2. PROBLEM FORMULATION

To set the stage for solving a problem of optimization, it's necessary first to formulate it in a manner not only reflecting the situation being modeled, but so as to be amenable to computational techniques. This raises a number of fundamental issues, which range from the problem format to be adopted to criteria for when a problem is "well posed."

Basic problem: Minimize a function $f_0 : \mathbb{R}^n \to \mathbb{R}$, the *objective function*, over a specified set $C \subset \mathbb{R}^n$, the *feasible set*.

Max versus min: Maximizing a function g is equivalent to minimizing $-g$, so there's no loss of generality in concentrating on minimization. This is the convention in much of optimization theory.

Solution concepts. Different things can be sought in a problem of optimization. The following terms identify the main concepts.

Feasible solution: Any point x that belongs to C, regardless of the value it gives to f_0. Just finding such a point could be difficult numerically in cases where the constraints are complicated or numerous, and indeed, the very existence of a feasible solution may sometimes be an open question. This is just a first-level solution concept, but important nevertheless.

Optimal solution: A point $\bar{x}$ furnishing the minimum value of f_0 over C, i.e., a feasible solution such that $f_0(\bar{x}) \leq f_0(x)$ for all other feasible solutions x. This is more specifically what is called a *globally* optimal solution, when contrast must be made with the next concept.

Locally optimal solution: A point $\bar{x} \in C$ such that, for some neighborhood U of $\bar{x}$, one has $f_0(\bar{x}) \leq f_0(x)$ for all $x \in C \cap U$. Optimality in this case is asserted not relative to C as a whole, but only relative to some sufficiently small ball around $\bar{x}$. In practice it may be very hard to distinguish whether a numerical method has produced a globally optimal solution or just a locally optimal one, if that much.

Optimal set: The set of all (globally) optimal solutions (if any).

Optimal value: The greatest lower bound to the values of $f_0(x)$ as x ranges over C. There may or may not be a point $\bar{x} \in C$ at which f_0 actually attains this value. Furthermore, although the optimal value is always well defined, it could fail to be finite. It is $-\infty$ when f_0 is not bounded below on C, and on the other hand, it is ∞ by convention if $C = \emptyset$.

Constraint manipulation: Constraints can be expressed in more than one way, and some forms of expression may be more convenient in one context than another.

Function constraints: In conditions like $f_i(x) = c_i$, or $f_i(x) \leq c_i$, or $f_i(x) \geq c_i$, f_i is called a *constraint function*.

(1) An equality constraint $f_i(x) = c_i$ can be expressed equivalently, if desired, as a pair of inequality constraints: $f_i(x) \leq c_i$ and $f_i(x) \geq c_i$.

(2) An inequality constraint $f_i(x) \geq c_i$ can be expressed also as $-f_i(x) \leq -c_i$.

(3) An inequality constraint $f_i(x) \leq c_i$ can be expressed also as $-f_i(x) \geq -c_i$.

(4) An inequality constraint $f_i(x) \leq c_i$ can be expressed as an equality constraint $f_i(x) + s_i = c_i$ involving an additional decision variable s_i, itself constrained to be nonnegative. Such a variable is called a *slack variable*.

(5) Any constraint $f_i(x) = c_i$, or $f_i(x) \leq c_i$, or $f_i(x) \geq c_i$, can be expressed in terms of $g_i(x) = f_i(x) - c_i$ as $g_i(x) = 0$, or $g_i(x) \leq 0$, or $g_i(x) \geq 0$.

Abstract constraints: For methodological purposes it's often convenient to represent only some of the constraints in a problem in terms of constraint functions f_i and to lump the rest together in the abstract form $x \in X$.

For instance, a requirement on $x = (x_1, \ldots, x_n)$ that $0 \leq x_1 \leq 1$ *could* be represented by two function constraints $g_1(x) \geq 0$ and $g_2(x) \leq 1$ with $g_1(x) = g_2(x) = 1{\cdot}x_1 + 0{\cdot}x_2 + \cdots + 0{\cdot}x_n$, but it could also be incorporated into the description of a set X to which x must belong.

Boxes: A set $X \subset \mathbb{R}^n$ is a *box* if it is a product $I_1 \times \cdots \times I_n$ of closed intervals $I_j \subset \mathbb{R}$. To require $x \in X$ is to require $x_j \in I_j$ for $j = 1, \ldots, n$. Here I_j could be bounded or even consist of just one point, or it could be unbounded or even $(-\infty, \infty)$.

Nonnegative orthant: the box $\mathbb{R}^n_+ = [0, \infty) \times \cdots \times [0, \infty)$.

Whole space: the box $\mathbb{R}^n = (-\infty, \infty) \times \cdots \times (-\infty, \infty)$.

Linear and affine functions: A function g on $\mathbb{R}^n$ is called *affine* if it can be expressed in the form $g(x_1, \ldots, x_n) = d_0 + d_1x_1 + \cdots + d_nx_n$ for some choice of constants $d_0, d_1, \ldots, d_n$. Many people simply refer to such a function as linear, and in this they are following a long tradition, but in higher mathematics the term *linear* is reserved for the special case of such a function where the constant term vanishes: $d_0 = 0$. Thus, g is linear when there's a vector $d = (d_1, \ldots, d_n) \in \mathbb{R}^n$ such that

$$g(x) = d{\cdot}x \qquad \text{(the inner product, or dot product, of two vectors)}$$

Linear constraints: Conditions $f_i(x) = c_i$, $f_i(x) \le c_i$ or $f_i(x) \ge c_i$ in which the function f_i is linear—or affine. Or, conditions $x \in X$ in which X is a box.

