Contents

Chapter 1

Models for Decision Making

This is Chapter 1 of Junior Level Web-Book for "Optimization Models for Decision Making" by Katta G. Murty.

1.1 Decision Making

Anyone who holds a technical, managerial, or administrative job these days is faced with making decisions daily at work. It may involve:

- determining which ingredients and in what quantities to add to a mixture being made so that it will meet specifications on its composition,
- selecting one among a small number of suppliers to order raw materials from,
- determining the quantities of various products to manufacture in the next period,
- allocating available funds among various competing agencies,
- determining how to invest savings among available investment options,
- deciding which route to take to go to a new location in the city,

- allocating available farm land to various crops that can be grown,
- determining how many checkout lanes to staff during a period of stable demand in a grocery store,

  etc., etc.

A situation such as one of these requiring some decisions to be made is known as a **decision making problem** or just **decision problem**. These problems arise in the operation of some system known as the **relevant system** for the problem. The person(s) responsible for making these decisions are called the **decision maker(s)** for the problem.

At one extreme, these decision making problems may be quite simple requiring the determination of the values of a small number of controllable variables with only simple conditions to be met; and at the other extreme they may be large scale and quite complex with thousands of variables and many conditions to be met.

Decision making always involves making a choice between various **possible alternatives**. Decision problems can be classified into two categories with very distinct features. It is important to understand the difference between these categories.

Two Categories of Decision Making Problems

Category 1: This category includes all decision problems for which the set of possible alternatives for the decision is a finite discrete set typically consisting of a small number of elements, in which each alternative is fully known in complete detail, and any one of them can be selected as the decision.

Even though many textbooks do not discuss these problems, these are the most common decision problems encountered in daily living, in school, at work, and almost everywhere. Some examples of this category are:

- A teenage girl knows four boys all of whom she likes, and has to decide who among them to go steady with.

- An automobile manufacturer has to decide whether to use a cast iron engine block, or an aluminum engine block in their new car line.

- A company has received merger offers from three other companies. It has to decide whether to accept any one of these offers, or to continue operating by itself.

Many more examples of this category can be seen in Chapter 2. Since all the alternatives for the decision are fully known in full detail, it is not necessary to construct a mathematical model to identify the set of all alternatives for the decision in this category. Instead, one can begin applying an algorithm for solving these problems directly. A specialized method known as the **scoring method** commonly used to handle these problems is discussed in Chapter 2.

Category 2: This category includes all decision problems for which each possible alternative for the decision is required to satisfy some restrictions and constraints under which the relevant system must operate. Even to identify the set of all possible alternatives for the decision, we need to construct a mathematical model of these restrictions and constraints in this category. An example of a decision problem in this category is discussed in Section 1.2.

Even when there are no constraints to be satisfied in a decision problem, if the number of possible alternatives is either infinite, or finite but very large; it becomes necessary to define the decision variables in the problem, and construct the objective function (the one to be optimized) as a mathematical function of the decision variables in order to find the best alternative to implement. Such decision problems also belong to Category 2.

So, the essential characteristic of a Category 2 decision problem is that in order to handle it we need to identify the decision variables in the problem and build a mathematical model of the objective function and/or the constraints in terms of the decision variables. The rest of this chapter deals only with this category of problems.

Quantitative Analysis for Solving Category 2 Decision Problems

In the past decisions were made exclusively on intuitive judgement based on hunches acquired from past experience. But to survive and thrive in this highly competitive technological world of today it is essential to make decisions on a rational basis. The most rational way for decision making is through **quantitative analysis** which consists of the following steps.

1. **Get a precise definition of the problem, all relevant data and information on it:** The initial statement of the problem may be vague or imprecise. Study the relevent system and develop an accurate and complete statement of the problem. Quite often the problem undergoes many changes in successive discussions until its final version is agreed upon by all the decision makers involved.

   Two types of factors or variables may be affecting the system. These are:

   > **Uncontrollable factors:** Factors such as environmental factors which are random variables not under the control of the decision makers.
   >
   > **Controllable inputs:** Factors whose levels can be controlled by the decision makers and set at desired values. These factors whose values the decision makers can manipulate are called **decision variables** in the problem. They may include other **ancillary variables** that are functions of the decision variables.

