SECTION C: CONTINUOUS OPTIMISATION
LECTURE 8: THE FUNDAMENTAL THEOREM OF LINEAR INEQUALITIES

HONOUR SCHOOL OF MATHEMATICS, OXFORD UNIVERSITY
HILARY TERM 2006, DR RAPHAEL HAUSER
WITH A FEW ADDITIONS FROM DR. NICK GOULD

1. Constrained Optimisation and the Need for Optimality Conditions.
In the remaining part of this course we will consider the problem of minimising objective functions over constrained domains. The general problem of this kind can be written in the form

$$\min_{x \in \mathbb{R}^n} f(x)$$
$$\text{s.t.} \quad g_i(x) = 0 \quad (i \in \mathcal{E}),$$
$$g_j(x) \geq 0 \quad (i \in \mathcal{I}),$$

where $\mathcal{E}$ and $\mathcal{I}$ are the finite index sets corresponding to the equality and inequality constraints, and where $f, g_i \in C^k(\mathbb{R}^n, \mathbb{R})$ for all $(i \in \mathcal{I} \cup \mathcal{E})$.

In unconstrained optimisation we found that we can use the optimality conditions derived in Lecture 1 to transform optimisation problems into zero-finding problems for systems of nonlinear equations. We will spend the next few lectures to develop a similar approach to constrained optimisation: in this case the optimal solutions can be characterised by systems of nonlinear equations and inequalities.

A natural by-product of this analysis will be the notion of a Lagrangian dual of an optimisation problem - called the primal - has a sister problem in the space of Lagrange multipliers - called the dual. In constrained optimisation it is often advantageous to think of the primal and dual in a combined primal-dual framework where each sheds light from a different angle on a certain saddle-point finding problem.

2. The Fundamental Theorem of Linear Inequalities. Before we proceed to developing these theories, we will take a closer look at systems of linear inequalities and prove a theorem that will be of fundamental importance in everything that follows:

THEOREM 2.1 (Fundamental theorem of linear inequalities).
Let $a_1, \ldots, a_m, b \in \mathbb{R}^n$ be a set of vectors. Then exactly one of the two following alternatives occurs:
(I) $\exists y \in \mathbb{R}^m_+$ such that $b = \sum_i^m y_i a_i$.
(II) $\exists d \in \mathbb{R}^n$ such that $d^{\mathrm{T}} b < 0$ and $d^{\mathrm{T}} a_i \geq 0$ for all $(i = 1, \ldots, m)$.

Note that Alternative (I) says that b lies in the convex cone generated by the vectors a_i:

$$b \in \text{cone}(a_1, \ldots, a_m) := \left\{ \sum_{i=1}^{n} \lambda_i a_i : \lambda_i \geq 0 \; \forall i \right\}.$$

Alternative (II) on the other hand says that the hyperplane $d^{\perp} := \{x \in \mathbb{R}^n : d^{\mathrm{T}} x = 0\}$ strictly separates b from the convex set $\text{cone}(a_1, \ldots, a_m)$. Thus, Theorem 2.1 is a result about convex separation: either b is a member of $\text{cone}(a_1, \ldots, a_m)$ or there exists

a hyperplane that strictly separates the two objects. See Figure 2.1 for an illustration of the two cases.
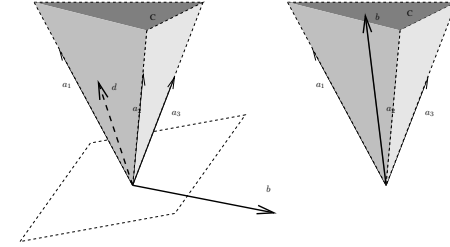

FIG. 2.1. *Illustration of the two alternative situations described by Theorem 2.1, where $C = \text{cone}(a_1, \ldots, a_m)$. On the left, Alternative (II) is depicted, and on the right, Alternative (I).*

Actually, our proof of Theorem 2.1 will reveal the slightly stronger property that when alternative (II) holds then d can be chosen so that $d^{\mathrm{T}} a_i = 0$ for at least $\text{rank}(a_1, \ldots, a_m) - 1$ indices i, a result that is useful in the theory of linear programming. We break down the proof of Theorem 2.1 into a series of lemmas:

LEMMA 2.2. *The two alternatives of Theorem 2.1 are mutually exclusive.*

Proof. If this is not the case then we find the contradiction

$$0 \leq \sum_{i=1}^{m} y_i (d^{\mathrm{T}} a_i) = d^{\mathrm{T}} \left(\sum_{i=1}^{m} y_i a_i \right) = d^{\mathrm{T}} b < 0.$$

□

LEMMA 2.3. *W.l.o.g. we may assume that $\text{span}\{a_1, \ldots, a_m\} = \mathbb{R}^n$.*

Proof. If $\text{span}\{a_1, \ldots, a_m\} \neq \mathbb{R}^n$ then either $b \in \text{span}\{a_1, \ldots, a_m\}$ and then we can restrict all arguments of the proof to the linear subspace $\text{span}\{a_1, \ldots, a_m\}$ of $\mathbb{R}^n$. Else, if $b \notin \text{span}\{a_1, \ldots, a_m\}$ then b cannot be written in the form $b = \sum_i^m \mu_i a_i$, so Alternative (I) does not hold. It remains to show that Alternative (II) applies in this case. Let π be the the orthogonal projection of $\mathbb{R}^n$ onto $\text{span}\{a_1, \ldots, a_m\}$, and let $d = \pi(b) - b$. Then $d \perp \text{span}\{a_1, \ldots, a_m\}$, so that

$$d^{\mathrm{T}} b = d^{\mathrm{T}} (b - \pi(b)) + d^{\mathrm{T}} \pi(b) = -\|d\|^2 + 0 < 0,$$
$$d^{\mathrm{T}} a_i = 0 \quad \forall i.$$

Therefore, Alternative (II) holds. □

Because of Lemma 2.3, we will henceforth assume that $\text{span}\{a_1, \ldots, a_m\} = \mathbb{R}^n$. We will next construct an algorithm that stops when a situation corresponding to either Alternative (I) or (II) is detected.

ALGORITHM 2.4.

S0 Choose $J^1 \subseteq \{1,\ldots,m\}$ such that $\operatorname{span}\{a_i\}_{J^1} = \mathbb{R}^n$, $|J^1| = n$.
S1 For $k = 1,2,\ldots$ repeat
    1. decompose $b = \sum_{i \in J^k} y_i^k a_i$
    2. if $y_i^k \geq 0\,\forall i \in J^k$ return y^k and stop.
    3. else begin
        let $j^k := \min\{i \in J^k : y_i^k < 0\}$
        let $\pi^k : \mathbb{R}^n \to \operatorname{span}\{a_i : i \in J^k \setminus \{j^k\}\}$ orthogonal projection
        let $d^k := \|a_{j^k} - \pi^k(a_{j^k})\|^{-1}(a_{j^k} - \pi^k(a_{j^k}))$
        if $(d^k)^{\mathrm{T}} a_i \geq 0$ for $(i = 1,\ldots,m)$ return d^k and stop.
    end
    4. let $l^k := \min\{i : (d^k)^{\mathrm{T}} a_i < 0\}$
    5. let $J^{k+1} := J^k \setminus \{j^k\} \cup \{l^k\}$
end.

Algorithm 2.4 is in fact the simplex algorithm for LP with Bland's rule in disguised form. Let us make a few remarks about the different stages of the algorithm:

- If the algorithm returns y^k in Step 2, then Alternative (I) holds: let $y_i = 0$ for $i \neq J^k$ and $y_i = y_i^k$ for $i \in J$. Then $y \in \mathbb{R}_+^m$ and $b = \sum_i y_i a_i$.
- If the algorithm enters Step 3, then $\{i \in J^k : y_i^k < 0\} \neq \emptyset$ because the condition of Step 2 is not satisfied.
- The vector d^k constructed in Step 3 satisfies

$$(d^k)^{\mathrm{T}} b = \sum_{i \in J^k} y_i^k (d^k)^{\mathrm{T}} a_i = y_{j^k}^k (d^k)^{\mathrm{T}} a_{j^k} < 0, \tag{2.1}$$

Therefore, if the algorithm returns d^k then Alternative (II) holds with $d = d^k$.