Conventional problem format in finite-dimensional optimization:

$$(\mathcal{P}) \qquad \begin{array}{l} \text{minimize } f_0(x) \text{ over all } x = (x_1, \ldots, x_n) \in X \subset \mathbb{R}^n \text{ satisfying} \\ f_i(x) \begin{cases} \le 0 & \text{for } i = 1, \ldots, s, \\ = 0 & \text{for } i = s+1, \ldots, m. \end{cases} \end{array}$$

The feasible set C for $(\mathcal{P})$ consists of the points $x \in X$ satisfying all the constraints $f_i(x) \le 0$ or $f_i(x) = 0$. Here in particular, X could be all of $\mathbb{R}^n$, in which case the condition $x \in X$ would impose no restriction whatever.

Unconstrained minimization: the case where $X = \mathbb{R}^n$ and "$m = 0$," i.e., no equality or inequality constraints are present, so that $C = \mathbb{R}^n$.

Linear programming: the case where a linear (or affine) function f_0 is minimized subject to linear constraints: the functions $f_1, \ldots, f_m$ are affine and the set X is a box (e.g. $X = \mathbb{R}^n$ or $X = \mathbb{R}^n_+$).

Quadratic programming: like linear programming, but the objective function f_0 is allowed to have quadratic terms, as long as it remains *convex*, as defined later. (Note: in quadratic programming the constraints are still only linear!)

Nonlinear programming: this term is used in contrast to linear programming, but a much more important watershed will eventually be seen in the distinction between *convex* programming and *nonconvex* programming.

Geometric considerations: In problems with a few, simple constraints, the feasible set C might be decomposable into a collection of "pieces," each of which could be inspected separately in an attempt to locate the minimum of the objective function. For instance, if C were a (solid) cube in $\mathbb{R}^3$, one could look at what happens at the 8 corners, along the 12 edges, on the 6 faces, and in the cube's interior.

For most problems of interest in modern applications, however, there is little hope in such an approach. The number of "pieces" would be astronomical, or there would be no easy organization or listing of them. A further difficulty would lie in identifying which of the constraints might be redundant. Then too, there could be problems of degeneracy, where the constraints line up in odd ways and spoil the possibility of a good description of the "pieces" of C. As if this weren't enough trouble, there is the real prospect that C might be disconnected. These considerations force a different perspective on the analyst, who must look instead for a new kind of geometric framework on which to base computational schemes.

Geometry of linear constraints: Initial insight into a kind of geometry that *does* provide important support in optimization can be gained through the following ideas, which will later be subsumed under the broader heading of "convexity."

Half-spaces and hyperplanes: Subsets of $\mathbb{R}^n$ of the form $\{x \mid d \cdot x = c\}$ for a vector $d = (d_1, \ldots, d_n) \neq (0, \ldots, 0)$ and some constant $c \in \mathbb{R}$ are called *hyperplanes*, while those of the form $\{x \mid d \cdot x \leq c\}$ or $\{x \mid d \cdot x \geq c\}$ are called *closed half-spaces*. (With strict inequality, the latter would be *open half-spaces*.) A linear equality or inequality constraint on x thus corresponds to making x belong to a certain hyperplane or closed half-space (unless the linear function is $\equiv 0$, in which case the set isn't a hyperplane or half-space but just $\emptyset$ or $\mathbb{R}^n$, depending on c).

Polyhedral sets: A set $C \subset \mathbb{R}^n$ is called *polyhedral* if it can be represented as the intersection of a collection of finitely many hyperplanes or closed half-spaces, or in other words, specified by a finite system of linear constraints. (The whole space $\mathbb{R}^n$ is regarded as fitting this description by virtue of being the intersection of the "empty collection" of hyperplanes. The empty set fits because it can be viewed as the intersection of two parallel hyperplanes with no point in common.)

Argument: When a set is specified by a collection of constraints, it is the intersection of the sets specified by each of these constraints individually.

Inequalities alone: In the definition of "polyhedral" it would be enough to speak just of closed half-spaces, inasmuch as any hyperplane is itself the intersection of the two closed half-spaces associated with it.

Boxes as a special case: Any box is in particular a polyhedral set, since it's determined by upper or lower bounds on coordinates x_j of $x = (x_1, \ldots, x_n)$, each of which could be expressed in the form $d \cdot x \leq c$ or $d \cdot x \geq c$ for a vector d lining up with some coordinate axis. A box is thus an intersection of certain closed half-spaces.

Linear subspaces as a special case: Polyhedral sets don't have to have "corners." For instance, any subspace of $\mathbb{R}^n$ is polyhedral, since by linear algebra it can be specified by finitely many homogeneous linear equations.

Geometric interpretation of linear programming: The feasible set C in any linear programming problem is a certain *polyhedral* set. The function f_0 being minimized over C is a *linear* function, so (unless $f_0 \equiv 0$) its "isosurfaces" $\{x \mid f_0(x) = \alpha\}$, as α ranges over $\mathbb{R}$, form a family of *parallel hyperplanes*; the gradient of f_0, which is the same everywhere, is a certain vector perpendicular to all these hyperplanes. One is seeking to find the "lowest" of the hyperplanes that still touches C.