   If there are no uncontrollable factors, or if the values of all the random variables among the uncontrollable factors are known exactly, the relevant system depends only on the values of the controllable decision variables and there is no uncertainty, i.e., all the relevent data in the decision problem is known with certainty.

In this case the decision problem is known as a **deterministic decision making problem**.

When the random variables among the uncontrollable inputs are subject to variation, the decision problem is known as a **stochastic** or **probabilistic decision making problem**. Here the outcome of the relevent system is uncertain even when the values of all the decision variables are fixed, because some importent variables will not have their values known before the decisions are finalized. This uncertainty must be incorporated into the decision making.

To solve a stochastic decision making problem, we need knowledge of the probability distributions of all the random variables among the uncontrollable inputs. Unless the decision problem is a very simple one, exact analysis of it using these probability distributions may become very complex. That's why very often stochastic decision problems are analyzed by studying appropriate deterministic approximations of them.

One commonly used hedging strategy to construct a deterministic approximation of a stochastic decision making problem is to replace each random variable by some location parameter of its probability distribution (mean, median, or some desirable percentile) plus some safety factor to account for the uncertainty in its value. This converts the problem into a deterministic decision making problem.

That is why studying techniques for solving deterministic decision making problems is of great importance. In this book we will discuss only deterministic decision making problems.

2. **Construct a mathematical model of the problem:** Construct a mathematical model that abstracts the essence of the decision problem. The model should express the various quantities in the problem including performance measures if any, in the form of mathematical functions of decision variables, and express the relationships among them using appropriate equations or inequalities, or objective functions to be optimized (maximized, or mini-

mized, as appropriate).

Real world problems are usually too complex to capture all the fine details of them in the form of simple mathematical models that we can analyze. Usually a model is a simplification that provides a sufficiently precise representation of the main features such that the conclusions obtained from it also remain valid to the original problem to a reasonable degree of approximation. Therefore, constructing a mathematical model usually involves making approximations, heuristic adjustments, and quite often ignoring (or putting aside or **relaxing**) features that are difficult to represent mathematically and handle by known mathematical techniques. When such relaxations are used, it may be necessary to make some manual adjustments to the final conclusions obtained from the model to incorporate the relaxed features.

It usually takes great skill to decide which features of the real problem to relax in constructing a model for it, this skill comes from experience. This is reflected in the word "mahaanubhaava" in Indian languages like Sanskrit, Telugu etc. for "great person or expert", which literally means "a person of vast experience". That's why many people consider modeling to be an art.

3. **Solve the model:** Solve the model to derive the solution, or conclusion for the decision problem.

   For some of the models we have efficient algorithms and high quality software systems implementing them. For some others we do not yet have efficient algorithms, and when the model is large, existing algorithms might take unduly long to solve it. In this case, one usually obtains approximate solutions using some heuristic approaches.

4. **Implement the solution:** In this final phase, the solution obtained is checked for practical feasibility. If it is found to be impractical for some reason, necessary modifications are carried out in the model and it is solved again; the same process is repeated as needed.

Often the output from the model is not implemented as is. It provides insight to the decision maker(s) who combine it with their practical knowledge and transform it into an implementable solution.

As an illustration, in the next section we develop a mathematical model for a very simple decision making problem of Category 2.

1.2 A Model for a Simple Decision Making Problem

Modern jogging shoes usually contain heel pads for cushioning to soften the impact when the foot hits the ground, and also to give some bounce. In some shoe brands, the heel pad is a sealed packet of plastic containing air under pressure. The following type of decision problem arises at companies making these heel pads.

Decision problem: There is 100 cc of a gas at 1500 mb of pressure in a closed container. Determine how much gas should be added to or expelled from the container to make sure that when the gas in the container is compressed to 3000 mb of pressure its volume will be 40 cc.

The only decision variable in this problem is:

$x =$ cc of gas to be added to or taken out of the present container

We adopt the convention that positive values of x indicate addition of gas to the container (i.e., for example, $x = 12$ means adding 12 cc of additional gas at 1500 mb of pressure to container), and negative values of x indicate expelling of present gas from container (i.e., for example, $x = -8$ means expelling 8 cc of gas from present container). A solution to this decision problem consists of obtaining a numerical value to the decision variable x. After implementing the solution x, the container will have $100 + x$ cc of gas at pressure 1500 mb.

Requirement to be met: When the contents $100 + x$ cc in the container are compressed to 3000 mb, the volume of gas in the container should be 40 cc.

