Minimizing Cost of Campus Food Subject to Nutritional Constraints

Anqi Cheng, Chuan Liu, Yangjun Sheng, Yan Zhao
Advisor: Alan Frieze

Table of Contents

Abstract

This report details our term research project for Operations Research II (21-393) at Carnegie Mellon University, advised by Professor Alan Frieze. Our project is about applying integer programming to solve a real-world problem: minimizing cost of campus food subject to nutritional constraints, such as calories consumption, macronutrients intake and different kinds of allergies. In this project, we aim to provide busy Carnegie Mellon students a one-stop solution for an affordable weekly dining menu when they are on campus, from Monday to Friday. Nutritional calculation takes into account inter-personal variations due to differences in height, weight, age and gender. We have constructed a web-based application to display our program. Users can obtain a customized lowest-cost weekly menu after they answer questions on our website. We invite all the readers of our report to try out our web application to gain deeper insights into the results of our model.

Introduction

Carnegie Mellon University has over 20 dining locations around campus, with over 30 dining concepts inspired by cuisines around the world, as well as many options that cater to students' various nutritional requirements. Students are often spoiled for choices of meals, because they have a huge variety of choices when it comes to picking their meal during the day.

In its efforts to make menu options more transparent, CMU dining service has published a nutrition calculator, that is both web and app-based, in order to inform students of the nutritional contents, ingredients and food allergens of the menu items. However, given their busy schedules, most CMU students do not have the time to pick the right items on the menu that would allow them to meet their nutritional and energy needs, while at the same time minimizing the cost of food.

We plan to solve this problem by providing a convenient, one-stop option for students by using a integer programming algorithm. Based on the student's personal information including height, weight, age and gender, we calculate the the nutritional needs of this student. We then provide a customized full-week's menu that minimizes the costs of campus food, while meeting the student's nutritional needs. One of our main deliverables for this project is a web application that will generate a personalized menu for a student based on his or her user inputs. In the project, we first use integer programming to get the total quantity of each type of food consumed in five days: Monday through Friday. We then solve an assignment problem afterwards to get a complete weekday menu. Our algorithm can accommodate special needs of the user when generating the menu, such as vegetarian-based diet and lactose intolerance.

Data

Our nutritional data come from Carnegie Mellon's web-based nutrition calculator. Its nutritional labels include calories (kcal), protein (gram), carbohydrate (gram) and fat (gram). However, due to minor discrepancy between calculated calories from the protein, carbohydrate, and fat contents of a food item and its nominal calorie label, we decide to use the calculated calories instead of the calorie label provided by this nutrition calculator. We calculate calories by the following widely accepted formula:

$$calories(kcal) = 4 * protein(in\ gram) + 4 * carbohydrate(in\ gram) + 9 * fat(in\ gram)$$

We also decide to include some special dining preference labels in our data: vegetarian, dairy-allergy, egg-allergy, wheat-allergy, fish-allergy, treenut-allergy, soy-allergy, peanut-allergy. We disregard some of the rarely used labels such as whether the food is healthy for heart.

Our price information come from menu published online, menu at the restaurant, or information from the cashiers. While we try to keep the prices as accurate as possible, there will still be minor differences between the price at the spot and the price from our application, due to the fact that some prices have a vague nature. For example, Zebra Lounge does not have a published food menu even at the restaurant and the employees are not familiar with all food items. In our dataset, we also include the location of the restaurant and at what specific time this food item is available: breakfast, lunch or dinner. Once again, the availability information is not perfectly accurate because some items are not offered daily. For example, while plain bagels are available daily, berry tarts at Tazza D' Oro are offered only occasionally.

Due to the our specific goal to provide a main dish suggestion for diners, our data mainly consist of entrees. While we do include some sides and drinks, those items are treated as entrees by our model to provide suggestions for people with lower nutritional needs. For example, for someone weighing 100 pounds, our model may suggest this user to consume a 16oz Fruit Punch(drink) for a meal.

We have 293 food items in total, from 16 restaurants. This, of course, only covers a portion of all food available. The complete list of restaurants is: Underground, Exchange, Carnegie Mellon Cafe, Zebra Lounge, City Grill, Taste of India, Entropy, Au Bon Pain, Entropy, Spinning Salad, Maggie Murph Cafe, EVGEFSTOS, Tazza D' Oro, Quick Piks, Stephanie's, Downtown Deli. One thing to notice is that Stephanie's is located at the Mellon Institute, which is a 15-minute walk from main campus.

