly available dataset, HUMMUS (Bölz et al., 2023), to recommend recipes for clinical nutrition. Code for our methods and experiments is accessible at our shared code repository¹, together with a working web application².

Institutional Review Board (IRB) This work does not require IRB approval.

1. Introduction

Meal prescriptions are an emerging form of dietary intervention for precision nutrition that provides patients with tailored meals aligned with clinical guidelines (Rodgers and Collins, 2020). They are increasingly integrated into clinical care and preventive medicine, particularly for patients managing chronic conditions such as type 2 diabetes or recovering from organ transplantation (Zeltzer et al., 2015). By offering concrete recommendations on what foods to eat and when, meal prescriptions translate dietary goals into actionable plans, supporting adherence, and enabling patients to meet clinical nutrition goals. Recent clinical trials show that these interventions outperform traditional single-nutrient interventions in improving diet quality and treating diet-related diseases (Chen et al., 2022; Cyrino et al., 2021).

1. <https://github.com/HumbleNutri/HumbleNutri>
2. <https://humblenutri.streamlit.app>

However, practical implementation in clinical settings remains a bottleneck. Current approaches either rely on Registered Dietitian Nutritionist (RDN) manually designing meal plans (a burdensome and unscalable approach), or employing off-the-shelf meal planning applications, which fail to support cultural dietary diversity. While these apps are generally able to account for dietary constraints, they often center on Western-normative diets, and their approaches do not offer customization to other sociocultural dietary patterns. This represents a major gap, given that the populations typically targeted in meal prescription interventions are of racial/ethnic minorities (Joo and Liu, 2021), limiting dietary adherence in the populations most at risk. This is especially concerning given the evidence that dietary changes are sustainably maintained and lead to nutritional benefits, only if the intervention works around individuals’ existing diets (Jinnette et al., 2021; Chapman, 2010).

In this paper, we introduce a culturally-tailored meal prescription plan recommender system, HUMBLE-NUTRI: Healthy and cUlturally-tailored Meal BundLE recommendation with NUTRItionist guidance, designed for nutritionists planning an individualized meal prescription program. Designed together with an RDN, the system build structured daily meal bundle recommendations of 3 meals a day that meet a specific set of nutritional guidelines for an individual patient based on their clinical profile. The major contribution of this work is the introduction of an integrated modular framework that generates meal prescriptions that jointly address clinical nutritional requirements, culturally relevant dietary practices, and patient food preferences in real-world settings.

2. Related Work

Prior work on food recommender systems is typically not designed to meet explicit nutritional guidelines, but has prioritized rating prediction or preference matching (Freyne and Berkovsky, 2010). Meanwhile, a few studies used optimization methods to develop meal plans that meet nutritional targets established by dietary guidelines or personalized nutritional profiles, without integrating patient preference information (Masset et al., 2009; Miow et al., 2025).

Culturally-matched recommendations remain underexplored in this space as well. Recent work in cultural matching models has focused primarily on general domains such as music or news (Casillo

et al., 2023), with limited application in food systems. Some recent dietary recommendation frameworks account for user taste preferences or dietary restrictions (Yang et al., 2017), but do not explicitly model or incorporate sociocultural dietary patterns, which are crucial for real-world adaptation in diverse clinical populations.

Our work bridges these gaps by integrating cultural matching recommendations for each patient, structured meal planning with nutritional optimization, and human expert guidance. Unlike prior systems, HumbleNutri is the first end-to-end system, to our knowledge, that generates culturally relevant weekly meal bundles tailored to individual dietary preferences and clinical needs, integrating personalization and nutritionist guidance in a unified pipeline.

3. Methods

Overview. HumbleNutri is a multimodule system that generates culturally tailored and clinically personalized meal recommendations for patients, guided by clinical nutrition principles and cultural relevance. The system works through the integration of four modules: (1) a semi-supervised learning method to infer cuisine and meal type labels at scale to identify culturally relevant recipes; (2) collaborative filtering to generate candidate recipes via meal type-specific recommenders, (3) a “recipe alignment” step to reflect nutritionist-specified constraints on ingredients and preparation methods, and (4) a structured optimization framework to generate personalized weekly meal bundles. Figure 1 illustrates an overview of the HumbleNutri system with the recommended weekly meal bundle structure. Following the guidance of the RDN, these modules are integrated to produce weekly plans consisting of daily meal bundles structured into breakfast, lunch with a main dish and 1 side dish (vegetables), and dinner with a main dish and 2 side dishes (vegetables and whole grains).

