

COPYRIGHT NOTICE:

Andrzej Ruszczyński: Nonlinear Optimization

is published by Princeton University Press and copyrighted, © 2006, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@pupress.princeton.edu

Chapter One

Introduction

A general optimization problem can be stated very simply as follows. We have a certain set X and a function f which assigns to every element of X a real number. The problem is to find a point $\hat{x} \in X$ such that

$$f(\hat{x}) \leq f(x) \quad \text{for all } x \in X.$$

The set X is usually called the *feasible set*, and the function f is called the *objective function*.

At this level of generality, very little can be said about optimization problems. In this book, we are interested in problems in which X is a subset of a finite dimensional real space \mathbb{R}^n , and the function f is sufficiently regular, for example, convex or differentiable. Frequently, the definition of X involves systems of equations and inequalities, which we call *constraints*.

The simplest case is when $X = \mathbb{R}^n$. Such a problem is called the *unconstrained optimization problem*. If X is a strict subset of the space \mathbb{R}^n , we speak about the *constrained problem*. The most popular constrained optimization problem is the *linear programming* problem, in which f is a linear function and the set X is defined by finitely many linear equations and inequalities.

If the objective function f or some of the equations or inequalities defining the feasible set X are nonlinear, the optimization problem is called the *nonlinear optimization* (or nonlinear programming) *problem*. In this case, the specific techniques and theoretical results of linear programming cannot be directly applied, and a more general approach is needed.

Nonlinear optimization problems have attracted the attention of science since ancient times. We remember classical geometrical problems like the problem of finding the largest area rectangle inscribed in a given circle, or the problem of finding the point that has the smallest sum of distances to the vertices of a triangle. We know that optimization occurs in nature: many laws of physics are formulated as principles of minimum or maximum of some scalar characteristics of observed objects or systems, like energy or entropy. Interestingly, bee cells have the optimal shape that minimizes the average amount of wall material per cell.

But the most important questions leading to optimization problems are associated with human activities. What is the cheapest way of reaching a

certain goal? What is the maximum effect that can be achieved, given limited resources? Which model reproduces the results of observations in the best way? These, and many other similar questions arise in very different areas of science, and motivate the development of nonlinear optimization models, nonlinear optimization theory, and numerical methods of optimization.

Let us examine several typical application problems. Of course, the versions of problems that we present here are drastically simplified for the purpose of illustrating their main features.

Image Reconstruction

One of the diagnostic methods used in modern medicine is *computer tomography*. Its idea is to send a beam through a part of the body at many angles, and to reconstruct a three-dimensional image of the internal organs from the results. Formally, this procedure can be described as follows. A part of the body (for example, the head) is represented as a set $S \subset \mathbb{R}^3$. For the purpose of analysis we assume that S is composed of very small cells S_i , $i = 1, \dots, n$ (for example, cubes). We further assume that the density of the matter is constant within each cell, similar to the darkness of each pixel in a digital picture. We denote by x_i the density within cell S_i . These quantities are unknown and the purpose of the analysis is to establish their values.

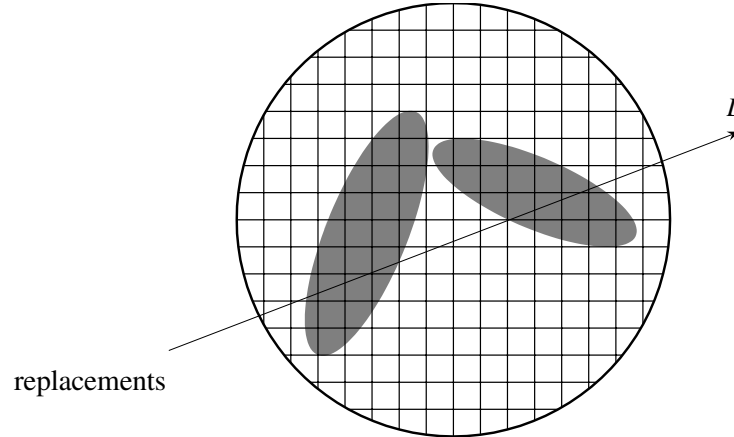


Figure 1.1. Beam L intersects a subset of cells.