General constraints in vector notation: In the conventional format and beyond, it may be useful to think of constraints as requiring the values of certain functions $f_1, \ldots, f_m$ to lie within certain ranges. For instance, $f_i(x)$ may be required to lie in a certain closed interval J_i of $\mathbb{R}$, which might or might not be bounded, and which might even consist of just a single number (meaning that $f_i(x)$ had to have a particular value—the case of an equality constraint). In vector notation such a constraint system $f_i(x) \in J_i$, $i = 1, \ldots, m$, comes out as

$$F(x) \in D, \quad \text{where} \quad F(x) = \big(f_1(x), \ldots, f_m(x)\big), \quad D = J_1 \times \cdots \times J_m.$$

Here D would be a box, but more generally one can envision cases where the feasible set over which $f_0(x)$ is to be minimized takes the form $C = \big\{x \in X \,\big|\, F(x) \in D\big\}$ for sets $X \subset \mathbb{R}^n$ and $D \subset \mathbb{R}^m$ that aren't necessarily boxes. Note that for a problem $(\mathcal{P})$ in conventional format one would have $D = [0, \infty) \times \cdots \times [0, \infty) \times [0, 0] \times \cdots \times [0, 0]$, the box formed by the product of s copies of $[0, \infty)$ followed by $m - s$ copies of $[0, 0]$.

Penalties and the choice of objectives: The conventional problem format suggests that a modeler should approach a situation looking for a family of functions f_i of certain decision variables x_j, one of these functions being the objective function, and the rest, constraint functions. But reality can be murkier. The distinction between what should be set up as a constraint and what should be incorporated into the expression to be minimized may be quite subtle and even in some cases just a matter of the notation being adopted.

Essential objective function: An extreme but important example for theoretical purposes is the device of translating a problem in the conventional format into the minimization over all $x \in \mathbb{R}^n$ of the function

$$f(x) = \begin{cases} f_0(x) & \text{if } x \in C, \text{ the feasible set,} \\ \infty & \text{if } x \notin C. \end{cases}$$

Here the condition $x \in C$ is incorporated into the objective by imposing an infinite penalty when it is violated.

Hard versus soft constraints: Some kinds of constraints are "hard" in the sense of representing intrinsic conditions that can't be violated. For instance, a vector $(x_1, \ldots, x_n)$ may give a system of probabilities or weights through the stipulation that $x_1 \geq 0, \cdots, x_n \geq 0$ and $x_1 + \ldots + x_n = 1$. It wouldn't make sense to consider the alternative where $x_1 \geq -.0001$ and $x_1 + \cdots + x_n = .9999$. Constraints of such type are often built into the specification of the set X in the conventional format.

Other constraints may have quite a different, "soft" character. For instance, in asking that a mechanism under design ought to have a strength coefficient of at least .78 the modeler may be expressing a general desire that could be changed a bit once the costs and trade-offs are better known. A coefficient value of .76 may be quite acceptable, once it is realized that the difference could cost a fortune.

Penalty expressions: In dealing with soft constraints $f_i(x) \leq 0$ or $f_i(x) = 0$, it may be better in many cases to introduce a penalty expression instead. Thus, instead of enforcing an exact constraint, a term $\varphi_i \circ f_i$ could be added to the objective where (for the inequality constraint) $\varphi_i(t) = 0$ when $t \leq 0$ but $\varphi_i(t) > 0$ when $t > 0$. A popular choice is

$$\varphi_i(f_i(x)) = \alpha_i \max\big\{0, f_i(x)\big\} \text{ with penalty parameter } \alpha_i > 0,$$

but many other possible choices are available. It's worthwhile sometimes to relax the requirement of φ_i to just $\varphi_i(t) \leq 0$ for $t \leq 0$, with a negative penalty interpreted as a reward (for satisfying the inequality with room to spare).

Multiple objectives: Contrary to what we hear every day, it is impossible to design something to be the quickest, the cheapest and the most convenient all at the same time. While a number of variables may be of keen interest in a situation, the best that can be done is to optimize one of them while keeping the others within reasonable ranges. As a compromise, one can look to minimizing an expression like a weighted combination of the variables or more generally, in the case of variables given by functions $f_1, \ldots, f_m$, an expression $\varphi\big(f_1(x), \ldots, f_m(x)\big)$.

Composite format in optimization: These considerations and others lead to a basic problem model for optimization problems which differs from the conventional one stated earlier, namely

$$\text{minimize } f(x) = \varphi\big(F(x)\big) \text{ over all } x \in X \subset \mathbb{R}^n$$

for a mapping $F : \mathbb{R}^n \to \mathbb{R}^m$ and a function $\varphi : \mathbb{R}^m \to \overline{\mathbb{R}}$, where $\overline{\mathbb{R}}$ denotes the *extended real numbers*, i.e., the interval $[-\infty, \infty]$, in contrast to $\mathbb{R} = (-\infty, \infty)$. The feasible set in this case is considered to be

$$C = \big\{x \in X \,\big|\, F(x) \in D\big\}, \text{ where } D = \big\{u \in \mathbb{R}^m \,\big|\, \varphi(u) < \infty\big\}.$$

Max functions and nonsmoothness: A "max function" is a function defined as the pointwise maximum of a collection of other functions, for instance

$$g(x) = \max\big\{g_1(x), \ldots, g_r(x)\big\}.$$

Here the "max" refers to the arithmetic operation of taking the highest of the r values $g_1(x), \ldots, g_r(x)$ (not necessarily all different) to be $g(x)$. Such a function g is generally *nonsmooth*; a function g is *smooth* if its first partial derivatives exist and are continuous. Sometimes it's preferable to deal with a nonsmooth function directly, but other times there are tricks that can be brought in to achieve a problem formulation with smooth functions only.

Example 1: Consider the problem of minimizing, over all of $\mathbb{R}^n$, a function g of the form just given. Suppose each of the functions g_k is itself smooth. By introducing an additional variable $u \in \mathbb{R}$, we can re-express the problem equivalently as

$$\begin{gathered}\text{minimize } f_0(x,u) := u \text{ over all } (x,u) \in \mathbb{R}^n \times \mathbb{R} \text{ satisfying} \\ f_k(x,u) := g_k(x) - u \le 0 \text{ for } k = 1, \ldots, r.\end{gathered}$$

The functions $f_0, \ldots, f_r$ are then smooth on $\mathbb{R}^n \times \mathbb{R}$.

