

BLOSSOMS - Optimizing Your Diet: What Linear Programming Can Tell You!

My name is Aysegul. I'm an operations researcher at Structured Decisions Corporation in Westminister, Massachusetts. I develop mathematical models to analyze complex statistical operations.

Today, I'm excited to tell you about linear programming and how it helps us make decisions in our daily lives. Before starting to explain to you what linear programming is, let's first look at what programming is before putting the word "linear" in front of it. In our work here, programming is not computer programming. For us, the word programming means planning. Programming problems is basically a systematic way of making decisions. It's like designing a program for a vacation day with your family or friends, which activities to pick, and how to schedule them. Let's face it, each morning we all encounter a programming problem: getting dressed. Let's assume I'm putting on four items of clothes: pants, shirt, shoes, and socks. In all the different ways I put on these clothes, which order will minimize the time for me to get dressed? That is my objective. Work with some other students in your class on this and come up with an answer. I'll see you in a few minutes.

Welcome back. So how many different ways did you come with? Well, there are 4 times 3 times 2 times 1 equals 24 different orders. However, can I wear socks over shoes? Not really. Or putting on my pants after putting my shoes on, will be really tough too. Therefore, many of these orderings are impractical programs, meaning that they either do not match with society's standards or are not efficient, which are my constraints. So let's eliminate some of these orderings. Well, hold on a second. I still have a large number of them. So how do you choose which one is the most efficient? In other words, which one is the optimum? And remember that each ordering has a degree of effectiveness that enables me to compare them so that I can select the optimum one. Also, do not forget that you might add other limits or, again in other terms, constraints to your problem, such as the number of drawers you open and close, let's say to minimize the noise.

Next, let's move to the term linear. If some things can be linear, that implies that others are then nonlinear. Here's an example. Let's say we are buying socks in a department store. If a pair of socks costs \$2.00 and there are no discounts, then we can say that our total cost is directly proportional to the number of pairs we purchased, meaning that we paid \$2.00 for a single pair, \$4.00 for two pairs, and \$20.00 for ten pairs. Here our total cost is linear with respect to the number of pairs we purchased. And looks like this. On the other hand, if there was a \$0.20 discount for the second pair, \$0.40 discount for the third pair-- up to the fifth pair, and \$1.00 afterwards, then our total cost will be nonlinear and looks like this.

In short, linear programming problems are really programming problems where all the mathematical relationships can be presented as linear functions. So what kind of mathematical relationships will we have in a linear programming problem? Well, we will have two. First, an objective function to maximize or minimize. And this objective function indicates what we want to do, like minimize our annual cost of travel or maximize profits in our business. Second, there will be a set of constraints that have to be satisfied in our solution. We will illustrate these concepts by an example in the next segment, but before that discuss with the person next to you and see if each of you can come up with examples of both linear and nonlinear relationships. And also determine some of the mathematical relationships in your own examples. See you in a couple minutes.

Welcome back. So were you able to come up it some examples? Well, the diet problem is one of the many classic applications of linear programming studied back in the 1940s. It basically involves finding a group of foods: a diet that meets all daily nutrition requirements at minimum cost.

One of the early researchers who studied the diet problem was George Stigler. That's why it's sometimes called the Stigler diet. George Stigler was a 1982 Nobel Laureate in Economics. And he posed the following problem: Let's assume a moderately active man who weighs 154 pounds. He has a list of 77 foods he can eat from. And his goal is to satisfy the minimum daily requirements of nine different nutrients while keeping his expense at minimum. He wants to figure out which of these 77 foods he should eat and in what quantities. Back then, Stigler wasn't able to use computers to solve such a problem. So he used his

intuition to find a solution. Here, his objective function was to minimize cost of food. And his constraints were all the dietary requirements to stay healthy. The first thing Stigler did was to eliminate through trial and error, 62 of the foods from the original 77. [UNINTELLIGIBLE] nutrients in comparison to remaining 15 anyway. Now from the reduced list, he calculated the required demands of each of the remaining 15 foods he needed to minimize his cost. So what do you think he came up with as his annual minimum cost? Well, in 1939 dollars, the annual cost of his solution was \$39.93. In today's dollars, this will be about \$600.00. But that is pretty inexpensive. Unfortunately, the Stigler diet has been criticized for its lack of variety and palatability. But the good news is that his methods have received praise and he was recognized by his peers.