A beam L passing through the set S intersects some of the cells (see Figure 1.1). Let $I(L)$ be this subset of cells. For every $i \in I(L)$ we know the length $a_i(L)$ of the path of the beam within cell i . The attenuation of the beam's energy depends on the total amount of matter on its way, $\sum_{i \in I(L)} a_i(L)x_i$. For every beam, given the measurements of its energy at the exit point, we

can estimate this sum by some quantity $b(L)$. Thus, we obtain the linear relation

$$\sum_{i \in I(L)} a_i(L)x_i = b(L).$$

If this analysis is repeated for a collection of beams L_j , $j = 1, \dots, m$, we obtain the system of linear equations

$$\sum_{i \in I(L_j)} a_i(L_j)x_i = b(L_j), \quad j = 1, \dots, m. \quad (1.1)$$

In addition, the physical nature of the problem dictates the inequalities:

$$x_i \geq 0, \quad i = 1, \dots, n. \quad (1.2)$$

Ideally, if all quantities were measured perfectly and the set S was actually built of small, uniform cells, the unknown density vector x would be a solution of this system. But several difficulties immediately arise. The number of beams, m , is not usually equal to the number of cells, n . The system of equations (1.1) may have too few equations to determine x in a unique way, or it may have too many or inconsistent equations. Its solution may violate inequalities (1.2), or it may have no solution at all. Moreover, the cell model is an approximation, the measurements have errors of their own, and it may be unreasonable to attempt to perfectly match their results.

For all these reasons, it is better to abandon the idea of perfectly satisfying equations (1.1) and to use instead some measure of fit. It may be the sum of the squares

$$f(x) = \sum_{j=1}^m \left(b(L_j) - \sum_{i \in I(L_j)} a_i(L_j)x_i \right)^2.$$

The smaller the values of f , the better, and therefore we formulate the following problem:

$$\text{minimize } \sum_{j=1}^m \left(b(L_j) - \sum_{i \in I(L_j)} a_i(L_j)x_i \right)^2 \quad (1.3)$$

subject to $x \geq 0$.

The above notation means that we want to find a point $\hat{x} \geq 0$ such that

$$f(\hat{x}) \leq f(x) \quad \text{for all } x \geq 0.$$

This problem will always have an optimal solution, because the function $f(x) \rightarrow \infty$ if any coordinate $x_i \rightarrow \infty$, and $f(\cdot)$ is a continuous function.

Our optimization problem (1.3) may still have some flaws. If the number of cells is larger than the number of beams used, the problem may have many optimal solutions. We would prefer a model that selects one of them in some reasonable way. One possibility is to modify the function $f(x)$ by adding to it the total amount of matter:

$$f(x) = \sum_{j=1}^m \left(b(L_j) - \sum_{i \in I(L_j)} a_i(L_j) x_i \right)^2 + \delta \sum_{j=1}^n x_j,$$

where $\delta > 0$ is a model parameter. The objective function of problem (1.3) and of the last modification is quadratic, and the constraints are linear. Such problems are called *quadratic programming problems*.

There are numerous other ways to formulate a meaningful objective function in this problem. They may involve nonlinear terms representing our knowledge about the properties of the image. We may formulate a stochastic model of attenuation and an appropriate maximum likelihood estimator. There is no need to discuss all these forms here. Most important for us is that the optimization model provides the flexibility that the system of equations did not have.

Portfolio Optimization

We want to invest our capital K in n assets. If we invest an amount z_i in asset i , after a fiscal quarter the value of this investment will become $(1 + R_i)z_i$, where R_i is the quarterly return rate of asset i . The return rate R_i is a random variable. Assume that all capital K is invested: $z_1 + z_2 + \cdots + z_n = K$. Introducing the variables $x_i = z_i/K$, $i = 1, \dots, n$, we can express the quarterly return rate of the entire portfolio as

$$\begin{aligned} R(x) &= \frac{1}{K} \left((1 + R_1)z_1 + (1 + R_2)z_2 + \cdots + (1 + R_n)z_n \right) - 1 \\ &= R_1x_1 + R_2x_2 + \cdots + R_nx_n. \end{aligned}$$

It is a random quantity depending on our asset allocations x_i , $i = 1, \dots, n$. By the definition of the variables x_i , they satisfy the equation

$$x_1 + x_2 + \cdots + x_n = 1.$$

We cannot sell assets that we do not own and thus we require that

$$x_i \geq 0, \quad i = 1, \dots, n.$$

We denote by X the set of asset allocations defined by the last two relations. Our idea is to decide about the asset allocations by an optimization problem.