Example 2: Suppose a problem is given in terms of penalty expressions as that of minimizing

$$f_0(x) + \sum_{i=1}^{s} \alpha_i \max\big\{0, f_i(x)\big\} + \sum_{i=s+1}^{m} \alpha_i \big|f_i(x)\big|$$

over all $x \in X \subset \mathbb{R}^n$, where the coefficients α_i are positive. With additional variables $u_i \in \mathbb{R}$ and the vector $(u_1, \ldots, u_m)$ denoted by $u \in \mathbb{R}^m$, we can write this as the problem of minimizing

$$\bar{f}_0(x,u) := f_0(x) + \sum_{i=1}^{m} \alpha_i u_i$$

over all $(x,u) \in \mathbb{R}^n \times \mathbb{R}^m$ satisfying $x \in X$ and

$$\begin{aligned} &f_i(x) - u_i \le 0 \text{ and } u_i \ge 0 \text{ for } i = 1, \ldots, s, \\ &f_i(x) - u_i \le 0 \text{ and } -f_i(x) - u_i \le 0 \text{ for } i = s+1, \ldots, m. \end{aligned}$$

Abstract format for optimization: Sometimes it is useful to think of an optimization problem with respect to $x \in \mathbb{R}^n$ as simply a problem of minimizing some function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ over all of $\mathbb{R}^n$. There is then a one-to-one correspondence between problems and functions. As seen earlier through the notion of the *essential* objective function for a problem in the conventional format, there is no real loss of generality in taking this view, just a certain suppression of details from the notation.

Epigraphical geometry: The geometry in this case, as suggested by the remarks just made about "max functions," centers not on the graph of f but on its *epigraph*, which is the set $\big\{(x, \alpha) \in \mathbb{R}^n \times \mathbb{R} \,\big|\, f(x) \leq \alpha\big\}$.

Overview of the different formats: While the conventional format has a long tradition behind it and has become almost synonymous with "optimization," the theoretical power and modeling advantages of the composite format are now encouraging a trend in that direction, at least among specialists. The abstract format is primarily a tool for thinking about problems in various ways for theoretical purposes, rather than an approach to modeling, but for that it is often very helpful. In this introductory course the conventional format will receive the main attention, but the impression should be resisted that all optimization models should be forced into that channel. In the long run it's better to have a broader perspective.

Aspects of good problem formulation: In many areas of mathematics, a problem targeted for the application of a numerical method is not considered to be well formulated unless the *existence* of a *unique* solution is assured, and the solution is *stable* in the sense of being affected only slightly when the data elements of the problem are shifted slightly. In optimization, however, the goals of uniqueness and stability are unrealistic in the way they are usually interpreted, and that of existence has to be adapted to the multiple notions of what may be meant by a solution.

Issues to consider:

Does a feasible solution exist?

Does an optimal solution exist (global optimality)?

Can there be more than one optimal solution?

What happens to the set of feasible solutions under perturbations in the problem?

What happens to the set of optimal solutions under perturbations in the problem?

What happens to the optimal value under perturbations in the problem?

Note: The optimal value always exists and is unique.

The role of sequences: These issues are all the more important in view of the fact that most problems have to be solved by a *numerical method*. Such methods don't just produce an answer, but instead (however this might be masked by the software) generate a sequence of solution *candidates* which, it is hoped, get closer and closer to something *perhaps acceptable in lieu of* a true optimal solution. Unless certain basic conditions are fulfilled, in particular ensuring the existence of an optimal solution, the candidates might not get progressively closer to anything meaningful at all. Anyway, they might not satisfy the problem's constraints exactly. Questions then arise as to the effects that could be encountered if constraints are perturbed.

Example: potential trouble in one-dimensional minimization. Even in the case of minimizing a function over an interval in $\mathbb{R}$, many pitfalls are apparent. An optimal solution can fail to exist because the function is unbounded below, or because the optimal value can be approached only in an asymptotic sense (getting arbitrarily close, but without attainment), or simply because the function lacks the requisite continuity properties. Gradual changes in the shape of the graph of the function, in the case of multiple dips and humps, can induce discontinuity and multiplicities in the behavior of the optimal solution set. All of these phenomena can make trouble for methods that are supposed to generate a sequence of points tending somehow toward a minimizing point.

Example: potential trouble in linear programming. As explained earlier, the feasible set C in a linear programming problem is a certain polyhedral set. It could be empty if the constraints are improperly chosen or even if they are perturbed only slightly from their proper values. Furthermore, in minimizing a linear function over such a set one can obtain as the set of optimal solutions a "corner," an entire "edge" of "face," or other such portion. Indeed a gradual, continuous change in the coefficients of the linear objective function can induce jumps in the answer.

Thus, even in the most elementary so-called linear cases of optimization, there can be difficulties under all three of the headings of existence, uniqueness and stability of solutions. In particular, two numerical formulations of a problem that differ only in roundoff in input data—the number of decimal points allocated to the representation of the various coefficients—could in principle have unique *optimal solutions* very different from each other. Because only *linear* programming is involved, the two *optimal values* would be close together, according to theoretical results we haven't discussed. But for more general classes of problems there can be discontinuities even in the behavior of optimal values unless special assumptions are invoked.