In 1947, eight years after Stigler made his initial estimates, George Dantzig came up with an algorithm called the simplex method to solve linear programming problems. He showed that Stigler's diet problem was indeed a linear programming problem, and using the mathematics and original 1939 data, he determined the true annual optimal cost using his methods to be \$39.69. Well, Stigler's guess for the optimal solution was off by only \$0.24 per year.

In the early 1950s, Dantzig decided to use Stigler's method to solve his own diet problem after his doctor advised him to go on a diet to lose weight. So he modeled the problem and let the computer decide what he would actually eat each day. His first solution called for various amounts of normal food, plus 500 gallons of vinegar. Well, imagine 2,000 of these. That's a lot of vinegar. It turned out that he had forgotten a constraint, so he reformulated the problem and his next solution called for 200 billion cubes per day. That's enough to make soup for your class. Something wasn't right again. This time he reformulated the problem with an upper limit on billion cubes. And his resulting solution this time called for two pounds of bran per day. That's more than a few plates of muffins. Still there was something right about the formulation. So he imposed an upper limit on bran, as well. But, his next diet called for two pounds of blackstrap molasses. Well, imagine the size of gingerbread house. At that point, it was clear to him that it wasn't easy to generate a test of values with the method he developed.

Your teacher is now handing out a list of all the foods that Stigler considered. If you had to restrict your diet to this set of foods, which ones would you prefer and which ones would you try to avoid? Do you think that you could come up with a diet that would satisfy your minimum daily requirements for nutrition? Discuss these questions with your neighbor, and I will see you in a few minutes.

Welcome back. Before moving further with solving a linear programming problem, let's first take a step back and formulate our own diet problem. We will be cooking two Mediterranean dishes. Our first dish is called green beans. In Turkish, it's called zeytinyagli taze fasulye. Green beans is traditionally served as cold, but it's perfect when it's warm too. One could enjoy the summer dish when it's really hot outside. The recipe for kisir varies from region to region in Turkey. However it's made, kisir is a perfect dish to be served with an afternoon tea or as an appetizer.

Well, we are done cooking, but before serving it, here comes our diet problem. Let's assume that I weigh around 120 pounds and I'm a moderately active person. So I require about 2,000 calories per day. My daily requirements are 271 grams of carbohydrates, 91 grams of protein, and 65 grams of fat. This dinner will cover half of my daily requirements, so the calorie intake should be at least 1,000. I should consume 435.5 grams of carbohydrates, 45.5 grams of protein, and 32.5 grams of fat from this dinner. One serving of green beans contains 12 grams of carbohydrates, 3 grams of protein, 9 grams of fat, and 171 calories, while one serving of kisir contains 33 grams of carbohydrates, 6 grams of protein, 1 gram of fat, and 150 calories. Our goal here is to have the right amounts of carbohydrates, fat, and protein from these two Mediterranean dishes, while minimizing the cost. You are now given all this information summarized in a handout. Work with some other students and come up with some reasonable combinations of these two foods.

Welcome back. What desirable levels of green beans and kisir did you come up with? We can usually say that we reach or even exceed our daily requirements by eating only either 16 servings of green beans or 33 servings of kisir. Let's look at these two options. First of all, what's the cost of the first option? If you were to have only 16 servings of green beans, it will cost us \$32.00. Second, what's the nutrient and calorie

intake? This option contains 192 grams of carbohydrates, 48 grams of protein, 144 grams of fat, and 2,736 calories. As seen here, even though we had the right amount of protein, we have too much of the others.

The second option, 33 servings of kisir, will cost us \$49.50 and contains 1,089 grams of carbohydrates, 198 grams of protein, 33 grams of fat, and 4,950 calories. In this case, we have the right amount of fat, but too much carbohydrates and protein. Also, who can eat 33 servings of kisir? We can't really lower it, since that will violate the daily requirement for fat. Similarly, if we were to have only green beans, we will have to eat all 16 servings. Again, that's a lot of green beans.

Next, we will look at how we can model this linear programming problem mathematically. And it's a linear program because the quantity of the nutrients consumed is exactly the quantity of the nutrients that your body uses. Basically, there are no interactions between kisir and green beans or no interactions among carbohydrates, fat, and protein. Let's continue with the mathematical formulation and note the quantity of green beans as x_1 servings and kisir as x_2 servings. And let's note the cost by z . Linear programming problems can involve both equality and inequality restrictions, and the latter is the case in our diet problem. To construct the limitations or constraints, let's start with the fact that we can't really have negative amounts of food. Hence, both x_1 and x_2 amounts of green beans and kisir should be greater than or equal to 0. Let's look at this from a plane geometric perspective: x_1 represents one dimension and x_2 represents a second dimension. So we are working in a two-dimensional space. The two dimensions, x_1 and x_2 , are represented by the right angle axis. If you graph these inequalities, for x_1 greater than or equal to zero, the solution space is any value of x_1 that's greater than 0. Basically, any value to the right of the point 0. Similarly, for x_2 greater than or equal to 0, the solution space is any value of x_2 straight up from the origin, perpendicular to the x_1 axis. Basically, these two constraints are used to make sure that x_1 and x_2 values are non-negative.