The expected return rate of the portfolio has the form

$$\mathbb{E}[R(x)] = \mathbb{E}[R_1]x_1 + \mathbb{E}[R_2]x_2 + \cdots + \mathbb{E}[R_n]x_n.$$

However, the problem

$$\underset{x \in X}{\text{maximize}} \mathbb{E}[R(x)]$$

has a trivial and meaningless solution: invest everything in the asset j^* having the highest expected return rate $\mathbb{E}[R_{j^*}]$. In order to correctly model our preferences, we have to modify the problem. We have to introduce the concept of *risk* into the model and to express the aversion to risk in the objective function. One way to accomplish this is to consider the variance of the portfolio return rate, $\mathbb{V}[R(x)]$, as the measure of risk. Then we can formulate a much more meaningful optimization problem

$$\underset{x \in X}{\text{maximize}} \mathbb{E}[R(x)] - \delta \mathbb{V}[R(x)].$$

The parameter $\delta > 0$ is fixed here. It represents the trade-off between the mean return rate and the risk. It is up to the modeler to select its value.

It is convenient to introduce a notation for the mean return rates,

$$\mu_i = \mathbb{E}[R_i], \quad i = 1, \dots, n,$$

for their covariances

$$c_{ij} = \text{cov}[R_i, R_j] = \mathbb{E}[(R_i - \mu_i)(R_j - \mu_j)], \quad i, j = 1, \dots, n,$$

and for the covariance matrix

$$C = \begin{bmatrix} c_{11} & c_{12} & \cdots & c_{1n} \\ c_{21} & c_{22} & \cdots & c_{2n} \\ \vdots & \vdots & & \vdots \\ c_{n1} & c_{n2} & \cdots & c_{nn} \end{bmatrix}.$$

Then we can express the mean return rate and the variance of the return rate as follows:

$$\begin{aligned} \mathbb{E}[R(x)] &= \sum_{i=1}^n \mu_i x_i = \langle \mu, x \rangle, \\ \mathbb{V}[R(x)] &= \sum_{i=1}^n \sum_{j=1}^n c_{ij} x_i x_j = \langle x, Cx \rangle. \end{aligned}$$

The optimization problem assumes the lucid form

$$\underset{x \in X}{\text{maximize}} \langle \mu, x \rangle - \delta \langle x, Cx \rangle. \quad (1.4)$$

The objective function of this problem is quadratic and the constraints are linear. Such an optimization problem is called a *quadratic optimization problem*.

Our model can be modified in various ways. Most importantly, the variance is not a good measure of risk, as it penalizes the surplus return (over the mean) equally as the shortfall. One possibility is to measure risk by the central semideviation of order p :

$$\sigma_p[R(x)] = \left(\mathbb{E}[\max(0, \mathbb{E}[R(x)] - R(x))]^p \right)^{1/p},$$

with $p \geq 1$. Then the optimization problem

$$\underset{x \in X}{\text{maximize}} \quad \langle \mu, x \rangle - \delta \sigma_p[R(x)]$$

becomes more difficult, because it involves a nonsmooth objective function.

We can also introduce additional constraints on asset allocations. And for all these models, we can change the trade-off coefficient $\delta > 0$ to analyze its influence on the solution to the problem.

Signal Processing

A transmitter sends in short time intervals I_k harmonic signals

$$s_k(t) = a_k \cos(\omega t), \quad t \in I_k, \quad k = 1, 2, \dots$$

The pulsation ω is fixed. The signals are received by n receivers. Because of their different locations, each of the receivers sees the signal with a different amplitude, and with a different phase shift. The incoming signals are equal to

$$u_{kj} = c_j a_k \cos(\omega t - \psi_j), \quad j = 1, \dots, n, \quad k = 1, 2, \dots,$$

where c_j and ψ_j denote the amplitude multiplier and the phase shift associated with receiver j .