Existence of feasible solutions: Generally speaking, there is no good criterion to apply to a system of constraints in order to ascertain on a theoretical level that there is as least one point satisfying the system. However, a numerical approach is possible. For a constraint system in the conventional problem format, a numerical method of optimization could be applied to the auxiliary problem of minimizing the function

$$g(x) := \sum_{i=1}^{s} \alpha_i \max\{0, f_i(x)\} + \sum_{i=s+1}^{m} \alpha_i \big| f_i(x) \big|$$

over all $x \in X$, where the introduced penalty coefficients α_i are all positive (e.g. $\alpha_i = 1$). Obviously, $g(x) = 0$ for each $x \in X$ satisfying the desired constraints, whereas $g(x) > 0$ for all other choices of $x \in X$. Thus if the optimal value in the auxiliary problem is 0 the optimal solutions to the auxiliary problem are precisely the feasible solutions to the original problem, but if the optimal value in the auxiliary problem is positive, there are no feasible solutions to the original problem.

Uniqueness of optimal solutions: The bad news is that there is no criterion, verifiable directly in terms of a problem's structure and data without going into computation, that can be imposed on a general *nonconvex* problem to ensure the existence of *at most one* optimal solution. This topic will be taken up later, after some theory of convexity has been built up.

Existence of optimal solutions: The good news is that readily verifiable criteria are available to ensure the existence of *at least one* optimal solution. The goal here will be to develop such a criterion, not just for the sake of a bare existence result, but in a form suited to the analysis of sequential approaches to finding solutions. This obliges us to work with the possibility that, in the course of calculations, constraints might only be satisfied approximately.

Approximate feasibility: Of special concern in connection with the stability of the feasible set C is what happens when constraints are only required to be satisfied to within a certain error bound. For any $\varepsilon > 0$, the set of *ε-feasible solutions* to a problem in conventional format is

$$C_\varepsilon := \Big\{ x \in X \;\Big|\; f_i(x) \le \varepsilon \text{ for } i = 1, \ldots, s;\ \big|f_i(x)\big| \le \varepsilon \text{ for } i = s+1 \ldots, m \Big\}.$$

Clearly C_ε includes C, so the minimum value of f_0 over C_ε is less than or equal to the minimum over C, the optimal value in the given problem.

Well posed problems: The key concept that we'll work with in connection with the existence and approximation of solutions is the following. An optimization problem in the conventional format will be deemed to be *well posed* when:

(a) the set $X \subset \mathbb{R}^n$ is closed as well as nonempty,

(b) the functions f_0, $f_1, \ldots, f_m$ on $\mathbb{R}^n$ are continuous,

(c) For some $\varepsilon > 0$, the set C_ε defined above has the property that, for every value $\alpha \in \mathbb{R}$, the set $\big\{x \in C_\varepsilon \,\big|\, f_0(x) \leq \alpha\big\}$ is bounded.

Easy special cases: Condition (a) is fulfilled when X is a nonempty box, or indeed any nonempty polyhedral set. Condition (b) is fulfilled when the functions are linear or affine, or are given by polynomial expressions in the variables x_j.

Condition (c) is certainly satisfied when X itself is bounded, or if for every $\alpha \in \mathbb{R}$ the set $\big\{x \in X \,\big|\, f_0(x) \leq \alpha\big\}$ is bounded. Also, (c) is sure to be satisfied if for some $\varepsilon > 0$ any one of the functions f_i for $i = 1, \ldots, s$ has the property that the set $\big\{x \in X \,\big|\, f_i(x) \leq \varepsilon\big\}$ is bounded, or one of the functions f_i for $i = s+1, \ldots, m$ is such that the set $\big\{x \in X \,\big|\, |f_i(x)| \leq \varepsilon\big\}$ is bounded.

Caution: This concept refers to the manner in which an application of optimization has been set up as a problem $(\mathcal{P})$ in conventional format: it's a property of the problem's *formulation* and depends on the specification of the functions f_i, the index s and the set X. A given application might be formulated in various ways as $(\mathcal{P})$, not only in the choice of decision variables but according to whether its requirements are taken as inequality constraints, equality constraints, or lumped into the abstract constraint $x \in X$. In some modes it could turn out to be well posed, but in others perhaps not. This is true in particular because the "perturbations" introduced in terms of ε affect the inequality and equality constraints differently, and don't affect the abstract constraint at all.

Review of terminology and notation for dealing with sequences: For the benefit of students wanting a refresher, we briefly go over some of the facts and ideas of advanced calculus coming into play here. Throughout these notes, we use superscript ν, the Greek letter "nu", as the running index for sequences so as to avoid conflicts with all the other indices that may appear in various situations in reference to coordinates, powers, etc. For instance, a sequence in $\mathbb{R}^n$ will be comprised of points x^ν for $\nu = 1, 2, \ldots$, where $x^\nu = (x_1^\nu, \ldots, x_n^\nu)$. (Note that "sequence" always means "infinite sequence," but there's nothing to preclude points from coinciding. As a special case, every x^ν could be the same point c, giving a *constant* sequence.)

Convergence: A sequence of points $x^\nu = (x_1^\nu, \ldots, x_n^\nu)$ in $\mathbb{R}^n$ is said to *converge* to a point $x = (x_1, \ldots, x_n)$ (and therefore be a *convergent sequence*) if for each coordinate index j one has $x_j^\nu \to x_j$ as $\nu \to \infty$, or equivalently

$$|x^\nu - x| \to 0, \text{ where } |x| := \sqrt{x_1^2 + \cdots + x_n^2} \text{ (Euclidean norm).}$$

Then x is the *limit* of the sequence; this is written as $x^\nu \to x$, or $x = \lim_{\nu\to\infty} x^\nu$.