We build on the HUMMUS dataset (Bölz et al., 2023), a recently released large-scale recipe dataset with user ratings and nutrient information curated for food recommendation tasks. Although this dataset provides a rich foundation for recommender systems, it lacks structured labels essential for downstream meal planning, such as meal type (e.g., breakfast, lunch, dinner) and cuisine type (e.g., Latin American). Hence, we extend the HUMMUS dataset by generating synthetic semantic labels using a semi-supervised learning approach.




Figure 1: Illustration of the overview of HumbleNutri.

Labeling Module. We utilize the subset of the HUMMUS dataset, $\sim 45\%$ of which includes human-labeled tags such as meal type, ingredients, and cuisine (e.g., ‘breakfast’, ‘wheat’, ‘ethiopian’). These 551 different tags are noisy and sparse, where many recipes are missing relevant tags even when the recipes clearly align with certain categories. We aim to expand this partial and incomplete labeling to full coverage for structuring personalized meal plans. Thus, we treat the subset of tagged recipes as a seed set of ground-truth examples and apply an iterative self-training framework inspired by prior work in semi-supervised learning and adaptive self-training to pseudo-label the unlabeled portion of the dataset (Yarowsky, 1995; Wang et al., 2021).

We train a fully-connected neural network classifier composed of a FastText-based (Marin et al., 2019) food-tailored language model called RecipeFT (Seo et al., 2023), and feed-forward layers. We create pseudo-labels for training by iteratively predicting class probabilities over a batch of unlabeled data and incorporates only the high-confidence predictions into the training set for subsequent rounds. We define prediction confidence based on class probability, retaining the top 90% most confident predictions (among those with probability > 0.5) and returning the rest to the unlabeled test dataset. We then retrain the

classifier on the expanded dataset, fine-tuning both model weights and hyperparameters using a hyperband tuner (Li et al., 2018) to adapt to the evolving label distribution. This iterative self-labeling procedure allows us to propagate semantic labels, including meal type and cuisine, throughout the data set.

In this paper, we test the application for Latin American and Hispanic patients, the primary sub-population that our RDN collaborator works with. Additional information about the classifier settings can be found in Appendix A.1.

Recipe Recommender. We initiate our recommendation module by filtering out nutritionally poor recipes using a simplified *NutriScore* system (Julia et al., 2017), as adopted in the HUMMUS dataset for healthiness-aware recommendation (Bölz et al., 2023). Specifically, we exclude all recipes with the lowest *NutriScore* out of five scores. We then proceed with recipe-level predictions using a collaborative filtering approach. We benchmark several established recommendation models, including BPR (Rendle et al., 2012), LightGCN (He et al., 2020), and BiVAE (Truong et al., 2021), and select BiVAE as the primary model based on its performance in offline recommendation metrics (see Appendix Table 2). To improve granularity and practical usability, we train independent models for each of the three meal types (from the labeling module): breakfast, side dish, and main dish. For breakfast and side dishes, we use all available recipes, whereas for main dishes, we restrict the training set to Latin American recipes to ensure cultural relevance for our target patient populations.

RDN Recipe Alignment. Next, we refine the recommended item for each patient by incorporating RDN-guided recipe alignment. Starting from individual recipes predicted from recipe recommenders, we then filter and rank candidate recipes using RDN-provided guidelines on ingredients and preparation methods. Specifically, our RDN defines two keyword sets: healthy (e.g., wholegrain) and unhealthy (e.g., fried), with the full set of words specified in the shared repository¹. We compute a composite alignment score using TF-IDF vectorization (Salton and Buckley, 1988) over the recipe text, including title, ingredients and preparation steps, and measure cosine similarity between each candidate recipe vector and the RDN-defined healthy and unhealthy keyword vectors. This score is detailed in Appendix A.2. While TF-IDF does not fully capture semantic meaning, this approach provides a transparent and interpretable way to apply expert-defined dietary guidelines, and

appears to successfully filter out many RDN-defined unhealthy recipes. In addition, these RDN-curated ingredient and preparation lists are fully customizable by the nutritionist user, allowing the user to fine-tune the alignment strategy based on individual patient requirements, evolving dietary guidelines, or their unique professional approach.

