opt.lin **Constrained linear optimization**

Consider a linear objective function $f : \mathbb{R}^n \to \mathbb{R}$ with variables x_i in vector $\mathbf{x}$ and coefficients c_i in vector $\mathbf{c}$:

$$f(\mathbf{x}) = \mathbf{c} \cdot \mathbf{x} \tag{1}$$

subject to the linear **constraints**—restrictions on x_i—

$$A\mathbf{x} \leqslant \mathbf{a}, \tag{2a}$$

$$B\mathbf{x} = \mathbf{b}, \text{ and} \tag{2b}$$

$$\mathbf{l} \leqslant \mathbf{x} \leqslant \mathbf{u} \tag{2c}$$

where A and B are constant matrices and $\mathbf{a}, \mathbf{b}, \mathbf{l}, \mathbf{u}$ are n-vectors. This is one formulation of what is called a **linear programming problem**. Usually we want to **maximize** f over the constraints. Such problems frequently arise throughout engineering, for instance in manufacturing, transportation, operations, etc. They are called **constrained** because there are constraints on $\mathbf{x}$; they are called **linear** because the objective function and the constraints are linear.

We call a pair $(\mathbf{x}, f(\mathbf{x}))$ for which $\mathbf{x}$ satisfies Eq. 2 a **feasible solution**. Of course, not every feasible solution is **optimal**: a feasible solution is optimal iff there exists no other feasible solution for which f is greater (assuming we're maximizing). We call the vector subspace of feasible solutions $S \subset \mathbb{R}^n$.

Feasible solutions form a polytope

Consider the effect of the constraints. Each of the equalities and inequalities defines a linear **hyperplane** in $\mathbb{R}^n$ (i.e. a linear subspace of dimension $n - 1$): either as a boundary of S (inequality) or as a restriction of S to the hyperplane. When joined, these hyperplanes are the boundary of S (equalities restrict S to

lower dimension). So we see that each of the boundaries of S is **flat**, which makes S a **polytope** (in $\mathbb{R}^2$, a polygon). What makes this especially interesting is that polytopes have **vertices** where the hyperplanes intersect. Solutions at the vertices are called **basic feasible solutions**.

Only the vertices matter

Our objective function f is linear, so for some constant h, $f(x) = h$ defines a **level set** that is itself a hyperplane H in $\mathbb{R}^n$. If this hyperplane intersects S at a point x, $(x, f(x) = h)$ is the corresponding solution. There are three possibilities when H intersects S:

1. $H \cap S$ is a vertex of S,
2. $H \cap S$ is a boundary hyperplane of S, or
3. $H \cap S$ slices through the interior of S.

However, this third option implies that there exists a level set G corresponding to $f(x) = g$ such that G intersects S and $g > h$, so solutions on $H \cap S$ are *not optimal*. (We have not proven this, but it may be clear from our progression.) We conclude that either the first or second case must be true for optimal solutions. And notice that in both cases, a (potentially optimal) solution occurs at at least one vertex. The key insight, then, is that **an optimal solution occurs at a vertex of** S.

This means we don't need to search all of S, or even its boundary: we need only search the vertices. Helpful as this is, it restricts us down to $\binom{n}{\#\,\text{constraints}}$ potentially optimal solutions—usually still too many to search in a naïve way. In Lec. opt.simplex, this is mitigated by introducing a powerful searching method.