It is convenient to represent the signal s_k as a complex number S_k with the real and imaginary parts $\Re(S_k) = a_k$ and $\Im(S_k) = 0$. Similarly, each amplitude and phase shift are represented by the complex number

$$H_j = c_j e^{-i\psi_j}, \quad j = 1, \dots, n.$$

Then the incoming signals can be represented in the complex domain as follows:

$$U_{kj} = H_j S_k, \quad j = 1, \dots, n, \quad k = 1, 2, \dots$$

Because of various noises and interferences, the receivers actually receive the signals

$$X_{kj} = U_{kj} + Z_{kj}, \quad j = 1, \dots, n, \quad k = 1, 2, \dots$$

Here Z_{kj} denotes the complex representation of the noise at receiver j in the k th interval. To facilitate the reconstruction of the transmitted signal $\{S_k\}$, we want to combine the received signals with some complex weights W_1, W_2, \dots, W_n to obtain the output signals:

$$Y_k = \sum_{j=1}^n W_j^* X_{kj}, \quad k = 1, 2, \dots$$

In the formula above, W_j^* denotes the complex conjugate of W_j .

Introduce the notation:

$$H = \begin{bmatrix} H_1 \\ \vdots \\ H_n \end{bmatrix}, \quad W = \begin{bmatrix} W_1 \\ \vdots \\ W_n \end{bmatrix}, \quad Z_k = \begin{bmatrix} Z_{k1} \\ \vdots \\ Z_{kn} \end{bmatrix}.$$

We have

$$Y_k = \sum_{j=1}^n W_j^* (H_j S_k + Z_{kj}) = (W^* H) S_k + W^* Z_k, \quad k = 1, 2, \dots,$$

where $W^* = [W_1^* \ W_2^* \ \dots \ W_n^*]$.

Our objective is to find the weights W_j to maximize the useful part of the signal Y_k , $k = 1, 2, \dots$, relative to the part due to the noise and interferences. To formalize this objective we consider the sequences $\{S_k\}$ and $\{Z_k\}$ as discrete-time stochastic processes with values in \mathbb{C} and \mathbb{C}^n , respectively.[†] The noise process $\{Z_k\}$ is assumed to be uncorrelated with the signal $\{S_k\}$ and to have the covariance matrix

$$\Theta = \mathbb{E}[Z_k Z_k^*].$$

The energy of the transmitted signal is denoted by

$$\sigma^2 = \mathbb{E}[S_k S_k^*].$$

The energy of the output signal $\{Y_k\}$ can be calculated as follows:

$$\mathbb{E}[Y_k Y_k^*] = \sigma^2 \|W^* H\|^2 + W^* \Theta W. \quad (1.5)$$

[†]The symbol \mathbb{C} denotes the space of complex numbers.

Here the symbol $\|A\|$ denotes the norm of the complex vector $A \in \mathbb{C}^n$,

$$\|A\| = \left(\sum_{j=1}^n |A_j|^2 \right)^{1/2} \quad \text{with} \quad |A_j|^2 = A_j A_j^*.$$

The first component on the right hand side of (1.5) represents the part of the output signal energy due to the transmitted signal $\{S_k\}$, while the second component is the energy resulting from the noise. The basic problem of minimizing the noise-to-signal ratio can now be formalized as follows:

$$\begin{aligned} & \text{minimize} \quad W^* \Theta W \\ & \text{subject to} \quad \|W^* H\| \geq 1. \end{aligned} \tag{1.6}$$

Unfortunately, in practice the vector H is not exactly known. We would like to formulate a problem that allows H to take any value of form $H_0 + \Delta$, where H_0 is some nominal value and $\|\Delta\| \leq \delta$, with the uncertainty radius $\delta > 0$ given. Problem (1.6) becomes:

$$\begin{aligned} & \text{minimize} \quad W^* \Theta W \\ & \text{subject to} \quad \|W^*(H_0 + \Delta)\| \geq 1 \quad \text{for all} \quad \|\Delta\| \leq \delta. \end{aligned}$$

It is a rather complex nonlinear optimization problem with a quadratic objective and with the feasible set defined by infinitely many nonlinear constraints. We can transform it to a simpler form as follows.