Modeling Assumptions

To ensure that our users enjoy a diversified array of food, we restrict the consumption of a single food item to be no more than twice.

We only consider monetary costs in the decision process. If the model tells the user to go to Stephenie's (a restaurant off-campus) for a cheaper meal, we do not care about the costs of time and inconvenience to walk for 15 minutes from the main campus and then come back for class. While this simplification is reasonable, as most restaurants are cramped together within a 10-minute walk circle, it can cause trouble for users that are sensitive to non-monetary costs.

We disregard other factors that may influence one's decision for food, such as satisfaction from food, need for vitamins and the like.

Based on the healthy diet assumption, we insist that the user consume three meals each day. Although Carnegie Mellon University offers late night meal as a 4th meal in a day, we disregard such options in our menu output. This 3-meal-a-day constraint also explains why our model cannot give solutions to users with high energy need. For example, for someone weighing 200 pounds and has a height of 180cm, it is impossible to satisfy his/her need in 3 meals a day and one entree per meal.

When assigning meals for each day, we make sure that there is one breakfast item, one lunch item and 1 dinner item. Beyond this, we do not distinguish menu for a specific day from that of the other days. In other word, the user may interchange Monday's menu with, say, Wednesday's menu at his/her will.

In addition, to simplify our model, we assume that it is acceptable for a user to have a deficit or an excess of nutrients for a specific day, as long as the total consumption of a week turns out to be the right amount. Motivated by this assumption, we decide to solve an integer programming problem and an assignment problem consecutively. If we must abide to daily nutritional needs strictly, we would have constructed five different integer programming problems.

For further simplification, we assume each meal consists of one entree, disregarding sides, drinks and combo meals such as picking several entrees from an Asian fast food restaurant like iNoodle. Although we include Taste of India, which provides combo meal, we assume that one meal consist of only one entree, motivated by this assumption. Although in reality, the user can have some curry chicken and some rice, in our model the user is allowed to pick only one to fill the whole plate.

Finally, we assume that Harris Benedict Equation[1] gives us an accurate guideline for nutritional needs. We assume that for a healthy diet, 45% -65% of the calories should come from carbohydrates, 20%-35% of the calories should come from fat, and 10%-35% of the calories should come from protein. These percentages are widely agreed by dietitians[2].

Model

Nutritional Constraints

We use Harris Benedict Equation to calculate the amount of calories an individual needs, given height, weight, age, gender and activity level. We first calculate the Basal Metabolic Rate (BMR) with the formula below. BMR is the minimum amount of calories of an individual to sustain his/her basic metabolic activities.

$Men: \ 66.5 + (13.75 * weight\ in\ kg) + (5.003 * height\ in\ cm) - (6.755 * \ age\ in\ years)$

$Women: \ 655.1 + (9.563 * \ weight\ in\ kg) + (1.850 * \ height\ in\ cm) - (4.676 * \ age\ in\ years)$

We then multiply this BMR with the user's activity level, which is a continuous user input on our web application, ranging from 1.2 to 1.9. According to Harris Benedict Equation, there are several cutoffs of the activity levels:

1.2 : the individual remains sedentary(little or no exercise).
1.75: the individual is lightly active (light exercise or sports for 1-3 days/week).
1.55: the individual is moderately active(exercise for 3-5 days/week).
1.725: the individual is very active(exercise for 6-7 days/week).
1.9: the maximum level when the individual has a physical job or exercise very hard each day.

Methodology

We decide to use integer programming for two reasons: First, this is naturally an integer programming problem, as students cannot buy half an entree at a restaurant. Second, although integer programming comes at a higher computational cost than the normal linear programming with fractions, the size of our data is very small compared to the capacity of the solver. We first translate the problem into Python codes and then use a library, PuLP, to do the actual calculation.

When solving the assignment problem after we get the total quantity of five weekdays, we exploited the fact that in our output there are 5 items for breakfast, lunch and dinner respectively. Thus we simply randomly pick one item from the 5 available and assign it to Monday, and pick one from the four left, assign it to Tuesday and so on.