Continuity: A function $f : \mathbb{R}^n \to \mathbb{R}$ is *continuous* if whenever $x^\nu \to x$ in $\mathbb{R}^n$ one has $f(x^\nu) \to f(x)$. The standard ways of verifying continuity involve such facts as the sum or product of continuous functions being continuous, along with the knowledge that certain elementary functions, for instance polynomial functions, trigonometric functions, etc., are continuous. For a mapping $F : \mathbb{R}^n \to \mathbb{R}^m$, continuity is defined similarly by the condition that $x^\nu \to x$ implies $F(x^\nu) \to F(x)$. In terms of a coordinate representation $F(x) = \big(f_1(x), \ldots, f_m(x)\big)$, this is equivalent to each of the component functions $f_i : \mathbb{R}^n \to \mathbb{R}$ being continuous.

Closedness: A set $S \subset \mathbb{R}^n$ is said to be *closed* if for every sequence of points $x^\nu \in S$ ($\nu = 1, 2, \ldots$) that converges to a point $x \in \mathbb{R}^n$, one has $x \in S$. A set is *open* if its complement in $\mathbb{R}^n$ is closed. In the extreme cases of $S = \mathbb{R}^n$ and $S = \emptyset$, S is both open and closed at the same time. The intersection of any family of closed sets is closed. The union of any family of *finitely many* closed sets is closed.

Example: level sets of continuous functions. If a function f is continuous, then for every choice of $c \in \mathbb{R}$ the set $\big\{x \,\big|\, f(x) \le c\big\}$ and the set $\big\{x \,\big|\, f(x) \ge c\big\}$ are both closed. So too is the set $\big\{x \,\big|\, f(x) = c\big\}$, which is their intersection.

Example: closedness of feasible sets. In an optimization problem in conventional format, the feasible set C is closed if the set X is closed and the functions $f_1, \ldots, f_m$ are continuous. This is because C is the intersection of X with m other sets of the form $\big\{x \,\big|\, f_i(x) \le 0\big\}$ or $\big\{x \,\big|\, f_i(x) = 0\big\}$, each of which is closed. Elementary cases where X is closed are those where X is all of $\mathbb{R}^n$ or where X is the product of closed intervals, i.e., X constrains each component x_j to lie in a certain closed interval of $\mathbb{R}$, either bounded or unbounded.

Example: boxes, polyhedral sets. These are the feasible sets for systems of linear constraints, so they are closed because affine functions are continuous.

Boundedness: A set $S \subset \mathbb{R}^n$ is called *bounded* if lies within some (large enough) ball, or in other words, if there exists $\rho \in (0, \infty)$ such that $|x| \le \rho$ for all $x \in S$. An equivalent characterization in terms of the coordinates x_j of x is that there exist

(finite) bounds α_j and β_j such that for every $x \in S$ one has $\alpha_j \leq x_j \leq \beta_j$ for $j = 1, \ldots, n$. As a special case, the empty subset $\emptyset$ of $\mathbb{R}^n$ is bounded.

Compactness: Closely related to closedness and boundedness is another property, which ultimately is crucial in any discussion of existence of solutions. A set S is called *compact* if every sequence $\{x^\nu\}_{\nu=1}^\infty$ of points in S has at least one subsequence $\{x^{\nu_\kappa}\}_{\kappa=1}^\infty$ that converges to a limit. The Heine-Borel Theorem asserts that a set $S \subset \mathbb{R}^n$ is compact if and only if S is both closed and bounded.

Cluster points: A point that is the limit of some *subsequence* of a given sequence, although not necessarily the limit of the sequence as a whole, is called a *cluster point* of the sequence. It follows from the theorem just quoted that *every bounded sequence in $\mathbb{R}^n$ has at least one cluster point* (since a bounded sequence can in particular be viewed as being in some large, closed ball, which by the theorem will be a compact set). This leads to the occasionally useful criterion that *a sequence $\{x^\nu\}_{\nu=1}^\infty$ in $\mathbb{R}^n$ converges (in its entirety) if and only if it is bounded and, because of circumstances, can't possibly have two different cluster points.*

Standard criterion for the attainment of a minimum or maximum: A basic fact of calculus related to optimization is the following. If a *continuous* function is minimized over a *nonempty*, *compact* set in $\mathbb{R}^n$, the minimum value is attained at some point (not necessarily unique). Likewise, the maximum value is attained.

Shortcomings for present purposes: This criterion could immediately be applied to optimization problems in conventional format by making assumptions that guarantee not only the closedness of the feasible set C (as already discussed) but also its boundedness. Then, as long as the objective function f_0 being minimized over C is continuous, an optimal solution will exist. But in many problems of optimization the feasible set *isn't* bounded, and yet we still want a criterion for existence. For instance, in "unconstrained" optimization we have $C = \mathbb{R}^n$. Therefore, we need a result that's more general.

THEOREM 1 (existence of optimal solutions). *Consider an optimization problem $(\mathcal{P})$ in conventional format, and assume it is well posed. Then the feasible set is closed. If the feasible set is also nonempty, then the optimal set is nonempty and the optimal value is finite. Furthermore, the optimal set is compact.*

Proof. The assumption that the problem is well posed entails (in conditions (a) and (b) of the definition of that property) the closedness of X and continuity of $f_1, \ldots, f_m$. These properties have already been seen above to imply that the feasible set C is closed.

Under the assumption now that C is also nonempty, let $\tilde{x}$ denote any point of C and let $\tilde{\alpha} = f_0(\tilde{x})$. The problem of minimizing f_0 over C has the same optimal solutions, if any, as the problem of minimizing f_0 over $\tilde{C} = \{x \in C \,|\, f_0(x) \leq \tilde{\alpha}\}$. The set $\tilde{C}$ is nonempty, because it contains $\tilde{x}$. It is closed by virtue of being the intersection of the closed set C and the set $\{x \,|\, f_0(x) \leq \tilde{\alpha}\}$, which is closed because f_0 is continuous by condition (b). Furthermore, it is bounded because of condition (c) in the definition of well posedness. Therefore, $\tilde{C}$ is compact. It follows from the standard criterion for the attainment of a minimum that the problem of minimizing f_0 over $\tilde{C}$ has an optimal solution. Hence the given problem, of minimizing f_0 over C, has an optimal solution as well.