First, we notice that the worst value of Δ for each W is $\Delta = -\delta W / \|W\|$. Thus all constraints are equivalent to the inequality

$$\|W^* H\| \geq 1 + \delta \|W\|.$$

Also, adding a phase shift to all components of W does not change the objective and the constraints, and therefore we can assume that the imaginary part of $W^* H$ is zero. This yields the problem

$$\begin{aligned} & \text{minimize} \quad W^* \Theta W \\ & \text{subject to} \quad \Re(W^* H) \geq 1 + \delta \|W\|, \\ & \quad \quad \quad \Im(W^* H) = 0. \end{aligned} \tag{1.7}$$

It has a convex objective function, one convex nonlinear constraint, and one linear constraint.

Classification

We have two sets: $A = \{a_1, \dots, a_k\}$, and $B = \{b_1, \dots, b_m\}$ in \mathbb{R}^n . Each vector a_i may be, for example, the set of results of various measurements or tests performed on patients for whom a particular treatment was effective,

while the vectors b_j may represent results of the same tests for patients who did not respond to this treatment.

Our goal is to find a *classifier*: ideally it would be a function $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $\varphi(a_i) > 0$ for all $i = 1, \dots, k$, and $\varphi(b_j) < 0$ for all $j = 1, \dots, m$. When a new vector of results $c \in \mathbb{R}^n$ appears, we can assign it to sets A or B , depending on the sign of $\varphi(c)$.

This problem is usually modeled as follows. We first determine a family of functions $\varphi(\cdot)$ in which we look for a suitable classifier. Suppose we decided to use linear classifiers of the form

$$\varphi(a) = \langle v, a \rangle - \gamma.$$

In this formula $v \in \mathbb{R}^n$ and $\gamma \in \mathbb{R}$ are unknown. To find the “best” classifier in this family means to determine the values of v and γ .




Figure 1.2. Linear classifier. The white points belong to the set A and the dark points belong to the set B .

When a fixed family of classifiers is used, we cannot expect to be able to separate the sets A and B with a classifier from this family. There is always a possibility of misclassification, and we have to determine the way in which we shall decide which classifier is “better.” Figure 1.2 illustrates this issue for two sets and for a linear classifier. One way to measure the quality of classification is to consider for each point a_i the classification error:

$$e_-(a_i; v, \gamma) = \begin{cases} 0 & \text{if } \langle v, a_i \rangle - \gamma \geq 0, \\ \gamma - \langle v, a_i \rangle & \text{if } \langle v, a_i \rangle - \gamma < 0. \end{cases}$$

Similarly, for each point b_j the classification error is

$$e_+(b_j; v, \gamma) = \begin{cases} 0 & \text{if } \langle v, b_j \rangle - \gamma \leq 0, \\ \langle v, b_j \rangle - \gamma & \text{if } \langle v, b_j \rangle - \gamma > 0. \end{cases}$$

Then our problem can be formulated as an optimization problem

$$\underset{v, \gamma}{\text{minimize}} \sum_{i=1}^k e_-(a_i; v, \gamma) + \sum_{j=1}^m e_+(b_j; v, \gamma). \quad (1.8)$$

We immediately notice a flaw in this model: setting $\gamma = 0$ and $v = 0$ we can make all classification errors zero, but the resulting classifier is useless. We need to restrict the set of parameters to ensure a discriminating power of the classifier. One possible condition is

$$\|v\| = 1. \quad (1.9)$$

Problem (1.8)–(1.9) is a nonlinear optimization problem with a piecewise linear objective function and a nonlinear constraint. By introducing *slack variables* $s \in \mathbb{R}^k$ and $z \in \mathbb{R}^m$, we can rewrite this problem as follows:

$$\begin{aligned} & \underset{v, \gamma, s, z}{\text{minimize}} \sum_{i=1}^k s_i + \sum_{j=1}^m z_j \\ & \text{subject to } \langle v, a_i \rangle - \gamma + s_i \geq 0, \quad i = 1, \dots, k, \\ & \quad \langle v, b_j \rangle - \gamma - z_j \leq 0, \quad j = 1, \dots, m, \\ & \quad s \geq 0, \quad z \geq 0, \\ & \quad \|v\| = 1. \end{aligned} \quad (1.10)$$