Objective Function

Minimize $\sum_{i=1}^{293} p_i(x_{iB} + x_{iL} + x_{iD})$

p_i represents the price for item i.

$1 <= i <= 293, \ j \in \{B, L, D\}, \ x_{ij} = 1$ if we consume food item i at meal time j.

$x_{ij} = 0$ if otherwise.

If the item is not offered at a specific meal time, then the corresponding variable is set to 0 in our constraints.

Nutritional Constraints

k_i = calories in kcal from 1 serving of food i

c_i = carbohydrates in gram from 1 serving of food i

f_i = fat in gram from 1 serving of food i

t_i = protein in gram from 1 serving of food i

K = needed calories in kcal

C_{lower} = lower bound of needed carbohydrate in gram = $\frac{K*45\%}{4}$

C_{upper} = upper bound of needed carbohydrate in gram = $\frac{K*65\%}{4}$

F_{lower} = lower bound of needed fat in gram = $\frac{K*20\%}{9}$

F_{upper} = upper bound of needed fat in gram = $\frac{K*35\%}{9}$

P_{lower} = lower bound of needed protein in gram = $\frac{K*10\%}{4}$

P_{upper} = upper bound of needed protein in gram = $\frac{K*35\%}{4}$

$$K \;=<\; total\ calories \;=\; \sum_{i=1}^{293} k_i(x_{iB} + x_{iL} + x_{iD})$$

$$C_{lower} \;=<\; total\ carbohydrate \;=\; \sum_{i=1}^{293} c_i(x_{iB} + x_{iL} + x_{iD}) \;<=\; C_{upper}$$

$$F_{lower} \;=<\; total\ fat \;=\; \sum_{i=1}^{293} f_i(x_{iB} + x_{iL} + x_{iD}) \;<=\; F_{upper}$$

$$P_{lower} \;=<\; total\ protein \;=\; \sum_{i=1}^{293} t_i(x_{iB} + x_{iL} + x_{iD}) \;<=\; P_{upper}$$

Meal Quantity Constraints

$$\sum_{i=1}^{293} x_{iB} = 5$$

$$\sum_{i=1}^{293} x_{iL} = 5$$

$$\sum_{i=1}^{293} x_{iD} = 5$$

Variety Constraints

$$\forall i, \quad x_{iB} + x_{iL} + x_{iD} <= 2$$

Dummy Variables

$x_{ij} = 0$ if item i is not offered at time j. $j \in \{B, L, D\}$

Assignment
Assume in the output from linear program we have
$x_{1j,} x_{2j}, x_{3j}, x_{4j}, x_{5j}$
Next we randomize the order. Assume that we get
$x_{2j,} x_{3j}, x_{1j}, x_{5j}, x_{4j}$
Then the menu for meal time j will look like:
Monday: food 2
Tuesday: food 3
Wednesday: food 1
Thursday: food 5
Friday: food 4

Demonstration and Analysis

Given the great variety of weekly menus that can be generated from the many permutations of user input parameters, it is impossible to show all of the results and analyze them. In this report, we will display and analyze the results of a few common cases, and invite our reader to use our web application to explore further.

To gain insight into the usability and practicability of our model, we pitch our web application against two students: the average CMU student vs the average CMU student with a lactose-intolerance.

The Average CMU Student

We will define the average CMU student to be:

- Of height 176cm and weight 88.3kg, as based on the height and weight of an average male in the US[3]
- Of sedentary activity level at 1.20. (The rigorous coursework in CMU is cause for long hours of studying/sitting)
- Of age 20, since most college students are of ages 18-22

Our results are displayed on the next page.