This requirement leads to a constraint that the decision variable x should satisfy. There are two important characteristics of gas in this problem, its pressure p in mb, and volume v in cc, which are ancillary variables. Determining the constraint on the decision variable needs an understanding of the relationship between p and v, i.e., as compression increases the value of p, how does the volume v vary?

We denote the volume of gas in the container at pressure p by $v(p)$. This relationship is provided in the form of an equation by **Boyle's Law** of physics which states that the pressure p and volume $v(p)$ of a fixed quantity (by weight say) of gas satisfy

$$pv(p) = \text{ a constant } a \text{ say.}$$

where the constant a in the RHS depends on the quantity of gas.

In reality Boyle's law does not hold exactly. But it offers a very good approximation to how $p, v(p)$ behave in the range of values of these variables encountered in this application, so we will use it.

So, $v(p) = a/p$. In our decision problem we know that $v(1500) = 100 + x$. Substituting, we find that $a = 1500(100 + x)$. So we have

$$v(p) = \frac{150,000 + 1500x}{p}$$

This provides the volume of gas in the container as a mathematical function of its pressure. The requirement is that $v(3000) = 40$. This can be expressed as the constraint $150,000 + 1500x = 3000 \times 40$, or

$$1500x = -30,000$$

This is a single linear equation in the decision variable x, it constitutes the **mathemtical model** for our decision problem. The left hand side of this constraint, $1500x$, is known as the **constraint function** and the right hand side constant $-30,000$ is abbreviated and called the **RHS constant** for the constraint.

This model has the unique solution $x = -30,000/1500 = -20$ which is the only value of the decision variable x satisfying the requirement. It is the **solution** for our decision problem, it corresponds to the action of releasing 20cc of gas from the container at original pressure of 1500 mb.

1.3 Optimization Models

The model of restrictions and constraints for the decision problem discussed in Section 1.2 is a single linear equation in one decision variable which has a unique solution. This is quite rare. Such models for most real world decision problems have many solutions. The question that arises then is how to select one of the many solutions of the model to implement?. This is usually done so as to optimize an objective function which is a measure of effectiveness of the relevent system.

Since prehistoric times, humans have had an abiding interest in optimizing the performance of systems that they use. Now-a-days all the decisions that we make at work, and those affecting our personal lives, usually have the goal of optimizing some desirable characteristic. If there are some objective functions to optimize in addition to satisfying the requirements on the decision variables, the resulting model is known as an **optimization model**.

Each of the objectives to optimize is typically a measure of effectiveness of performance of the relevant system, and should be expressed as a mathematical function of the decision variables.

If higher values of a measure of performance are more desirable (such a measure could be considered as a **profit measure**) we seek to attain the maximum or highest possible value for it. If lower values of a measure of performance are more desirable (such a measure could be interpreted as a **cost measure**) we seek to attain the minimum or the lowest possible value for it. The various measures of performance are usually called **the objective function(s)** in the mathematical model for the system. To **optimize** an objective function means to either maximize or minimize it as desired.

If there is only one measure of performance (such as yearly total profit, or production cost per unit, etc.) the model will be a **single ob-**

jective model. When there are several measures of performance, we get a **multiobjective model** in which two or more objective functions are required to be optimized simultaneously.

In optimization models the requirements come from the relationships that must hold among the decision variables and the various static or dynamic structural elements by the nature of system operation. Each requirement leads to a constraint on the decision variables that will be expressed as a mathematical equation or inequality in the model for the problem. The model also includes any bounds (lower and/or upper) that the decision variables or some functions of them must satisfy in order to account for the physical limitations under which the system must operate.

In some problems, in addition to all these requirements, there may be others that specify that the values of some decision variables must come from specified sets (for example, if the decision variable x_1 is the diameter of pipe used in designing a component, and this pipe is available in diameters 1", or 1.5", or 2" only; then the value of x_1 must come from the set $\{1'', 1.5'', 2''\}$).

We know that if an objective function is a cost function (profit function) we would like to minimize (maximize) it. Fortunately, it is not necessary to consider minimization and maximization problems separately, since any minimization problem can be transformed directly into a maximization problem and vice versa. For example, to maximize a function $f(x)$ of decision variables x, is equivalent to minimizing $-f(x)$ subject to the same system of constraints, and both these problems have the same set of optimum solutions. Also, we can use

$$\begin{pmatrix} \text{Maximum value of } f(x) \\ \text{subject to some constraints} \end{pmatrix} = - \begin{pmatrix} \text{Minimum value of } -f(x) \\ \text{subject to the same con-} \\ \text{straints} \end{pmatrix}$$

For this reason, we will discuss algorithms for minimization only in this book.