- If the algorithm enters Step 4 then $\{i : (d^k)^{\mathrm{T}} a_i < 0\} \neq \emptyset$ because the condition of the last "if" statement of Step 3 is not satisfied. Moreover, since

$$(d^k)^{\mathrm{T}} a_{j^k} = 1,$$
$$(d^k)^{\mathrm{T}} a_i = 0 \quad (i \in J^k \setminus \{j^k\}),$$

we have $\{i : (d^k)^{\mathrm{T}} a_i < 0\} \cap J^k = \emptyset$. This shows that $l^k \notin J^k$.
- We have $\operatorname{span}\{a_i : i \in J^{k+1}\} = \mathbb{R}^n$, because $(d^k)^{\mathrm{T}} a_{l^k} \neq 0$ and $(d^k)^{\mathrm{T}} a_i = 0$ $(i \in J^k \setminus \{j^k\})$ show that $a_{l^k} \notin \operatorname{span}\{a_i : i \in J^k \setminus \{j^k\}\}$. Moreover, $|J^{k+1}| = n$.

LEMMA 2.5. *In Algorithm 2.4 it can never occur that $J^k = J^t$ for $k < t$.*

Proof. Let us assume to the contrary that $J^k = J^t$ for some iterations $k < t$. Let $j^{\max} := \max\{j^s : k \leq s \leq t-1\}$. Then there exists $p \in \{k, k+1, \ldots, t-1\}$ such that $j^{\max} = j^p$. Since $J^k = J^t$, there also exists $q \in \{k, k+1, \ldots, t-1\}$ such that $j^{\max} = l^q$. In other words, the index must have once left J and then reentered, or else it must have entered and then left again.

Now $j^{\max} = j^p$ implies that for all $i \in J^p$ such that $i < j^{\max}$ we have $y_i^p \geq 0$. Likewise, for the same indices i we have $(d^q)^{\mathrm{T}} a_i \geq 0$, as

$$i < j^{\max} = l^q = \min\{i : (d^q)^{\mathrm{T}} a_i < 0\}.$$

Furthermore, we have $y_{j^{\max}}^p = y_{j^p}^p < 0$ and $(d^q)^{\mathrm{T}} a_{j^{\max}} = (d^q)^{\mathrm{T}} a_{l^q} < 0$. And finally, since $J^s \cap \{j^{\max}+1, \ldots, m\}$ remains unchanged for $s = k, \ldots, t$ we have $(d^q)^{\mathrm{T}} a_i = 0$ for all $i \in J^p$ such that $i > j^{\max}$. Therefore,

$$(d^q)^{\mathrm{T}} b = \sum_{i \in J^p} y_i^p (d^q)^{\mathrm{T}} a_i \geq 0. \tag{2.2}$$

On the other hand, (2.1) shows $(d^q)^{\mathrm{T}} b < 0$, contradicting (2.2). Thus, indices $k < t$ such that $J^k = J^t$ do not exist. ∎

We are finally ready to prove Theorem 2.1:

Proof. Since $J^k \subseteq \{1,\ldots,m\}$ and there are finitely many choices for these index sets and Lemma 2.5 shows that there are no repetitions in the sequence $J^1, J^2, \ldots$, the sequence must be finite. But this is only possible if in some iteration k Algorithm 2.4 either returns y^k, detecting that Alternative (I) holds, or d^k, detecting that Alternative (II) holds. ∎

3. The Implicit Function Theorem. Another fundamental tool we will need is the implicit function theorem. This is a standard results from multivariate analysis and can be proven via a rather technical fixed point argument. For a proof see any advanced calculus book. Before stating the theorem, let us illustrate it with an example:

EXAMPLE 3.1. *The function $f(x_1, x_2) = x_1^2 + x_2^2 - 1$ has a zero at the point $(1,0)$ and $\frac{\partial}{\partial x_1} f(1,0) = 1 \neq 0$. In a neighbourhood of this point the level set $\{(x_1, x_2) : f(x_1, x_2) = 0\}$ can be explicitly parameterised in terms of x_2, that is, there exists a function $h(t)$ such that $f(x_1, x_2) = 0$ if and only if $(x_1, x_2) = (h(t), t)$ for some value of t.*