Let $\bar{x}$ denote an optimal solution, not necessarily the only one, and let $\bar{\alpha} = f_0(\bar{x})$. Then $\bar{\alpha}$ is the optimal value in the problem, and because f_0 is a real-valued function this optimal value is finite. The optimal set is $\{x \in C \,|\, f_0(x) = \bar{\alpha}\}$, and this is the same as $\{x \in C \,|\, f_0(x) \leq \bar{\alpha}\}$ because strict inequality is impossible. The same argument applied to the set $\tilde{C}$ tells us that this set, like $\tilde{C}$, is compact. □

Example: problems with only an abstract constraint. As a simple case to which the existence criterion in Theorem 1 can be applied, consider the problem

$$\text{minimize } f_0(x) \text{ over all } x \in X \subset \mathbb{R}^n,$$

where there are no side conditions of the form $f_i(x) \leq 0$ or $f_i(x) = 0$. The basic criterion for good formulation comes down in this case to

(a) the set $X \subset \mathbb{R}^n$ is closed as well as nonempty,

(b) the function f_0 is continuous,

(c) the set $\{x \in X \,|\, f_0(x) \leq \alpha\}$ is bounded in $\mathbb{R}^n$ for all values of $\alpha \in \mathbb{R}$.

Under these conditions, the conclusion is that the f_0 does attain a minimum value over X, and the set of minimizing points is compact. In particular (a) and (c) are satisfied when X is compact. On the other hand, the result covers also the case of *unconstrained* optimization, where $X = \mathbb{R}^n$. The basic criterion in that special situation is that f_0 should be a continuous function whose level sets are bounded.

Example: the gap distance between two sets. For two nonempty, closed sets C_1 and C_2 in $\mathbb{R}^n$, the *gap distance* between C_1 and C_2 is the optimal value in the problem of minimizing $|x_1 - x_2|$ (Euclidean distance) over all $x_1 \in C_1$ and $x_2 \in C_2$, or in other words, over all pairs (x_1, x_2) in the set $X = C_1 \times C_2$. *While this optimal value is well defined, as always, are there optimal solutions to this problem?* That is, do there exist

pairs $(\bar{x}_1, \bar{x}_2) \in C_1 \times C_2$ for which $|\bar{x}_1 - \bar{x}_2|$ is the minimum? A sufficient condition for the existence of such a pair, as provided by Theorem 1 through the preceding example, is the boundedness of the sets $\big\{(x_1, x_2) \in C_1 \times C_2 \,\big|\, |x_1 - x_2| \le \rho\big\}$. This obviously holds if both C_1 and C_2 are bounded, but one can show actually that the boundedness of just one of the sets is enough.

If the gap distance between C_1 and C_2 is 0, does that mean these sets necessarily have at least one point in common? Again, this hinges on the existence of a solution to the optimization problem described. If the $(\bar{x}_1, \bar{x}_2) \in C_1 \times C_2$ yields the minimum value $|\bar{x}_1 - \bar{x}_2|$, and this is 0, then obviously $\bar{x}_1 = \bar{x}_2$, and this is a point in $C_1 \cap C_2$. An example where the gap distance is 0 but $C_1 \cap C_2 = \emptyset$ is furnished in $\mathbb{R}^2$ by taking C_1 to be a hyperbola having C_2 as one of its asymptotes. There are pairs of points arbitrarily near to each other in these sets, yet the sets don't meet.

Convergence to a solution: As groundwork for the consideration of numerical methods, it's important to broaden Theorem 1 cover sequences such as could be generated by such methods. The ε provision in "well-posedness" will be utilized in this. We keep the discussion focused on a problem in conventional format and notation.

Feasible sequence: A sequence of points x^ν all belonging to the feasible set, or in other words, satisfying

$$x^\nu \in X,\ f_i(x^\nu) \le 0 \text{ for } i \in [1, s],\ f_i(x^\nu) = 0 \text{ for } i \in [s+1, m].$$

Optimal sequence: A feasible sequence of points x^ν such that, for the optimal value $\bar{\alpha}$ in the problem, $f_0(x^\nu) \to \bar{\alpha}$. (Note that this property says nothing about the points x^ν themselves converging or even remaining bounded as $\nu \to \infty$!)

Asymptotically feasible sequence: A sequence of points $x^\nu \in X$ with the property that $\max\big\{0, f_i(x^\nu)\big\} \to 0$ for $i = 1, \ldots, s$, and $f_i(x^\nu) \to 0$ for $i = s+1, \ldots, m$. An equivalent description of this property is that $x^\nu \in C_{\varepsilon^\nu}$ for some choice of shrinking "error bounds" $\varepsilon^\nu \to 0$.

Asymptotically optimal sequence: An asymptotically feasible sequence of points x^ν with the additional property that, in terms of the optimal value $\bar{\alpha}$ in the problem, one has $\max\big\{\bar{\alpha}, f_0(x^\nu)\big\} \to \bar{\alpha}$.

Comments: For some kinds of unconstrained problems or problems with linear constraints only, methods can be devised that generate an optimal sequence. Usually, however, the most one can hope for is an asymptotically optimal sequence. The properties of such sequences are therefore of fundamental interest. Note that every optimal sequence is in particular an asymptotically optimal sequence.