Let $x = (x_1, \ldots, x_n)^T$ denote the vector of decision variables. A typical single objective optimization model has the following form:

$$\begin{aligned}
\text{Minimize} \quad & \theta(x) & (1.3.1)\\
\text{subject to} \quad & g_i(x) \quad \begin{cases} = b_i, & i = 1, \ldots, m \\ \leq b_i, & i = m+1, \ldots, m+p \end{cases} & (1.3.2)\\
\ell_j \leq x_j \quad & \leq \quad u_j, \quad j = 1, \ldots, n & (1.3.3)\\
x_j \quad & \in \quad \Delta_j, \quad j \in J \subset \{1, \ldots, n\}. & (1.3.4)
\end{aligned}$$

where all the functions are assumed to be continuous and differentiable, and for each $j \in J$, Δ_j is a specified set within which the value selected for the variable x_j is required to lie. The function $g_i(x)$, constant b_i are respectively the **constraint function, RHS constant** respectively for the ith constraint in (1.3.2).

Any "$\geq$" inequality constraint can be transformed into a "$\leq$" constraint by multiplying both sides of it by -1. That's why we listed all the inequality constraints in the "$\leq$" form.

ℓ_j, u_j are the **upper and lower bounds** on the decision variable x_j. In many problems $\ell_j = 0, u_j = \infty$ is common (i.e., x_j is required to be nonnegative) because economic activities can only be carried out at nonnegative levels. But in general ℓ_j, u_j can have any real values satisfying $\ell_j \leq u_j$, in fact we can have $\ell_j = -\infty$ and $u_j = +\infty$ (in this case x_j is called an **unrestricted variable**.

Constraints like those in (1.3.4) mainly arise in **discrete problems** where some variables are required to assume only values from specified discrete sets.

For (1.3.1)$-$(1.3.4), a numerical vector x is said to be a **feasible solution** if it satisfies all the constraints (1.3.2)$-$(1.3.4). A feasible solution $\bar{x}$ satisfying $\theta(\bar{x}) \leq \theta(x)$ for all feasible solutions x is said to be an **optimum solution** or **optimum feasible solution** for (1.3.1) to (1.3.4), because it has the smallest value for the objective function among all feasible solutions.

The typical multiobjective problem is of the form

Minimize $\theta_i(x)$; $i = 1$ to k simultaneously
subject to constraints of form (1.3.2)$-$(1.3.4).

If constraint (1.3.4) is absent, the above models are said to be **continuous variable optimization models** since each variable can assume any value within its bounds subject to the other constraints. If constraints (1.3.4) are there, and Δ_j are discrete sets (like the set of integers, or the set $\{0, 1\}$ etc.) the models are said to be **discrete optimization models**.

Single Versus Multiobjective Models

Mathematical theory of single objective models is well developed. In contrast, for multiobjective optimization models, we do not even have the concept of an optimum solution. Often, the various objective functions conflict with each other (i.e., optimizing one of them usually tends to move another towards undesirable values), for solving such models one needs to know how many units of one function can be sacrificed to gain one unit of another, but this trade-off information is not available. In other words, one is forced to determine the best compromise that can be achieved. Since trade-off information among the various objective functions is not given, multi-objective optimization problems are not precisely stated mathematical problems. Techniques for handling them usually involve trial and error using several degrees of compromises among the various objective functions until a consensus is reached that the present solution looks reasonable from the point of view of all the objective functions.

We restrict the scope of this book to single objective optimization models.

Static Versus Dynamic Models

Models that deal with a one-shot situation are known as **static models**. These include models which involve determining an optimum solution for a one period problem. For example, consider the production planning problem in a company making a variety of products. To determine the optimum quantities of each product that this company should produce in a single year, leads to a static model.

However, planning does involve the time element, and if an application is concerned with a situation that lasts over several years, the same types of decisions may have to be made in each year. In the production planning problem discussed above, if a planning horizon of 5 years is being considered, it is necessary to determine the optimum quantities of each product to produce, in each year of the planning horizon. Models that involve a sequence of such decisions over multiple periods are called **multi-period or dynamic models**.

