

The Diet Problem:
A WWW-based Interactive Case Study
in Linear Programming
(<http://www.mcs.anl.gov/home/otc/Guide/CaseStudies/diet/>)¹

Joseph Czyzyk

Timothy J. Wisniewski

Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, IL 60439-4844

July 1996

¹This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Computational and Technology Research, U.S. Department of Energy, under Contract W-31-109-Eng-38.

Abstract

The Diet Problem is an interactive WWW-based case study that introduces users (particularly students and practitioners) to the mathematics of optimization. Users select foods for their menus, edit a set of nutritional constraints, and solve the linear program simply by clicking some buttons and making simple entries. A detailed analysis of the diet, complete with graphs and tables, is returned to the user.

The Diet Problem utilizes a simple, yet interesting, linear program to introduce the concepts of model planning, data gathering, optimal solutions, and dual variables. The power of the World Wide Web makes the case study accessible to a wide variety of users in many different countries.

1. Introduction

The dramatic growth in the use of the World Wide Web (WWW) is creating a new and unique opportunity to make educational tools available to large numbers of users. Given the interactive nature of Netscape and other Web browsers, it is now possible to create interesting exercises for students to learn about new fields of study and application. *The Diet Problem* is an interactive case study created by the Optimization Technology Center (OTC) to introduce users to a simple mathematical model, problem formulation, solution strategies, and the use of optimization techniques in general.

Section 2 will describe the diet problem as a linear program and some issues of formulating the mathematical model. Section 3 describes how this case study takes advantage of the Web interface and how users can access the model for themselves. Section 4 makes some concluding remarks about the Web and mathematics.

2. The Diet Problem

The Diet Problem was chosen because of its historical significance, instructional value, and ease of understanding.

The diet problem can be easily stated as follows:

Minimize the cost of food eaten during one day
Subject to the requirements that the diet satisfy a person's nutritional
 requirements and that not too much of any one food be eaten.

This formulation is included in *The Diet Problem* Web page where the user is shown the steps to convert this verbal formulation of the problem into a mathematical problem (a linear program) and then into a form that is computer readable. This process is summarized here for completeness.

The first step is to assign variables for the amounts of food that we are solving for and parameters for the data that we know.

Suppose that we have three foods and two nutritional requirements to satisfy:

Food	Cost / Serving	Vitamin A	Calories	Maximum Servings
Corn (C)	\$0.18	107	72	10
2% milk (M)	\$0.23	500	121	10
Wheat bread (W)	\$0.05	0	65	10

Nutrient	Minimum Amount in Diet	Maximum Amount in Diet
Calories (Cal)	2000	2250
Vitamin A (VA)	5000	50000

We can define our variables as follows:

$x[\text{food}]$	=	amount of <i>food</i> to eat
$c[\text{food}]$	=	cost of 1 serving of <i>food</i>
$A[\text{food}]$	=	amount of Vitamin A in 1 serving of <i>food</i>
$Cal[\text{food}]$	=	amount of calories in 1 serving of <i>food</i>
$MinF[\text{food}]$	=	minimum number of servings for <i>food</i>
$MaxF[\text{food}]$	=	maximum number of servings for <i>food</i>
$MinN[\text{nutrient}]$	=	minimum amount of <i>nutrient</i> required
$MaxN[\text{nutrient}]$	=	maximum amount of <i>nutrient</i> required

In this problem, we will be solving for the values of $x[\text{food}]$. The other values (C, A, Cal, etc.) are provided as input.

The problem can then be written as follows.

Minimize	$cost[C] * x[C] + cost[M] * x[M] + cost[W] * x[W]$		
Subject to	$MinN[VA] \leq$	$A[C] * x[C] + A[M] * x[M] + A[W] * x[W]$	$\leq MaxN[VA]$
	$MinN[Cal] \leq$	$Cal[C] * x[C] + Cal[M] * x[M] + Cal[W] * x[W]$	$\leq MaxN[Cal]$
	$MinF[C] \leq$	$x[C]$	$\leq MaxF[C]$
	$MinF[M] \leq$	$x[M]$	$\leq MaxF[M]$
	$MinF[W] \leq$	$x[W]$	$\leq MaxF[W]$

The objective function is the value that we are trying to minimize. In this problem, the objective function is the cost of the entire menu. Since the objective function and the constraints are both linear functions of the variables, $x[\text{food}]$, this problem is a linear program.