Let's continue with carbohydrates. One serving of green beans contains 12 grams of carbohydrates, whereas one serving of kisir contains 33 grams of carbohydrates. We should then have the following: $12x_1$ plus $33x_2$ greater than or equal to 135.5. So we are searching for a solution space in which every point, x_1 and x_2 , yields a sum greater than or equal to 135.5.

The first step is to draw the barrier line. Remember, this is a linear program so we know that the constraints can be expressed as linear combinations of x_1 and x_2 and the barrier must be a straight line. In order to do this, we find two points, one on the x_1 axis and the other on the x_2 axis. If x_1 is equal to 0, we must have x_2 equal to 4.1. That's our first point. If x_2 is equal to 0, now we have x_1 as 11.3. And this is our second point. These two points lie on the barrier line. Since we are searching for a sum greater than or equal to 135.5, all the points in the shaded area and its boundaries satisfy this inequality.

Similarly for protein, one serving of green beans contains 3 grams of protein and one serving of kisir contains 6 grams of protein. And it should be less than 45.5. Hence, the limitation on protein should be $3x_1$ plus $6x_2$ greater than or equal to 45.5. Here we marked the barrier line for this inequality too.

For fat, it is $9x_1$ plus x_2 greater than or equal to 32.5. As seen here, this inequality will eliminate even more points in our solution space.

We also have a limitation on the total number of calories we intake. One serving of green beans has 171 calories, whereas one serving of kisir contains 150 calories. Since the minimum number of calories is 1,000, the cost change looks like $171x_1$ plus $150x_2$ greater than or equal to 1,000.

Before moving further, one last but important thing I want to point out is that the constraints on protein and fat are binding, meaning that changing any one of them also changes the optimum solution.

We are now finished formulating our own diet problem. Here comes the fun part, how do we solve it? In other words, how do we determine the required amounts of x_1 and x_2 which simultaneously satisfy these six inequalities. What do you think? Any ideas? Talk with your neighbors and with your teacher. See you soon.

Do you have any ideas on how to solve this problem? One approach is a graphical method. The blue area here represents the combined solution space or intersection of all these inequalities. Each of the extreme points of this region corresponds to the binding of a pair of constraints. Binding here means these constraints are tied: the inequality becomes an equality. Let's further elevate this by looking at points A, B, and C as seen here. We could have considered other points such as point V, but that's not a physical extreme point. On the other hand, points A, B, and C are physical.

Point A corresponds to the binding of the constraints $9x_1 + x_2 \geq 32.5$, and $x_2 \geq 0$. Elevating our objective, which was $V = 2x_1 + 1.5x_2$ at point A with x_1 equal to 0 and x_2 equal to 32.5, gives us a total cost of \$48.75.

Let's move to point B, which corresponds to the binding of the constraints $3x_1 + 6x_2 \geq 45.5$ and $9x_1 + x_2 \geq 32.5$. Elevating our objective at point B with x_1 equal to 2.93 and x_2 equal to 6.12 results in \$15.04.

Our last extreme point is point C, which corresponds to the binding of the constraints $3x_1 + 6x_2 \geq 45.5$ and $x_1 \geq 0$. Again elevating our objective at point C with x_1 equal to 15 and x_2 equal to 0 provides us with a total cost value of \$30. That's higher than point B. Since point B results in the smallest objective value, we will pick point B, which is the optimum extreme point.

The optimum solution is to have 2.93 servings of green beans and 6.12 servings of kisir, with a minimum cost of \$15.04.

Now we are done solving our diet problem. Here comes the challenge. Most linear programming problems in real life involve thousands of physical extreme points. It will be really impractical to find all physical extreme points and evaluate each and every one of them. That's why as I mentioned at the beginning of the lesson, George Dantzig developed this method called the simplex. It's one of the fundamental methods used to solve linear programming problems. Simplex methods somehow systematically search a small set of such extreme points and find the optimum by always improving the value of the objective function or at least staying at its current level.