When planning for a multi-period horizon, if there is no change in the data at all from one period to the next, then the optimum solution for the first period determined from a static model for that period, will continue to be optimal for every period of the planning horizon. Thus multi-period problems in which the changes in the data over the various periods are small, can be handled through a static one period model, by repeating the same optimum solution in every period. Even when changes in the data from one period to the next are significant, many companies find it convenient to construct a static single period model for their production planning decisions, which they solve at the beginning of every period with the most current estimates of data for the optimum plan for that period. This points out the importance of static models, even though most real world problems are dynamic.

In most multi-period problems, data changes from one period to the next are not insignificant. In this case the optimum decisions for the various periods may be different, and the sequence of decisions will be interrelated, i.e., a decision taken during a period may influence the state of the system for several periods in the future. Optimizing such a system through a sequence of single period static models solved one at a time, may not produce a policy that is optimal over the planning horizon as a whole. However, constructing a dynamic model with the aim of finding a sequence of decisions (one for each period) that is optimum for the planning horizon as a whole, requires reasonably accurate estimates of data for each period of the planning horizon. When such data is available, a dynamic model tries to find the entire sequence of interrelated decisions that is optimal for the system over the planning horizon.

In this book we will discuss both static and dynamic models. We begin with techniques for finding optimum solutions to static models, and then discuss how to extend these to handle dynamic models. We discuss the basic approach known as dynamic programming at an elementary level for solving problems that are posed as multi-period or multistage decision problems.

Even though the theory for handling dynamic models in a multi-period setting is well developed, practitioners find it difficult to use this theory in applications, due to lack of reliable information on how conditions might change in future periods.

Stochastic Versus Deterministic Models

An optimization model in which there is no uncertainty (i.e., all the data elements are known with certainty) is known as a **deterministic optimization model**.

In a single objective static optimization model, the objective function can be interpreted as the yield or profit that is required to be maximized. The objective function expresses the yield as a function of the various decision variables. In real world applications, the yield is almost never known with certainty, typically it is a random variable subject to many random fluctuations that are not under our control. For example the yield may depend on the unit profit coefficients of the various goods manufactured by the company (these are the **data elements** in the model) and these things fluctuate randomly. To analyze the problem treating the yield as a random variable requires the use of complicated **stochastic optimization (programming) models**. Instead, one normally analyses the problem using a deterministic model in which the random variables in the yield function are replaced by either their most likely values, or expected values etc. The solution of the deterministic approximation often gives the decision maker an excellent insight for making the best choice. We can also perform sensitivity analysis on the deterministic model. This involves a study of how the optimum solution varies as the data elements in the model vary within a small neighborhood of their current values. Decision makers combine all this information with their human judgment to come up with the best decision to implement.

Some people may feel that even though it is more complicated, a stochastic programming model treating the data elements as random variables (which they are), leads to more accurate solutions than a deterministic approximation obtained by substituting expected values and the like for the data elements. In most cases this is not true. To analyze the stochastic model one needs the probability distributions of the random data elements. Usually, this information is not available. One constructs stochastic models by making assumptions about the nature of probability distributions of random data elements, or estimating these distributions from past data. The closeness of the optimum solution obtained from the model may depend on how close the selected probability distributions are to the true ones. In the world of today, economic conditions and technology are changing constantly, and probability distributions estimated in a month may no longer be valid in the next. Because of this constant change, many companies find it necessary almost in every period to find a new optimum solution by solving the model with new estimates for the data elements. In this mode, an optimum solution is in use only for a short time (one period), and the solution obtained by solving a reasonable deterministic approximation of the real problem is quite suitable for this purpose. For all these reasons most real world optimization applications are based on deterministic models. In this book we discuss only methods for solving deterministic optimization models.

1.4 Optimization in Practice

Optimization is concerned with decision making. Optimization techniques provide tools for making optimal or best decisions. To maintain their market position, or even to continue to exist in business these days, businesses everywhere have to organize their operations to deliver products on time and at the least possible cost, offer services that consistently satisfy customers at the smallest possible price, and introduce new and efficient products and services that are cheaper and faster than competitors. These developments indicate the profound importance of optimization techniques. The organizations that master these techniques are emerging as the new leaders. All the countries in the world today that have a thriving export trade in manufactured goods have achieved it by applying optimization techniques in their manufacturing industries much more vigorously than the other countries.