Bundle nutritional optimization. Finally, we model daily meal bundle generation through an Integer Linear Programming (ILP) using the candidate recipes. The objective maximizes the recipe recommendation scores, subject to per-day patient-specific nutritional requirements, along with structural and practical constraints. We gather patient clinical information, full list detailed in Appendix A.3, to calculate the Body mass index (BMI) and Ideal Body Weight (IBW) (Devine, 1974; Peterson et al., 2016) and compute Resting Metabolic Rate (RMR) using the Mifflin-St. Jeor Equation (Mifflin et al., 1990) that will define personalized nutritional constraints. Additionally, we add structural and practical constraints to increase the usability in the real-world. Each bundle includes breakfast, lunch, and dinner, with three bundles per weekly plan, considering the possible leftover meals. See Appendix A.3 for a detailed formulation of this optimization module.

Together, these components form an end-to-end system that combines collaborative filtering and clinical nutrition modeling with domain expert guidance.

4. Results

In this section, we summarize our experiment results for choosing a primary recommendation model and an ablation study on HumbleNutri system modules. Appendix Table 2 shows average offline recommendation metric scores across benchmarked collaborative filtering models. Our model demonstrates competitive performances, while not a direct 1:1 comparison, compared to the HUMMUS paper (Bölz et al., 2023) reporting their best baseline to be BPR (Rendle et al., 2012) with $NDCG@10 = 0.0627$. Moreover, they identify a healthiness-accuracy trade-off when attempting to achieve healthiness-aware recommendation through similar thresholded data-processing as ours, with $NDCG@10$ decreased by 32%. In contrast, our approach achieves a 71% improvement in *NutriScore* as well as a 51% increase in $NDCG@10$, without compromising predictive performance. This highlights the benefit of our structured meal type-

specific recommender design and focus on a single cuisine type for more targeted recommendations.

In Appendix Table 3, we present an ablation study on the HumbleNutri system, where we incrementally combine the system modules and evaluate their impact using quantitative metrics. Specifically, we test the advantages of the RDN-guided recipe aligning module and the bundle optimization module in the system. We compare the system with (WGA-RB) and without guided recipe alignment and random bundling (WOGA-RB) instead of bundle optimization, with bundle optimization but without guided recipe alignment (WOGA-BO), and the full HumbleNutri (HN) system. We employ several metrics to test the healthiness and personalization of the recommended weekly plan. We use *NutriScore* (Julia et al., 2017), a rating system similar to the World Health Organization (WHO) score (Amine et al., 2003) that is based on how much nutrients and essential food groups (e.g., vegetables) are included in the food, and *RRR*, a ratio of recommended to restricted nutrients quantifying overall healthiness of the food (Scheidt and Daniel, 2004). To assess personalization, we use the average patient-specific nutrient deviation of each bundle from the required daily intake by calculating the Root Mean Squared Deviation (RMSD). The required daily intake for each patient is set as the median of its prescribed range, derived from patient clinical information used in our experiments. This metric measures how closely the generated meal plans adhere to individual patients’ nutritional requirements, with lower deviations indicating better personalization.

Our ablation study experiments show a clear and consistent improvement across the metrics as we incrementally add modules, demonstrating the value of each component. We specifically generate twelve bundles (4 weeks) and measure the average value of these three metrics. In particular, the RDN-guided recipe aligning module improved *NutriScore* by 51% on average compared to those without it, and the bundle optimization model improving *RRR* by 52% and reducing *BundleRMSD* by 77% on average compared to random bundles. The complete HN system achieves the best performance in all metrics, indicating that these modules work synergistically to produce healthier and more personalized meal plans. This shows the advantage of integrating expert knowledge with optimization techniques to address the challenge of personalized dietary recommendations for precision nutrition.

5. Discussion

In this work, we present HumbleNutri, a novel, open source recommendation system built for nutritionists designing personalized meal prescriptions. This synergistic modular system generates semantic labels that allow culturally relevant recipe recommendations, aligning recipes for target patients based on clinical nutritional guidance, and constraint-based optimization for creating weekly meal bundles. Our experiments show the contribution of each module to the generation of healthy, personalized meal plans.

Future work will incorporate dynamic feedback loops to adapt recommendations to user preferences and adherence, integrating nutritionist and patient input on factors such as cooking skills and cost. We also plan to develop an evaluation pipeline with qualitative studies alongside RDNs to assess clinical impact, positioning HumbleNutri as a foundation for adaptive precision nutrition systems. Moreover, our work could be improved with further validation of our pseudo-labels within the labeling module and exploring other recipe alignment methods through context-aware embeddings to better capture semantics of ingredients and preparation steps.