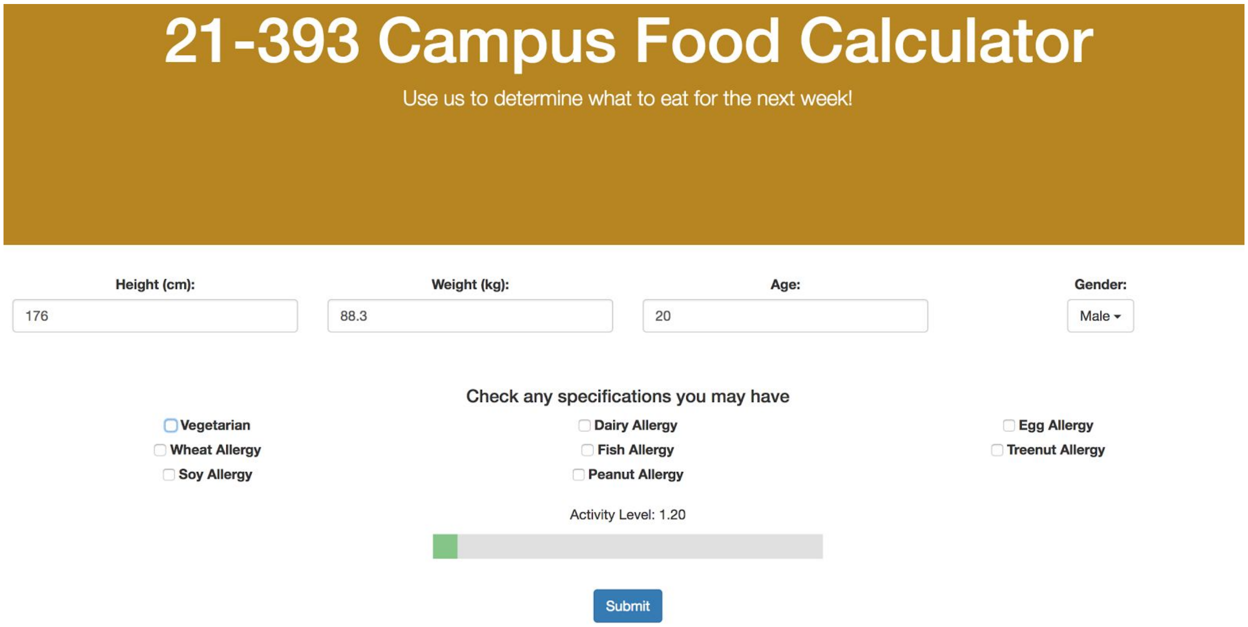

Fig 1.1: User input parameters for the average CMU student

Menu for a Week

Weekday	Breakfast	Lunch	Dinner
Monday	French Toast ($4.5)	Asiago Cheese Bagle ($1.29)	Farro ($2.15)
Tuesday	Chocolate Chip Pancakes ($5.28)	Asiago Cheese Bagle ($1.29)	Farro ($2.15)
Wednesday	Chocolate Chip Pancakes ($5.28)	Cinnamon Danish ($1.99)	Mediterranean Barley Salad ($2.15)
Thursday	Raisin Scone ($2.5)	Smoked Turkey, Bacon, Swiss, on Ciabatta ($7.29)	Mediterranean Barley Salad ($2.15)
Friday	Raisin Scone ($2.5)	Smoked Turkey, Bacon, Swiss, on Ciabatta ($7.29)	Cinnamon Danish ($1.99)

	Calories	Protein(g)	Fat(g)	Carbohydrates(g)	Cost($)
Menu	10151	316 (12%)	391 (35%)	1342 (53%)	49.80
Requirements	10130	10-35%	20-35%	45-65%	

Fig 1.2: Generated menu for one week for the average CMU student
Click on the item to see detailed nutritional contents and location.

Here, we observe that the average CMU student will spend $49.80. We also observe that all menu items appear twice except French Toast. When we click on the items to display their detailed nutritional contents, we can see that these selected items are of high energy values .

The CMU Student with Dietary Restrictions

To test our model further, we will now add additional constraints. Below is a menu for an average CMU student that is lactose-intolerant.

Menu for a Week

Weekday	Breakfast	Lunch	Dinner
Monday	Plain Bagel ($1.25)	Jamaican Chicken Wrap ($7.25)	BBQ Wings ($8.29)
Tuesday	Plain Bagel ($1.25)	Chicken Curry ($7.25)	BBQ Wings ($8.29)
Wednesday	McCann's steel Cut Oatmeal ($3.59)	Chicken Curry ($7.25)	Chicken Caesar Wrap ($6.25)
Thursday	McCann's steel Cut Oatmeal ($3.59)	Large Roll ($8.99)	Chicken Curry ($9.5)
Friday	Supergrain Bagel ($2.5)	Large Roll ($8.99)	Chicken Curry ($9.5)

	Calories	Protein(g)	Fat(g)	Carbohydrates(g)	Cost($)
Menu	10196	757 (30%)	396 (35%)	901 (35%)	93.74
Requirements	10130	10-35%	20-35%	45-65%	

Fig 1.3: Generated menu for one week for the average lactose-intolerant CMU student

There are no intersection of menu items with that for the average CMU student. This is due to the fact that most commonly seen entrees contain cheese or other dairy contents. The cost is substantially higher than that of an average student because the lactose-intolerant CMU student must choose more expensive foods such as sushi and chicken wings that have no dairy contents. In reality the cost can be somewhat lower because it is usually an option to order food with dairy contents and customize it. For example, the lactose-intolerant student can order a cheese burger, insisting no cheese to be added to the burger.