1.5 Various Types of Optimization Models

Chapter 2 discusses a commonly used approach for handling Category 1 decision problems.

When constraints (1.3.4) are not there, the optimization model (1.3.1) − (1.3.3) is said to be a **linear programming model (LP)** if all the functions $\theta(x), g_i(x)$ are linear functions (i.e., each of them is of the form $a_1x_1 + \ldots + a_nx_n$, where $a_1, \ldots, a_n$ are given constants). LP is a very important model because it has many applications in a wide variety of areas. Also, many other models are solved by algorithms that have subroutines which require the solution of LP subproblems. The rich mathematical theory of LP is in a very highly developed and beautiful state, and many efficient algorithms have been developed for solving LPs. High quality software implementations of these algorithms are also widely available. Chapters 3 to 6 discuss examples of LP applications, and algorithms for solving and analyzing LP models including specialized LP models with special properties.

The optimization model (1.3.1)−(1.3.4) is said to be a **linear integer programming model (ILP)** or commonly **IP** if all the functions $\theta(x), g_i(x)$ are linear functions, and all the sets Δ_j are the set of integers. Often the word "linear" is omitted and the model is referred to as an integer program or IP. IP is even more widely applicable than LP since combinatorial choices found in many applications and combinatorial optimization problems can be modeled using binary and integer variables. IP theory is well developed, but more so in a negative way. For many IP models existing algorithms can only handle problems of moderate size within a reasonable time. So, the development of clever and ingenious heuristic approaches to obtain reasonable solutions to large IP models fast is a highly thriving area of research. Chapters 7, 8, 9 discuss IP models, and approaches for handling them.

Mathematical models for some multiperiod decision problems can be expressed in a form similar to (1.3.1)−(1.3.4), but Chapter 10 discusses the recursive approach that can be applied on simple problems posed in the multiperiod format directly without using such models. The application of this recursive approach to solve simple project planning problems without any complicating constraints is the subject of Chapter 11.

Finally, when constraints (1.3.4) are not there, and at least one of the functions $\theta(x), g_i(x)$ is nonlinear, (1.3.1)−(1.3.3) is known as a **continuous variable nonlinear program (NLP)**. Development of the theory of nonlinear programming has been going on ever since Newton and Lebnitz discovered calculus in the 17th century. We do not discuss NLP models in this book.

The subject that includes linear, integer, and nonlinear programming problems under its umbrella is called **mathematical programming**.

1.6 Background Needed

The most important background necessary for studying this book is knowledge of the **Gauss-Jordan (GJ) method** for solving systems of linear equations, the concepts of linear independence and bases from linear algebra, and the fundamental concepts of n-dimensional geometry. An excellent way to acquire this is to study Chapters 1, 4 for the GJ method and background in linear algebra, and Chapter 3 for background in n-dimensional geometry, in the self-study webbook **Sophomore Level Self-Teaching Webbook for Computational and Algorithmic Linear Algebra and n-Dimensional Geometry** [1.1].

1.7 Exercises

1.1: We have discussed classifications of decision problems into several types:

Categories 1, 2;

Deterministic, stochastic;

Single period, multiperiod;

Static, dynamic;

Single objective, multiobjective.

Explain these classifications clearly. Think of some examples of your own for each type and explain them in complete detail.

1.2: Discuss some strategies used in practice for handling stochastic decision problems, explaining why they may be preferred to others.

1.3: Explain the practical difficulties in applying the many nice approaches developed in theory to handle multiperiod decision problems, on problems involving many periods.

1.4: Think of some decision problems involving optimization, and state them clearly in your own words. Explain what data you will need to solve them. Discuss how you will handle these problems using your present state of knowledge wothout looking at the rest of this book. Keep these with you. Later after you have studied the book completely, review these notes and see if studying this book has helped you reach better decisions for these problems.

1.8 References

[1.1] K. G. MURTY, *Sophomore Level Self-Teaching Webbook for Computational and Algorithmic Linear Algebra and n-Dimensional Geometry*, available at the website:
http://ioe.engin.umich.edu/people/fac/books/murty/algorithmic_linear_algebra/

Index

For each index entry we provide the section number where it is defined or discussed first, and number of some sections where it appears prominently.