References

- EK Amine, NH Baba, M Belhadj, M Deurenberg-Yap, A Djazayery, T Forrestre, DA Galuska, Sofie Herman, WPT James, JR M’Buyamba Kabangu, et al. Diet, nutrition and the prevention of chronic diseases. *World Health Organization technical report series*, (916), 2003.
- Felix Bölz, Diana Nurbakova, Sylvie Calabretto, Armin Gerl, Lionel Brunie, and Harald Kosch. Hummus: A linked, healthiness-aware, user-centered and argument-enabling recipe data set for recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 1–11, 2023.
- Mario Casillo, Francesco Colace, Dajana Conte, Marco Lombardi, Domenico Santaniello, and Carmine Valentino. Context-aware recommender systems and cultural heritage: a survey. *Journal of Ambient Intelligence and Humanized Computing*, 14(4):3109–3127, 2023.
- Katarzyna Chapman. Can people make healthy changes to their diet and maintain them in the long term? a review of the evidence. *Appetite*, 54(3):433–441, 2010.
- Aleda MH Chen, Juanita A Draime, Sarah Berman, Julia Gardner, Zach Krauss, and Joe Martinez. Food as medicine? exploring the impact of providing healthy foods on adherence and clinical and economic outcomes. *Exploratory Research in Clinical and Social Pharmacy*, 5:100129, 2022.
- L Goldfarb Cyrino, Jennie Galpern, Lori Moore, Lea Borgi, and Leonardo V Riella. A narrative review of dietary approaches for kidney transplant patients. *Kidney International Reports*, 6(7):1764–1774, 2021.
- Ben J Devine. Gentamicin therapy. 1974.
- Jill Freyne and Shlomo Berkovsky. Recommending food: Reasoning on recipes and ingredients. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 381–386. Springer, 2010.
- Steven Haussmann, Oshani Seneviratne, Yu Chen, Yarden Ne’eman, James Codella, Ching-Hua Chen, Deborah L McGuinness, and Mohammed J Zaki. Foodkg: a semantics-driven knowledge graph for food recommendation. In *The Semantic Web—ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part II 18*, pages 146–162. Springer, 2019.
- Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 639–648, 2020.
- Rachael Jinnette, Ai Narita, Byron Manning, Sarah A McNaughton, John C Mathers, and Katherine M Livingstone. Does personalized nutrition advice improve dietary intake in healthy adults? a systematic review of randomized controlled trials. *Advances in Nutrition*, 12(3):657–669, 2021.
- Jee Young Joo and Megan F Liu. Culturally tailored interventions for ethnic minorities: a scoping review. *Nursing open*, 8(5):2078–2090, 2021.