Comparing the menus for Fig 1.2 and Fig 1.3, we observe that the total amount of calories derived for a lactose-intolerant student is also higher than an average CMU student, with the protein percentage being higher as well by 18%, while carbohydrates percentage is lower by

18%. This points to a greater number of dairy foods with high carbohydrate content than that with protein content.

Another interesting observation is that the minimum cost for students with lactose intolerance is much higher. Such observation makes intuitive sense because we are searching for optimum in a much smaller dataset given additional constraints. The stark jump in price from $49.80 to $93.74 reflects the lack of cheap and good nutritional foods for lactose-intolerant people. This identifies as a potential gap to be addressed by CMU dining, since lactose intolerance impacts 33% of the American population[4].

Conclusion

From our demonstration, it is clear that our application does meet our initial goal of providing a lowest-cost menu for students, while meeting their nutritional constraints. However, it is also apparent that the menu exhausts options that are cheap and high in energy contents, and it also repeats item quite often to save money. The reason why we end up getting menu we might not enjoy is that we disregard many important non-monetary factors in our modeling assumptions.

The most salient of all disregarded factors is the satisfaction we get from eating. Some menu items, while having high nutritional contents, may not have a correspondingly high tastiness level. For example, sushi rolls are healthy, but they might not be as attractive as french fries. This fact may impact the utilization rate of our application: users who receive their personally generated menus may not get options that are suited to their palates and thus do not follow the menu. One possible solution is that when asking for user input, we can let them describe what types of food they like. Then we calculate a satisfaction score for each item, based on how many ingredients in the food agree with the user's preference. For example, a naive way of doing so is as follows: A user likes "fried", and "chicken" as keywords. Food whose name has "fried" gets 1 point. Food with "chicken" gets 1 points. Food with "fried" and " chicken" in its name gets 2 points. While minimizing costs, we make sure the satisfaction level is acceptable by adding one more constraint.

Another problem is that sometimes a food item appear twice in a row. In Fig 1.2 and 1.3, we observe that despite restricting the number of repeated menu options to two, the menu options frequently occur consecutively for the same meal-time. For example, according to the menu in Fig 1.2, the average CMU student will have the Mediterranean Barley Salad for both Wednesday and Thursday dinner. Given that people generally do not enjoy eating the same food for dinner twice in a row, this may also deter people from using our model. A simple solution to this problem is when we assign items to daily menus in the assignment problem, we add some more constraints so that even if the same menu item is selected twice for a weekly menu, it can be interspersed among other menu items.

Although we can make significant improvement in solving these two major problems and fine-tune our models in many other ways, the program currently is a functional tool that can be put into practical use for anyone whose primary aim is to tighten their budget for campus food.

Citation

1. Mifflin, MD, ST St Jeor, LA Hill, BJ Scott, SA Daugherty, and Yo Koh. "A new predictive equation for resting energy expenditure in healthy individuals." *Am J Clin Nutr*51, no. 2 (February 1990): 241-47. https://www.ncbi.nlm.nih.gov/pubmed/2305711.
2. Penner, Elle. "Ask the Dietitian: What's the Best Carb, Protein and Fat Breakdown for Weight Loss?" Myfitnesspal. June 29, 2016. Accessed November 20, 2017. http://blog.myfitnesspal.com/ask-the-dietitian-whats-the-best-carb-protein-and-fat-breakdown-for-weight-loss/.
3. National Center for Health Statistics. (2017, May 03). Accessed December 19, 2017, from https://www.cdc.gov/nchs/fastats/body-measurements.htm
4. Lactose Intolerance Statistics. Accessed December 19, 2017, from https://www.statisticbrain.com/lactose-intolerance-statistics/