The Diet Problem formulation page continues this derivation to show a more mathematical formulation and the AMPL modeling language [1] formulation so that we can solve the problem by using a computer.

The optimal solution for our small sample problem is shown here.

Variables:

Food	Servings
Corn	1.94
Milk	10.0
Wheat Bread	10.0

Dual Variables:

Variable	Dual Cost for Upper Bound (\$)
Corn	0.0
Milk	-0.073
Bread	-0.11

Constraint	Dual Cost for Lower Bound (\$)	Dual Cost for Upper Bound (\$)
Calories	0.0025	0.0
Vitamin A	0.0	0.0

Simply stated, the dual variables represent the cost of including a constraint in the model. Without any constraints (other than those that do not allow negative amounts of food), one can see how the optimal menu has zero cost and includes no foods. From the above tables we can make the following simple interpretations:

- Since the amounts of milk and wheat bread in the diet are at their upper bounds, a cost is associated with these bounds. The dual variable associated with the upper bound on milk is -0.073 . If we increased the upper bound on milk servings from 10 to 11, the cost of the diet would decrease by \$0.073 and similarly for wheat bread.
- Since the amount of corn in the diet is not at the upper bound, the constraint $x[\text{Corn}] \leq 10$ is not active and has a dual cost of 0.0. Changing the upper bound on corn will have no effect on the objective value.
- A similar argument can be made about the Calories constraint. If we are willing to accept one less calorie in the diet, the cost can be decreased by \$0.0025.
- Since neither the upper or lower bounds on the Vitamin A constraint are active, the corresponding dual cost is 0.0.

3. Taking Advantage of the Web

Linear programs are not difficult to formulate or understand. In fact, many high-school textbooks have sections on linear programming. Simple two-variable problems can be solved by using graphical techniques, and even larger problems can be solved by hand. Realistic problems with more than a few variables and constraints, however, quickly become too difficult to solve by hand. Although good LP software exists, it is unlikely that students will correctly formulate their problems for the

software package in order to play with larger or more complicated models. By using the remote execution features of the Web, students no longer need to correctly encode their problems or download software to solve the diet problem. They can simply select the foods of their choice, modify the constraints, solve the problem, and obtain tables and graphs of the solution—all by simply clicking some buttons.

The Diet Problem has a database of 64 foods that users can select to be in their menus. A small portion of the input form in Figure 1 shows some of the foods in the model. They also have the option of editing the constraints or using the defaults. The information from these forms is collected and sent to the CGI² program that runs *The Diet Problem*. The program extracts the data, formulates the linear program, and solves it. If the linear program is infeasible, that is, the nutritional requirements cannot be met with the foods selected, the user is notified and given another chance. If the problem is feasible, that is, a solution exists, a results page is created that shows the optimal menu and cost (see Figure 2). The results page also has pointers to other pages that show the cost breakdown by food (see Figure 3) and how much of each nutrient was contributed by each food. This information is shown in pie charts. The information about the dual variables is also shown along with an explanation.

<input type="checkbox"/>	Hotdog, Plain	\$0.31	1 Hotdog
<input type="checkbox"/>	White Rice	\$0.08	1/2 Cup
<input type="checkbox"/>	Peanut Butter	\$0.07	2 Tbsp
<input type="checkbox"/>	Sardines in Oil	\$0.45	2 Sardines
<input type="checkbox"/>	Popcorn, Air-Popped	\$0.04	1 Oz

Figure 1. A portion of the food selection input form.

Additional links to further information are available or those people who are more interested in linear programming or in additional applications. There are pointers into the remainder of the NEOS Guide³ for information on modeling languages, software for linear programming, solution techniques, and other applications. There also are pointers to the NEOS Server,⁴ an OTC facility for solving users' mathematical problems (linear programs, stochastic linear programs, nonlinear programs, network problems, and complementarity problems) over the Internet. Following these and a host of other links, the user can quickly learn about a variety of optimization models and solution techniques.