Indeed, this level set is nothing else but the unit circle S^1, and for (x_1, x_2) with $x_1 > 0$ we have $f(x_1, x_2) = 0$ if and only if $x_1 = h(x_2)$ where $h(t) = \sqrt{1 - t^2}$. Thus,

$$S^1 \cap \{x \in \mathbb{R}^2 : x_1 > 0\} = \{(h(t), t) : t \in (-1, 1)\},$$

as claimed. Another way to say this is that S^1 is a differentiable manifold with local coordinate map

$$\varphi : S^1 \cap \{x \in \mathbb{R}^2 : x_1 > 0\} \to (-1, 1),$$
$$x \mapsto x_2.$$

The parameterisation in terms of x_2 was only possible because $\frac{\partial}{\partial x_1} f(1,0) \neq 0$. To illustrate this, note that we also have $f(0,1) = 0$, but now $\frac{\partial}{\partial x_1} f(0,1) = 0$ and we cannot parameterise S^1 by x_2 in a neighbourhood of $(0,1)$. In fact, in a neighbourhood of $x_2 = 1$, there are two x_1 values $\pm\sqrt{1 - x_2^2}$ such that $f(x_1, x_2) = 0$ when $x_2 < 1$ and none when $x_2 > 0$.

These observations can be generalised. For another 2D example, see Figure 3.1. To describe the general case, for $f \in C^k(\mathbb{R}^{p+q}, \mathbb{R}^p)$, let $f_B'(x)$ be the leading $p \times p$ block of the Jacobian matrix $f'(x) = [\,f_B'(x)\ f_N'(x)\,]$, and $f_N'(x)$ the trailing $p \times q$ block. Let x_B be the first $p \times 1$ block of the vector x and x_N the trailing $q \times 1$ block. The

following theorem generalises the above made observations:

THEOREM 3.2. *[Implicit Function Theorem]*
Let $f \in C^k(\mathbb{R}^{p+q}, \mathbb{R}^p)$ and let $\bar{x} \in \mathbb{R}^{p+q}$ be such that $f(\bar{x}) = 0$ and $f'_B(\bar{x})$ nonsingular. Then there exist open neighbourhoods $U_B \subset \mathbb{R}^p$ of $\bar{x}_B$ and $U_N \subset \mathbb{R}^q$ of $\bar{x}_N$ and a function $h \in C^k(U_N, U_B)$ such that for all $(x_B, x_N) \in U_B \times U_N$,

i) $f(x_B, x_N) = 0 \Leftrightarrow x_B = h(x_N)$,
ii) $f'_B(x)$ is nonsingular,
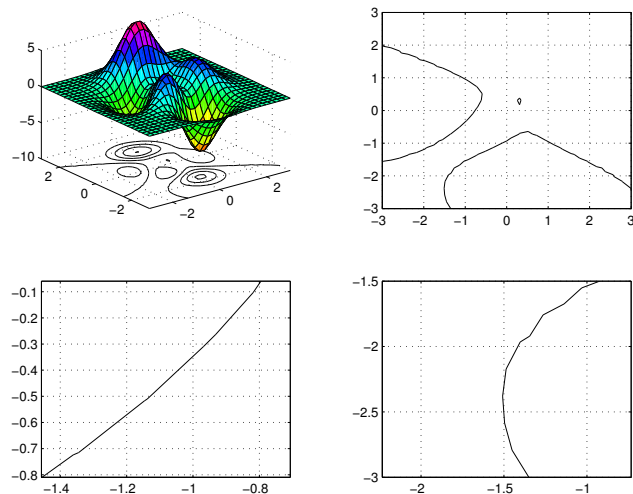*iii) $h'(x_N) = -\left(f'_B(x)\right)^{-1} f'_N(x)$.*

FIG. 3.1. *The first window shows the Matlab built-in sample function peaks(30) and some of its level sets. The second window shows the level sets corresponding to the level zero. Windows 3 and 4 are zooms of window 2. In the domain of window 3 the x-coordinate can be used to locally express the level set as the graph of a function $y = y(x)$. In the domain of window 4 this is not possible.*