Another way to guarantee a discriminating power of the classifier is to enforce a “buffer zone” between the two sets, by requiring that

$$\begin{aligned} \langle v, a_i \rangle - \gamma + s_i &\geq 1, \quad i = 1, \dots, k, \\ \langle v, b_j \rangle - \gamma - z_j &\leq -1, \quad j = 1, \dots, m. \end{aligned}$$

The width of the buffer zone is equal to $2/\|v\|$. The condition $\|v\| = 1$ is no longer needed, because $v = 0$ and $\gamma = 0$ are no longer attractive. However, we have to prevent $\|v\|$ from becoming too large (otherwise the buffer zone has little meaning). One way to accomplish this is to add to the objective

function a term which quickly grows with $\|v\|$, as in the problem below:

$$\begin{aligned} & \underset{v, \gamma, s, z}{\text{minimize}} && \sum_{i=1}^k s_i + \sum_{j=1}^m z_j + \delta \|v\|^2 \\ & \text{subject to} && \langle v, a_i \rangle - \gamma + s_i \geq 1, \quad i = 1, \dots, k, \\ & && \langle v, b_j \rangle - \gamma - z_j \leq -1, \quad j = 1, \dots, m, \\ & && s \geq 0, \quad z \geq 0. \end{aligned} \quad (1.11)$$

Here $\delta > 0$ is a fixed parameter of the problem. This quadratic optimization problem, under the name *support vector machine*, is commonly used in computer science for data classification.

Optimal Control

A robot has to move an object of mass M from position x^{ini} to position x^{fin} within time $[0, T]$ avoiding an obstacle (see Figure 1.3).




Figure 1.3. The object has to be moved from x^{ini} to x^{fin} avoiding the obstacle Y .

One way to approach this problem is to split the time interval $[0, T]$ into n equal pieces

$$I_k = \left[\frac{kT}{n}, \frac{(k+1)T}{n} \right), \quad k = 0, \dots, n-1,$$

and to apply a constant force to the object within each interval. Denote by F_k the force used in interval k , and by v_k the velocity of the object at the beginning of the interval k . Both F_k and v_k are three-dimensional vectors. The velocity at the end of interval k equals

$$v_{k+1} = v_k + \frac{T}{nM}(F_k + G), \quad k = 1, \dots, n. \quad (1.12)$$

Here G denotes the gravitation force. We ignore, for simplicity, the mass of the robot arm and the mechanics of the arm.

Denoting by x_k the three-dimensional vector representing the location of the object at the beginning of interval k , we can write the equation for the

position at the end of interval k :

$$x_{k+1} = x_k + \frac{T}{n}v_k + \frac{T^2}{2n^2M}(F_k + G), \quad k = 1, \dots, n. \quad (1.13)$$

Equations (1.12)–(1.13) are an example of state equations for a discrete-time dynamical system. The six-dimensional vector (x_k, v_k) represents the *state* of the system at the beginning of the k th interval. The three-dimensional vector F_k is the *control* applied in the k th interval.

The condition that the object cannot hit the obstacle can be modeled as follows. We describe the obstacle as a certain closed set $Y \subset \mathbb{R}^3$ and we introduce the distance function[†]

$$d(x, Y) \triangleq \min_{y \in Y} \|x - y\|.$$

Our requirement can now be formulated as follows:

$$d(x_k, Y) \geq \delta, \quad k = 0, 1, \dots, n,$$

where $\delta > 0$ is some minimum distance to the obstacle allowed. For the purpose of practical tractability we may replace this condition with a slightly weaker one, formulated for a sufficiently dense net of points $y^j \in Y$, $j = 1, \dots, J$:

$$\|x^k - y^j\| \geq \delta, \quad k = 0, 1, \dots, n, \quad j = 1, \dots, J. \quad (1.14)$$