²CGI, The Common Gateway Interface, provides a programming interface for Web pages

³<http://www.mcs.anl.gov/home/otc/Guide/>

⁴<http://www.mcs.anl.gov/home/otc/Server/>

The Optimized Menu

The cost of this *optimal* diet is \$1.63 per day.

The Solution and Cost Breakdown by Food

Food	Servings	Cost(\$)
Wheat Bread	10.00	0.50
Frozen Broccoli	0.51	0.08
Carrots, Raw	0.07	0.01
2% Lowfat Milk	1.79	0.41
Peanut Butter	2.82	0.20
White Rice	5.48	0.44

Figure 2. The optimal menu

The World Wide Web supports a close interaction between the users of the diet problem and the others. The “comments and suggestions” page encourages users to submit their comments and suggestions for improvements. Because of the overwhelming interest in the case study and because of the helpful comments that were received, many improvements to the case study were made: many of the pages were reworded for clarity, a section on dual variables was added so that professors could better use the case study in their classes, and additional foods were added to the menu.

The World Wide Web makes the distribution of the latest version of the software painless. Since one version of the software resides on the Web Server at Argonne, anyone accessing the diet problem is interacting with the latest version. In addition, there is no need to support software for many different machines because all the different Web browsers understand the same Web protocol. Designing information for the World Wide Web promotes quick development and easy user feedback.

The *Diet Problem* has proven to be extremely popular. In the five months that the case study has been on the Web, the site has solved more than 2,600 instances of the diet problem. The users come from all over the world, and many are students in colleges and universities.

4. Conclusion

The growth of the World Wide Web is enabling many innovations in education technology. With each new update of the HTML standard and the advent of the Java programming language, new possibilities arise for creating interesting and interactive

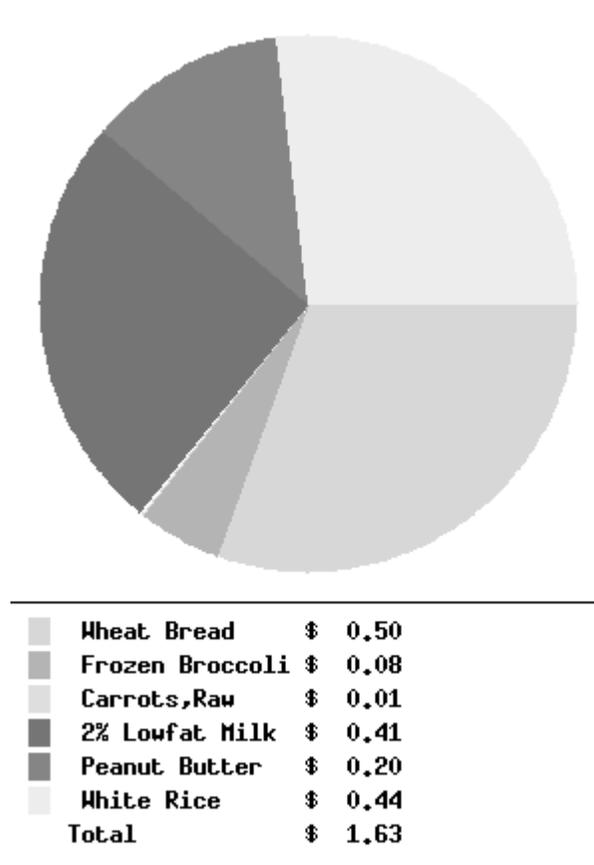


Figure 3. The cost breakdown by food.

Web pages. As more and more people become connected to the Internet and have access to the World Wide Web, the audience for Web applications is becoming extremely large. We are seeking to take advantage of this growth to introduce new users to the applicability and power of optimization models through the use of interesting case studies. *The Diet Problem* provides a simple introduction to linear programming and model construction. We enable users to interact with a mathematical model without having to use mathematical notation. We also enable them to view the solution by using easy-to-understand tables and charts.

Bibliography

- [1] R. FOURER, D. M. GAY, and B. W. KERNIGHAN, 1993. *AMPL: A Modeling Language for Mathematical Programming*. Scientific Press, San Francisco, California.