THEOREM 2 (optimality from sequences). *Consider an optimization problem* $(\mathcal{P})$ *in conventional format, and suppose it is well posed. Assume also that the feasible set in* $(\mathcal{P})$ *is nonempty. Then any asymptotically optimal sequence* $\{x^\nu\}_{\nu=1}^\infty$ *(hence any optimal sequence) is bounded, and all of its cluster points (at least one exists) are optimal solutions to* $(\mathcal{P})$*. Furthermore, the sequence* $\{f_0(x^\nu)\}_{\nu=1}^\infty$ *converges to the optimal value in* $(\mathcal{P})$*.*

If in fact $(\mathcal{P})$ *has a unique optimal solution* $\bar{x}$*, any asymptotically optimal sequence (hence any optimal sequence) must converge (as a whole) to* $\bar{x}$*.*

Proof. Let ε be an error bound value for which condition (c) in the definition of well posedness is satisfied. Denote the optimal value by $\bar{\alpha}$, and consider any number $\alpha \in (\bar{\alpha}, \infty)$. For any asymptotically optimal sequence $\{x^\nu\}_{\nu=1}^\infty$, we have $x^\nu \in X$ and there is an index $\bar{\nu}$ such that, for all $\nu \geq \bar{\nu}$, we have

$$f_i(x^\nu) \leq \varepsilon \text{ for } i \in [1,s], \quad |f_i(x^\nu)| \leq \varepsilon \text{ for } i \in [s+1,m], \quad f_0(x^\nu) \leq \alpha.$$

In other words, all the points x^ν with $\nu \geq \bar{\nu}$ lie in the set $\{x \in C_\varepsilon \,|\, f_0(x) \leq \alpha\}$. This set is bounded because of condition (c). Therefore, the sequence $\{x^\nu\}_{\nu=1}^\infty$ is bounded.

Let $\bar{x}$ denote any cluster point of the sequence; $\bar{x} = \lim_{\kappa\to\infty} x^{\nu_\kappa}$ for some subsequence $\{x^{\nu_\kappa}\}_{\kappa=1}^\infty$. Because X is closed by condition (a), and $x^{\nu_\kappa} \in X$, we have $\bar{x} \in X$. From the continuity of the functions f_i in condition (b), we have $f_i(x^{\nu_\kappa}) \to f_i(\bar{x})$, and through the asymptotic optimality of the sequence $\{x^\nu\}_{\nu=1}^\infty$ this implies that

$$f_i(\bar{x}) \leq 0 \text{ for } i \in [1,s], \quad f_i(\bar{x}) = 0 \text{ for } i \in [s+1,m], \quad f_0(\bar{x}) \leq \bar{\alpha}.$$

Since no feasible solution can give the objective function a value lower than the optimal value $\bar{\alpha}$, we conclude that $\bar{x}$ is an optimal solution.

To justify the final assertion of the theorem, we note that in the case described the asymptotically optimal sequence, which is bounded, can't have two different cluster points, so it must converge. The only candidate for the limit is the unique optimal solution. □

Comment: It merits emphasis that the case at the end of Theorem 2 is generally the *only* one in which a proposed numerical method of solution that is truly able to generate an asymptotically optimal sequence (some don't even claim to do that) is *guaranteed* to produce points closer and closer to a particular optimal solution. Aside from some special situations dependent on convexity, or ones where a locally rather than globally optimal solution is acceptable, the cluster point property is the best that can be hoped for in computation.

Comparison of possible solution methods: Different numerical techniques for trying to solve an optimization problem can be quite different in their behavior. Some general questions to ask are these:

Is the method valid, i.e., are the assumptions on which its justification depends satisfied for the problem being addressed? If not, the output may be worthless. Here the most frequent abuse is the neglect of convexity or differentiability assumptions.

What does the method actually claim to find, and how reliable is the claim? Methods that are said to "solve" a problem often just try to approximate a point for which various conditions associated mathematically with optimality, but not necessarily guaranteeing optimality, are fulfilled. Some methods generate more information about a problem and its potential solutions than do others, and this auxiliary information could be useful on the side.

How robust is the method? Here the issue is whether the technique works dependably or is liable to get stuck on problems for which the "flavor" isn't quite to its liking.

Rates of convergence: An important issue in comparing different methods is their speed. Real judgment can only come from computer experiments with well selected test problems. But there's also a theoretical discipline which tries to make comparisons on the basis of the rate at which the distance of $f_0(x^\nu)$ to the optimal value $\bar{\alpha}$ decreases, or the distance of x^ν from an optimal solution $\bar{x}$ decreases. For instance, one can try to ascertain the existence of a constant c such that

$$|x^{\nu+1} - \bar{x}| \;\leq\; c|x^\nu - \bar{x}| \text{ for all indices } \nu \text{ beyond some } \bar{\nu}.$$

Then the method is said to *converge linearly at the rate c*. If two methods both guarantee convergence if this type, but one method typically yields a lower value of c than the other, then that method can be expected to find solutions faster—at least if computations are continued until the convergence behavior takes over. (Rates of convergence usually apply only to some undetermined "tail portion" of the sequence generated.) Modes of convergence can also differ more profoundly. A method with

$$|x^{\nu+1} - \bar{x}| \;\leq\; c|x^\nu - \bar{x}|^2 \text{ for all indices } \nu \text{ beyond some } \bar{\nu}$$

is said to *converge quadratically*. Such a method is likely to be much quicker than one that converges only linearly, but it may carry additional overhead or be applicable only in special situations. A method may exhibit *finite convergence* in the sense of being certain to terminate after only finitely many iterations (no infinite sequence then being involved at all), as is often true for example in linear programming. For such cases still other forms of theoretical comparison have been developed.