Other constraints are obvious from the problem formulation:

$$\begin{aligned} (x_0, v_0) &= (x^{\text{ini}}, 0), & (x_n, v_n) &= (x^{\text{fin}}, 0), \\ x_{k3} &\geq 0, & k &= 1, \dots, n-1. \end{aligned} \quad (1.15)$$

Finally, to formulate an optimization problem, we need to define a suitable objective function. Many reasonable formulations are possible. We can be concerned with the comfort of the travel by minimizing the variation of the force:

$$f_1(F) = \|F_0 + G\| + \sum_{k=1}^{n-1} \|F_k - F_{k-1}\| + \|F_{n-1} + G\|.$$

We can minimize the total work:

$$f_2(x, F) = \sum_{k=1}^{n-1} \langle F_k, x_{k+1} - x_k \rangle,$$

[†]The symbol \triangleq means equal by definition.

a combination of the two functions, or some other expression representing our preferences. The decision variables in the model are (x_k, v_k, F_k) , $k = 1, \dots, n$, and the state equations (1.12)–(1.13), as well as additional conditions (1.14)–(1.15), are constraints of the problem.

Approximation of Functions

We have a certain space of functions \mathcal{S} defined on a domain $D \subset \mathbb{R}^m$, and a function $\psi \in \mathcal{S}$. The space \mathcal{S} may be, for example, the space $\mathcal{C}(D)$ of continuous functions on a compact set D , or the space of integrable functions $\mathcal{L}_p(D, \mathcal{B}, \mu)$, where $p \in [1, \infty]$, D is a Lebesgue measurable set in \mathbb{R}^m , μ is the Lebesgue measure, and \mathcal{B} is a σ -algebra of measurable subsets of D . The space \mathcal{S} is equipped with a metric $\text{dist}(\cdot, \cdot)$.

We are also given a mapping $A : \mathbb{R}^n \rightarrow \mathcal{S}$. For example, it may have the form

$$A(x) = \sum_{j=1}^n x_j \varphi_j,$$

where φ_j , $j = 1, \dots, n$, are given functions in \mathcal{S} . The last relation is understood as follows: $A(x)$ is a function with values

$$[A(x)](u) = \sum_{j=1}^n x_j \varphi_j(u), \quad u \in D.$$

The *approximation problem* is simply the optimization problem:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \text{dist}(\psi, A(x)).$$

Several special cases are of interest. If $\mathcal{S} = \mathcal{C}(D)$ and the distance between two continuous functions is measured by the maximum norm, we obtain the *Chebyshev approximation problem*

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \max_{u \in D} |\psi(u) - [A(x)](u)|.$$

If $\mathcal{S} = \mathcal{L}_p(D, \mathcal{B}, \mu)$, with some $p \in [1, \infty)$, then our problem has the form

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \int_D |\psi(u) - [A(x)](u)|^p d\mu.$$

A special case of practical importance arises when the domain D is finite: $D = \{u_1, u_2, \dots, u_m\}$. Then the approximated function ψ can be viewed as a finite collection of data $y_i = \psi(u_i)$, $i = 1, \dots, m$. The approximation




Figure 1.4. Chebyshev approximation.

problem is then the problem of constructing a model which is close to the data in the sense of the selected metric $\text{dist}(\cdot, \cdot)$.

Figure 1.4 illustrates the solution of a Chebyshev approximation problem with a finite set of data.

Optimization problems arise in engineering, economics, statistics, business, medicine, physics, chemistry, biology, and in mathematics itself. New optimization problems appear all the time, and scientists analyze their properties and look for their solutions. Frequently, as in the examples discussed above, the models have to be adjusted or modified, to reflect the intentions of their authors and the specificity of the application. Because of these reasons, almost no applied nonlinear optimization problem is exactly the same as models discussed in textbooks. Therefore, it is our intention to provide readers with tools that will help them to analyze their models, to choose the best solution methods, and to improve the models if the results are not appropriate. Such tools necessarily involve solid mathematical foundations.