- Chantal Julia, Serge Hercberg, World Health Organization, et al. Development of a new front-of-pack nutrition label in france: the five-colour nutriscore. *Public health panorama*, 3(4):712–725, 2017.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2018.
- Javier Marin, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusuf Aytar, Ingmar Weber, and Antonio Torralba. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- Gabriel Masset, Pablo Monsivais, Matthieu Maillot, Nicole Darmon, and Adam Drewnowski. Diet optimization methods can help translate dietary guidelines into a cancer prevention food plan. *The Journal of nutrition*, 139(8):1541–1548, 2009.
- Mark D Mifflin, Sachiko T St Jeor, Lisa A Hill, Barbara J Scott, Sandra A Daugherty, and Young O Koh. A new predictive equation for resting energy expenditure in healthy individuals. *The American journal of clinical nutrition*, 51(2):241–247, 1990.
- Yee Xuen Miow, Wilfred Kok Hoe Mok, Wan Ying Gan, Poh Ying Lim, Geeta Appannah, and Siti Nur ‘Asyura Adznam. The use of linear programming approach in diet optimization among children under five: a scoping review. *BMC Public Health*, 25(1):1279, 2025.
- Courtney M Peterson, Diana M Thomas, George L Blackburn, and Steven B Heymsfield. Universal equation for estimating ideal body weight and body weight at any bmi. *The American journal of clinical nutrition*, 103(5):1197–1203, 2016.
- Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*, 2012.
- Griffin P Rodgers and Francis S Collins. Precision nutrition—the answer to “what to eat to stay healthy”. *Jama*, 324(8):735–736, 2020.
- Gerard Salton and Christopher Buckley. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523, 1988.
- Douglas M Scheidt and Eileen Daniel. Composite index for aggregating nutrient density using food labels: ratio of recommended to restricted food components. *Journal of nutrition education and behavior*, 36(1):35–39, 2004.
- DongHyeon Seo, Abigail Horn, Andrés Abeliuk, and Keith Burghardt. What’s on the menu? towards predicting nutritional quality of food environments. *medRxiv*, 2023.
- Quoc-Tuan Truong, Aghiles Salah, and Hady W Lauw. Bilateral variational autoencoder for collaborative filtering. In *Proceedings of the 14th ACM international conference on web search and data mining*, pages 292–300, 2021.
- Yaqing Wang, Subhabrata Mukherjee, Haoda Chu, Yuancheng Tu, Ming Wu, Jing Gao, and Ahmed Hassan Awadallah. Meta self-training for few-shot neural sequence labeling. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1737–1747, 2021.
- Longqi Yang, Cheng-Kang Hsieh, Hongjian Yang, John P Pollak, Nicola Dell, Serge Belongie, Curtis Cole, and Deborah Estrin. Yum-me: a personalized nutrient-based meal recommender system. *ACM Transactions on Information Systems (TOIS)*, 36(1):1–31, 2017.
- David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics*, pages 189–196, 1995.
- Stuart M Zeltzer, David O Taylor, and WH Wilson Tang. Long-term dietary habits and interventions in solid-organ transplantation. *The Journal of Heart and Lung Transplantation*, 34(11):1357–1365, 2015.

Appendix A.

A.1. Labeling Module

HUMMUS dataset (Bölz et al., 2023) includes 507,335 recipes, 302,412 users, and 1,916,424 user-item interactions sourced from *Food.com*, and *FoodKG* (Hausmann et al., 2019). The self-labeled HUMMUS dataset extension through a semi-supervised, iterative self-training process is a contribution of this work and can facilitate future research and practical use for culturally-tailored and structured meal recommendations beyond the application of Latin American and Hispanic patients that is showcased in this paper. Researchers and nutritionists can readily extend HumbleNutri by leveraging cuisine-specific tags (e.g., Asian, Mediterranean) tailored to their applications or patient populations.

HumbleNutri employs two distinct classifiers, a binary sigmoid classifier for breakfast recipes and a multi-class softmax classifier for meal types: appetizer, main-dish, dessert, drink, and sauce/condiment. Each recipe is assigned two labels, a breakfast indicator label and a primary meal type label corresponding to its most likely category. HumbleNutri specifically utilizes appetizer and main-dish meal type label to use for side dishes and main dishes included in the lunch and dinner in the daily meal bundle.

Additionally, HumbleNutri train a cuisine type classifier to facilitate culturally-tailored meal prescriptions. We leverage available tags relevant to Latin American cuisines and employ a binary sigmoid classifier to label all recipes. See Table 1 for the dataset statistics after this labeling module for Latin American application.

The complete list of utilized tags and the source code for these trained classifiers are available at our shared code repository¹.

A.2. RDN-Guided Recipe Alignment Module

In the HumbleNutri system, RDN-guided recipe alignment leverages a composite score derived from TF-IDF vectorization (Salton and Buckley, 1988), as detailed in the main text. The score is computed as: $\text{Score} = \cos(\mathbf{r}, \mathbf{h}) - k \cdot \cos(\mathbf{r}, \mathbf{u})$, where $\mathbf{r}$ is the TF-IDF vector of the recipe text, $\mathbf{h}$ is the healthy keyword vector, $\mathbf{u}$ is the unhealthy keyword vector, and k is a constant penalty weight. We retain recipes with $\text{Score} > 0$.

Table 1: Data statistics after pseudo-labeling meal types and Latin American (Lat.Am.) cuisine.

Data	Counts
Breakfast recipes	79,056 (15.6%)
Side-dish recipes	95,484 (18.8%)
Lat.Am. recipes	74,466 (14.7%)
Lat.Am. main-dish recipes	35,292 (7.0%)
Lat.Am. user-item interactions	252,794 (13.2%)

A full array of the ingredients and preparation constraint keywords provided by RDN are included in our shared code repository¹.