You now know about linear programming and how to solve linear programming problems using the graphical method.

In our lesson today, we formulated and solved a diet problem for people. Just as a note, the diet problem for chickens, cattle, and pigs has been a notable linear programming example too. In fact, it's an expansion of the simple diet problem we face in our daily lives. It's used daily to purchase feed for cattle in real-time marketing costs, which is done next by special linear programming formulations.

Here's a challenging problem for you. Why don't you formulate a nice diet problem for your cat or dog. It might be a little challenging since the simple model we just developed is one that doesn't take into account the dietary needs of a pet. Thus, it will probably require much more thought and modification to reflect the real world situation for them. This even holds true for any person who needs to modify the simple formulation to reflect his own specific dietary needs.

I hear some of you asking whether we can apply these methods to solve problems other than the diet problem in our daily lives. Yes. And of course so many of them. There are many industries that use linear programming in their decisions they make. One example is from the entertainment industry. Theme parks such as Universal Studios use linear programming to decide on queue lines. Or most of us have probably been in one of those retail stores in a theme park. They use linear programming to determine which and how much of a product to order, which constitutes a significant portion of parks' profits. Another industry is manufacturing. For example, how do you think that a manufacturer of school bags decides how many from each style and color to make in order to maximize its profits. Most of you [UNINTELLIGIBLE] in the class own a school bag, right? You buy a new one because it's either old or out of style or some other

reason. Based on the constraints set by the market trends, which you probably have an impact on too, the manufacturer decides on how many to produce to maximize its profits. The airline industry is another one which uses linear programming to schedule their pilots, plane routes, direct and indirect flights, or to determine their ticket prices.

I hope you enjoyed learning about linear programming. I also hope that you will use these methods to optimize your own daily lives. Thank you.

I hope you enjoyed this lesson since it has been a joy for me to create it. Hopefully, your students will enjoy it too. I have a couple of things to mention regarding some of the segments.

In segment one, you might want to ask students to give other examples of programming problems they face in their daily lives and what their approaches were to solve them.

In segment two, students need to have some prior knowledge of permutations. If your curriculum doesn't cover it, one way might be to assign some homework problems before the lessons. Or you could work through some examples in the classroom the day before this lesson. There's actually an excellent source of lecture notes on the fine art of counting. It's available online on MIT's OpenCourseWare website. You could also further explain the term "linear" with some additional examples, since most students in your classroom might have never seen it. A potential source might be Hillier and Lieberman's book called Introduction to Operations Research.

In segment three, a great idea to enhance a student's knowledge on the history of linear programming could be handing out the articles by Stigler, Dantzig, and [? Saltz ?] for further reading, following this lesson. For these articles, please refer to the additional online resources section.

Regarding the material in segment four, some of your students might actually enjoy doing additional reading on nutrition facts and dietary guidelines. This information could be found on the websites of the U.S. Department of Agriculture, the U.S. Department of Health and Human Services, and also the Food and Drug Administration.

Segments six and seven explain solving our diet problem using the graphical method. Involves a lot of geometry. To make sure that your students feel comfortable with graphing equations, you could assign homework problems before this lesson. Also students probably need graphing paper while working in groups in the classroom.

In segment eight, there's some time knowledge being used, which your students might have never seen before. So to better prepare them, you might want to explain the terms such as "extreme point" and "feasible" in further detail. John Chinneck's notes on linear programming explains some of these [UNINTELLIGIBLE] through the prototype example of the Acme Bicycle Company. If you want to further explore linear programming in the classroom, you could work on other prototype examples, such as product mix, resource allocation, workforce scheduling, blending or mixing problems, and aggregate planning. In the additional online resources section, we provided some useful links to those prototype problems. You could also go into greater depth on the solution methods for linear programming. We have included in the supplementary material how to solve our diet problem using the simplex method. You could either spend another class period on the simplex method or provide the supplementary material as a handout for students to read after the lesson. For other methods, such as the [UNINTELLIGIBLE] simplex method or the transportation problem algorithm, John Chinneck's notes on Linear Programming in Practice might be a good source. Further information is again provided in the additional online resources section.

Before we end our lesson today, I would like to take Professor Saul Gass for all that he has done in teaching linear programming and also for letting us use his ideas from his book An Illustrated Guide to Linear Programming in this video.

Thank you for your interest in this lesson. I hope it helps making the topic of linear programming exciting to your students. I wish you the best of luck.