A.3. Bundle Optimization Module

The bundle optimization module employs an Integer Linear Programming (ILP) approach to integrate the patient’s clinical profile, generating recipe combinations that satisfy individualized nutritional requirements and form a daily meal bundle comprising breakfast, lunch, and dinner. The clinical information that we consider includes: ‘Gender’, ‘Height’, ‘Weight’, ‘Age’, ‘Post surgery recovery phase (T/F)’, ‘Activity level (sedentary, lightly-active, moderately-active, active, very-active)’, ‘Prediabetes (T/F)’, ‘High Cholesterol (T/F)’, and ‘Hypertension (T/F)’, which we use to calculate the Body mass index (BMI), Ideal Body Weight (IBW) (Devine, 1974; Peterson et al., 2016), and Resting Metabolic Rate (RMR) using the Mifflin-St. Jeor Equation (Mifflin et al., 1990) that will define personalized nutritional constraints. These formulas are detailed in this shared document³.

We formulate the optimization problem as follows:
Sets: M : meal types; I_m : candidate recipes for $m \in M$; C : nutritional constraint parameters.

Parameters: $s_{i,m}$: recommendation score for $i \in I_m$; $v_{i,m,c}$: nutrient value of i for $c \in C$; $t_{i,m}$: preparation time of i ; T_m^{upper} : preparation time upper bound for m ; N_c^{lower}, N_c^{upper} : patient-specific bounds for c .

Decision Variables: $x_{i,m} \in \{0,1\}$: 1 if recipe i selected for meal type m , 0 otherwise.

$$\text{Maximize: } \sum_{m \in M} \sum_{i \in I_m} s_{i,m} \cdot x_{i,m} \quad (1)$$

3. <https://shorturl.at/Uo5Qg>

We maximize our objective of recipe recommendation score with six types of constraints. First, constraint at (2) ensures that there is one recipe per meal type. Constraint at (3) defines upper bound and lower bound for a patient-specific daily nutrient value requirements, namely ‘*Calories*’, ‘*Carbohydrate*’, ‘*Total fat*’, ‘*Saturated fat*’, ‘*Total sugar*’, ‘*Sodium*’, ‘*Protein*’, and ‘*Fiber*’. Practical constraint at (4) limits the preparation time for each meal to address patients’ weekly routines, and structural constraints at (5-7) prevent the repetition of recipes within and across the bundles.

$$\text{subject to: } \sum_{i \in I_m} x_{i,m} = 1, \quad \forall m \in M \quad (2)$$

$$N_c^{lower} \leq \sum_{m \in M} \sum_{i \in I_m} v_{i,m,c} \cdot x_{i,m} \leq N_c^{upper}, \quad \forall c \in C \quad (3)$$

$$\sum_{i \in I_m} t_{i,m} \cdot x_{i,m} \leq T_m^{upper}, \quad \forall m \in M \quad (4)$$

$$x_{k,\text{LunMain}} + x_{k,\text{DinMain}} \leq 1, \quad \forall k \in I_{\text{LunMain}} \cap I_{\text{DinMain}} \quad (5)$$

$$x_{k,\text{LunSide}} + x_{k,\text{DinSide-VG}} \leq 1, \quad \forall k \in I_{\text{LunSide}} \cap I_{\text{DinSide-VG}} \quad (6)$$

$$x_{i,m} = 0, \quad \forall \text{ previously selected recipe } i \text{ for meal type } m \quad (7)$$

Table 2: Recommender model comparison.

Avg. score:	NDCG@10 $\uparrow$	MAP $\uparrow$	Recall@10 $\uparrow$
BPR (Rendle et al., 2012)	0.0960	0.0815	0.1297
LightGCN (He et al., 2020)	0.0404	0.0267	0.0730
BiVAE (Truong et al., 2021)	0.0989	0.0839	0.1305

Table 3: Ablation study (with $\pm$ standard deviation).

Modules	NutriScore $\uparrow$	RRR $\uparrow$	BundleRMSD $\downarrow$
WOGA-RB	0.39 $\pm$ 0.27	1.54 $\pm$ 0.13	372.51 $\pm$ 24.50
WGA-RB	0.53 $\pm$ 0.26	2.55 $\pm$ 0.07	212.63 $\pm$ 26.29
WOGA-BO	0.37 $\pm$ 0.27	2.64 $\pm$ 0.11	71.09 $\pm$ 17.88
HN	0.61 $\pm$ 0.26	3.36 $\pm$ 0.09	54.96 $\pm$ 7.56